WHEN DOES LEARNED POSTERIOR GEOMETRY RECOVER WEIGHTED MANIFOLD FORMS?

Anonymous authors
Paper under double-blind review

ABSTRACT

When does accurate local posterior geometry yield accurate weighted function energies? We study a fixed ambient trial space, separating learning error, finite-noise geometric bias, and numerical integration. A gradient-dependent bound connects local matrix error to the weighted form, while circle examples, including an exact Gaussian observation posterior, show how small normal components produce large form errors through ambient extensions. On controlled image-translation manifolds, we compare matched covariance models trained with sample residuals or conditional second moments. These objectives have the same ideal population minimizers but different conditional label randomness. With oracle-assisted centering and known-dimensional filtering, conditional supervision reduces mean linear-form error from 24.80% to 3.73%, with improvements numerically distinguished in all four image–initialization pairs. Same-form attribution identifies learning-induced normal error, rather than finite-noise posterior variance, as the dominant raw discrepancy in this cohort. Improving a learned mean instead improves tangent directions, worsens restricted metric scales, and produces large opposing fitting and centering contributions. Together, these findings connect extension-sensitive posterior forms, same-objective supervision effects, and signed geometric errors. They help researchers choose supervision and ambient probes by their effect on weighted energies, beyond aggregate local matrix accuracy.

1 INTRODUCTION

Gaussian denoising connects posterior covariance to data geometry. Under appropriate small-noise conditions, normalized covariance approaches the tangent projector, enabling conditional regression of local geometry without an explicit neighborhood graph (Manor & Michaeli, 2024; Bamberger et al., 2026; Kharitenko et al., 2026). We study the next question: *when does a learned local matrix field give accurate weighted energies of functions?*

Local matrix accuracy does not settle this question. Function energy weights matrix directions through ambient derivatives, the density, and mass whitening. Small normal components can dominate the energy of functions with large normal derivatives. Circle examples, including a genuine posterior covariance, expose this dependence while preserving fixed-extension small-noise convergence. Our reference is $\mathcal{E}_\nu(f, g) = \mathbb{E}_\nu[\nabla f^\top P \nabla g]$, compressed on a fixed ambient trial space. We separate learning, finite-noise and integration errors; signed tangent, cross and normal blocks reveal amplification and cancellation. Eigenvalue accuracy follows from form control; stable eigenspace recovery additionally requires spectral separation.

Controlled image-translation orbits supply matched training distributions and known geometry. Initial oracle-assisted models recover local projectors to about 3% yet incur large form errors. Frozen-field quadrature checks resolve the discrepancies and filtering effects. Matched cold-start retraining then compares old, improved and oracle centers under a common budget.

The strongest intervention changes *conditional label randomness* while preserving the ideal population regression objective. With oracle centering, conditional second-moment supervision reduces mean rank-two linear-form error from 24.80% to 3.73% at concentrated density, with numerically resolved improvement in all four pairs (two images, two covariance repeats). Cubic functions give mixed outcomes. Relative-energy and extension diagnostics clarify what each finite-space comparison recovers.

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.

Signed attribution identifies learning-induced normal error as the dominant raw discrepancy, with small genuine posterior bias. After conditional supervision and filtering, cross contributions become relatively important. Improving a learned mean instead helps tangent directions but worsens metric scales, while large fitting and centering terms partially cancel.

Together, the exact-posterior example, matched intervention and signed attribution identify how local-to-form transfer depends on ambient extensions, supervision and error direction. They distinguish learning-induced normal and cross energy from posterior bias on a common geometric reference.

2 Related work

Learned posterior geometry and conditional regression. Posterior moments and score derivatives underlie denoising-based geometric estimation (Efron, 2011; Vincent, 2011; Meng et al., 2021; Manor & Michaeli, 2024). Riemannian Metric Matching formulates conditional geometric regression with mean-centered PSD models and analyzes its manifold limit (Bamberger et al., 2026). Rao–Blackwellization reduces conditional randomness in stochastic gradients (Liu et al., 2019). We study how matched sample and conditional supervision affect weighted-form recovery under a fixed training budget, and localize the resulting differences through a common geometric reference. Landing with the Score uses denoising derivatives for tangent projection and Riemannian optimization (Kharitenko et al., 2026).

Ambient extensions and geometric operators. Closest-point methods prescribe extensions whose ambient derivatives recover intrinsic surface derivatives (März & Macdonald, 2012). We analyze how learned posterior fields act on fixed ambient extensions, then measure the resulting weighted-form discrepancy. Graph Laplacians and diffusion maps supply established geometric limits (Belkin & Niyogi, 2003; Coifman & Lafon, 2006; Singer, 2006). Neural eigensolvers (Pfau et al., 2019; Deng et al., 2022), operator-residual methods (Colbrook et al., 2023), and residual-guided dictionary learning (Coote & Colbrook, 2026) address function-space approximation or spectral reliability. Our spectral analysis uses subspace perturbation bounds (Davis & Kahan, 1970). The geometric transfer bound gives sufficient conditions in terms of local matrix error, ambient gradients and numerical accuracy; direct form comparisons measure the experimental error.

Local spectra and image-specific models. Score smoothing (Chen, 2026), pointwise Hessian filtering (Osada, 2026), and weighted tangent-measure identification (Zhao & Ni, 2026) concern different aspects of local geometry. Our filtering assumes known dimension and targets function energies. Image-specific networks (Shocher et al., 2018) and manifold chart models (Schonsheck et al., 2019; Alberti et al., 2024) provide approaches to local learning. Our separate image orbits supply controlled geometric references for comparing supervision interventions and attributing their weighted-form errors.

3 From local geometry to a weighted manifold form

Let ν be a probability measure on a smooth embedded d-dimensional manifold $\mathcal{M} \subset \mathbb{R}^D$, and let $P(x)$ be its orthogonal tangent projector. For specified ambient C^1 functions $\Phi : \mathbb{R}^D \to \mathbb{R}^r$, fix $M = \mathbb{E}_\nu[\Phi\Phi^\top] \succ 0$, $\Psi = M^{-1/2}\Phi$, and $J = D\Psi$. Write $\mathcal{E}_C(f, g) = \mathbb{E}_\nu[\nabla f^\top C \nabla g]$. The reference and estimated compressed forms are

$$H_P = \mathbb{E}_\nu[JPJ^\top], \qquad H_C = \mathbb{E}_\nu[JCJ^\top]. \tag{1}$$

Here $C(x)$ is a symmetric learned matrix field. On a compact manifold without boundary, with smooth positive density q relative to Riemannian volume, H_P compresses the form of $-q^{-1}\operatorname{div}_{\mathcal{M}}(q\nabla_{\mathcal{M}})$. We compare these finite matrices; convergence to the full operator also requires function-space approximation.

3.1 The amplification factor belongs to the whitened functions

Proposition 1 (Local error to geometric form error)**.** *Suppose* $\kappa_J = (\mathbb{E}_\nu\|J\|_{\mathrm{op}}^4)^{1/2} < \infty$ *and* $\mathbb{E}_\nu\|C - P\|_F^2 < \infty$. *Then*

$$\|H_C - H_P\|_{\mathrm{op}} \le \kappa_J(\mathbb{E}_\nu\|C - P\|_F^2)^{1/2}. \tag{2}$$

For $e_{\mathrm{loc}} = (\mathbb{E}_\nu \|C - P\|_F^2/d)^{1/2}$ *and* $\|H_P\|_{\mathrm{op}} > 0$*, the relative error is at most* $\mathfrak{a}_\Phi e_{\mathrm{loc}}$*, where* $\mathfrak{a}_\Phi = \sqrt{d}\,\kappa_J/\|H_P\|_{\mathrm{op}}$.

This Cauchy–Schwarz bound exposes the density-weighted gradients of *whitened ambient functions*. Mass conditioning alone does not determine $\mathfrak{a}_\Phi$; the bound applies per orbit and dictionary, not to substituted averages across different geometries.

For the canonical Gaussian posterior covariance $A_t(x) = t^{-1}\,\mathrm{Cov}(Z \mid Y_t = x)$ (Appendix A), with $Y_t = Z + \sqrt{t}\epsilon$, the geometric comparison follows

$$\widehat{H} \;\longrightarrow\; H_C \;\longrightarrow\; H_{A_t} \;\longrightarrow\; H_P. \tag{3}$$

Here $\widehat{H} = M^{-1/2}\widehat{K}M^{-1/2}$ uses a numerical stiffness matrix and exact mass. The errors are integration, clean-query estimation, and finite-noise geometric bias. Writing $\|F\|_{2,F} = (\mathbb{E}_\nu\|F\|_F^2)^{1/2}$, the latter two contribute at most $\kappa_J(\|C - A_t\|_{2,F} + \|A_t - P\|_{2,F})$. We assume any asserted small-noise convergence of A_t; this bound supplies neither its rate nor control from noisy-query regression alone. Appendix B adds integration and relative-mass errors and gives a gap-dependent projector bound with reference H_P.

3.2 Normal leakage can dominate despite a correct tangent direction

Put $N = I - P$. The exact signed decomposition is

$$H_C - H_P = \underbrace{\mathbb{E}[JP(C-P)PJ^\top]}_{D_T} + \underbrace{\mathbb{E}[J(PCN + NCP)J^\top]}_{D_X} + \underbrace{\mathbb{E}[JNCNJ^\top]}_{D_N}. \tag{4}$$

These are tangent, cross, and normal contributions; their norms do not add because the matrices can cancel. The appendix bounds each contribution using the corresponding tangent and normal gradient moments.

Proposition 2 (Ambient-extension obstruction on a circle)**.** *On the uniform unit circle, let* $P(x) = I - xx^\top$ *and* $C_\delta = P + \delta xx^\top$, $0 < \delta < 1$*. Its leading eigenspace is the exact tangent space, and* $\|C_\delta - P\|_F/\|P\|_F = \delta$ *everywhere. For the ambient functions*

$$\psi_a(x) = \sqrt{2}\,x_1[1 + \tfrac{a}{2}(\|x\|^2 - 1)], \qquad a \in \mathbb{R}, \tag{5}$$

all restrictions equal $\sqrt{2}\cos\theta$*, with unit* $L^2(\nu)$ *norm, but*

$$\mathcal{E}_P(\psi_a, \psi_a) = 1, \qquad \mathcal{E}_{C_\delta}(\psi_a, \psi_a) = 1 + \delta(a+1)^2. \tag{6}$$

Thus no extension-independent guarantee follows from local projector error alone. Fixed-a convergence as $\delta \to 0$ is preserved. The next statement shows that this issue also occurs for an actual posterior covariance.

Corollary 3 (Exact-posterior extension obstruction)**.** *Let* Z *be uniform on* S^1 *and* $Y_t = Z + \sqrt{t}\epsilon$*,* $t > 0$*. At a clean query* x*, let* u *be the posterior angle relative to* x*, with density proportional to* $e^{\cos u/t}$*. Put* $m_t = \mathbb{E}\cos u$ *and* $n_t = t^{-1}\,\mathrm{Var}(\cos u) > 0$*. Then*

$$A_t(x) = m_t P(x) + n_t xx^\top, \qquad \mathcal{E}_{A_t}(\psi_a, \psi_a) = m_t + n_t(a+1)^2. \tag{7}$$

With $\chi(x) = \sqrt{2}(x_1^2 - x_2^2)(2 - \|x\|^2)$*, the orthonormal restrictions of* (ψ_a, χ) *give* $H_P = \mathrm{diag}(1, 4)$ *and* $H_{A_t} = \mathrm{diag}(m_t + n_t(a+1)^2, 4m_t)$*. Their lowest modes are orthogonal if* $n_t(a+1)^2 > 3m_t$*. Yet* $\|A_t - P\|_F = O(t)$ *as* $t \to 0$*.*

The posterior has arbitrarily large extension-dependent energy at fixed t; for every fixed a, its form error still tends to zero. The proof is in Appendix B.3. Neither example attributes the observed image errors to posterior variance; our signed analysis tests that separately.

4 Supervision interventions and same-form attribution

Let $Z \sim \nu$, $Y_t = Z + \sqrt{t}\epsilon$, $\epsilon \sim N(0, I)$, and $\mu_t(y) = \mathbb{E}[Z \mid Y_t = y]$. A frozen mean m defines the residual label and its conditional target

$$R_m = \frac{(Z - m(Y_t))(Z - m(Y_t))^\top}{t}, \quad B_m(y) = \mathbb{E}[R_m \mid Y_t = y] = A_t(y) + \frac{b_m(y)b_m(y)^\top}{t}, \tag{8}$$

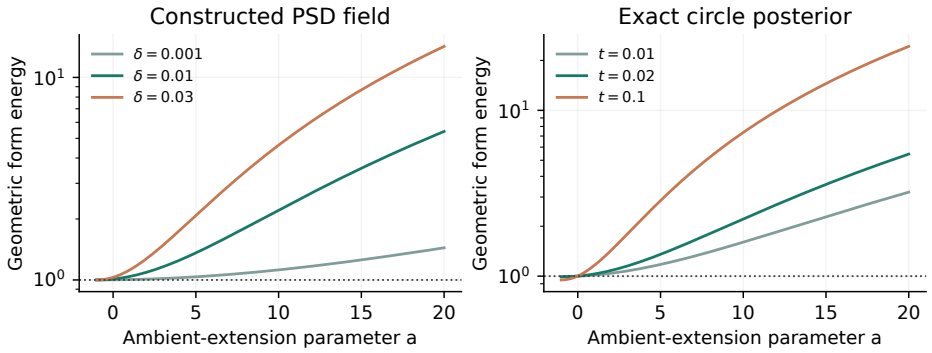

Figure 1: **Same intrinsic function, different ambient extensions.** *Left*: constructed PSD field. *Right*: exact circle posterior, with noise variance t. Dotted line: intrinsic energy. Fixed-extension small-noise convergence is preserved.

where $b_m = \mu_t - m$. Thus mean error introduces a positive semidefinite rank-one contamination. Pointwise, its trace and operator norm both equal $\|b_m\|^2/t$. Their averages are not the operator norm of the averaged contamination, nor the resulting geometric form error.

4.1 Conditional labels preserve the population objective

Proposition 4 (Matched residual supervision)**.** *For a fixed m, a matrix predictor $C(Y_t)$, and finite displayed second moments in Frobenius norm,*

$$\mathbb{E}\|C(Y_t) - R_m\|_F^2 = \mathbb{E}\|C(Y_t) - B_m(Y_t)\|_F^2 + \mathbb{E}\|R_m - B_m(Y_t)\|_F^2. \tag{9}$$

Consequently the two regression objectives have identical population minimizers on any common hypothesis class.

The mixed term vanishes after conditioning (Bamberger et al., 2026). Under oracle centering, both objectives have population regression target A_t. Numerical teachers have separately checked approximation error. Oracle means construct labels, never additional network inputs; the frozen mean input is deterministic given Y_t.

Conditioning also removes conditional raw-gradient randomness at a fixed parameter without changing its expectation (Liu et al., 2019); Appendix B.5 states the covariance identity. Our matched runs test the final supervision effect under Adam, clipping and EMA. The identity does not determine these transformed updates or their trajectory.

Target fit and geometry require separate evaluation: even $C = A_t$ has own-target error $-b_m b_m^\top/t$, whose relative norm can approach one when contamination dominates. Conversely, reducing this contamination need not improve the fitted tangent scale under a fixed budget.

4.2 A common matrix exposes amplification and cancellation

Fix ν, M, and J as in Section 3, and define $\mathcal{T}(E) = \mathbb{E}_\nu[JEJ^\top]$. All fields below are evaluated at clean queries $x \sim \nu$, not at $Y_t \sim p_t$. The raw field obeys the exact identity

$$\widehat{H}_C - H_P = \underbrace{\widehat{H}_C - H_C}_{\text{integration}} + \underbrace{\mathcal{T}(C - B_m)}_{\text{own-target fitting}} + \underbrace{\mathcal{T}(b_m b_m^\top/t)}_{\text{centering}} + \underbrace{\mathcal{T}(A_t - P)}_{\text{finite-noise geometry}}. \tag{10}$$

Each term admits the signed blocks in Equation (4): $\mathcal{T}(N(C - A_t)N)$ is normal learning error, while $\mathcal{T}(NA_tN)$ is genuine posterior normal variance. Their common-direction quadratic forms expose cancellation; separate norms are not causal percentages.

Let $\mathcal{F}_d$ retain the leading d pointwise eigenpairs, including their eigenvalues. The corresponding processed comparison is

$$H_{\mathcal{F}_d(C)} - H_P = \big(H_{\mathcal{F}_d(C)} - H_{\mathcal{F}_d(A_t)}\big) + \big(H_{\mathcal{F}_d(A_t)} - H_P\big). \tag{11}$$

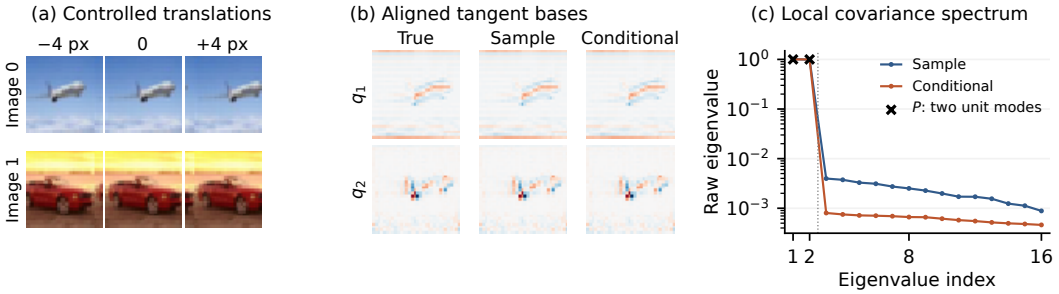

Figure 2: Controlled images and learned local geometry, $\kappa = 8$, $\sigma = .02$. Left: translations of both paired images. Middle: true and learned leading two-dimensional bases at image 0, repeat 0, zero shift; luminance maps use one scale and orthogonal alignment only for display. Right: raw covariance eigenvalues. Both learned models use oracle-centered training. These pointwise objects do not establish weighted-form spectral recovery.

This separates processed learning error and processed finite-noise bias. Filtering is nonlinear: the first term is not $H_{\mathcal{F}_d(C-A_t)}$. Proposition 1 applies to processed fields only with their corresponding processed errors. Comparing A_t with $\mathcal{F}_d(A_t)$ also reveals whether filtering removes appreciable genuine posterior content.

Numerical attribution. A numerical posterior's error cancels between fitting and finite-noise terms, so exact reconstruction cannot validate it. We vary posterior order and spatial quadrature for each component, with fixed analytic mass and reference. Grid changes are sensitivity measures, not rigorous integration bounds.

5 Controlled Recovery and Matched Learning Interventions

5.1 Design: Known Geometry, Matched Distribution, Fixed Functions

Each CIFAR base image defines a smooth periodic translation orbit in $\mathbb{R}^{3072}$. The Fourier renderer removes Nyquist components and performs no clipping; analytic derivatives supply a two-dimensional tangent frame Q, metric, and $P = QQ^\top$. Shifts follow the known density $q_\kappa(\theta) \propto \exp\{\kappa(\cos\theta_1 + \cos\theta_2)\}$, with $\kappa \in \{0, 8\}$. Training uses this same renderer, prior, and Gaussian noise, with primary $\sigma = \sqrt{t} = .02$. Density is given, not estimated. Separate 12-dimensional linear and cubic ambient dictionaries have analytic mass and reference stiffness matrices, held fixed throughout each comparison. Linear gradients are constant, so $H_C = J\mathbb{E}_\nu[C]J^\top$: this space isolates averaged-field errors, not spatially resolved field recovery. Cubic gradients vary with the query, testing a different spatially weighted interaction. Appendix C.7 specifies architectures, image indices and trial functions. Figure 2 shows the orbits and local matrix geometry.

At this query both models have two near-unit leading eigenvalues; conditional supervision reduces the remaining displayed eigenvalues.

The initial study uses 16 images and one initialization per image–density setting. The paired study uses images 0 and 1, both densities and two covariance initializations: four center/input arms (old/old, improved/improved, improved/old, oracle/old), each with sample or conditional labels. Within each cell all eight arms share initial parameters, draws, rank-16-plus-diagonal architecture and 9,000-update budget. Covariance models start cold; each improved mean is one frozen continuation checkpoint. The 64 models are a factorial design, not 64 independent geometries. Inputs are noisy pixels and a frozen learned mean, never true shifts or geometry. Oracle centering only constructs labels (Appendices C and E).

We evaluate a common set of 512 noisy queries and a separate set of 512 clean queries, with additional fresh-query checks. Noisy-query fit to B_m and A_t is distinct from clean tangent and metric recovery. Form errors use $\|\widehat{H}_C - H_P\|_{\mathrm{op}}/\|H_P\|_{\mathrm{op}}$, always on the same whitened functions. Rank-

Table 1: Matched two-image cohort, $\kappa = 8$: mean rank-two form-error ratios over two images and two covariance initializations. Label types share the ideal target within each row; centers change it across rows. The full center/input/label design is reported in Appendix E.

Center	Linear: sample	Conditional	Cubic: sample	Conditional
Old mean	62.76	21.80	4.18	2.13
Improved mean	10.66	7.66	1.30	1.12
Oracle center	0.24795	0.03730	0.02303	0.00677

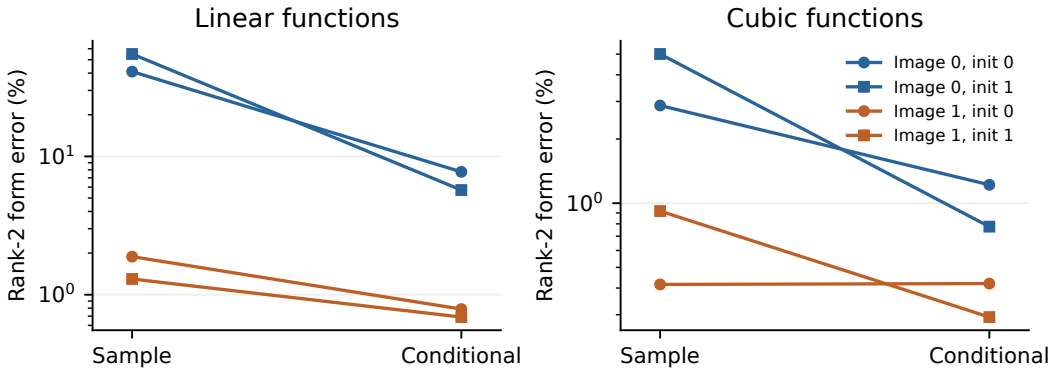

Figure 3: Oracle-centered, rank-two comparisons at $\kappa = 8$ on the same four image–initialization pairs. Lines connect matched models on a log axis. The post-run numerical screen passes 4/4 linear and 2/4 cubic units; these are two images with two repeats.

two processing retains the leading two eigenvalues and directions and uses the known dimension. It is not an intrinsic-dimension estimator.

5.2 Local geometry and weighted-form accuracy

Oracle-assisted models achieve about 3% full-projector error, measured as weighted RMS within each orbit and then averaged across images. Yet the concentrated linear form error is 56.38 times the reference norm (Table 9). Rank-two filtering reduces it to 0.23235; unit leading eigenvalues give 0.23528, whereas true-tangent projection PCP gives 0.00814. The last control preserves the predicted tangent block; it does not replace the prediction by P.

Raw normal error dominates; filtering leaves normal and cross contributions. These initial blocks combine learning and posterior error. The conservative gradient-moment bound supplies sufficient conditions, while direct form comparisons establish experimental accuracy (Appendix D).

Frozen-field integration. With fixed mass and functions, we use $64 \to 128 \to 256$ grids, shifted grids, and an independent reference-control-variate random check. Concentrated linear rank-two mean changes are 0.005253 (128-to-256), 0.000139 (shifted) and 0.004372 (random), below the 0.23235 error. All 16 concentrated cases pass the filtering-improvement screen in each space; one uniform linear case worsens. Appendix D details these numerical sensitivity checks, which are not rigorous quadrature bounds.

5.3 Same-target supervision yields controlled positive recovery

With oracle centering, replacing sample labels by conditional second moments reduces mean rank-two linear error from 0.24795 to 0.03730 (Table 1). Noisy-query covariance error falls $0.03393 \to 0.01377$, and leading-tangent-projector RMS error $e_{\tan}$ falls $0.02498 \to 0.00641$ (Appendix E). The new fields' mean integration sensitivity is 0.000668. In each of the four paired linear units, improvement exceeds ten times the sum of both fields' measured sensitivities. This post-run descriptive screen was not preregistered; it is neither a significance test nor a rigorous error bound.

Table 2: Oracle-centered two-image cohort at $\kappa = 8$, linear functions. Mean operator norms relative to $\|H_P\|_{\mathrm{op}}$. T, X, N denote tangent, cross, and normal blocks of learning error $C - A_t$, or of $\mathcal{F}_2(C) - \mathcal{F}_2(A_t)$ for Rank 2. Posterior N is the normal block of A_t or $\mathcal{F}_2(A_t)$. Norms cannot be summed.

Field	Labels	Total error	Learn: T	Learn: X	Learn: N	Posterior: N
Raw	Sample	45.03	0.00848	0.03522	45.03	6.45×10^{-6}
Raw	Conditional	10.83	0.00038	0.03149	10.83	6.45×10^{-6}
Rank 2	Sample	0.24795	0.00848	0.03526	0.24378	3.19×10^{-10}
Rank 2	Conditional	0.03730	0.00038	0.03150	0.01145	3.19×10^{-10}

Cubic mean error improves $0.02303 \to 0.00677$, with new-field sensitivity 0.000103, but only two of four improvements pass the same screen (Figure 3). A uniform-density image–initialization case has larger errors in both spaces. Appendix E reports all pairs. The intervention removes conditional label randomness while preserving the ideal objective (Proposition 4); the measured effect concerns final accuracy under the matched training budget.

5.4 Mean refinement separates direction and scale

Across the eight image–density–initialization cells with sample supervision, mean bias energy falls $35.05 \to 15.54$ and error to noisy-query A_t falls $1.869 \to 0.983$. Relative own-target error changes $0.9951 \to 0.9871$. Leading-tangent-projector RMS improves $0.3704 \to 0.1732$, while restricted cometric RMS worsens $0.2323 \to 0.3334$ and inverse-metric RMS $0.1744 \to 0.2294$; each direction of change holds in all eight cells. Concentrated linear form error decreases $62.76 \to 10.66$. Figure 8 visualizes the direction–scale contrast. Table 1 compares the corresponding conditional-label results for the same learned centers.

Changing the center alone, with the old mean input, gives concentrated sample linear error 10.27, versus 10.66 when both change. The information-preserving input change $(Y, m) \to (Y, (Y - m)/\sigma)$ increases errors $10.66 \to 11.47$ (sample) and $7.66 \to 8.01$ (conditional), distinguishing input parameterization from centering.

Fresh-query diagnosis finds about 99.3% of conditional learned-center own-target MSE along the bias-outer-product direction; this is matrix-regression error, not an operator-error percentage. Clean restricted target contamination falls while fitted scale worsens (Appendix F). The following same-form analysis tests which terms enter the geometric error.

5.5 Signed attribution distinguishes learning from finite-noise bias

For the two-image cohort we integrate the posterior, its rank-two version and both mean-bias outer products on the same four quadrature rules. Orders 9, 13 and 17 independently probe reference sensitivity for every term of Equations (10) and (11).

At concentrated density, the linear finite-noise form bias is 9.76×10^{-6} raw and 6.98×10^{-6} after rank-two processing. The raw genuine posterior normal contribution is 6.45×10^{-6}, versus learning-normal contributions 45.03 and 10.83 for oracle sample and conditional supervision (Table 2). The filter changes the genuine posterior form by only 6.45×10^{-6}. Thus in this cohort it mainly removes learned extra normal energy, rather than correcting a large finite-noise posterior mismatch. After conditional supervision and filtering, learning-normal error is 0.01145, cross error 0.03150, and tangent error 0.000379. Cross terms remain substantial; these separate norms are descriptive, not additive causal shares.

For raw improved-mean sample fields, quadratic forms along each case's largest-absolute-error direction average -39854.61 for fitting, $+44468.10$ for centering, and $+4.16 \times 10^{-6}$ for finite-noise bias, giving $+4613.49$ (Figure 8). Fitting and centering sensitivities average 10.39 and 10.54; each term is checked before cancellation. All values use $\|H_P\|_{\mathrm{op}}$ normalization. Across components and settings, the largest posterior-order 13-to-17 matrix change is 3.68×10^{-9}. Appendix G retains the signed components and both numerical axes.

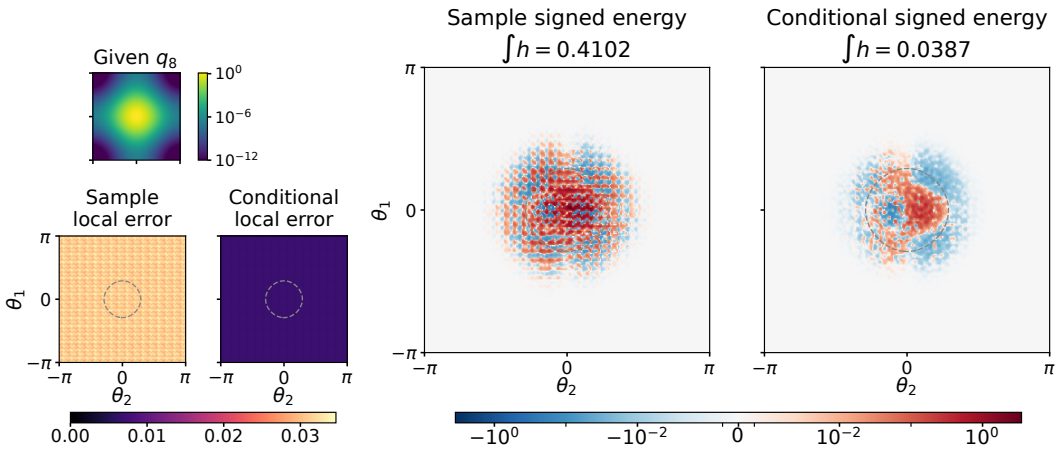

Figure 4: Spatial local-to-form diagnostic, image 0, initialization 0, oracle centering, rank-two fields. Both models use the same baseline-worst probe. Local: $\|C - P\|_F/\sqrt{2}$. Right: $q_8 g^\top (C - P)g/\|H_P\|_{\mathrm{op}}$, with $g = D\Psi^\top v$; its integral is the signed normalized energy error. Paired maps share scales; energy colors are symmetric-logarithmic. Dashed contours enclose 95% of quadrature mass. Full periodic domains are shown.

5.6 Where local discrepancies enter a function's energy

Figure 4 fixes image 0, initialization 0 and a common unit-L^2 linear probe: the largest-absolute-error eigenvector of the saved sample/rank-two comparison. This post-result diagnostic evaluates both models on that same function. Its signed normalized energy error falls $0.41017 \to 0.03873$; the integral of its pointwise absolute discrepancy falls $0.57245 \to 0.12395$. Local matrix error varies little across this orbit, while weighted energy error has positive and negative regions. Density and the ambient gradient mediate transfer; a small integral need not mean uniformly small energy discrepancy. Appendix I retains signed blocks, all four units, and grid/shift sensitivity.

5.7 Energy distortion and ambient-extension sensitivity

For a unit-L^2 function $f = v^\top \Psi$, e_C gives the additive bound $|\widehat{\mathcal{E}}_C(f, f) - \mathcal{E}_P(f, f)| \le e_C \|H_P\|_{\mathrm{op}}$. The worst relative energy distortion, when $H_P \succ 0$, is

$$e_{\mathrm{energy}} = \|H_P^{-1/2}(\widehat{H}_C - H_P)H_P^{-1/2}\|_{\mathrm{op}}, \tag{12}$$

All tested references are numerically positive definite. At concentrated density, oracle rank-two sample-to-conditional mean e_{energy} falls $35.75\% \to 7.02\%$ (linear) and $4.74\% \to 1.72\%$ (cubic). The four conditional linear values span 1.38%–14.64%. Mean energy-whitened integration changes are 0.141% (linear) and 0.0320% (cubic).

To vary only the ambient extension, use the known closest-point projection π and set $\Phi_s = \Phi \circ \pi + s(\Phi - \Phi \circ \pi)$. Intrinsic functions, masses and reference forms stay fixed, while $J_s = JP + sJN$. For the chosen quadrature $\mathcal{Q}$, define $D_{T,\mathcal{Q}} = \mathcal{Q}[JPCPJ^\top] - H_P$; let $D_{X,\mathcal{Q}}$ and $D_{N,\mathcal{Q}}$ replace the expectations in Equation (4) by $\mathcal{Q}$. Then exactly

$$\widehat{H}_C^{(s)} - H_P = D_{T,\mathcal{Q}} + sD_{X,\mathcal{Q}} + s^2 D_{N,\mathcal{Q}}. \tag{13}$$

For fixed concentrated oracle conditional linear fields, mean error is 0.000357 at $s = 0$, 0.03730 at $s = 1$, and 0.08851 at $s = 2$. Figure 5 shows all four units and both dictionaries. The $s = 0$ endpoint uses true tangents and is an oracle diagnostic. Appendix H retains all 320 comparisons and checks both numerical normalizations.

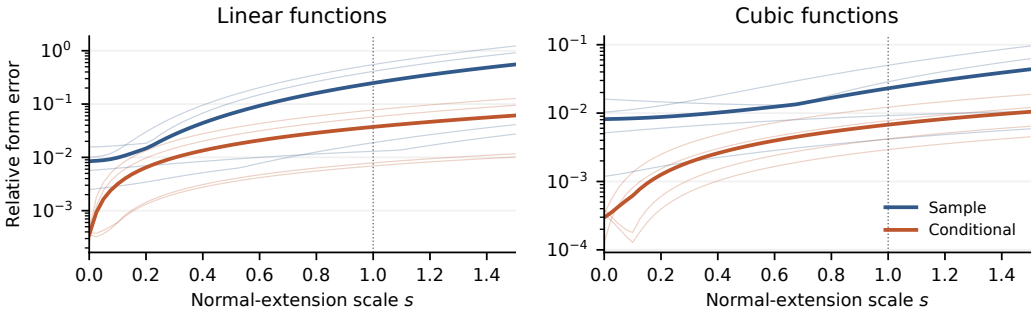

Figure 5: Extension sensitivity of concentrated oracle rank-two fields. Thin curves: four paired units; thick: their mean errors. $s = 1$ recovers the original comparison; $s = 0$ uses known tangents. Intrinsic functions and fields are fixed.

6 From form recovery to spectral claims

Eigenvalues without a gap. With analytic mass, put $\varepsilon = e_{\text{energy}}$ from Equation (12). For $H_P \succ 0$, congruence gives

$$(1-\varepsilon)H_P \preceq \widehat{H}_C \preceq (1+\varepsilon)H_P. \tag{14}$$

For $\varepsilon < 1$, the min–max principle therefore bounds every ordered eigenvalue:

$$(1-\varepsilon)\lambda_j(H_P) \le \lambda_j(\widehat{H}_C) \le (1+\varepsilon)\lambda_j(H_P). \tag{15}$$

No neighboring gap is needed. The four concentrated oracle-conditional linear rank-two matrices have $\varepsilon \le 0.1463842$, giving a 14.64% relative eigenvalue bound for these computed matrices. Integration uncertainty remains separate (Appendix H).

Stable eigenspaces. For symmetric stiffness $\widehat{K}$ and mass $\widehat{M} \succ 0$, put $K_C = \mathbb{E}_\nu[D\Phi C D\Phi^\top]$ and

$$\eta_Q = \|M^{-1/2}(\widehat{K} - K_C)M^{-1/2}\|_{\text{op}}, \quad \rho = \|M^{-1/2}\widehat{M}M^{-1/2} - I\|_{\text{op}} < 1. \tag{16}$$

Set $\eta = \eta_Q + \kappa_J\|C - P\|_{2,F}$.

Corollary 5 (Finite-space geometric projector transfer)**.** *For the lowest k modes of H_P, suppose $\gamma_P = \lambda_{k+1}(H_P) - \lambda_k(H_P) > 0$. Set*

$$\varepsilon_P = \frac{\eta + \rho\|H_P\|_{\text{op}}}{1-\rho}, \qquad b_M = \sqrt{1+\rho}\big((1-\rho)^{-1/2} - 1\big).$$

If $\varepsilon_P < \gamma_P/2$, the true $L^2(\nu)$ projectors onto the reference and estimated generalized-eigenfunction subspaces satisfy

$$\|\widehat{\Pi}_k - \Pi_{P,k}\|_{\text{op}} \le \min\{1, \varepsilon_P/(\gamma_P - \varepsilon_P) + b_M\}. \tag{17}$$

Appendix B proves the perturbation and mass-coordinate transport steps. Analytic mass gives $\rho = 0$, but grid sensitivities do not certify η_Q. The reference gap controls eigenspace stability; the eigenvalue comparison above uses relative form control alone.

7 Conclusion

Finite-noise posterior forms can depend strongly on ambient extensions. Matched supervision preserving the ideal objective substantially changes form accuracy; signed attribution separates learning-induced normal error, residual cross contributions and cancellation from posterior bias. These connected findings help researchers evaluate supervision and probe design through weighted energies alongside local matrix accuracy.

The empirical evidence concerns translation orbits with constant positive-definite metrics, supplied densities, fixed ambient functions and the tested budget. The 16 images give different embeddings and metrics within this intrinsically flat family. The strongest recovery uses oracle-assisted centering and known dimension. These experiments establish finite-form accuracy in that setting; the sufficient inequalities apply to general smooth embedded manifolds.

Reproducibility statement

The appendix specifies the data indices, network architectures, training protocols and ambient trial functions. Proofs, algebra checks and numerical audits document the reported comparisons and their limitations; these checks are not formal proof verification.

AI use statement

Generative AI assisted grammar checking.

References

Giovanni S. Alberti, Johannes Hertrich, Matteo Santacesaria, and Silvia Sciutto. Manifold learning by mixture models of VAEs for inverse problems. *Journal of Machine Learning Research*, 25(202): 1–35, 2024. URL https://www.jmlr.org/papers/v25/23-0396.html.

Jacob Bamberger, Adam Gosztolai, Pierre Vandergheynst, Michael Bronstein, and Iolo Jones. Riemannian metric matching for scalable geometric modeling of distributions. In *Proceedings of the 43rd International Conference on Machine Learning*, 2026. URL https://icml.cc/virtual/2026/poster/64721.

Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15(6):1373–1396, 2003. doi: 10.1162/089976603321780317. URL https://doi.org/10.1162/089976603321780317.

Zhengdao Chen. On the interpolation effect of score smoothing in diffusion models. In *International Conference on Learning Representations*, 2026. URL https://proceedings.iclr.cc/paper_files/paper/2026/hash/71f7154547c748c8041505521ca433ab-Abstract-Conference.html.

Ronald R. Coifman and Stéphane Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1):5–30, 2006. doi: 10.1016/j.acha.2006.04.006. URL https://doi.org/10.1016/j.acha.2006.04.006.

Matthew J. Colbrook, Lorna J. Ayton, and Máté Szőke. Residual dynamic mode decomposition: Robust and verified Koopmanism. *Journal of Fluid Mechanics*, 955:A21, 2023. doi: 10.1017/jfm.2022.1052. URL https://doi.org/10.1017/jfm.2022.1052.

George Coote and Matthew J. Colbrook. Residual-guided dictionary learning for spectrally accurate Koopman approximation. *arXiv preprint arXiv:2606.29083*, 2026. URL https://arxiv.org/abs/2606.29083.

Chandler Davis and W. M. Kahan. The rotation of eigenvectors by a perturbation. III. *SIAM Journal on Numerical Analysis*, 7(1):1–46, 1970. doi: 10.1137/0707001. URL https://doi.org/10.1137/0707001.

Zhijie Deng, Jiaxin Shi, and Jun Zhu. NeuralEF: Deconstructing kernels by deep neural networks. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 4976–4992, 2022. URL https://proceedings.mlr.press/v162/deng22b.html.

Bradley Efron. Tweedie's formula and selection bias. *Journal of the American Statistical Association*, 106(496):1602–1614, 2011. doi: 10.1198/jasa.2011.tm11181. URL https://doi.org/10.1198/jasa.2011.tm11181.

Andrey Kharitenko, Zebang Shen, Riccardo De Santi, Niao He, and Florian Dörfler. Landing with the score: Riemannian optimization through denoising. In *International Conference on Learning Representations*, 2026. URL https://proceedings.iclr.cc/paper_files/paper/2026/hash/d1a6fe144bdff7b95138bca8097d4e69-Abstract-Conference.html.

Runjing Liu, Jeffrey Regier, Nilesh Tripuraneni, Michael Jordan, and Jon McAuliffe. Rao-Blackwellized stochastic gradients for discrete distributions. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 4023–4031, 2019. URL https://proceedings.mlr.press/v97/liu19c.html.

Hila Manor and Tomer Michaeli. On the posterior distribution in denoising: Application to uncertainty quantification. In *International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=adSGeugiuj.

Thomas März and Colin B. Macdonald. Calculus on surfaces with general closest point functions. *SIAM Journal on Numerical Analysis*, 50(6):3303–3328, 2012. doi: 10.1137/120865537. URL https://doi.org/10.1137/120865537.

Chenlin Meng, Yang Song, Wenzhe Li, and Stefano Ermon. Estimating high order gradients of the data distribution by denoising. In *Advances in Neural Information Processing Systems*, volume 34, pp. 25359–25369, 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/hash/d582ac40970f9885836a61d7b2c662e4-Abstract.html.

Genki Osada. Local Hessian spectral filtering for robust intrinsic dimension estimation. In *Proceedings of the 43rd International Conference on Machine Learning*, 2026. URL https://icml.cc/virtual/2026/poster/61325.

David Pfau, Stig Petersen, Ashish Agarwal, David G. T. Barrett, and Kimberly L. Stachenfeld. Spectral inference networks: Unifying deep and spectral learning. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=SJzqpj09YQ.

Stefan Schonsheck, Jie Chen, and Rongjie Lai. Chart Auto-Encoders for manifold structured data. *arXiv preprint arXiv:1912.10094*, 2019. URL https://arxiv.org/abs/1912.10094.

Assaf Shocher, Nadav Cohen, and Michal Irani. Zero-Shot super-resolution using deep internal learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3118–3126, 2018. URL https://openaccess.thecvf.com/content_cvpr_2018/html/Shocher_Zero-Shot_Super-Resolution_Using_CVPR_2018_paper.html.

Amit Singer. From graph to manifold Laplacian: The convergence rate. *Applied and Computational Harmonic Analysis*, 21(1):128–134, 2006. doi: 10.1016/j.acha.2006.03.004. URL https://doi.org/10.1016/j.acha.2006.03.004.

Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural Computation*, 23(7):1661–1674, 2011. doi: 10.1162/NECO_a_00142. URL https://doi.org/10.1162/NECO_a_00142.

Ziqi Zhao and Qingjian Ni. Recovering weighted tangent geometry from a single-scale score field. *arXiv preprint arXiv:2608.22334*, 2026. URL https://arxiv.org/abs/2608.22334.

A NOTATION, ASSUMPTIONS, AND GUIDE TO THE EVIDENCE

The appendix follows the main argument: proofs for fixed weighted forms; controlled orbit construction; frozen-field integration; matched supervision and centering; input/bias diagnostics; and same-form attribution. The tables below identify the reference measure and the strength of each conclusion. Theorems apply to declared functions and fields. Numerical audits test their relevance in a controlled setting, without turning finite grid differences into mathematical error bounds.

A.1 OBJECTS AND QUERY DOMAINS

Symbol	Meaning and domain
D, d, r	Ambient dimension, positive manifold dimension, and finite trial-space dimension. The orbit experiments use $D = 3072, d = 2, r = 12$.
ν, q, p_t	Clean probability measure; density relative to manifold volume when used; law of noisy $Y_t = Z + \sqrt{t}\epsilon$.
t, σ	Gaussian variance $t = \sigma^2$ and standard deviation σ.
P, N, Q	Clean tangent projector, $N = I - P$, and a local orthonormal tangent frame with $QQ^\top = P$.
μ_t, A_t	Canonical posterior mean and covariance divided by t, defined at every query by the Gaussian Bayes kernel below.
m, b_m, B_m	Frozen learned mean; $b_m = \mu_t - m$; residual second-moment target $B_m = A_t + b_m b_m^\top/t$.
R_m, C	Random sample residual label and a deterministic matrix predictor. The contamination $b_m b_m^\top/t$ is not the random label R_m.
Φ, M, Ψ, J	Specified ambient functions, clean mass, whitened functions, and $J = D\Psi$. Ambient extensions are part of the comparison.
$\mathcal{T}, H_C, H_P$	$\mathcal{T}(E) = \mathbb{E}_\nu[JEJ^\top]$, estimated and tangent-reference compressed forms.
$\mathcal{F}_d$	Pointwise leading-d eigenpair retention, including eigenvalues; dimension is supplied.
η_Q, ρ, γ_P	True whitened stiffness integration error, relative mass error, and reference compressed spectral gap.

Canonical posterior values on clean queries. Conditional moments specified only p_t-almost everywhere cannot by themselves specify values on a lower-dimensional manifold. We use the Gaussian Bayes version at every $y \in \mathbb{R}^D$:

$$\pi_t(dz \mid y) = \frac{\exp(-\|y - z\|^2/(2t))\,\nu(dz)}{\int \exp(-\|y - z'\|^2/(2t))\,\nu(dz')},$$
$$\mu_t(y) = \int z\,\pi_t(dz \mid y), \qquad A_t(y) = t^{-1}\,\mathrm{Cov}_{\pi_t(\cdot|y)}(Z). \tag{18}$$

The denominator is positive for any probability prior. Under the compact orbit priors used here, all moments and derivatives needed for these posterior fields exist. Thus $A_t(x)$ at a clean point is a specified field, not an arbitrary representative of a conditional expectation. The neural predictor likewise defines a field at clean points. Training error under p_t alone does not bound error under singular ν; no such measure-change step is assumed in our transfer argument.

Whitening convention. The main text uses symmetric $M^{-1/2}$. Implementations may use $W = V\Lambda^{-1/2}$ when $M = V\Lambda V^\top$, so $W^\top M W = I$ and $\Psi = W^\top \Phi$. More generally $W = M^{-1/2}O$ for an orthogonal O. The resulting compressed matrices are common orthogonal conjugates; operator norms, spectral gaps and angles are unchanged. Signed comparisons must use the same chosen coordinates within each orbit, space and intervention. No matrix averaging across unrelated per-orbit whitening coordinates is used.

A.2 Assumption and proof map

Result	Required conditions	Proof location
Proposition 1	Fixed ν, ambient C^1 functions, $M \succ 0$, finite displayed gradient and field moments; symmetric measurable C. PSD is unnecessary for the bound.	Appendix B.1
Proposition 2	Uniform unit circle, $0 < \delta < 1$, the explicit ambient extensions. No learned network assumption.	Appendix B.3
Corollary 3	Uniform circle prior, Gaussian observation, canonical clean posterior and specified extensions; fixed-a small-noise limit.	Appendix B.3
Proposition 4	Fixed center and common hypothesis class; finite squared Frobenius moments. The gradient identity additionally requires differentiability and finite gradient second moments at the fixed parameter.	Appendix B.5
Corollary 5	Symmetric estimated stiffness, positive definite mass, $\rho < 1$, positive reference gap and $\varepsilon_P < \gamma_P/2$.	Appendix B.2

Interpreting H_P as a weighted Laplace form uses the compact, smooth, boundaryless manifold and positive smooth volume density stated in the main text. The finite-matrix inequalities themselves need only their displayed moments and mass condition. They do not impose an unreported volume, reach, or curvature class, supply a general small-noise rate for $A_t \to P$, or assert consistency over expanding trial spaces. The explicit circle posterior has a separate small-noise calculation. If a pointwise leading-eigenvalue tie occurs, $\mathcal{F}_d$ requires a fixed measurable choice; no Lipschitz stability of that filter is asserted without a gap. The processed comparison evaluates its own processed fields, not a filtered raw-error difference.

A.3 Which evidence addresses which question?

The initial 16-image experiment asks whether local geometry can be learned under matched priors and assisted centering. The four paired units at a fixed density are two images with two covariance training initializations, not four independent manifolds. The mean model remains fixed across those repeats. The improved-mean and target-only arms distinguish centering from the frozen mean input. The later input intervention is a separate matched follow-up. Every paired outcome is included in its corresponding comparison.

Numerical posterior order checks vary the teacher approximation; grid and shifted-grid checks vary learned-field integration. Multi-start mode checks probe a different issue: a stable local quadrature may still miss remote posterior mass. The reference uses the generated latent only as an initializer; it never enters the predictor. Finite checks are not global tail certificates, and numerical teachers are not assumed identical to the ideal conditional expectations in Proposition 4.

The training design and final checkpoint rule were frozen before training. The tenfold paired-effect screen was introduced in the post-run analysis: if e_s, e_c are the sample/conditional grid-256 errors, a positive improvement is called numerically resolved when

$$e_s - e_c > 10(\Delta_s + \Delta_c), \qquad \Delta_a = \frac{\max\{\|H_{a,256} - H_{a,128}\|_{\mathrm{op}}, \|H_{a,256} - H_{a,256,\mathrm{shift}}\|_{\mathrm{op}}\}}{\|H_P\|_{\mathrm{op}}}.$$

This descriptive screen gives the reported four-of-four linear and two-of-four cubic counts. It was not preregistered and is neither a statistical test nor a rigorous quadrature bound. A positive form comparison does not require the loose sufficient bound to be numerically certified; a spectral claim still requires a separate gap and numerical-precision assessment.

B Geometric transfer: proofs, mass error, and diagnostic scope

This section concerns the geometric reference H_P of Equation (1). All bounds use the specified ambient trial functions and the fixed geometric reference measure.

B.1 LOCAL AND BLOCKWISE FORM BOUNDS

Proof of Proposition 1. Write $E = C - P$. Jensen's inequality, submultiplicativity, and $\|E\|_{\mathrm{op}} \le \|E\|_F$ give

$$\|\mathbb{E}[JEJ^\top]\|_{\mathrm{op}} \le \mathbb{E}[\|J\|_{\mathrm{op}}^2 \|E\|_F] \le (\mathbb{E}\|J\|_{\mathrm{op}}^4)^{1/2} (\mathbb{E}\|E\|_F^2)^{1/2}.$$

The assumed moments make this difference integrable; H_P is integrable because $\|P\|_{\mathrm{op}} = 1$ and ν is a probability measure. Substitution of e_{loc} gives the relative statement. Inserting A_t and applying Minkowski's inequality proves the two-term geometric error budget. □

For a more directional bound, let $J_T = JP$ and $J_N = JN$, and define

$$e_T^2 = \mathbb{E}\|P(C-P)P\|_F^2, \qquad e_X^2 = \mathbb{E}\|PCN\|_F^2, \qquad e_N^2 = \mathbb{E}\|NCN\|_F^2,$$
$$\kappa_T^2 = \mathbb{E}\|J_T\|_{\mathrm{op}}^4, \qquad \kappa_X^2 = \mathbb{E}[\|J_T\|_{\mathrm{op}}^2 \|J_N\|_{\mathrm{op}}^2], \qquad \kappa_N^2 = \mathbb{E}\|J_N\|_{\mathrm{op}}^4.$$

When these moments are finite, the three terms in Equation (4) satisfy

$$\|H_C - H_P\|_{\mathrm{op}} \le \kappa_T e_T + 2\kappa_X e_X + \kappa_N e_N. \tag{19}$$

Indeed, $C - P = P(C-P)P + PCN + NCP + NCN$ for symmetric C; each block follows from the same Cauchy–Schwarz argument, and the cross term includes its transpose. This sum is an upper bound, not an attribution identity for norms. The signed contribution matrices must be retained to inspect cancellation. If $J_N = 0$ on the manifold, the cross and normal contributions vanish. True PCP removes those contributions while preserving predicted tangent eigenvalues; it is not the same control as replacing C by P.

B.2 INTEGRATION, MASS ESTIMATION, AND A GEOMETRIC SPECTRAL BOUND

Let $K_C = \mathbb{E}[D\Phi C D\Phi^\top]$, with empirical or quadrature stiffness $\widehat{K}$ and mass $\widehat{M}$. Define

$$\eta_Q = \|M^{-1/2}(\widehat{K} - K_C)M^{-1/2}\|_{\mathrm{op}}, \qquad \rho = \|M^{-1/2}\widehat{M}M^{-1/2} - I\|_{\mathrm{op}} < 1. \tag{20}$$

Set $\eta = \eta_Q + \kappa_J \|C - P\|_{2,F}$, or replace the second term by either the two-term A_t budget or the blockwise bound above. Then $\|M^{-1/2}\widehat{K}M^{-1/2} - H_P\|_{\mathrm{op}} \le \eta$.

Proof of Corollary 5. Let $S = M^{-1/2}\widehat{M}M^{-1/2}$ and $R = S^{-1/2}$. The relative-mass congruence $R(M^{-1/2}\widehat{K}M^{-1/2})R$ differs from H_P by at most $\eta/(1-\rho) + \rho\|H_P\|_{\mathrm{op}}/(1-\rho)$. For the reference part this follows by inserting $R - I$ on either side, using $\|R - I\|_{\mathrm{op}} \le (1-\rho)^{-1/2} - 1$ and $\|R\|_{\mathrm{op}} \le (1-\rho)^{-1/2}$. Weyl's inequality preserves the separated prefix, and the separated-prefix Davis–Kahan bound is $\varepsilon_P/(\gamma_P - \varepsilon_P)$. Transporting its basis through R back to the true whitened function coordinates adds at most $\|R - I\|_{\mathrm{op}}/s_{\min}(R) \le b_M$. Indeed, for an orthonormal basis V, the distance from RV to its original span is at most $\|(I - VV^\top)(R-I)V\|_{\mathrm{op}}/s_{\min}(RV)$, which is bounded by b_M. The triangle inequality for orthogonal projectors then gives the stated bound, capped by one. □

The stated condition is sufficient, not necessary. Grid-to-grid differences are not upper bounds on η_Q without additional regularity and quadrature analysis. Accurate integration of P and the mass does not establish accurate integration of the learned field C. Even a valid finite-space certificate would require a separate Galerkin approximation argument before it could imply recovery of an infinite-dimensional operator or all its eigenfunctions.

B.3 PROOF AND LIMITS OF THE CIRCLE EXAMPLE

Proof of Proposition 2. On the circle write $n = x = (\cos\theta, \sin\theta)$ and $\tau = (-\sin\theta, \cos\theta)$. Then $P = \tau\tau^\top$ and $C_\delta = \tau\tau^\top + \delta nn^\top$, so its eigenvalues are $1, \delta$. The gradient of Equation (5) restricted to the circle is $\sqrt{2}(e_1 + ax_1 x)$. Its tangent and normal components have scalar coefficients $-\sqrt{2}\sin\theta$ and $\sqrt{2}(a+1)\cos\theta$. Uniform integration gives energies 1 and $1 + \delta(a+1)^2$, respectively. Also $C_\delta - P = \delta nn^\top$ and $\|P\|_F = 1$. □

The manifold function, mass, and intrinsic energy remain fixed as a varies; only its ambient extension changes. The local form using a leaking PSD field therefore need not be well-defined on equivalence classes of manifold functions without specifying the extension. The example supplies no claim that C_δ is a normalized posterior covariance of a particular prior. It does not preclude convergence for a fixed, gradient-controlled dictionary.

A gapped geometric mode reversal. Pair ψ_a with the ambient polynomial $\chi(x) = \sqrt{2}(x_1^2 - x_2^2)(2 - \|x\|^2)$. Its restriction is $\sqrt{2}\cos 2\theta$, its normal derivative vanishes on the circle, and its tangent derivative is $-2\sqrt{2}\sin 2\theta$. The restricted functions are centered and orthonormal. Orthogonality of the first and second harmonics therefore gives, in the same trial basis,

$$M = I_2, \qquad H_P = \operatorname{diag}(1, 4), \qquad H_{C_\delta} = \operatorname{diag}\big(1 + \delta(a+1)^2, 4\big). \tag{21}$$

The reference gap is 3, and its lowest mode is orthogonal to the estimated lowest mode whenever $\delta(a+1)^2 > 3$. For example, $\delta = .03, a = 10$ gives $H_{C_\delta} = \operatorname{diag}(4.63, 4)$ with projector error one. Local matrix error can be arbitrarily small by choosing the extension parameter accordingly; fixing the parameter restores convergence as $\delta \to 0$. These are finite compressions with specified ambient extensions, not a claim that a leaking field defines an intrinsic operator on all manifold functions.

An exact posterior variant. We prove Corollary 3, illustrated in Figure 1. The extension obstruction also occurs for a genuine Gaussian observation model, without asserting that its covariance equals the constructed C_δ. Let Z be uniform on S^1 and $Y_t = Z + \sqrt{t}\epsilon$. At a clean query $x \in S^1$, the posterior angle u relative to x has density proportional to $\exp(\cos u/t)$. Write $m_t = \mathbb{E}[\cos u \mid Y_t = x]$. Rotational symmetry and integration by parts give

$$A_t(x) = m_t P(x) + n_t x x^\top, \qquad n_t = t^{-1}\operatorname{Var}(\cos u \mid Y_t = x) > 0. \tag{22}$$

Indeed, the cross covariance vanishes by symmetry, and integrating $\partial_u[\sin u \exp(\cos u/t)]$ over a period gives $t^{-1}\mathbb{E}\sin^2 u = \mathbb{E}\cos u = m_t$. Hence, for the same ambient extensions,

$$\mathcal{E}_{A_t}(\psi_a, \psi_a) = m_t + n_t(a+1)^2. \tag{23}$$

For every fixed $t > 0$, the geometric form error is unbounded as $|a|$ grows, although the intrinsic function and its mass remain unchanged. On the two-function space above, $H_{A_t} = \operatorname{diag}(m_t + n_t(a+1)^2, 4m_t)$, so a geometric mode reversal occurs when $n_t(a+1)^2 > 3m_t$.

The local matrix error nevertheless tends to zero. On $[-\pi, \pi]$, $2u^2/\pi^2 \le 1 - \cos u \le u^2/2$. The normalizing integral of $\exp[(\cos u - 1)/t]$ is bounded below by a constant times $\sqrt{t}$ for small t. Gaussian domination of the numerator then gives $\mathbb{E}|u|^{2r} = O(t^r)$ for fixed r. Consequently, $1 - m_t \le \mathbb{E}u^2/2 = O(t)$ and $n_t \le \mathbb{E}u^4/(4t) = O(t)$. Thus $\|A_t - P\|_F = O(t)$ and the leading direction is exactly tangent for all sufficiently small t. For fixed a, the form error remains $O(t)$; there is no contradiction to the small-noise limit. This exact-posterior control is a geometric-extension example, not an attribution of the CIFAR form errors.

B.4 Mean contamination and query distributions

For $b = \mu_t - m$, Equation (8) implies pointwise

$$\operatorname{tr}(B_m - A_t) = \|B_m - A_t\|_{\mathrm{op}} = \|b\|^2/t. \tag{24}$$

Thus the noisy-query mean metric near 51 in Table 3 is also the average pointwise operator norm of target contamination. It is not $\|\mathbb{E}(B_m - A_t)\|_{\mathrm{op}}$ and is not a geometric form error. A clean-query L^2 budget for this rank-one term would instead use $(\mathbb{E}_\nu\|b(x)\|^4/t^2)^{1/2}$, together with the appropriate gradients. Noisy-query and clean-query measures cannot silently be interchanged.

The path $H_C \to H_{A_t} \to H_P$ separates estimation and finite-noise geometric bias. Appendix G evaluates these components on the same form and clean-query measure in the paired cohort.

B.5 Conditional supervision and gradient variance

For Proposition 4, condition on Y_t and expand $C - R_m = (C - B_m) + (B_m - R_m)$. The cross inner product has zero conditional expectation because C and B_m are measurable functions of Y_t and $\mathbb{E}[R_m - B_m \mid Y_t] = 0$. This proves the loss identity for any common hypothesis class.

The corresponding Rao–Blackwell conditional-variance identity (Liu et al., 2019) concerns gradients at a fixed parameter. For the loss $\frac{1}{2}\|C_\theta(Y_t) - R_m\|_F^2$, write $L_\theta = D_\theta \operatorname{vec} C_\theta(Y_t)$ and let g_R, g_B be its parameter gradients using sample and exact conditional labels. At fixed θ, assuming the required derivative moments exist,

$$\mathbb{E}[g_R \mid Y_t] = g_B, \qquad \operatorname{Cov}(g_R) - \operatorname{Cov}(g_B) = \mathbb{E}\big[L_\theta^\top \operatorname{Cov}(\operatorname{vec} R_m \mid Y_t) L_\theta\big] \succeq 0. \tag{25}$$

This removes conditional randomness from the raw gradient without changing its expectation. Our optimizer uses Adam, clipping, and EMA: equality of expected raw gradients does not imply equality of expected clipped or adaptively transformed updates. We test the final effect of matched supervision; we do not infer the optimizer trajectory or a convergence guarantee from this identity.

For Equation (25), write $g_R = g_B - L_\theta^\top \operatorname{vec}(R_m - B_m)$. Conditional centering eliminates the cross covariance, and the law of total covariance gives the stated PSD difference. Both statements assume the displayed moments exist. With a numerical teacher $\widetilde{B}_m$, the population target is instead $\widetilde{B}_m$; the ideal identity does not remove its approximation error. The order checks in the experiments assess that numerical difference rather than asserting exact equality for a numerical teacher.

C Distribution-matched image-orbit geometry and repair diagnostics

This study provides the controlled geometric reference for the manifold transfer analysis. It asks whether learning on a specified image distribution recovers its normalized posterior covariance, tangent projector, and a finite weighted geometric form. The subsequent filtering and continuation experiments were designed after observing the initial form discrepancies; their protocols were frozen before their own outputs, not before the original study. All results are descriptive. Checkpoints, trial functions and spectral prefixes are held fixed across the reported comparisons. The coarse-grid form estimates below document the initial study. For $\sigma = .02$, the refined evaluation in Appendix D supersedes them; matched retraining and attribution follow in Appendices E and G. The $\sigma = .16$ results remain coarse-grid estimates.

C.1 Matching the training prior and the reference geometry

For each of 16 previously exposed CIFAR base images, construct a smooth periodic translation orbit $F_i : \mathbb{T}^2 \to \mathbb{R}^{3072}$. Scale pixel values to $[-1, 1]$, remove the Nyquist row and column in the two-dimensional Fourier transform, and apply phase shifts before inverse transformation. The images are *not* clipped after rendering. For angular frequencies n_1, n_2, the multiplier is $\exp[-\mathrm{i}(n_1\theta_1 + n_2\theta_2)]$; a 2π shift corresponds to 32 pixels. This filtering and continuous rendering change the original CIFAR distribution. They provide an explicit geometric test, not an identified semantic manifold of natural images.

For each image and $\kappa \in \{0, 8\}$, generate

$$\theta \sim q_\kappa(\theta) \propto \exp[\kappa(\cos\theta_1 + \cos\theta_2)], \qquad X = F_i(\theta), \qquad Y = X + \sigma\epsilon. \tag{26}$$

Training and reference integration share the renderer, density, scaling, clipping policy, and noise levels. In contrast, an unconditional CIFAR denoiser and the posterior of one specified translation orbit generally have different priors. Prior mismatch therefore enters the comparison alongside local estimation error. The same 16 base images are conditional training priors for both densities, not 32 independent images or a held-out-image evaluation. Only continuous parameters and noise are newly sampled at evaluation.

Let $J_i(\theta) = \partial_\theta F_i(\theta)$ and $G_i = J_i^\top J_i$. Translation is orthogonal in the filtered ambient representation, so G_i is constant and positive definite for these images. With $Q_i = J_i G_i^{-1/2}$, the tangent projector and induced weighted form are

$$P_i(\theta) = Q_i(\theta) Q_i(\theta)^\top, \qquad J_i^+ P_i (J_i^+)^\top = G_i^{-1}, \tag{27}$$

$$\mathcal{E}_{q_\kappa}(u, v) = \int_{\mathbb{T}^2} (\nabla_\theta u)^\top G_i^{-1} \nabla_\theta v \, q_\kappa \, d\theta. \tag{28}$$

Here q_κ is a coordinate density, supplied rather than learned. The corresponding nonnegative operator is $-q_\kappa^{-1}\nabla_\theta \cdot (q_\kappa G_i^{-1}\nabla_\theta)$. For nonconstant metrics the Riemannian volume factor must also be

Table 3: Matched orbit study at $\sigma = .02$: descriptive means over 16 base images, one initialization per density. Noisy covariance errors average per-query ratios; clean projector error is a per-image weighted RMS ratio. All entries are ratios, not percentages. Oracle centering is oracle-assisted training. The mean Jacobian's normal response is a signed quadratic diagnostic.

Metric (C_A: oracle center; C_B: learned center)	$\kappa = 0$	$\kappa = 8$
Noisy mean bias / σ^2	50.877	51.150
Noisy C_A error to A_t	0.03172	0.03151
Noisy C_B error to B_m	0.99781	0.99748
Clean C_A full-projector error	0.02987	0.02967
Clean C_B full-projector error	1.80734	1.80822
Clean C_A restricted cometric error	0.01406	0.01377
Clean C_A inverse-metric error	0.01405	0.01387
Mean-Jacobian normal response	-0.18499	-0.18748

included when converting between coordinate and volume densities. A restricted tangent check alone is insufficient: ambient identity gives $Q_i^T I Q_i = I_2$ but does not suppress any normal direction.

C.2 Training, numerical posteriors, and error definitions

Each of the 32 image–density cases trains one mean network for 12,000 updates and two covariance networks for 9,000 updates, with batch size 64 and fresh online inputs. Thus a mean model sees 768,000 sampled pairs and each covariance model 576,000 pairs, not that many distinct source images. There is one initialization per case. Mean training minimizes pixel MSE divided by σ^2. The two covariance arms share initialization, batches, input features, optimizer schedule, and final EMA selection. Both use a rank-16-plus-diagonal PSD parameterization; rank 16 is an approximation at finite noise, not an asserted posterior rank.

The input consists of the noisy image, σ, and, for the covariance networks, the same frozen learned mean. The separate per-case weights encode image identity and density. No network receives θ, J_i, G_i, P_i, or renderer coordinates. The arm-specific residual centers are the numerical posterior mean or the learned mean, giving targets A_t or B_m in Equation (8). The oracle-centered arm therefore uses an oracle *training target*, although the oracle mean is not an additional network input. Covariance regression uses the factorized Frobenius objective divided by the fixed common constant 2, with target-only terms removed.

Noise standard deviations are $\{.02, .04, .08, .16, .32, .64\}$. Adam learning rates start at 5×10^{-4} for the mean and 10^{-3} for covariance, with 100-step warmup, cosine decay, and final EMA decay 0.995. The diagonal initialization is 10^{-9} with floor 10^{-12}; a larger isotropic floor would force normal leakage on this two-dimensional benchmark. These choices were fixed before the main cases. Independent calibration streams of 4,096 examples per scale and arm supply full constant covariance baselines; calibration does not use evaluation observations.

The posterior reference integrates the orbit likelihood times q_κ. A local Gauss–Hermite importance grid with proposal covariance $\sigma^2 G_i^{-1}$ is centered by six Gauss–Newton updates. The generated latent initializes the reference only, never the learned model. Training uses order 9, checked against 13; evaluation uses 13, checked against 17 on fixed queries. A global 64×64 scan and four local starts audit alternative modes. These finite diagnostics are not global tail-integral certificates. Earlier clean-query references for images 00 and 12 at $\kappa = 8$, $\sigma = .64$ remain flagged by the tail gate and excluded from certified posterior claims. All 32 matched-study numerical diagnostic suites passed.

At each scale, 256 fresh noisy queries measure $\mathbb{E}\|\mu_t(Y) - m(Y)\|^2/\sigma^2$ and mean per-query relative Frobenius errors $\mathbb{E}[\|C - A_t\|_F/\|A_t\|_F]$, and likewise for B_m. Clean-orbit evaluation uses weighted 32×32 periodic quadrature at all scales, and 64×64 refinement at $\sigma = .02$. For each image, full projector error is $[\mathbb{E}_{q_\kappa}\|C(F_i(\theta)) - P_i(\theta)\|_F^2/2]^{1/2}$. Restricted cometric error averages $\|Q_i^T C Q_i - I_2\|_{\mathrm{op}}$; inverse-metric error uses the inverse restricted matrix only where its minimum eigenvalue exceeds 10^{-8} and retains the valid mass. Invalid or nearly singular inverses are not silently dropped. Tables average the resulting per-image metrics over the same 16 images.

C.3 Comparing local and weighted-form accuracy

Oracle-assisted training yields about 3% clean full-projector error and about 1.4% restricted cometric and inverse-metric errors at $\sigma = .02$ (Table 3). Learned centering has mean bias around 51 and relative error around 0.998 against its own residual target B_m. These measurements distinguish target contamination from fit to that target. The protocol first learns a mean and then fits covariance around the frozen center. The local metrics average over each orbit. A separate mean-Jacobian comparator has small tangent restriction error but negative mean normal quadratic responses, about $-0.185/-0.187$ for the two densities. It is not a PSD covariance field. For the learned-centered image-11 uniform case, the restricted matrix becomes nearly singular; the inverse-metric error is about 2.59×10^6. The large average is retained rather than replaced by the much smaller median.

For fixed ambient functions ϕ with clean mass M, let

$$H_C = M^{-1/2}\mathbb{E}_q[D\phi(F_i)C(F_i)D\phi(F_i)^\top]M^{-1/2}, \qquad e_C = \frac{\|H_C - H_P\|_{\mathrm{op}}}{\|H_P\|_{\mathrm{op}}}, \tag{29}$$

where H_P substitutes the analytic tangent projector. This tests a compression of Equation (28). Comparisons to the finite-noise reference A_t are retained separately. Neither noisy-query fitting nor clean local comparison alone controls the integrated form error or every term of Corollary 5.

The initial 12 ambient functions are linear Fourier pixel probes, with frequencies $(1,0)$, $(0,1)$, $(1,1)$, $(1,-1)$, $(2,0)$, and $(0,2)$, and sine/cosine pairs. Their ambient gradients are constant. If R contains the probe directions and T is their fixed mass-normalization map, the measured matrix is

$$H_C = T^T R^T \mathbb{E}_q[C(F_i(\theta))]RT. \tag{30}$$

Thus this test observes a density-averaged field, not its spatial variation. An oracle-centered constant baseline has full-projector errors about $0.996/0.959$, yet linear-form errors only $0.0789/0.2884$ for $\kappa = 0/8$. The learned oracle-centered field instead has linear-form errors $0.0521/56.3456$, despite substantially better local geometry. The latter value is a multiple of the reference norm, not a percentage. Proposition 1 identifies a direct derivative-weighted control of this comparison; unwhitened mass conditioning alone is insufficient.

C.4 Historical coarse-grid filtering and nonlinear probes

The numerical values in this subsection are the historical 32/64-grid estimates. The completed $\sigma = .02$ refinement in Appendix D resolves the principal filtering effects; the coarse sensitivities below motivated that study. The repair study freezes the original models and infers the leading two vectors $\widehat{Q}$ of their full predicted covariance, including its learned diagonal. It compares raw C, retained-eigenvalue filtering $\widehat{P}C\widehat{P}$, and the unit projector $\widehat{P} = \widehat{Q}\widehat{Q}^T$. Rank two is supplied by the known benchmark dimension, not estimated. True PCP is an oracle diagnostic. No true tangent vectors enter $\widehat{Q}$. Spectral truncation is a standard baseline; setting eigenvalues to one targets a small-noise geometric limit and is not a general identity for finite-noise covariance.

The same models are evaluated at $\sigma = .02, .16$, using 32×32 clean grids and 256 new noisy queries per case/scale; $\sigma = .02$ is refined to 64×64. New 12-function cubic probes are frozen before these repair outputs. Writing $z(x) = x^T RB^{-1}$ for the orbit-normalized Fourier coordinates, their unwhitened functions are $z_j(x)^3/3$; B here is a deterministic coordinate matrix, unrelated to the residual target B_m. Their gradients vary with x. Analytic centering, mass, and geometric stiffness use Fourier identities for first and third harmonics. No candidate uses oracle tangent projection in its ambient function gradients.

Filtering improves the reported oracle-centered form errors in both trial spaces (Table 4). In the concentrated linear case, 56.3456 falls to 0.2267; unit eigenvalues give 0.2294, while true PCP gives 0.00809. The latter preserves learned tangent eigenvalues. These comparisons prioritize directional/cross/normal diagnostics, without yet establishing a unique cause. The density is supplied analytically and held fixed throughout these comparisons. The cubic estimate falls from 3.4067 to 0.0252. However, the mean rank-2 matrix changes under grid refinement are 0.2554 and 0.0420, respectively. Even under uniform density the change is about 0.0100, comparable to the rank-2 error. These historical estimates alone did not resolve the form improvements; the subsequent frozen-field refinement provides the current values. Local full-projector errors change only modestly:

Table 4: Historical filtering estimates at $\sigma = .02$, 64×64 clean grid: mean $\|H_C - H_P\|_{\mathrm{op}}/\|H_P\|_{\mathrm{op}}$ over 16 images. Rank 2 retains predicted eigenvalues; Unit sets them to one; true P uses oracle tangent directions. $\Delta_{32,64}$ is the rank-2 matrix change divided by $\|H_P\|_{\mathrm{op}}$, not an error bar. These estimates are superseded by the refined Appendix D evaluation. Every entry is a ratio.

κ	Functions	Raw	Rank 2	Unit	True P	$\Delta_{32,64}$
Oracle-centered training						
0	linear	0.05205	0.01008	0.00479	0.00844	0.01002
0	cubic	0.16583	0.01041	0.00539	0.00843	0.01003
8	linear	56.346	0.22673	0.22941	0.00809	0.25540
8	cubic	3.40671	0.02515	0.02554	0.00772	0.04199
Learned-centered training						
0	linear	71.292	0.88089	0.53262	0.33772	0.06772
0	cubic	198.019	1.31133	0.89985	0.33683	0.14267
8	linear	64510.31	234.109	192.800	0.30776	55.709
8	cubic	4018.45	15.115	12.241	0.31161	3.70846

Table 5: Historical oracle-centered constant-per-scale control and trial-space conditioning at $\sigma = .02$. Constant-form errors use the 64-grid estimate. This control was added after the repair results, reusing the original frozen calibration matrices. Condition numbers refer to the declared unwhitened angular basis; different function spaces are not interchangeable tests.

κ	Constant: linear	Constant: cubic	Cond.: linear	Cond.: cubic
0	0.07893	2.06980	1.00	1.00
8	0.28839	1.46179	23442.91	117.98

approximately 0.030 raw, 0.029 rank 2, and 0.027 unit projector for the oracle-centered models. Removing normal leakage can therefore matter much more to selected ambient functions than to an aggregate local metric.

Cubic functions expose the limitation of a constant field (Table 5), but they also change the tested space: at $\kappa = 8$ its unwhitened mass condition number is about 118, versus 23,443 for the linear space. We therefore report both dictionaries separately, including the filtered learned-center values at concentrated density. The full $\sigma = .16$ comparison is retained in Table 6; its 32-grid estimates have no 64-grid refinement in this study. Clean-form improvements do not establish better noisy-query fitting to A_t or B_m.

Independent checks cover 160 archives and recompute 1,536 matrix errors. All eigensolver and reference-order gates pass; the maximum relative eigen residual is below 10^{-6}. Raw predictions reproduce the earlier study to relative error below 10^{-12}. Analytic reference mass and stiffness agree with quadrature to below 5×10^{-9}. Those reference-function checks do not imply equally accurate integration of a learned field; the refinement column explicitly measures the latter sensitivity.

C.5 Training pilots and the two uses of 16

Pilots use only images 00 and 01, both densities, and two continuation random streams per case: eight runs, not eight images or independent initializations. At $\sigma = .02$, two mean branches start from the same original checkpoint and receive sample-X or numerical posterior-mean targets for 3,000 steps, batch 64, learning rate 10^{-4}. Separately, two covariance branches start from the original learned-centered checkpoint and use sample residuals or numerical exact-B_m targets for 3,000 steps, learning rate 2×10^{-4}. Both covariance branches keep the *original* mean fixed. Paired branches share online inputs; each stage uses 192,000 fresh parameter/noise pairs. Final EMA decay is 0.995, with no evaluation-based checkpoint selection. Oracle teacher labels are diagnostic assistance, not a reference-free method.

Mean bias decreases from 35.186 to 15.627 with sample targets and 15.217 with oracle means (Table 7), still well above one. The starting value differs from the 16-image mean of about 51 because the pilot uses a smaller cohort. Covariance error against its own target barely changes: 0.995049

Table 6: Historical repair comparisons at $\sigma = .16$, 32×32 grid, mean over 16 images. These are geometric-reference errors, not finite-noise posterior-regression errors, and have no 64-grid convergence check.

κ	Functions	Raw	Rank 2	Unit	True P
Oracle-centered training					
0	linear	0.04802	0.03420	0.03030	0.00921
0	cubic	0.16002	0.03359	0.02969	0.00925
8	linear	57.580	0.43842	0.44441	0.00977
8	cubic	3.47858	0.06133	0.05984	0.01006
Learned-centered training					
0	linear	6.79355	0.75828	0.18800	0.55977
0	cubic	20.290	0.77945	0.18972	0.55867
8	linear	7067.03	31.194	21.989	0.54030
8	cubic	423.904	2.71127	1.64603	0.54769

Table 7: Continuation pilots: descriptive means of eight runs on two base images. Mean and covariance branches are separate ablations. Covariance branches keep the old mean fixed and do not test retraining around either improved mean. Errors against B_m are relative Frobenius ratios.

Metric	Original	Sample targets	Exact targets
Mean bias / σ^2	35.185954	15.626843	15.216886
Covariance error to B_m	0.995049	0.994897	0.994874

to 0.994897 or 0.994874. The two supervision choices give similar own-target errors under this initialization and budget. Architecture, factor parameterization and optimizer are held fixed in this comparison. The late-continuation gradient-clipping fraction is zero.

A separate diagnostic freshly initializes the covariance network and fits *16 fixed noisy examples* with exact-B_m labels for 1,200 updates at learning rate 10^{-3}. Unlike the continuations, it evaluates final raw weights, not EMA. Mean training relative error reaches 0.04912, but error on fresh queries is 1.33761. This contrast measures fitting to the fixed subset and transfer to fresh queries from the same orbit. The *16 base images* of the main study define continuous online training distributions. The original pilots vary mean and covariance training separately; the matched cold-start experiment in Appendix E evaluates covariance retraining around the improved means.

C.6 Independent image orbits and mixture geometry

The current implementation aggregates diagnostics over independently trained orbits. It does not learn overlaps, transitions, a joint density, or a global operator coupling those orbits. Image-specific learning and local chart models have precedents (Shocher et al., 2018; Schonsheck et al., 2019; Alberti et al., 2024). The contribution of this setup is a matched reference for distinguishing centering, target fitting, local geometry, and weighted-form errors.

For a mixture indexed by I, define $\alpha_i(y) = \Pr(I = i \mid Y = y)$, $\mu_i(y) = \mathbb{E}[X \mid Y = y, I = i]$, and $C_i(y) = \mathrm{Cov}(X \mid Y = y, I = i)$. Conditioning on the component index gives

$$\mathrm{Cov}(X \mid y) = \sum_i \alpha_i C_i + \sum_i \alpha_i (\mu_i - \mu)(\mu_i - \mu)^T, \qquad \mu = \sum_i \alpha_i \mu_i. \tag{31}$$

A prior-weighted average of local covariances omits posterior responsibilities and the between-component term. The different image orbits have not been shown to overlap as charts of a shared semantic manifold, so this mixture formula alone does not provide the transition maps needed for an atlas. An error-controlled mixture-operator assembly would require a new implemented estimator, analysis, and baselines; it remains future work.

The subsequent studies have completed learned-field integration checks, matched conditional-target controls, and covariance retraining around the improved mean (Appendices D and E). Appendix G adds the signed form-level attribution. These separate targets, fitting and numerical sensitivity without

Table 8: Exact shared body and output specifications. Every padded convolution uses circular padding. Upsampling is nearest-neighbor; concatenation uses the same-resolution encoder feature.

Stage	Operation
Input / resolution 32	3×3 convolution, 3 (mean) or 6 (covariance) input channels to 32; residual block width 32.
Resolution 16	3×3 stride-2 convolution $32 \to 64$; residual block width 64.
Resolution 8	3×3 stride-2 convolution $64 \to 128$; two residual blocks width 128.
Decoder resolution 16	Upsample, concatenate 64-channel skip; 3×3 convolution $192 \to 64$; residual block width 64.
Decoder resolution 32	Upsample, concatenate 32-channel skip; 3×3 convolution $96 \to 64$; residual block width 64; GroupNorm(8), SiLU.
Mean head	1×1 convolution $64 \to 3$, zero weights/bias; scalar preconditioning below. Total 1,236,483 parameters.
Covariance heads	1×1 convolutions $64 \to 48$ and $64 \to 3$; factor rank 16 and diagonal. Total 1,240,467 parameters.

proving a capacity or optimization impossibility. A shared model and genuinely unseen base images would be needed before expanding the generalization claim. The original per-case results for both function spaces are retained.

C.7 Implementation and trial specification

This section specifies the data identities, network architectures, training settings and ambient trial construction used for the runs. As an internal implementation check, standalone definitions extracted from the training code give identical initialized state dictionaries and CPU forward outputs for both networks, and identical trial-construction arrays for all 16 images and both densities. This consistency check uses the existing frozen training runs.

Data identity. We use the CIFAR-10 Python archive `cifar-10-python.tar.gz` from `https://www.cs.toronto.edu/~kriz/cifar.html`, with MD5 `c58f30108f718f92721af3b95e74349a`. Internal images 00–15 are zero-based entries in the original `test_batch`, in the following order:

$$(8795, 6915, 2806, 7283, 2075, 488, 8792, 5247, 9015, 6289, 2976, 3074, 8383, 6283, 7628, 7540).$$

Their labels are $(0, 1, \ldots, 9, 0, 1, \ldots, 5)$. They are previously exposed conditional base images, not an unseen-image test split. The paired cohort uses original indices 8795 and 6915. Original unsigned-byte arrays are scaled by $u \mapsto u/127.5 - 1$ before Nyquist filtering. Flattening is channel–row–column. The training-only smoke image is training entry 0 and does not enter reported geometry results.

Architecture and initialization. Both models use the same circular-convolution encoder–decoder body. Table 8 specifies its blocks. The noise embedding is $(\sin(\omega_j \log \sigma), \cos(\omega_j \log \sigma))_{j=1}^{16}$, with 16 linearly spaced $\omega_j \in [.25, 4]$, followed by linear layers $32 \to 128 \to 128$ with SiLU between them. Each residual block of width c uses GroupNorm(8), SiLU, a 3×3 convolution, a second GroupNorm(8), embedding-conditioned affine scale/shift, SiLU, and a second 3×3 convolution; residual addition is divided by $\sqrt{2}$. GroupNorm uses affine parameters and $\epsilon = 10^{-5}$. The embedding maps $128 \to 2c$ and second block convolutions start at zero. Other body layers use PyTorch's default initializers. No dropout or attention is used.

Writing the mean head output as h, with $s_d = .5$, the network is

$$m(y, \sigma) = \frac{s_d^2}{\sigma^2 + s_d^2} y + \frac{\sigma s_d}{\sqrt{\sigma^2 + s_d^2}} h\left(y/\sqrt{\sigma^2 + s_d^2}, \sigma\right).$$

The covariance body takes the unrescaled concatenation $(y, m(y, \sigma))$; the specified input-intervention arm alone uses $(y, (y - m)/\sigma)$. Its prediction is $C = UU^\top + \mathrm{diag}(d)$. Reshape the 48 factor channels as channel–rank–row–column, then flatten channel–row–column, and divide by $\sqrt{3072}$ to

obtain $U \in \mathbb{R}^{3072\times 16}$. Factor weights start i.i.d. $N(0, .03^2)$ with zero bias. Diagonal-head weights are zero, its bias is $\log(\mathrm{expm1}(10^{-9}))$, and $d = \mathrm{softplus}(h_d) + 10^{-12}$.

Network forward/backward arithmetic is float32, with no mixed precision or TF32; posterior geometry, rendering and form accumulation use float64. Numerical training targets are cast to float32 for regression. Adam uses $(\beta_1, \beta_2) = (.9, .999)$, $\epsilon = 10^{-8}$, no weight decay, and global gradient-norm clipping at 10. At zero-based update j of T, the learning rate is $\eta_0 \min(1, (j+1)/100)[.1 + .9(1 + \cos(\pi j/T))/2]$. EMA decay is .995 and only the final EMA is evaluated. The original mean and covariance models use 12,000 and 9,000 updates, with initial learning rates 5×10^{-4} and 10^{-3}, respectively, and batch size 64. Paired covariance retraining and the input intervention use 9,000 updates, batch size 64 and initial learning rate 10^{-3}. Their initialization and stream base seeds are 202609240101 and 202609240201; the evaluation base seed is 202609240301. For a paired case $c = 2i + \mathbf{1}_{\kappa=8}$ and covariance repeat r, initialization/stream seeds are their respective base seeds plus $100c + r$; the stream's Gaussian generator adds 1,000,000. All matched arms and label types share these seeds, draws and initial tensors. Mean improvements are frozen continuation models, not an additional independent mean replication.

Complete ambient trial construction. Let w_j run, in order, through $((1,0),(0,1),(1,1),(1,-1),(2,0),(0,2))$, and set $\ell = (.299, .587, .114)/\|(.299, .587, .114)\|$. At pixel (h, l) use $\alpha_{hl} = 2\pi(h, l)/32$. The 3072×12 probe matrix R has first six columns $\ell_c \sin(w_j^\top \alpha_{hl})/\sqrt{512}$ and last six $\ell_c \cos(w_j^\top \alpha_{hl})/\sqrt{512}$, with channel–row–column flattening. For the Nyquist-filtered zero-shift image x_0, write $x_0^\top R = (a^\top, b^\top)$ and define

$$B = \begin{pmatrix} \mathrm{diag}(b) & -\mathrm{diag}(a) \\ \mathrm{diag}(a) & \mathrm{diag}(b) \end{pmatrix}, \quad Z = RB^{-1}, \quad z(x) = Z^\top x.$$

On the orbit, z consists of the six sines followed by the six cosines of $w_j^\top \theta$. The unwhitened ambient functions are $z - \mathbb{E}_\nu z$ (linear) and $z^{\odot 3}/3 - \mathbb{E}_\nu[z^{\odot 3}/3]$ (cubic). Mass and stiffness follow exactly from $\mathbb{E}e^{iv^\top \theta} = \prod_{a=1}^{2} I_{|v_a|}(\kappa)/I_0(\kappa)$ and the constant G^{-1}. The cubic construction uses frequencies $w_j, 3w_j$ and the elementary third-power trigonometric identities. For each mass $M = V\Lambda V^\top$, multiply functions on the right by $W = V\Lambda^{-1/2}$. Thus the gradient columns are ZW for linear functions and $Z\,\mathrm{diag}(z^2)W$ for cubic functions. These formulas determine every per-image B, B^{-1}, G, whitening matrix and H_P. A fresh eigendecomposition can choose equivalent eigenbasis signs or rotations; comparisons of saved matrices must first align these coordinates.

Fixed-query visual check. Figure 6 expands the main-text illustration to all four image–initialization pairs. Each query is the clean zero-shift image at $\kappa = 8$ and $\sigma = .02$, using the saved 9,000-update final EMA. The main panel uses image 0 and repeat 0 by fixed index, without selection on a geometric score. The model inputs remain pixels and the frozen learned mean; true tangents are used only in evaluation and display alignment. If $\widehat{Q}$ spans the leading two eigenvectors, we display $\widehat{Q}R$, where R minimizes $\|\widehat{Q}R - Q\|_F$ over orthogonal matrices. This preserves the predicted projector but the aligned columns are not individually eigenvectors. Signed luminance maps share one color scale; eigenvalues are shown separately without rescaling. These clean-query images illustrate local geometry, not orbit-average form accuracy or cross-image geometric compatibility.

D FROZEN LEARNED-FIELD INTEGRATION AND DIRECTIONAL DIAGNOSTICS

This section checks the numerical integration of the original 16-image learned fields. The mass and tangent reference being analytic does not imply that neural matrix fields are integrated accurately. This study keeps the fields, dictionaries, and whitening fixed; subsequent matched retraining and same-form posterior attribution are separate studies.

D.1 COMPLETED FROZEN-FIELD INTEGRATION FOLLOW-UP

The follow-up freezes all 32 image–density cases, both learned covariance arms, $\sigma = .02$, and both ambient dictionaries. It retains the same analytic mass, centering, whitening, and reference stiffness while evaluating unshifted $64, 128, 256$ grids, half-cell shifted $128, 256$ grids, and an independent 65,536-sample density-weighted random integration. It saves the complete signed matrices

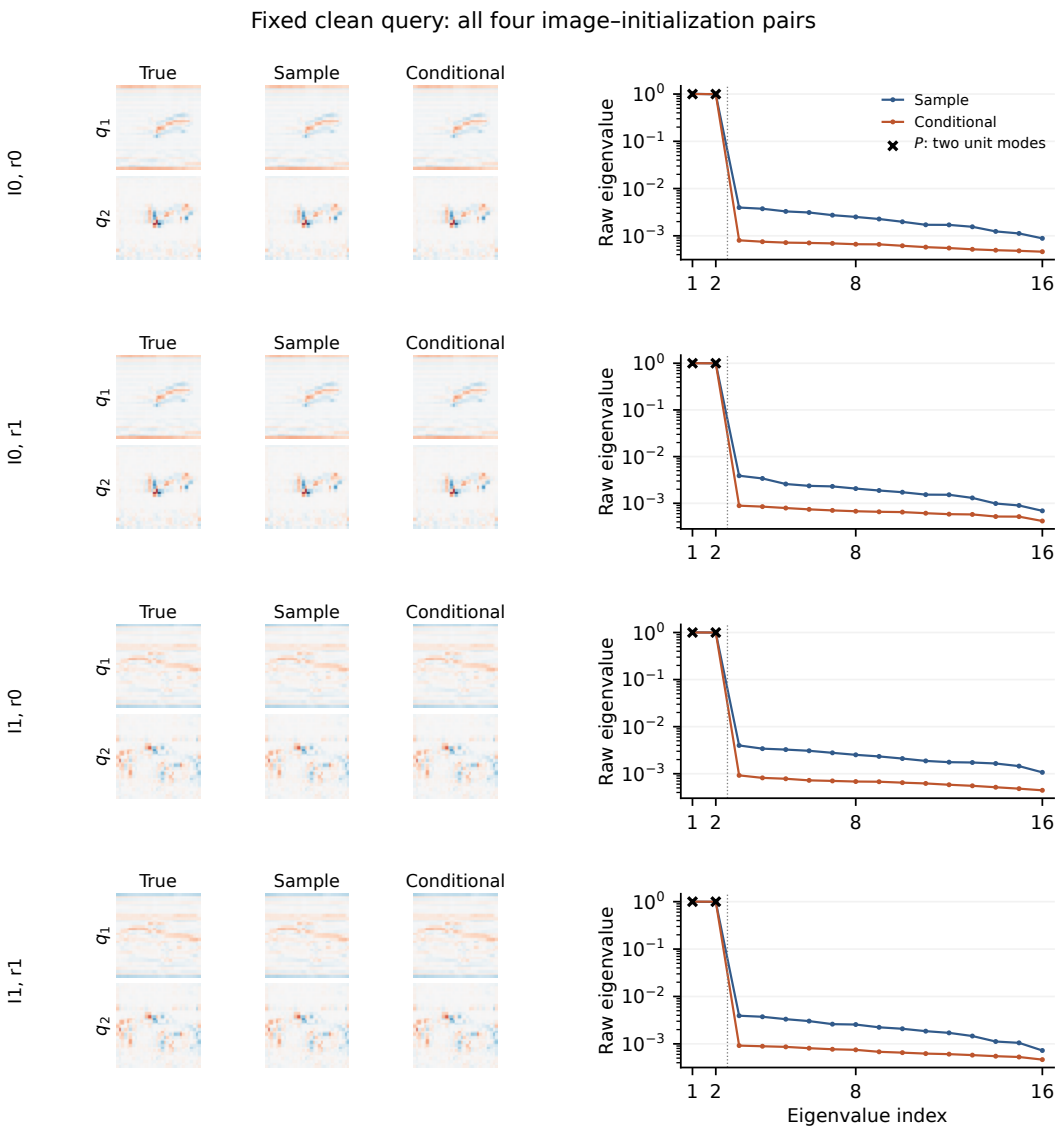

Figure 6: Pointwise local geometry for both base images and both covariance initializations. Each row compares the analytic frame with oracle-centered sample and conditional models, followed by their raw local spectra. Orthogonal basis alignment affects display only; no individual eigenvalue or projector is altered. These local spectra are distinct from the spectra of the weighted form matrices.

Table 9: Initial oracle-assisted cohort, $\sigma = .02$, 16-image means. Grid-256 errors and refinement/shift differences are normalized by $\|H_P\|_{\mathrm{op}}$. Pass counts resolved filtering improvements under the numerical screen.

κ	Functions	Raw	Rank 2	$\Delta_{128,256}$	Shift	Pass
0	Linear	0.05176	0.01055	0.000101	2.68×10^{-6}	15/16
0	Cubic	0.1653	0.01093	0.0001101	2.82×10^{-6}	16/16
8	Linear	56.38	0.2323	0.005253	0.0001391	16/16
8	Cubic	3.411	0.02567	0.0006373	1.13×10^{-5}	16/16

Table 10: Frozen-model integration follow-up, retained-rank-two field, $\sigma = .02$: means over 16 images. Error uses the 256-grid geometric reference comparison; Refine is the 128-to-256 matrix difference, Shift compares the two 256 grids, and Random compares 65,536 independent draws with the unshifted 256 grid. All are normalized by $\|H_P\|_{\mathrm{op}}$. Pass counts resolved improvements: cases whose maximum Refine/Shift is below 10% of both the remaining error and the raw-to-filter error reduction. This declared numerical tolerance is not a rigorous quadrature certificate.

κ	Functions	Center	Error	Refine	Shift	Random	Pass
0	Linear	Oracle	0.01055	0.000101	2.68×10^{-6}	0.008415	15/16
0	Linear	Learned	0.8807	0.000172	8.08×10^{-6}	0.01408	15/16
0	Cubic	Oracle	0.01093	0.0001101	2.82×10^{-6}	0.01063	16/16
0	Cubic	Learned	1.309	0.0003696	1.59×10^{-5}	0.02079	15/16
8	Linear	Oracle	0.2323	0.005253	0.0001391	0.03972	16/16
8	Linear	Learned	230	0.2526	0.00843	0.4699	16/16
8	Cubic	Oracle	0.02567	0.0006373	1.13×10^{-5}	0.01396	16/16
8	Cubic	Learned	14.84	0.01623	0.0005387	0.05382	16/16

in Equation (4), the gradient moments, and per-orbit amplification bounds, rather than inferring amplification from unwhitened mass conditioning. These are numerical sensitivity diagnostics, not rigorous integration certificates. The protocol was fixed after the initial form discrepancies and before this follow-up's results.

The follow-up reduces the earlier grid ambiguity without asserting a rigorous integration bound. For oracle-assisted rank-two fields at $\kappa = 8$, the 256-grid linear/cubic errors are $0.23235/0.02567$; the 128-to-256 changes are $0.005253/0.000637$, and the shifted-256 differences are $0.000139/0.0000113$. All 16 cases in each space satisfy the declared resolved-improvement tolerance. Under uniform density the linear/cubic means are $0.01055/0.01093$. One uniform linear case (image 01) worsens from raw error 0.006108 to 0.009135 after rank-two filtering; its refinement change is only 0.0000287. The increase in error exceeds the measured integration change, so the numerical screen records this case as a resolved deterioration. All cases are included.

Plain random integration has appreciable sampling variation (Table 10). After observing this, we additionally compute the known-reference control variate

$$\widetilde{H}_C^{\mathrm{MC}} = H_P + \frac{1}{n}\sum_{j=1}^{n} J(X_j)(C(X_j) - P(X_j))J(X_j)^\top.$$

This is unbiased for H_C, with the same independent queries, functions, and whitening. It is a post-hoc numerical diagnostic, not a learned-model repair. Its mean discrepancy from grid 256 for rank-two oracle fields is $0.000201/0.000361$ (uniform linear/cubic) and $0.004372/0.000862$ (concentrated linear/cubic). The original Monte Carlo estimates and declared flags are retained. No confidence interval or rigorous integration certificate is inferred.

In the concentrated linear case, raw normal, cross, and tangent contribution norms are 56.3824, 0.07621, and 0.00814, against total error 56.3793. After rank-two filtering they are 0.20882, 0.07622, and 0.00814, against total error 0.23235 (Table 11). Unit eigenvalues give 0.23528, whereas true PCP gives about 0.00814. This locates the dominant raw geometric discrepancy in the normal block and retains directional/cross error after truncation. It does not separate the true finite-noise normal

Table 11: Oracle-assisted fields at the 256 grid: norms of the tangent, cross (including transpose), and normal contribution matrices, and total error, each divided by $\|H_P\|_{\mathrm{op}}$ and averaged over 16 images. Contribution norms must not be added: the signed matrices, their exact sum, and per-image values are retained in the artifact.

κ	Functions	Field	Tangent	Cross	Normal	Total error
0	Linear	Raw	0.008439	0.004542	0.05112	0.05176
0	Linear	Rank 2	0.008439	0.004542	0.0002965	0.01055
0	Linear	Unit	0.0003906	0.004548	0.000297	0.004684
0	Cubic	Raw	0.008425	0.004861	0.1652	0.1653
0	Cubic	Rank 2	0.008425	0.00486	0.0009173	0.01093
0	Cubic	Unit	0.0003904	0.004873	0.0009184	0.00522
8	Linear	Raw	0.00814	0.07621	56.38	56.38
8	Linear	Rank 2	0.00814	0.07622	0.2088	0.2323
8	Linear	Unit	0.0003888	0.07639	0.2093	0.2353
8	Cubic	Raw	0.007771	0.01852	3.411	3.411
8	Cubic	Rank 2	0.007771	0.01852	0.0134	0.02567
8	Cubic	Unit	0.0003911	0.01855	0.01342	0.0261

Table 12: Oracle-assisted raw-field derivative diagnostics at the 256 grid. Local RMS, amplification factor, and relative bound values are computed *per image* before averaging. Direct and block bounds use Equations (2) and (19); they are empirical integral diagnostics, not population certificates. Their looseness is retained rather than interpreting them as accurate error predictions. Complete cases and fixed-prefix spectral errors are exported.

κ	Functions	e_{loc}	$\mathfrak{a}_\Phi$	Direct bound	Block bound	Error
0	Linear	0.02987	1740	52.24	14.04	0.05176
0	Cubic	0.02987	1.06×10^4	318.9	73.86	0.1653
8	Linear	0.02965	1.32×10^6	3.98×10^4	8680	56.38
8	Cubic	0.02965	8.76×10^4	2642	591.9	3.411

variance of A_t from covariance estimation error, or uniquely diagnose an optimization cause. The density is supplied.

The measured amplification factor averages about 1.32×10^6 for the concentrated linear space, versus 1.74×10^3 for the uniform one. The resulting bounds are loose (Table 12); they explain why local RMS alone supplies no useful guarantee here, not the precise error magnitude or a successful spectral certificate. Signed-block reconstruction has maximum relative residual 2.75×10^{-10}, the eigensolver residual is below 10^{-6}, and analytic reference stiffness agrees with grid 256 to 5.38×10^{-14}. These checks concern different errors. Batch-64 inference changes old-grid matrices by at most 1.47×10^{-6} relative to their own norms compared with the original batch-16 evaluation; the follow-up uses batch 64 throughout. All source checkpoints remain frozen.

The matched training controls are now complete (Appendix E): old, improved, and oracle residual centers are crossed with sample and conditional supervision, with cold starts, matched initialization, budget, and evaluation queries. A new-target/old-input arm separates target changes from the frozen-mean input representation. Appendices F and G add the input intervention and signed form attribution. Shared models and unseen base images concern broader generalization; they are not prerequisites for this controlled-manifold analysis or evidence of chart compatibility.

E Paired covariance retraining after mean improvement

This follow-up executes the previously specified centering repair, after the frozen-field integration audit. All 64 final-EMA models and their evaluations completed. It supplies a measured chain from mean bias to own-target fitting, true posterior covariance, clean geometry, and compressed weighted forms. Each link is measured separately, allowing target fitting, local geometry, and weighted forms to be compared within each training arm.

Table 13: Paired retraining: mean over two images and two covariance replications per density. Bias is $\mathbb{E}\|\mu_t - m\|^2/\sigma^2$ for the target center; Oracle has zero by construction although its input remains the old learned mean. Own target is B_m for learned centers and A_t for Oracle. Covariance columns are mean per-query relative Frobenius errors. Clean A_t uses independent clean queries. Ratios, not percentages.

Target	κ	Center/input	Bias	Own target	A_t	Clean A_t
Sample	0	Old	35.059	0.99501	1.8752	1.8744
Sample	0	New	15.524	0.98691	0.97297	0.9728
Sample	0	New/old input	15.524	0.98691	0.97352	0.97315
Sample	0	Oracle	0	0.033382	0.033382	0.027318
Sample	8	Old	35.049	0.99512	1.8638	1.8629
Sample	8	New	15.557	0.98725	0.9924	0.992
Sample	8	New/old input	15.557	0.98726	0.99102	0.99048
Sample	8	Oracle	0	0.033935	0.033935	0.028085
Conditional	0	Old	35.059	0.99504	1.8747	1.8741
Conditional	0	New	15.524	0.98681	0.98724	0.9872
Conditional	0	New/old input	15.524	0.98681	0.98733	0.98728
Conditional	0	Oracle	0	0.013635	0.013635	0.0062942
Conditional	8	Old	35.049	0.99489	1.8902	1.8899
Conditional	8	New	15.557	0.98711	1.0062	1.0066
Conditional	8	New/old input	15.557	0.98711	1.0064	1.0067
Conditional	8	Oracle	0	0.013772	0.013772	0.0066881

E.1 Matched interventions and inference scope

We use base images 00 and 01, both $\kappa = 0, 8$, and $\sigma = .02$. The improved mean is the existing sample-supervised 3,000-update continuation from pilot repeat zero, fixed by index in every case rather than selected by geometry. The oracle-mean-supervised pilot is not used. Both old and improved means are frozen and identical across the two new covariance replications; these do not replicate mean training or introduce held-out base images.

For each image–density case and covariance initialization, eight arms cross sample-residual or conditional-second-moment supervision with four center/input settings. Old uses the old mean for both; New uses the improved mean for both; New/old input changes the center but retains the old mean input; Oracle uses the numerical posterior mean as center and the old learned mean as input. Thus no oracle mean or tangent coordinates enter a network as features. Conditional targets are numerical A_t for Oracle and $B_m = A_t + (\mu_t - m)(\mu_t - m)^\top/\sigma^2$ otherwise. Sample targets use $(X - m)(X - m)^\top/\sigma^2$ with the corresponding center.

All covariance networks start from scratch: rank 16 plus diagonal, Adam with base learning rate 10^{-3}, 100-step warmup and cosine-to-0.1 decay, 9,000 updates of 64 fresh pairs, gradient clipping at 10, and final EMA decay 0.995. Within a replication the eight arms share initialization and input stream; initial tensor hashes and first/last batch hashes verify pairing across workers. There are two independent covariance initializations and streams per case, 64 networks in total. Historical multiscale covariance checkpoints are not matched comparators. All comparisons are descriptive: eight case/seed cells are only two base images, not eight independent image replications.

E.2 Noisy fitting and clean geometry are separate measurements

Each case uses 512 fresh noisy queries, common across all arms and replications, and a separate 512-query clean-orbit sample from the same known density. We retain per-query errors, target squared norms, absolute MSE, mean relative ratios, and aggregate RMS ratios. Training reference order 9 is checked against 13, and evaluation order 13 against 17 on fixed diagnostic queries; these remain numerical reference checks rather than rigorous tail bounds.

For sample supervision, pooling the two densities reduces mean bias from 35.05368 to 15.54035, and noisy covariance error to A_t from 1.86947 to 0.98268. But error to the arm's own target only changes from 0.995065 to 0.987081. Absolute own-target MSE falls from 1424.535 to 275.714, while target

Table 14: Clean geometry after paired retraining. Columns use the distinct RMS definitions in the text. Full P includes predicted normal energy; Tangent uses the leading rank-two projector. Improving directions can worsen restricted metric scales. Each row averages two images and two replications.

Target	κ	Center/input	Full P	Tangent	Cometric	Inverse
Sample	0	Old	1.8743	0.38075	0.25647	0.19055
Sample	0	New	0.97255	0.16951	0.34822	0.23543
Sample	0	New/old input	0.9729	0.16968	0.34851	0.23553
Sample	0	Oracle	0.027343	0.024204	0.011388	0.011545
Sample	8	Old	1.8628	0.36002	0.20808	0.15821
Sample	8	New	0.9917	0.17695	0.31849	0.22332
Sample	8	New/old input	0.99021	0.1766	0.31894	0.22358
Sample	8	Oracle	0.028106	0.024981	0.011008	0.011032
Conditional	0	Old	1.874	0.36594	0.26295	0.19322
Conditional	0	New	0.98696	0.16518	0.35627	0.23963
Conditional	0	New/old input	0.98705	0.16617	0.35678	0.23983
Conditional	0	Oracle	0.006295	0.0060253	0.00070398	0.00070401
Conditional	8	Old	1.8898	0.33524	0.24334	0.18399
Conditional	8	New	1.0062	0.16934	0.32477	0.22673
Conditional	8	New/old input	1.0063	0.17005	0.32505	0.22691
Conditional	8	Oracle	0.0066874	0.0064139	0.00066095	0.0006612

squared norm falls from 1434.287 to 280.749. Relative to the zero-matrix predictor, aggregate MSE is still $0.99320/0.98207$ for old/new centers. The absolute MSE and target norm decrease together, whereas their ratio changes less. Conditional supervision gives own-target errors $0.994964/0.986958$, close to the sample-supervised values under the same learned centers and training budget.

For clean geometry, define $S = Q^\top C Q$ and let $\widehat{P}$ project onto the two leading eigenvectors of the full predicted covariance, including its diagonal. In this follow-up tangent, cometric, and inverse errors are $e_{\tan} = [\mathbb{E}\|\widehat{P} - P\|_F^2/2]^{1/2}$, $[\mathbb{E}\|S - I_2\|_F^2/2]^{1/2}$, and $[\mathbb{E}\|S^{-1} - I_2\|_F^2/2]^{1/2}$, respectively. The last two are RMS Frobenius metrics, unlike the mean operator-norm restrictions reported for the original study. The minimum observed restricted eigenvalue is 0.96949; no singular inverse is discarded. Full P error instead compares the entire C to P on the weighted 256 grid and is not the leading-subspace error.

Sample-supervised tangent error improves from 0.37039 to 0.17323 in all 8 paired cells, but cometric error worsens from 0.23228 to 0.33335 and inverse error from 0.17438 to 0.22938, also in all 8. Conditional supervision has the same directions of change. Thus improved centering yields a partial geometric repair, not simultaneous recovery of directions and scales. New/old input gives nearly the same local results as New; these improvements do not require changing the mean input in this tested architecture, without establishing equivalence of the learned models.

E.3 Paired form comparisons and numerical resolution

All new fields are integrated on grids 64, 128, 256 and half-cell-shifted 256, with the original analytic mass, stiffness, functions and whitening unchanged. Write Δ for the larger normalized matrix difference between 128/256 and regular/shifted 256. Signed tangent/cross/normal matrices and spectral prefixes 2, 4, 8 are retained separately; reference quadrature does not stand in for learned-field convergence.

For sample supervision at $\kappa = 8$, New reduces the raw linear error from 27647.4 to 4613.49 and rank-two error from 62.7558 to 10.6592; the cubic rank-two error decreases from 4.1783 to 1.29947. Under uniform weighting, the linear rank-two mean increases from 0.80006 to 0.82612. Target-only concentrated linear error is 10.2693, close to New's 10.6592.

Conditional targets are useful when the center is accurate. For Oracle at $\kappa = 8$, noisy error to A_t drops from 0.03393 to 0.01377, tangent error from 0.02498 to 0.00641, and rank-two linear/cubic form errors from $0.24795/0.02303$ to $0.03730/0.00677$. The corresponding conditional-field integration changes are $0.000668/0.000103$. This rules out the broad claim that target sampling noise

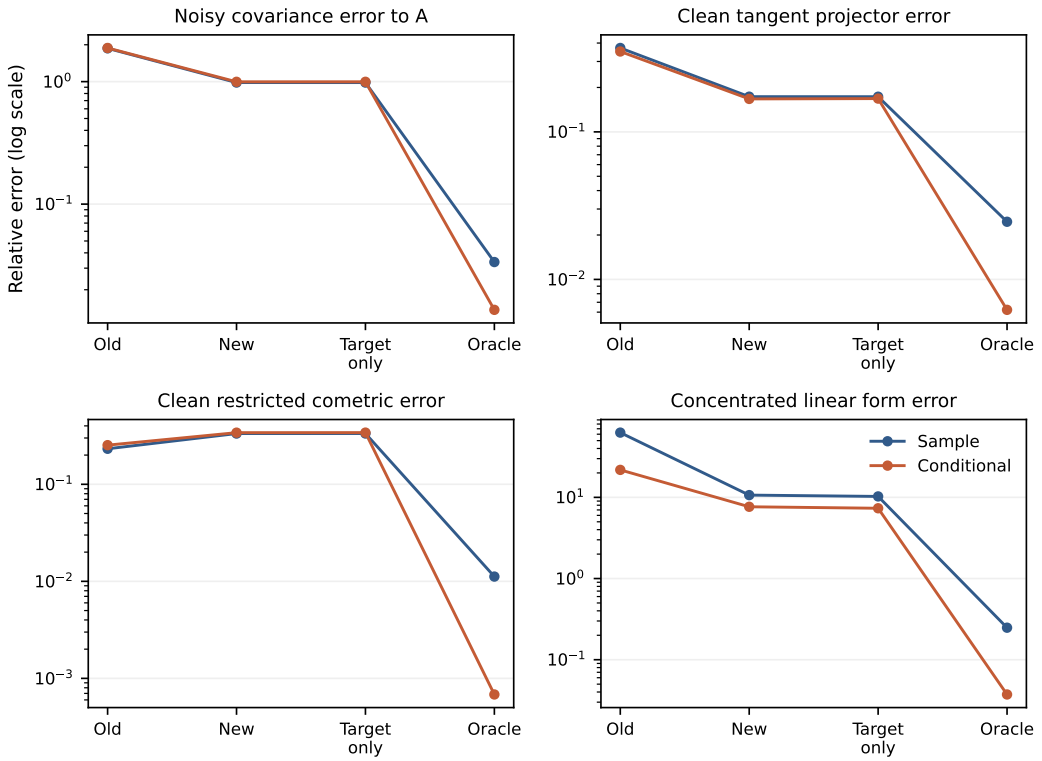

Figure 7: Completed paired interventions. Curves are descriptive means, with no confidence bands: the first three panels pool both densities and two covariance replications on two images; the last uses only $\kappa = 8$ and rank-two filtering. All vertical axes are logarithmic. The Oracle branches remain assisted controls. Per-case outcomes, including reversals of average improvements, are retained in the artifact.

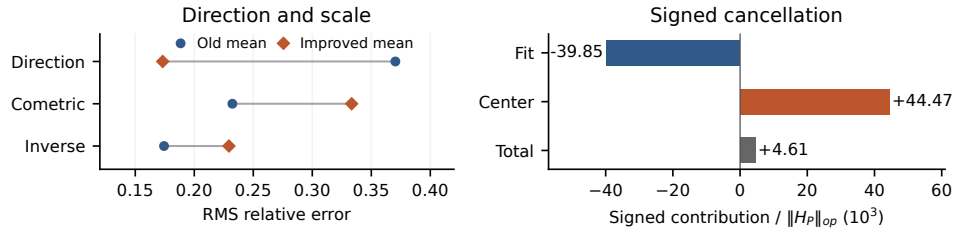

Figure 8: Left: old/improved mean comparison over eight paired cells; all improve direction and worsen both scales. Right: raw improved-mean sample fields, concentrated linear space; signed contributions along each case's common worst-error direction. Finite-noise bias is 4.16×10^{-6}.

is irrelevant; its observed impact depends on the center and fitting regime. However, uniform image01/replication0 worsens in both form spaces, and its concentrated cubic contrast is slightly harmful and unresolved at this numerical precision.

As a post-run numerical diagnostic, compare each paired effect with ten times the sum of the two empirical integration changes. Oracle sample-to-conditional linear improvements at $\kappa = 8$ clear this screen in all 4 cells, but cubic improvements clear it in only 2. This is neither a statistical test nor a quadrature certificate. For New with sample supervision, concentrated linear Δ averages 0.05441, much smaller than the residual 10.6592, yet its ratio to the two-mode reference gap reaches 3.66 in one cell. Thus numerical resolution of form discrepancies is distinct from resolution at the scale of

Table 15: Paired weighted-form errors at grid 256, normalized by the same reference norm in each trial space. R2 retains the predicted leading two eigenvalues. Δ_L, Δ_C are empirical integration changes for the linear/cubic R2 fields, not confidence intervals or rigorous error bounds. The two function spaces are different tests.

Target	κ	Center/input	Linear raw	Linear R2	Cubic R2	Δ_L	Δ_C
Sample	0	Old	21.395	0.80006	0.91642	0.00023821	0.00027985
Sample	0	New	3.0802	0.82612	0.90125	0.00011682	0.00011175
Sample	0	New/old input	2.9396	0.83113	0.89873	0.00014834	0.00017217
Sample	0	Oracle	0.02741	0.014582	0.013613	8.0553e-05	9.0767e-05
Sample	8	Old	27647	62.756	4.1783	0.2148	0.015903
Sample	8	New	4613.5	10.659	1.2995	0.054407	0.0034952
Sample	8	New/old input	4541.9	10.269	1.2502	0.02881	0.0024983
Sample	8	Oracle	45.027	0.24795	0.02303	0.01309	0.0011221
Conditional	0	Old	22.951	0.82741	0.92042	0.00033206	0.00042274
Conditional	0	New	3.1448	0.84306	0.93009	6.3821e-05	4.3571e-05
Conditional	0	New/old input	3.1927	0.84163	0.92825	8.9304e-05	8.5131e-05
Conditional	0	Oracle	0.0072432	0.0050145	0.0047107	2.2534e-05	2.4276e-05
Conditional	8	Old	35626	21.799	2.1286	0.070831	0.0071941
Conditional	8	New	5142.6	7.6557	1.1235	0.022818	0.0013635
Conditional	8	New/old input	5788.6	7.3404	1.1451	0.020974	0.0014997
Conditional	8	Oracle	10.83	0.037301	0.0067697	0.00066798	0.00010253

the reference spectral gap. The original 16-image study's small-gap cases are not resolved by this two-image experiment.

Independent recomputation verifies all 64 checkpoint/source hashes, shared noisy evaluation queries, and scalar summaries from saved arrays. Signed-block reconstruction has relative residual below 1.65×10^{-14}, reference quadrature error below 4.95×10^{-14}, and posterior covariance order changes below 3×10^{-8}. None supplies a population spectral guarantee. The comparisons separate residual centering bias, fitting to B_m, and the direction/scale tradeoff on the same controlled image orbits.

Appendix F locates the remaining error directions and reports a matched input-representation intervention.

F Locating the remaining bias and tangent-scale errors

F.1 Frozen-checkpoint audit

We retain all 64 final models from Appendix E, replay its original evaluation queries, and independently draw 512 new noisy and 512 new clean queries per case. All arms and covariance replications share these queries. No models are trained or selected in this audit. The same two base images and fixed mean checkpoints limit inference; covariance seeds are not independent mean or image replications.

Let $b = (\mu_t - m)/\sigma$, $R = bb^\top$, $B_m = A_t + R$ and $\Delta = C - A_t$. For $R \neq 0$, Frobenius orthogonal projection gives

$$\alpha = \frac{\langle \Delta, R \rangle_F}{\|R\|_F^2}, \qquad \|C - B_m\|_F^2 = (\alpha - 1)^2 \|R\|_F^2 + \|\Delta - \alpha R\|_F^2. \tag{32}$$

This locates error in matrix space. It is not an identifiable learned split $C = \widehat{A} + \widehat{R}$; neither Δ nor its orthogonal residual must be positive semidefinite. We do not interpret α when $R = 0$. On the new noisy queries under conditional supervision, the first term accounts for $99.503\%/99.286\%$ of the pooled own-target MSE with old/improved means. The $\|R\|_F^2$-weighted coefficients are $0.004958/0.007210$. Sample supervision gives $99.510\%/99.307\%$. These are error shares, not fractions of the covariance successfully learned.

Table 16: Independent clean-query audit with conditional supervision. Each entry averages the per-cell RMS Frobenius error normalized by $\sqrt{2}$ over two images, two densities, and two covariance initializations. All eight cells exhibit decreasing target contamination but increasing fitting and total scale errors.

Tangent-scale quantity	Old mean	Improved mean
Total $S_C - I_2$	0.25306	0.34059
Fitting $S_C - S_B$	0.23865	0.32979
Contamination S_R	0.03830	0.02578
Finite-noise $S_A - I_2$	0.00001896	0.00001896

For clean-query tangent frame Q, write $S_C = Q^\top C Q$ and analogously S_A, S_R, S_B. The signed identity

$$S_C - I_2 = \underbrace{S_C - S_B}_{\text{target fitting}} + \underbrace{S_R}_{\text{center contamination}} + \underbrace{S_A - I_2}_{\text{finite-noise geometry}} \tag{33}$$

retains all quadratic cross inner products when squared. Component norms cannot be added as an error identity. The clean target's average largest tangent eigenvalue falls from 1.03791 to 1.02564, whereas the prediction's rises from 1.33391 to 1.45539. Thus the scale deterioration cannot be assigned simply to greater clean target contamination. Fitting error is a separate substantial contribution; this does not identify a unique architectural or optimization cause. Sample supervision has the same direction: fitting error $0.21872 \to 0.32254$.

Noisy and clean statements must remain separate. On noisy queries, tangent bias energy instead increases $0.50322 \to 0.59385$, while normal bias decreases $34.68487 \to 15.01588$. These noisy-query components use the known tangent at the generating clean X, not a tangent inferred at Y. Total mean-error improvement need not improve every directional component or input domain.

All source/checkpoint hashes remain fixed. Scalar summaries recomputed from saved arrays agree exactly; original-query covariance squared errors replay to 3.56×10^{-15}, signed matrix reconstruction to 1.23×10^{-15}, and the relative residual of (32) is below 4.18×10^{-16}. Posterior covariance order-13/17 changes are below 2.98×10^{-8}. These are numerical checks, with no new form integration or spectral certificate supplied by the frozen audit.

F.2 Matched input reparameterization

Motivated by the bias-direction audit, we change only the covariance input from (Y, m_{new}) to $(Y, (Y - m_{\text{new}})/\sigma)$. Given Y and σ, this is invertible in exact arithmetic and adds no oracle information. The actual transformation is in float32, matching the input dtype. The six-channel backbone, rank-16-plus-diagonal output, parameter shapes, initial tensors, improved mean, training pairs, labels, optimizer, schedule and 9,000-update budget are fixed. Equal initial tensors do not mean equal initial functions after the feature transformation; this intervention tests input representation and conditioning, not added expressivity.

We train 16 new models: two images, two densities, two covariance initializations, and sample/conditional supervision. Comparators are the frozen New-arm checkpoints of Appendix E, not newly trained controls. Initial parameter hashes, frozen-mean hashes and first/last training-stream hashes match exactly. Only final EMA checkpoints are used. The numerical conditional labels remain an assisted diagnostic; the sample residual labels require no oracle posterior. Evaluation reuses the original paired queries and adds 512 independently drawn noisy and 512 clean queries with a new seed fixed before training. All case/seed contrasts are retained.

The simple residual representation does not repair own-target fitting within this fixed budget. Sample-supervised error to B_m is $0.987081 \to 0.987107$; conditional error is $0.986958 \to 0.987008$. Conditional noisy error to A_t changes $0.99671 \to 0.99339$, while clean cometric error changes $0.34052 \to 0.33940$. Sample-supervised tangent error worsens $0.17323 \to 0.17794$. On the additional independently drawn noisy queries, the bias-direction shares remain $99.314\%/99.294\%$ for original-input sample/conditional models and $99.314\%/99.299\%$ after the transformation. The bias-direction shares are close across the two input parameterizations on these fixed and fresh query sets. The paired comparison holds the optimizer, training budget and mean estimator fixed.

Table 17: Input intervention on the original paired evaluation queries. Means over both densities and eight case/initialization cells per row. Own B_m and noisy A_t are mean per-query relative Frobenius errors; clean tangent and cometric use the RMS definitions in Appendix E. Both inputs use the same improved mean.

Target	Input	Own B_m	Noisy A_t	Tangent	Cometric
Sample	Original	0.98708	0.98268	0.17323	0.33335
Sample	Residual	0.98711	0.98322	0.17794	0.33104
Conditional	Original	0.98696	0.99671	0.16726	0.34052
Conditional	Residual	0.98701	0.99339	0.16732	0.33940

Table 18: Input intervention at $\kappa = 8$ in the same fixed trial spaces. R2 retains the two leading predicted eigenvalues. Δ is the larger normalized matrix change between grids 128/256 and regular/shifted 256; it is an empirical sensitivity, not a certified error bound.

Target	Input	Linear R2	Cubic R2	Δ_L	Δ_C
Sample	Original	10.65915	1.29947	0.05441	0.00350
Sample	Residual	11.47262	1.37704	0.01855	0.00186
Conditional	Original	7.65572	1.12347	0.02282	0.00136
Conditional	Residual	8.01384	1.15587	0.02992	0.00374

All 16 new fields are re-integrated on grids 64, 128, 256 and shifted 256, without changing analytic mass, reference stiffness, functions or whitening. Concentrated linear rank-two mean errors worsen from 10.6592 to 11.4726 for sample supervision and 7.6557 to 8.0138 for conditional supervision. Each comparison has one improving and three worsening case/seed contrasts. Uniform linear means also worsen, $0.82612 \to 0.84523$ and $0.84306 \to 0.84973$. All raw/rank-two/unit fields, signed blocks and individual contrasts remain in the artifact.

The post-run tenfold paired-effect screen from Appendix E resolves all four concentrated sample linear contrasts, but only two conditional linear and two conditional cubic contrasts. It is neither a significance test nor a quadrature certificate. Residual form discrepancies exceed the grid changes, while the new concentrated linear sensitivity divided by the two-mode reference gap reaches 0.762 for sample and 1.762 for conditional supervision. Thus there is no uniform compressed spectral guarantee, and this intervention does not resolve the original sixteen-image study's other small-gap cases.

All 16 checkpoint, initialization, mean and training-stream hashes pass. Raw scalar recomputation agrees exactly; signed-form relative reconstruction residual is below 1.56×10^{-14} and reference quadrature error below 4.95×10^{-14}. The minimum observed clean restricted eigenvalue is 1.08782, with no singular inverse removed. The result establishes the insufficiency of this particular feature transformation in the tested setting, not the impossibility of learning the missing component or recovering a manifold operator with another estimator.

G Signed geometric-form attribution on frozen models

G.1 Scope and reference calculation

This post-training analysis reuses every saved learned-field integration matrix from the 64-model paired study and the 16-model input intervention. It adds no training. Its units are two base images, densities $\kappa = 0, 8$, two covariance initializations, two label types, and five input/center arms. Raw and rank-two fields on both 12-dimensional dictionaries give 320 reported comparisons. References are shared between covariance initializations and arms, not independent replications. Attribution does not extend to the other 14 images of the initial study.

The fixed protocol integrates A_t, $\mathcal{F}_2(A_t)$, and $R(x) = (\mu_t(x) - m(x))(\mu_t(x) - m(x))^\top / t$ for both frozen means. It evaluates clean queries on grids 64, 128, 256, and half-cell-shifted 256, using analytic mass, stiffness, and whitening throughout. Numerical posterior orders 9, 13, and 17 are evaluated on every rule, with order 17 primary. The learned covariance is not reevaluated: its stored

Table 19: Numerical posterior geometric bias, mean over the two images. All entries are norms relative to $\|H_P\|_{\mathrm{op}}$. The order-change column compares 13 and 17 on the full component. The same posterior reference is shared by all learned arms and covariance initializations.

κ	Space	Field	Full	Tangent	Cross	Normal	Order change
0	linear	raw	2.21×10^{-5}	2.34×10^{-5}	4.43×10^{-5}	1.30×10^{-6}	6.83×10^{-14}
0	linear	rank2	2.08×10^{-5}	2.34×10^{-5}	4.43×10^{-5}	5.47×10^{-10}	5.87×10^{-14}
0	cubic	raw	2.73×10^{-5}	2.34×10^{-5}	4.43×10^{-5}	6.51×10^{-6}	5.09×10^{-14}
0	cubic	rank2	2.08×10^{-5}	2.34×10^{-5}	4.43×10^{-5}	5.47×10^{-10}	4.60×10^{-14}
8	linear	raw	9.76×10^{-6}	2.70×10^{-5}	2.92×10^{-5}	6.45×10^{-6}	4.58×10^{-13}
8	linear	rank2	6.98×10^{-6}	2.70×10^{-5}	2.92×10^{-5}	3.19×10^{-10}	4.58×10^{-13}
8	cubic	raw	1.67×10^{-5}	2.81×10^{-5}	3.53×10^{-5}	9.90×10^{-6}	1.12×10^{-13}
8	cubic	rank2	1.15×10^{-5}	2.81×10^{-5}	3.53×10^{-5}	4.18×10^{-10}	1.12×10^{-13}

matrices use the same queries and rules. Frozen-mean inference uses batch 64, matching those stored calculations. Two nodes and 16 distinct GPUs compute the new references; reference code, mean checkpoints, protocol, and saved input matrices are hashed before execution. The numerical posterior uses the known orbit model only in the reference and supervision calculations, never as an inference input to the covariance networks.

For reference covariance factors $A_t = VV^\top$, a two-column subspace iteration and Rayleigh–Ritz extraction computes $\mathcal{F}_2(A_t)$. The analytic tangent frame initializes this reference eigensolver, not a learned estimator. Six iterations give maximum relative residual 8.37×10^{-15}; a PSD trace-remainder check separates the selected pair from the remaining eigenvalues at every query. Dense independent algebra controls compare the extraction with a full eigendecomposition. All numerical posterior mode iterations have final step below 4.12×10^{-16}; this is a solver check, not proof that numerical posterior error vanishes. Order changes probe the same local quadrature family. They do not themselves exclude missing remote posterior mass; the earlier multi-start checks and their limitations are described in Appendix C.2.

G.2 Matrices, directions, and numerical sensitivity

For raw fields we store $\mathcal{T}_{\mathcal{Q}}(C - B_m)$, $\mathcal{T}_{\mathcal{Q}}(R)$, and $\mathcal{T}_{\mathcal{Q}}(A_t) - H_P$, where $\mathcal{Q}$ denotes the chosen quadrature and H_P remains analytic. This last expression includes the reference quadrature residual $R_{\mathcal{Q}} = \mathcal{Q}[JPJ^\top] - H_P$; its relative norm is below 4.84×10^{-14}. Each term has tangent, symmetric cross, and normal blocks. The posterior-minus-reference tangent block includes $R_{\mathcal{Q}}$. Thus the full-field tangent contribution is $D_{T,\mathcal{Q}} = \mathcal{Q}[JP(C-P)PJ^\top] + R_{\mathcal{Q}}$, equivalently $\mathcal{Q}[JPCPJ^\top] - H_P$. This is the convention used in Equation (13) and Appendix H. For rank-two fields we store $\mathcal{T}_{\mathcal{Q}}(\mathcal{F}_2(C) - \mathcal{F}_2(A_t))$ and $\mathcal{T}_{\mathcal{Q}}(\mathcal{F}_2(A_t)) - H_P$. We also retain the posterior content removed by filtering, $\mathcal{T}_{\mathcal{Q}}(A_t - \mathcal{F}_2(A_t))$. The quadrature symbol $\mathcal{Q}$ is distinct from the tangent frame Q.

For a fixed model, processing rule, and dictionary, the primary signed comparison uses the largest-absolute-eigenvalue eigenvector of its total error at grid 256. This direction stays fixed across all component, posterior order, and integration comparisons. Archives also contain quadratic forms on every analytic H_P eigenvector and every fixed whitened coordinate axis. A numerical basis in a degenerate eigenspace is not a unique geometric observable. We never select a different worst direction for each component. Averaged signed quantities average per-case quadratic forms, not matrices expressed in unrelated per-orbit coordinates.

Define the reported spatial sensitivity for a matrix component D by

$$\Delta_Q(D) = \frac{\max\{\|D_{256} - D_{128}\|_{\mathrm{op}}, \|D_{256} - D_{256,\mathrm{shift}}\|_{\mathrm{op}}\}}{\|H_P\|_{\mathrm{op}}}.$$

The archive separately retains $64 \to 128$ changes. The maximum order $13 \to 17$ component change is 3.68×10^{-9}, and the maximum $9 \to 17$ change is 4.02×10^{-9}. Exact signed reconstruction has maximum relative residual 8.04×10^{-16}. The latter is an algebra check only: an inaccurate posterior reference can cancel between fitting and finite-noise terms. Order changes and spatial changes test different issues, and neither is a rigorous continuum error certificate.

Table 20: Raw learned-center fields, $\kappa = 8$, linear functions. Components are signed quadratic forms in each total error's common worst direction, then averaged over four paired units. Last columns report component spatial sensitivities, not confidence limits. All values use $\|H_P\|_{\mathrm{op}}$ units.

Center	Labels	Fit	Centering	Finite	Total	Δ_Q fit	Δ_Q center
old	sample	7913.63	19733.80	4.27×10^{-6}	27647.43	6.93	5.60
old	conditional	16239.70	19386.24	4.23×10^{-6}	35625.94	5.03	5.60
new	sample	-39854.61	44468.10	4.16×10^{-6}	4613.49	10.39	10.54
new	conditional	-39252.83	44395.47	4.20×10^{-6}	5142.64	10.42	10.54

The raw finite-noise full-form errors range from 9.76×10^{-6} to 2.74×10^{-5} across the four density/dictionary settings (Table 19). For concentrated linear functions, filtering removes posterior content of norm 6.45×10^{-6}, while it removes much larger extra energy from the learned oracle-centered fields. For conditional rank-two models, the remaining learning-cross norm exceeds the learning-normal norm on average. These statements concern the specified spaces at $\sigma = .02$, not all finite-noise scales or ambient extensions.

For improved means, fitting and centering strongly oppose one another (Table 20); conditional supervision does not remove that phenomenon. On these clean-query weighted forms, the centering term can be large despite reduced noisy-query mean bias. Changed distribution and whitened gradient weighting prevent replacing one quantity by the other. Large negative fitting contribution means underfitting relative to B_m in that direction; it is not itself proof of better or worse geometry. The raw and filtered attributions answer different questions and cannot be interchanged through a linear decomposition of the filter.

G.3 Complete grouped results

The following tables retain every arm and both label types. "Fit" denotes own-target fitting for raw fields, and processed learning error for rank-two fields; centering is absorbed into processed learning error for the latter, not separately decomposed. Dashes in those rows mark this distinction; they do not assert that centering has vanished. The underlying calculations retain individual signed matrices, cross/normal blocks, common-direction evaluations, and separate posterior-order and spatial-grid comparisons. These aggregate tables average two images and two covariance initializations, not four independent images.

Table 21: All arms, uniform density, linear functions. Ratios relative to $\|H_P\|_{\mathrm{op}}$; component norms do not sum.

Arm	Labels	Field	Total	Fit	Center	Δ_Q fit	Δ_Q center
old	samp	raw	21.40	14.29	14.01	0.00152	0.00136
old	samp	rank2	0.80006	0.80004	—	0.00024	—
old	cond	raw	22.95	17.42	14.01	0.00190	0.00136
old	cond	rank2	0.82741	0.82739	—	0.00033	—
new	samp	raw	3.08	13.40	13.65	0.00566	0.00567
new	samp	rank2	0.82612	0.82610	—	0.00012	—
new	cond	raw	3.14	13.07	13.65	0.00566	0.00567
new	cond	rank2	0.84306	0.84304	—	6.38×10^{-5}	—
target_only	samp	raw	2.94	13.33	13.65	0.00569	0.00567
target_only	samp	rank2	0.83113	0.83111	—	0.00015	—
target_only	cond	raw	3.19	13.06	13.65	0.00567	0.00567
target_only	cond	rank2	0.84163	0.84161	—	8.93×10^{-5}	—
oracle	samp	raw	0.02741	0.02740	0	7.84×10^{-5}	0
oracle	samp	rank2	0.01458	0.01460	—	8.06×10^{-5}	—
oracle	cond	raw	0.00724	0.00724	0	2.27×10^{-5}	0
oracle	cond	rank2	0.00501	0.00504	—	2.25×10^{-5}	—
input_residual	samp	raw	3.42	13.23	13.65	0.00579	0.00567
input_residual	samp	rank2	0.84523	0.84521	—	9.39×10^{-5}	—
input_residual	cond	raw	3.40	13.34	13.65	0.00574	0.00567
input_residual	cond	rank2	0.84973	0.84971	—	5.25×10^{-5}	—

Table 22: All arms, uniform density, cubic functions; same conventions.

Arm	Labels	Field	Total	Fit	Center	Δ_Q fit	Δ_Q center
old	samp	raw	68.71	45.03	53.17	0.02140	0.02091
old	samp	rank2	0.91642	0.91640	—	0.00028	—
old	cond	raw	73.13	50.19	53.17	0.01958	0.02091
old	cond	rank2	0.92042	0.92040	—	0.00042	—
new	samp	raw	8.70	46.74	50.20	0.02437	0.02462
new	samp	rank2	0.90125	0.90123	—	0.00011	—
new	cond	raw	9.04	46.70	50.20	0.02457	0.02462
new	cond	rank2	0.93009	0.93007	—	4.36×10^{-5}	—
target_only	samp	raw	8.28	46.78	50.20	0.02447	0.02462
target_only	samp	rank2	0.89873	0.89871	—	0.00017	—
target_only	cond	raw	9.15	46.87	50.20	0.02458	0.02462
target_only	cond	rank2	0.92825	0.92822	—	8.51×10^{-5}	—
oracle	samp	raw	0.08748	0.08746	0	8.57×10^{-5}	0
oracle	samp	rank2	0.01361	0.01363	—	9.08×10^{-5}	—
oracle	cond	raw	0.02605	0.02604	0	2.65×10^{-5}	0
oracle	cond	rank2	0.00471	0.00473	—	2.43×10^{-5}	—
input_residual	samp	raw	9.91	46.21	50.20	0.02498	0.02462
input_residual	samp	rank2	0.91176	0.91174	—	9.11×10^{-5}	—
input_residual	cond	raw	10.14	46.76	50.20	0.02502	0.02462
input_residual	cond	rank2	0.94262	0.94260	—	6.12×10^{-5}	—

Table 23: All arms, concentrated density, linear functions; same conventions.

Arm	Labels	Field	Total	Fit	Center	Δ_Q fit	Δ_Q center
old	samp	raw	27647.43	9990.74	20886.60	6.93	5.60
old	samp	rank2	62.76	62.76	—	0.21480	—
old	cond	raw	35625.94	17729.55	20886.60	5.03	5.60
old	cond	rank2	21.80	21.80	—	0.07083	—
new	samp	raw	4613.49	41179.05	45489.92	10.39	10.54
new	samp	rank2	10.66	10.66	—	0.05441	—
new	cond	raw	5142.64	40862.03	45489.92	10.42	10.54
new	cond	rank2	7.66	7.66	—	0.02282	—
target_only	samp	raw	4541.91	41271.84	45489.92	10.19	10.54
target_only	samp	rank2	10.27	10.27	—	0.02881	—
target_only	cond	raw	5788.60	40222.05	45489.92	10.66	10.54
target_only	cond	rank2	7.34	7.34	—	0.02097	—
oracle	samp	raw	45.03	45.03	0	0.02577	0
oracle	samp	rank2	0.24795	0.24795	—	0.01309	—
oracle	cond	raw	10.83	10.83	0	0.00545	0
oracle	cond	rank2	0.03730	0.03730	—	0.00067	—
input_residual	samp	raw	4534.80	41316.39	45489.92	9.92	10.54
input_residual	samp	rank2	11.47	11.47	—	0.01855	—
input_residual	cond	raw	4490.75	41548.12	45489.92	11.16	10.54
input_residual	cond	rank2	8.01	8.01	—	0.02992	—

Table 24: All arms, concentrated density, cubic functions; same conventions.

Arm	Labels	Field	Total	Fit	Center	Δ_Q fit	Δ_Q center
old	samp	raw	1590.05	612.77	1385.22	0.60138	0.78191
old	samp	rank2	4.18	4.18	—	0.01590	—
old	cond	raw	2096.37	1031.52	1385.22	0.67023	0.78191
old	cond	rank2	2.13	2.13	—	0.00719	—
new	samp	raw	256.18	2787.27	3019.76	0.81147	0.81545
new	samp	rank2	1.30	1.30	—	0.00350	—
new	cond	raw	292.01	2772.07	3019.76	0.80896	0.81545
new	cond	rank2	1.12	1.12	—	0.00136	—
target_only	samp	raw	253.31	2793.45	3019.76	0.80057	0.81545
target_only	samp	rank2	1.25	1.25	—	0.00250	—
target_only	cond	raw	332.10	2732.12	3019.76	0.82370	0.81545
target_only	cond	rank2	1.15	1.15	—	0.00150	—
oracle	samp	raw	3.09	3.09	0	0.00234	0
oracle	samp	rank2	0.02303	0.02303	—	0.00112	—
oracle	cond	raw	0.72562	0.72561	0	0.00045	0
oracle	cond	rank2	0.00677	0.00677	—	0.00010	—
input_residual	samp	raw	259.93	2792.18	3019.76	0.77358	0.81545
input_residual	samp	rank2	1.38	1.38	—	0.00186	—
input_residual	cond	raw	254.08	2806.95	3019.76	0.85712	0.81545
input_residual	cond	rank2	1.16	1.16	—	0.00374	—

H Relative-Energy and Ambient-Extension Diagnostics

These bounded analyses reuse the saved small matrices from all 320 attribution comparisons. They add no training, new observations, selected checkpoints or optimized extension parameter. Their protocol was written after review and before the computations, not before the underlying training experiment. The primary rule is grid 256; grid 64, grid 128 and half-cell-shifted grid 256 are retained. All analytic masses, references and original model fields stay fixed.

H.1 Two different accuracy statements

For $f = v^\top \Psi$, $\|f\|_{L^2(\nu)} = \|v\|$, and the symmetric error $D = \widehat{H}_C - H_P$ satisfies

$$|v^\top D v| \le \|D\|_{\mathrm{op}} \|v\|^2 = e_C \|H_P\|_{\mathrm{op}} \|f\|^2_{L^2(\nu)}.$$

When $H_P \succ 0$, substitute $u = H_P^{1/2} v$ to obtain

$$\sup_{v \neq 0} \frac{|v^\top D v|}{v^\top H_P v} = \|H_P^{-1/2} D H_P^{-1/2}\|_{\mathrm{op}} = e_{\mathrm{energy}}.$$

This is a deterministic statement about the computed finite matrices; their integration uncertainty remains separate. For a positive-semidefinite reference with a nullspace, the same statement applies on the positive-energy subspace only, not to arbitrary functions with a null component.

Gap-free relative eigenvalue control. Writing $\varepsilon = e_{\mathrm{energy}}$ gives $-\varepsilon I \preceq H_P^{-1/2} D H_P^{-1/2} \preceq \varepsilon I$. Congruence by $H_P^{1/2}$ proves Equation (14). When $0 \le \varepsilon < 1$, positive scalar multiplication preserves eigenvalue order; monotonicity under the Loewner order, by the min–max principle, yields Equation (15). The reference and estimate need not commute, and repeated reference eigenvalues are allowed. For a singular reference the same argument applies to both matrices compressed onto its positive-energy subspace, not automatically to their full spectra. Across the four concentrated oracle-conditional linear rank-two comparisons, $\max e_{\mathrm{energy}} = 0.14638418$. Thus every ordered eigenvalue of each of these computed matrices has relative error at most approximately 14.64%. This bound does not include continuum integration or statistical uncertainty.

We add no ridge. The numerical positive threshold is $512\epsilon_{\mathrm{mach}}\|H_P\|_{\mathrm{op}}$. All 12 eigenvalues exceed it in every case. At concentrated density the reference condition numbers range from 16.72 to 17.16 (linear) and 23.20 to 24.22 (cubic).

For each comparison, transform the grid-128 and shifted-grid differences by the same $H_P^{-1/2}$ and report the largest norm as Δ_{energy}. Reusing an additive-normalization grid sensitivity would be incorrect. Neither this sensitivity nor a small difference between two rules is a rigorous integration bound. The tables report means of four per-model worst-case relative errors, not one guarantee for their average field or a uniform guarantee across cases. For the concentrated oracle conditional linear models, per-case relative-energy errors span 0.01380–0.14638; the mean is 0.07021 and mean sensitivity 0.001412. The corresponding cubic mean is 0.01723 with sensitivity 0.000320.

H.2 Changing extensions while holding the intrinsic problem fixed

For a smooth embedded orbit and its nearest-point projection π in a tubular neighborhood, $D\pi = P$ on the orbit. The family $\Phi_s = \Phi \circ \pi + s(\Phi - \Phi \circ \pi)$ therefore has identical restrictions, centering, mass and intrinsic gradients for every s, but $D(M^{-1/2}\Phi_s) = JP + sJN$. Population integration gives $H_C^{(s)} - H_P = D_T + sD_X + s^2 D_N$, with the blocks of Equation (4). For a computed form, use the explicit quadrature convention

$$\begin{aligned} D_{T,\mathcal{Q}} &= \mathcal{Q}[JPCPJ^\top] - H_P, \\ D_{X,\mathcal{Q}} &= \mathcal{Q}[J(PCN + NCP)J^\top], \\ D_{N,\mathcal{Q}} &= \mathcal{Q}[JNCNJ^\top]. \end{aligned} \tag{34}$$

Expanding $J_s C J_s^\top$ proves Equation (13) exactly for that rule. The symmetric cross block includes both transposes. In particular, the tangent contribution includes the analytic reference subtraction. If

an archive instead stores $\widetilde{D}_{T,\mathcal{Q}} = \mathcal{Q}[JP(C-P)PJ^\top]$, convert it using

$$D_{T,\mathcal{Q}} = \widetilde{D}_{T,\mathcal{Q}} + R_{\mathcal{Q}}, \qquad R_{\mathcal{Q}} = \mathcal{Q}[JPJ^\top] - H_P.$$

Our diagnostic matrices apply this correction, consistently with Appendix G; its small size does not justify omitting it from an exact identity. Only on-manifold first derivatives are needed. We do not construct or query a numerical nearest-point solver. We use closest-point extensions (März & Macdonald, 2012) to vary the ambient normal derivative while holding the intrinsic problem fixed.

Every raw and rank-two case is evaluated on $s = -1, -.975, \ldots, 2$, without selecting an optimal s. Figure 5 displays $0 \le s \le 1.5$; each curve is computed directly from the same signed component matrices. At $s = 1$ the original error is reconstructed. At $s = 0$, the form uses the predicted tangent block, retaining its scale; it does not replace C by P. In the concentrated oracle linear comparison, sample-supervised mean errors at $s = 0, 1, 2$ are $0.008485, 0.24795, 0.97925$; conditional values are $0.000357, 0.03730, 0.08851$. These means average per-case norms in their own whitened bases, never matrices in different bases. The different curve shapes reflect the remaining linear cross term and quadratic normal term; individual signed contributions can also cancel.

Numerical consistency checks. All 135 source files were hashed before calculation. An independent dense synthetic field reproduces the extension polynomial to 6.16×10^{-16} relative error. Direct generalized Rayleigh quotients at the extreme energy-whitened eigenvectors agree to 3.78×10^{-13} relative error. The saved block sums and original full errors agree within 2.14×10^{-10} relative to $\|H_P\|_{\mathrm{op}}$, including the large raw learned-center cases. These are algebra checks, not additional quadrature certificates. The following tables summarize per-case errors and sensitivities; the common reference and positive-energy restriction are unchanged throughout.

Table 25: Relative-energy and extension diagnostics, uniform density, linear space. Four-unit means; raw and rank-two cases remain separate. $e_C^{(0)}$ uses the known-tangent extension, and Δ_{energy} measures integration sensitivity.

Arm	Labels	Field	e_C	e_{energy}	Δ_{energy}	$e_C^{(0)}$
old	samp	raw	21.395	97.228	0.0037749	0.33765
old	samp	rank2	0.80006	1.0737	0.00057545	0.30229
old	cond	raw	22.951	103.4	0.0029988	0.3468
old	cond	rank2	0.82741	0.99175	0.0019881	0.31455
new	samp	raw	3.0802	14.738	0.00037072	0.46519
new	samp	rank2	0.82612	0.91525	0.00024053	0.45891
new	cond	raw	3.1448	14.838	0.00014358	0.4761
new	cond	rank2	0.84306	0.90818	8.9847e-05	0.46999
target_only	samp	raw	2.9396	14.127	0.00056748	0.4657
target_only	samp	rank2	0.83113	0.92084	0.00030888	0.45935
target_only	cond	raw	3.1927	14.962	0.00014547	0.47685
target_only	cond	rank2	0.84163	0.90567	0.00021786	0.47069
oracle	samp	raw	0.02741	0.091457	0.0001415	0.0094758
oracle	samp	rank2	0.014582	0.019554	0.00012237	0.0094758
oracle	cond	raw	0.0072432	0.023458	7.0533e-05	0.00052189
oracle	cond	rank2	0.0050145	0.0089662	6.7806e-05	0.00052189
input_residual	samp	raw	3.4209	16.174	0.00061304	0.46402
input_residual	samp	rank2	0.84523	0.9144	0.00030638	0.45782
input_residual	cond	raw	3.4042	16.156	0.00035641	0.47421
input_residual	cond	rank2	0.84973	0.90375	0.00019813	0.4683

Table 26: Relative-energy and extension diagnostics, uniform density, cubic space; definitions as above.

Arm	Labels	Field	e_C	e_{energy}	Δ_{energy}	$e_C^{(0)}$
old	samp	raw	68.712	312.45	0.014975	0.33805
old	samp	rank2	0.91642	1.3929	0.00075075	0.3034
old	cond	raw	73.131	330.72	0.012108	0.34648
old	cond	rank2	0.92042	1.2558	0.0032148	0.31463
new	samp	raw	8.6978	40.971	0.0016935	0.46409
new	samp	rank2	0.90125	1.0118	0.00026002	0.458
new	cond	raw	9.0433	40.486	0.00069095	0.47437
new	cond	rank2	0.93009	1.0241	9.3102e-05	0.46854
target_only	samp	raw	8.2819	39.409	0.0016461	0.46454
target_only	samp	rank2	0.89873	1.0103	0.0002794	0.45839
target_only	cond	raw	9.1501	40.751	0.00043146	0.47499
target_only	cond	rank2	0.92825	1.0203	0.00025378	0.4691
oracle	samp	raw	0.087483	0.25556	0.0002368	0.0094069
oracle	samp	rank2	0.013613	0.018641	0.00015005	0.0094069
oracle	cond	raw	0.026053	0.10172	7.6963e-05	0.00051847
oracle	cond	rank2	0.0047107	0.0089683	6.8454e-05	0.00051847
input_residual	samp	raw	9.9144	45.21	0.0014657	0.46297
input_residual	samp	rank2	0.91176	1.0111	0.00028459	0.45695
input_residual	cond	raw	10.143	45.972	0.0016512	0.47243
input_residual	cond	rank2	0.94262	1.0225	0.00022103	0.46678

Table 27: Relative-energy and extension diagnostics, concentrated density, linear space; definitions as above.

Arm	Labels	Field	e_C	e_{energy}	Δ_{energy}	$e_C^{(0)}$
old	samp	raw	27647	38826	7.8263	0.26654
old	samp	rank2	62.756	97.547	0.30363	0.24386
old	cond	raw	35626	49134	2.3742	0.31802
old	cond	rank2	21.799	40.915	0.13807	0.28982
new	samp	raw	4613.5	6738.3	0.56467	0.42159
new	samp	rank2	10.659	18.036	0.07941	0.41709
new	cond	raw	5142.6	7355.6	0.41332	0.43133
new	cond	rank2	7.6557	12.863	0.034421	0.42621
target_only	samp	raw	4541.9	6619.2	0.63666	0.42233
target_only	samp	rank2	10.269	17.494	0.046907	0.4178
target_only	cond	raw	5788.6	8208.6	0.60977	0.43194
target_only	cond	rank2	7.3404	12.412	0.029028	0.42678
oracle	samp	raw	45.027	59.482	0.033383	0.0084847
oracle	samp	rank2	0.24795	0.35747	0.017804	0.0084847
oracle	cond	raw	10.83	14.403	0.0076243	0.0003569
oracle	cond	rank2	0.037301	0.070208	0.001412	0.0003569
input_residual	samp	raw	4534.8	6619.1	1.1647	0.41634
input_residual	samp	rank2	11.473	19.443	0.027714	0.41187
input_residual	cond	raw	4490.7	6478	1.0045	0.42904
input_residual	cond	rank2	8.0138	13.419	0.051414	0.42403

Table 28: Relative-energy and extension diagnostics, concentrated density, cubic space; definitions as above.

Arm	Labels	Field	e_C	e_{energy}	Δ_{energy}	$e_C^{(0)}$
old	samp	raw	1590.1	3685.5	0.71103	0.26762
old	samp	rank2	4.1783	9.6252	0.030786	0.24475
old	cond	raw	2096.4	4562.5	0.28818	0.31854
old	cond	rank2	2.1286	5.7644	0.020201	0.29051
new	samp	raw	256.18	611.87	0.059745	0.4226
new	samp	rank2	1.2995	2.7331	0.0078891	0.41803
new	cond	raw	292.01	657.69	0.046809	0.43226
new	cond	rank2	1.1235	2.2343	0.0033586	0.42709
target_only	samp	raw	253.31	602.81	0.053896	0.42366
target_only	samp	rank2	1.2502	2.6593	0.0059468	0.41904
target_only	cond	raw	332.1	729.15	0.052352	0.43284
target_only	cond	rank2	1.1451	2.2264	0.0029185	0.42763
oracle	samp	raw	3.0906	5.4111	0.003734	0.0081649
oracle	samp	rank2	0.02303	0.04743	0.0019963	0.0081649
oracle	cond	raw	0.72562	1.2605	0.00082416	0.00029731
oracle	cond	rank2	0.0067697	0.017235	0.00031992	0.00029731
input_residual	samp	raw	259.93	604.07	0.10298	0.41763
input_residual	samp	rank2	1.377	2.8771	0.0037772	0.41309
input_residual	cond	raw	254.08	597.5	0.094871	0.43018
input_residual	cond	rank2	1.1559	2.2644	0.007411	0.42512

I SPATIAL MAPS OF LOCAL ERROR AND WEIGHTED ENERGY DISCREPANCY

This follow-up evaluates the frozen oracle-centered checkpoints on spatial grids; it does not retrain a model, change a target, or add base images. The cohort is exactly the two images and two covariance initializations at $\kappa = 8$, $\sigma = .02$. Both sample and conditional supervision are retained, with raw and rank-two fields kept separate. Image 0, initialization 0 is the main example, chosen by index before spatial evaluation.

I.1 ONE DECLARED FUNCTION FOR BOTH SUPERVISION TYPES

For each pair, let v be a unit eigenvector associated with the largest-absolute eigenvalue of the saved grid-256 sample/rank-two error matrix $\widehat{H}_C - H_P$. Its sign is fixed by making its largest-magnitude coordinate positive. We hold v, the original linear dictionary and analytic mass whitening fixed across both models, both field processings and all four quadrature rules. Consequently $f = v^\top \Psi$ has unit $L^2(\nu)$ norm and the same intrinsic restriction in every comparison within a pair. This selection is a post-result diagnostic of the baseline's worst direction, not a randomly chosen function or an independent test. For the sample/rank-two field its absolute energy discrepancy equals the reported form norm by construction; for the other fields it need not equal their worst-direction discrepancy. We never average coefficient vectors or matrices from different image bases.

Write $g = D\Psi^\top v$, $g_T = Pg$ and $g_N = (I - P)g$. The linear probe has constant ambient g, while its tangent and normal parts vary around the orbit. With $Z = \|H_P\|_{\mathrm{op}}$, define the following signed densities with respect to the parameter area $d\theta_1 d\theta_2$:

$$\begin{aligned} h_T(\theta) &= q_8(\theta) g_T^\top (C - P) g_T / Z, \\ h_X(\theta) &= 2 q_8(\theta) g_T^\top C g_N / Z, \\ h_N(\theta) &= q_8(\theta) g_N^\top C g_N / Z, \\ h(\theta) &= q_8(\theta) g^\top (C - P) g / Z = h_T + h_X + h_N. \end{aligned} \tag{35}$$

Thus $\int h_b \, d\theta = v^\top D_b v / Z$ for the signed matrices in Equation (4). The factor two includes both cross transposes. These blocks decompose $C - P$, not $C - A_t$; this spatial diagnostic is distinct from the posterior attribution in Appendix G. Normal density is nonnegative because C is PSD. Tangent and cross densities can have either sign, and contributions can cancel both between blocks and between locations. We therefore report both $\int h$ and $\int |h|$.

The local map is $\|C - P\|_F / \sqrt{2}$, using the stated raw or filtered field. It is not the leading-projector error $e_{\tan}$. All figures display the entire periodic parameter square. The dashed contour encloses the highest-density grid cells carrying at least 95% of quadrature mass; no tail cells are dropped from any integral. Density uses a logarithmic color scale with a display floor 10^{-12}, local errors a common linear scale, and signed energies a symmetric-logarithmic scale. These display transforms do not change the stored values or numerical integration.

I.2 NUMERICAL RECONSTRUCTION AND SPATIAL CANCELLATION

We use grids 64, 128, 256 and half-cell-shifted 256, with the same analytic density, mass and reference stiffness. Float32 network inference is followed by float64 geometry and eigenspace calculations; TF32 is disabled. The same final-EMA checkpoints, input means and rank-two eigensolver settings as the paired experiment are retained. Each scalar integral is compared with the corresponding previously saved signed matrix, not only with the sum of its newly computed spatial blocks. Probe normalization and reference energy are checked separately. The sensitivity Δ_v below is the larger absolute change from grid 128 to 256 or from grid 256 to shifted 256, in normalization Z. It is an empirical sensitivity for this fixed function, not an operator norm, confidence interval or rigorous quadrature bound.

All 256 scalar block comparisons with previously saved matrices agree to 1.42×10^{-13} in normalization Z; the largest pointwise additive residual is 9.53×10^{-12}. The largest relative leading-space residual is 9.73×10^{-15} and the largest probe-mass error is 1.15×10^{-12}. These are reconstruction and solver checks, not additional integration bounds.

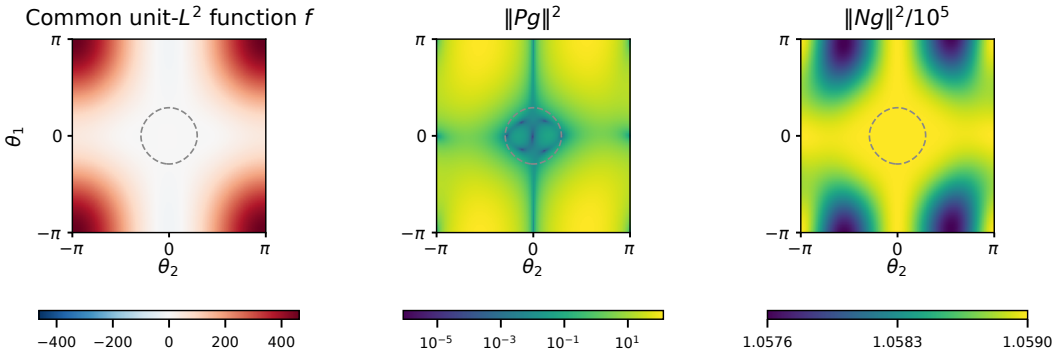

Figure 9: The common function and its ambient-gradient decomposition for the main spatial example. The function is mass-normalized; the two gradient maps are unweighted squared norms, not error contributions. Dashed contours mark the same density region as Figure 4.

Table 29: Rank-two spatial diagnostics. Unit is image/initialization. Local RMS is density-weighted full-field Frobenius error divided by $\sqrt{2}$. T, X, N and Total are signed integrals in Equation (35); $\int |h|$ retains spatial cancellation and Δ_v measures quadrature sensitivity. Each pair uses its own common baseline-worst function.

Unit	Labels	Local RMS	T	X	N	Total	$\int \lvert h \rvert$	Δ_v
0/0	Sample	0.03119	-0.008857	0.02769	0.3913	0.4102	0.5724	0.002969
0/0	Cond.	0.006973	0.000166	0.01529	0.02326	0.03873	0.124	0.000204
0/1	Sample	0.0273	0.005814	0.006057	0.5379	0.5498	0.7284	0.03854
0/1	Cond.	0.007662	1.275e-05	0.002505	0.01946	0.02198	0.1058	0.001233
1/0	Sample	0.0257	-0.0002305	0.003139	0.01595	0.01885	0.08142	0.0001437
1/0	Cond.	0.00552	0.0002065	-0.005346	0.000718	-0.004421	0.01612	1.481e-05
1/1	Sample	0.02543	-0.003448	-0.01225	0.002705	-0.013	0.03321	2.197e-05
1/1	Cond.	0.005651	0.0002054	-0.002482	7.976e-05	-0.002197	0.006767	5.667e-05

For image 0, initialization 0, rank-two local RMS falls $0.03119 \to 0.00697$. The signed energy integral falls $0.41017 \to 0.03873$, while its absolute integral falls $0.57245 \to 0.12395$. Consequently $1 - |\int h| / \int |h|$ changes from 28.35% to 68.76%: there is a genuine decrease in spatial absolute discrepancy, and the remaining signed integral also benefits from stronger cancellation. This ratio is a description of this one function, not a causal percentage of recovery. The highest-density region has mass 0.95023 and covers 6.52% of the parameter square. It contains 83.36% and 79.65% of the sample and conditional absolute discrepancies, respectively. Thus low-mass tails cannot simply be omitted from the function-energy comparison.

The conditional rank-two signed tangent, cross and normal contributions are 0.000166, 0.01529 and 0.02326. Their sum is the reported total; these are common-direction quadratic forms, not the block operator norms in Table 2. Raw and processed fields are both retained in Tables 30 and 29 and Figure 10. Across the four pairs, 4/4 reductions in absolute signed error exceed ten times the sum of both fields' measured scalar sensitivities. This descriptive check applies only to the declared functions; the pre-existing matrix-norm comparisons remain the form-level evidence. Figure 11 retains every paired outcome.

Table 30: Raw-field spatial diagnostics on exactly the same functions as Table 29. The probe is selected from the sample/rank-two field, not reselected to maximize any raw-field discrepancy.

Unit	Labels	Local RMS	T	X	N	Total	$\int \lvert h \rvert$	Δ_v
0/0	Sample	0.03182	-0.008857	0.02764	83.29	83.31	83.31	0.02375
0/0	Cond.	0.007162	0.000166	0.01529	14.89	14.91	14.91	0.004272
0/1	Sample	0.02786	0.005814	0.006059	81.21	81.23	81.23	0.05464
0/1	Cond.	0.00786	1.275e-05	0.002502	24.57	24.57	24.57	0.01616
1/0	Sample	0.02651	-0.0002305	0.003142	2.846	2.849	2.849	6.225e-06
1/0	Cond.	0.005789	0.0002065	-0.005346	1.052	1.047	1.047	2.634e-05
1/1	Sample	0.02623	-0.003448	-0.01225	0.5522	0.5365	0.5377	5.539e-05
1/1	Cond.	0.005938	0.0002054	-0.002482	0.1171	0.1148	0.1149	8.243e-05

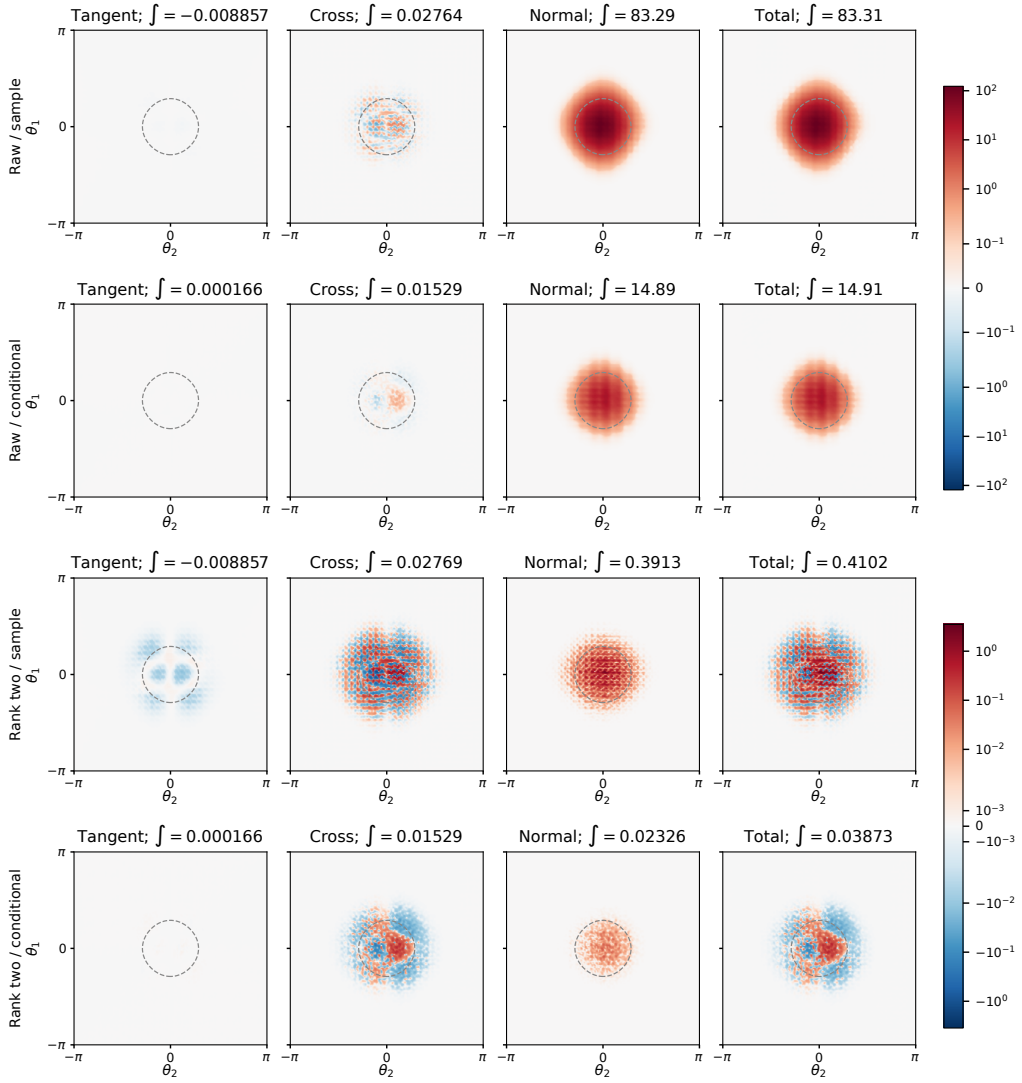

Figure 10: Signed tangent, cross, normal and total energy densities for image 0, initialization 0. Integrals in the titles retain signs; contribution norms must not be substituted for them. The raw and rank-two row pairs use separate symmetric-logarithmic scales, shared within each pair, because their magnitudes differ substantially. All panels use the same function, density and reference.

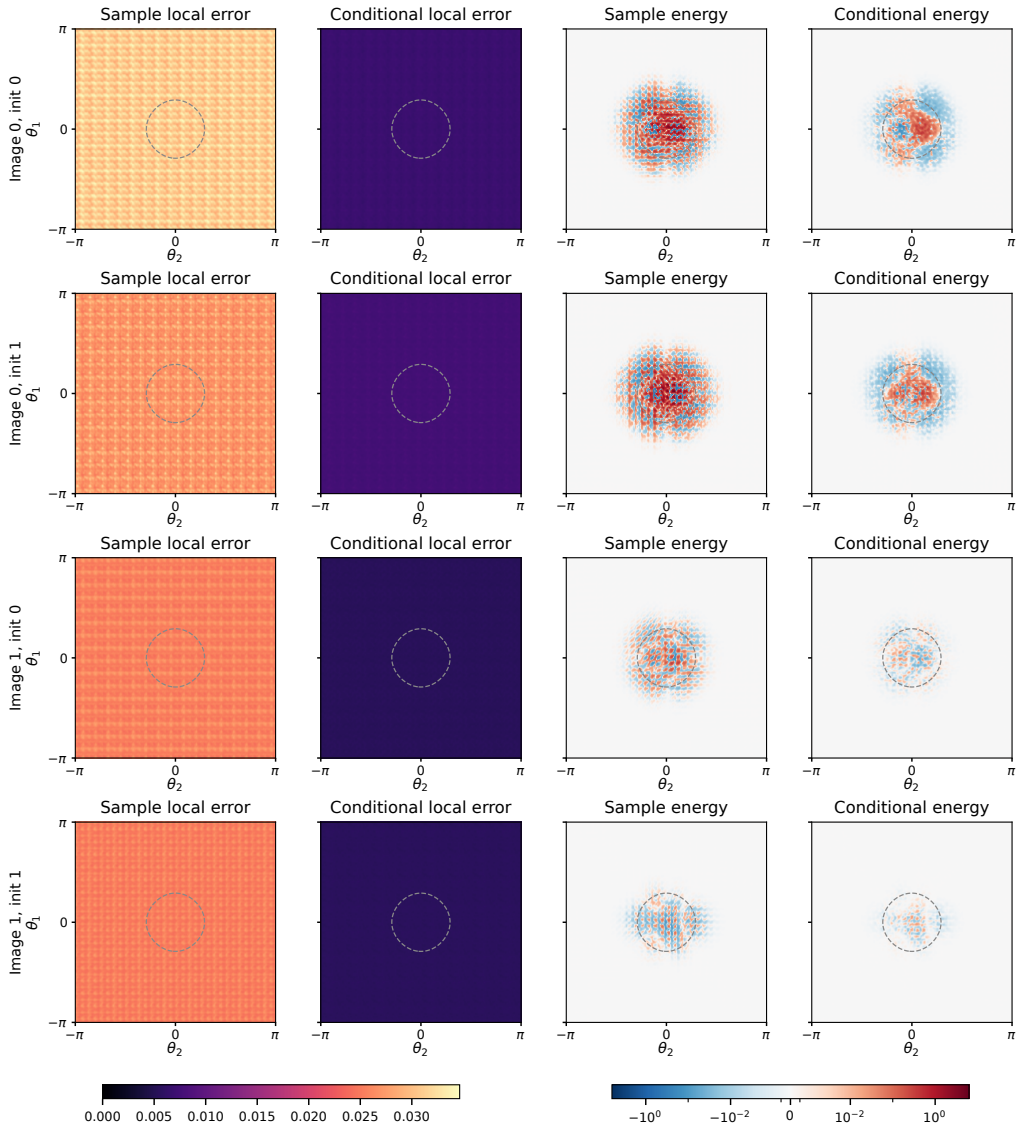

Figure 11: All four image–initialization pairs, rank-two fields. Local-error maps share one scale across all rows; signed energy maps share another. Each row holds its baseline-selected function fixed between supervision types. The common direction is different between rows, so spatial patterns are within-pair diagnostics, not a comparison of one function across images.