Coarse information and finite rank obstructions for flat lattice Hodge forms

Lluis Eriksson

31 October 2026 — candidate revision v4, not deposited

Abstract

On a finite periodic lattice with unit spacing, a flat Hodge form cannot have a positive coercivity constant uniform under refinement if it is supplemented only by linear observations in a fixed finite-dimensional real space and linear fine-space repairs of uniformly bounded rank. The linear observations may vary arbitrarily with the refinement, and their penalty need only vanish at zero. We prove the same obstruction directly on kernels of uniformly bounded rank constraints. The proof constructs a real physical Fourier subspace of dimension $2(2R+1)^d$ with Hodge energy at most $4\pi^2 dR^2/n^2$ times its squared norm. Its intersection with the unobserved directions gives a necessary rank budget. We also explain why this budget is not sufficient, using a two-dimensional cross-coupling example and a classical complete criterion for a fixed map. The Fourier construction, finite budget, and uniform obstructions have Lean proofs. We also construct derivative repairs, prove a differentiable-observation tangent obstruction, and derive an exact nonlinear regular-sector obstruction using the implicit function theorem. The operator, full-space tangent and later regular-sector endpoints have separately pinned Lean evidence; the fixed-map criterion remains analytic here. We derive output-rank divergence on the actual divergence-free space under a subquadratic directional spatial mean of nonnegative repair weights, allowing their pointwise maximum to be large. A constructed quadratic-weight mean observation and explicit penalty counterexamples delimit this result. These are finite flat statements, with no interacting or continuum mass-gap claim.

1 Purpose and relation to the existing record

This text consolidates a continuation of the exact-kernel obstruction in [1]. That paper exhibits fields in the kernel of a particular line-integral block map. Changing the map can remove a particular witness. Here the obstruction depends instead on the amount of coarse information: a whole low-energy subspace is too large to be seen by a fixed-dimensional measurement together with a bounded-rank repair. The map need not factor through the original block variable, be local, or obey a norm bound; it is real-linear in the unrestricted-penalty results. The penalty applied to its output may be nonlinear. Later sections treat differentiable observations with a locally upper-quadratic penalty and exact nonlinear regular sectors. The restricted-sector result is proved separately and does not follow merely from failure of an estimate on the full space.

Fourier diagonalization, rank-nullity and finite-dimensional Schur complements are classical. For the singular Schur complement and pseudoinverse background, see Albert and Carlson–Haynsworth–Markham [3, 4]. Finite-difference Hodge theory has a long history; Dodziuk's work [2] provides historical context, rather than this paper's specific lattice normalization. We make no claim to a new general spectral principle or an exhaustive priority search. The contribution is

a concrete, reproducibly checked constraint on proposed repair routes in this programme. This consolidation is intended for the existing paper dossier, rather than an additional deposit for each lemma.

Gauge-fixing and block averaging in renormalization-group constructions require their own domain and scaling analysis. Dimock [5] develops covariant axial gauge for an abelian model and discusses its relation to Balaban's programme. His introduction distinguishes a finite unit lattice from a torus with small spacing, equivalently a large unit-spacing volume. Our unrestricted, unit-spacing flat form is not identified with every constrained fluctuation form in those constructions. This negative diagnostic does not refute their gauge-fixed estimates or establish an interacting bound.

2 Spaces and the actual flat form

Fix integers $d \geq 1$, $n \geq 1$ and $N_c \geq 2$. Set $X_n = (\mathbb{Z}/n\mathbb{Z})^d$ and $m = N_c^2 - 1$. The real one-cochain space is

$$E_n = \{A : X_n \times \{1, \dots, d\} \longrightarrow \mathbb{R}^m\}, \qquad \|A\|^2 = \sum_{x \in X_n} \sum_{i=1}^{d} \|A_i(x)\|^2.$$

The internal coordinates are those of the repository's Euclidean Lie-coordinate model. No identification with a matrix Lie algebra is required here. For positive unordered plaquettes $i < k$, define

$$\begin{aligned} (\mathrm{d}A)_{ik}(x) &= A_k(x + e_i) - A_k(x) - A_i(x + e_k) + A_i(x), \\ (\delta A)(x) &= \sum_i (A_i(x) - A_i(x - e_i)). \end{aligned}$$

The flat operator $K_{0,n}$ is the actual repository Hodge operator, with

$$\langle A, K_{0,n} A \rangle = \|\mathrm{d}A\|^2 + \|\delta A\|^2. \tag{1}$$

The plaquette norm sums each pair once. All differences have unit spacing; there is no hidden n^2 or block-side squared factor.

Choose a direction j and two distinct internal unit coordinates u, v. They exist since $m \geq 3$. For $f : X_n \to \mathbb{C}$, put

$$(E_j f)_i(x) = \begin{cases} (\Re f(x))u + (\Im f(x))v, & i = j, \\ 0, & i \neq j. \end{cases}$$

This is a real-linear injection and $\|E_j f\|^2 = \sum_x |f(x)|^2$. Writing $D_i f(x) = f(x + e_i) - f(x)$, the actual curl and divergence give

$$\langle E_j f, K_{0,n} E_j f \rangle = \sum_{i \neq j} \sum_x |D_i f(x)|^2 + \sum_x |f(x) - f(x - e_j)|^2 = \sum_i \sum_x |D_i f(x)|^2. \tag{2}$$

The last equality is finite translation invariance. The exclusive two plaquette incidences and the divergence term account for every direction exactly once, with no extra factor two.

3 The original block kernel and its repair budget

We retain the principal results of the historical paper [1] before explaining the new information-budget obstruction. The original Appendix A (an even-block square-wave witness with quotient $8/\ell$) remains in the preserved v1; it is not reproduced here. Use positive integers ℓ, N', set $n = \ell N'$, and for a coarse site y put $B_y = \{\ell y + r : 0 \le r_i < \ell\}$. The original line-integral map is

$$(Q_\ell A)_i(y) = \ell^{-d} \sum_{x \in B_y} \sum_{k=0}^{\ell-1} A_i(x + ke_i).$$

For a constant field $A_i(x) = v_i$ and $S = \sum_i \|v_i\|^2 > 0$,

$$\|A\|^2 = \ell^d (N')^d S, \qquad \|\ell^\sigma Q_\ell A\|^2 = \ell^{2\sigma+2} (N')^d S.$$

A uniform Poincare estimate must therefore have $\sigma \ge (d-2)/2$ if it is to survive this sector. Equality alone makes the two powers identical. Larger exponents pass this necessary test too; neither equality nor the inequality is a sufficient full-space condition. This follows by cancelling $(N')^d S$ to obtain $\ell^{d-2-2\sigma} \le C_P$. The real-exponent comparison is analytic here.

For $d \ge 2$ choose distinct directions p, q and the real isometry $\iota(z) = (\Re z)u + (\Im z)v$ from Section 2. For $\ell \ge 2$ define

$$A_i^{(\ell)}(x) = \mathbf{1}_{i=p}\, \iota\big(e^{2\pi \mathrm{i} x_q/\ell}\big), \qquad \lambda_\ell = |e^{2\pi \mathrm{i}/\ell} - 1|^2 = 4\sin^2(\pi/\ell).$$

The character is well defined on $X_{\ell N'}$. The sum over the q-offset in every block is a nontrivial root-of-unity sum and vanishes. Shifts in the line direction p do not change the profile. Hence

$$Q_\ell A^{(\ell)} = 0, \quad \delta A^{(\ell)} = 0, \quad \|A^{(\ell)}\|^2 = (\ell N')^d, \quad \langle A^{(\ell)}, K_{0,n} A^{(\ell)} \rangle = (\ell N')^d \lambda_\ell. \tag{3}$$

Only the plaquette pair containing p, q contributes to the curl energy; its difference has constant squared norm λ_ℓ at every site. This also explains the exact energy in (3), including periodic boundary jumps.

Consequently any scale-wise estimate

$$\|A\|^2 \le C_P(\ell)\big(\langle A, K_{0,n} A\rangle + \Psi_\ell(Q_\ell A)\big), \qquad \Psi_\ell(0) = 0,$$

forces $C_P(\ell) \ge \lambda_\ell^{-1} \ge \ell^2/(4\pi^2)$. No fixed $C_P > 0$ before all block sides can work. This includes $\Psi_\ell(y) = \|T_\ell y\|^2$ for arbitrary linear postconditioners, and all critical or faster scalar rescalings. There is no regularity or sign assumption on the scalar function Ψ_ℓ.

The original spectral strengthening is retained analytically. For $1 \le r \le R \le \lfloor (\ell-1)/2 \rfloor$, replace the profile by $\iota(ce^{2\pi \mathrm{i} r x_q/\ell})$, $c \in \mathbb{C}$. These are mutually orthogonal real two-dimensional spaces in $\ker Q_\ell$. Derivative orthogonality must be checked too: every forward difference multiplies a character by its frequency multiplier, so the same root-of-unity Gram sum annihilates cross terms between distinct frequencies. Thus for $H_\ell = K_{0,n} + (T_\ell Q_\ell)^*(T_\ell Q_\ell)$ the sum of these spaces has dimension $2R$ and maximum Rayleigh quotient at most $4\sin^2(\pi R/\ell) \le 4\pi^2 R^2/\ell^2$. Finite-dimensional min–max gives $\mu_{2R}(H_\ell) \le 4\pi^2 R^2/\ell^2$. Taking $R = \lfloor \ell^\alpha \rfloor$, $0 < \alpha < 1$, for sufficiently large ℓ, yields a growing cluster below $4\pi^2 \ell^{-2(1-\alpha)}$. Eigenvalue language is restricted to this linear self-adjoint operator, not applied to a general nonlinear penalty.

Finally, for any linear changed measurement $\widetilde{Q}_\ell$, linear postconditioner T_ℓ and positive self-adjoint fine term V_ℓ, put

$$\varepsilon_\ell = \frac{\|T_\ell \widetilde{Q}_\ell A^{(\ell)}\|}{\|A^{(\ell)}\|}, \qquad \nu_\ell = \frac{\langle A^{(\ell)}, V_\ell A^{(\ell)}\rangle}{\|A^{(\ell)}\|^2}.$$

Set $\widetilde{H}_\ell = K_{0,n} + (T_\ell \widetilde{Q}_\ell)^*(T_\ell \widetilde{Q}_\ell) + V_\ell$. Expansion gives the exact repaired quotient

$$\frac{\langle A^{(\ell)}, \widetilde{H}_\ell A^{(\ell)}\rangle}{\|A^{(\ell)}\|^2} = \lambda_\ell + \varepsilon_\ell^2 + \nu_\ell.$$

A full-space constant $c > 0$ therefore requires $\varepsilon_\ell^2 + \nu_\ell \geq c - \lambda_\ell$ at every scale, hence $\liminf_{\ell\to\infty}(\varepsilon_\ell^2 + \nu_\ell) \geq c$. On the single witness line, uniform coercivity is equivalent to $\inf_{\ell\geq 2}(\lambda_\ell + \varepsilon_\ell^2 + \nu_\ell) > 0$; this is not full-space sufficiency. In particular if $\widetilde{Q}_\ell = Q_\ell + R_\ell$, $T_\ell = \ell^{(d-2)/2} I$ and $\eta_\ell = \|R_\ell A^{(\ell)}\|/\|A^{(\ell)}\|$, then $\varepsilon_\ell^2 = \ell^{d-2}\eta_\ell^2$. In dimension four and with $V_\ell = 0$, leakage $o(\ell^{-1})$ cannot even repair this line. The liminf, real-exponent, spectral multiplicity and witness-line equivalence are analytic statements. The older exact kernel, norm-form energy, finite budget and dimension-four factorized no-go have their own preserved formal source and reproductions, separately anchored below.

4 A constructed low band

For a nonnegative integer R with $2R < n$, let $\Lambda_R = \{-R, \ldots, R\}^d$ and $B = |\Lambda_R| = (2R+1)^d$. Define

$$\chi_k(x) = \exp(2\pi \mathrm{i}\, k \cdot x/n), \qquad f_a = \sum_{k\in\Lambda_R} a_k \chi_k, \qquad S_R = \{E_j f_a : a \in \mathbb{C}^{\Lambda_R}\} \subset E_n.$$

Lemma 4.1. *The real subspace S_R has dimension $2B$, and for every $A \in S_R$,*

$$\langle A, K_{0,n} A\rangle \leq b_R(n)\|A\|^2, \qquad b_R(n) = \frac{4\pi^2 d R^2}{n^2}. \tag{4}$$

Proof. For $k, \ell \in \Lambda_R$, the product root-of-unity sum is $\sum_x \overline{\chi_k(x)}\chi_\ell(x) = n^d \mathbf{1}_{k=\ell}$. The no-alias condition ensures that coordinate differences, of absolute value at most $2R < n$, cannot be nonzero multiples of n. Thus synthesis is injective and its complex image has dimension B. As a real space that image has dimension $2B$; applying the real-linear injection E_j preserves dimension. The factor two represents independent real and imaginary coefficients stored in two real internal coordinates, not two copies of a real scalar band.

For completeness, energy orthogonality is needed as well as mass orthogonality. For each i,

$$D_i f_a = \sum_k a_k (e^{2\pi \mathrm{i} k_i/n} - 1)\chi_k.$$

Applying the same Gram sum to these derivative amplitudes gives

$$\|E_j f_a\|^2 = n^d \sum_k |a_k|^2,$$
$$\langle E_j f_a, K_{0,n} E_j f_a\rangle = n^d \sum_k |a_k|^2 \sum_i |e^{2\pi \mathrm{i} k_i/n} - 1|^2.$$

Since $|e^{\mathrm{i}t} - 1| \leq |t|$ and $|k_i| \leq R$, each sum of multipliers is at most $4\pi^2 d R^2/n^2$. This proves the bound for every member of the constructed image, not only its individual generators. □

5 The finite information budget

Let F be a finite-dimensional real vector space, $C : E_n \to F$ a linear measurement and $V : E_n \to E_n$ a linear operator. No symmetry or positivity of V is needed. Let $\Phi : F \to \mathbb{R}$ be any function with $\Phi(0) = 0$.

Theorem 5.1. *If $c > b_R(n)$ and*

$$c\|A\|^2 \le \langle A, K_{0,n}A\rangle + \Phi(CA) + \langle A, VA\rangle \quad \textit{for all } A \in E_n, \tag{5}$$

then $2(2R+1)^d \le \dim F + \operatorname{rank} V$.

Proof. If the budget is smaller than $\dim S_R$, restrict the map (C, V) to S_R. Its image has dimension at most $\dim F + \operatorname{rank} V$. Rank-nullity yields a nonzero $A \in S_R$ with $CA = VA = 0$. Equation (5) and Lemma 4.1 imply $c\|A\|^2 \le b_R(n)\|A\|^2$, a contradiction. □

The codomain dimension is a conservative measurement budget. We do not label this formal endpoint as a sharper theorem with $\operatorname{rank} C$ in its place. No regularity, sign, polynomial form or growth restriction on Φ enters.

Proposition 5.2. *If $P : E_n \to E_n$ is linear, $c > b_R(n)$ and $\dim F + \operatorname{rank} P < 2(2R+1)^d$, the inequality*

$$c\|A\|^2 \le \langle A, K_{0,n}A\rangle + \Phi(CA) \qquad (A \in \ker P)$$

is false.

Proof. Apply the same rank-nullity argument to $(C, P)|_{S_R}$. Its nonzero kernel vector belongs to the stipulated sector as well as the unobserved band. □

For example, $d = 4, n = 16, R = 1$ gives dimension 162 and $b_R(n) = \pi^2/16 < 1$. Twelve measurement coordinates and a repair of rank at most 149 cannot yield (5) with $c = 1$. This illustrates a necessary budget, not the programme's chosen target constant.

6 Uniform refinement and sector exclusion

Fix a positive integer N' and put $n_L = (L+1)N'$ for nonnegative integers L. The finite-dimensional space F is fixed. Allow arbitrary families of linear maps $C_L : E_{n_L} \to F$, arbitrary functions $\Phi_L : F \to \mathbb{R}$ with $\Phi_L(0) = 0$, and linear maps $V_L : E_{n_L} \to E_{n_L}$ with $\operatorname{rank} V_L \le r$ for one fixed nonnegative integer r.

Theorem 6.1. *There is no $c > 0$ for which (5) holds for every L and every $A \in E_{n_L}$ with these families. Likewise, for linear P_L of rank at most r, there is no positive uniform c for the sector estimates of Proposition 5.2 on all $\ker P_L$.*

Proof. Set $M = \dim F$ and $R = M + r + 1$. Since $d \ge 1$, $2(2R+1)^d > M + r$. Choose L so large that $2R < n_L$ and $b_R(n_L) < c$. Here is a fully finite choice argument, without an assumed limit theorem: with $a = 4\pi^2 dR^2$, take an integer $L > \max\{2R, a/c, 1\}$. Since $N' \ge 1$, $n_L \ge L + 1 > L > 1$, and $a < cL \le cn_L^2$. Theorem 5.1 now contradicts the rank bound. Use Proposition 5.2 for the separate sector conclusion. □

The constants c, r precede the quantifier over L. Arbitrary rescalings of C_L preserve its fixed codomain and do not escape the theorem. Removing harmonic or constant directions is covered only when a concrete linear constraint describing the proposed sector has uniformly bounded rank. A quotient or its norm is not silently identified with such a sector. Linearity is essential: if nonlinear observations were allowed, the single real observation $C_L(A) = \|A\|^2$, with $\Phi_L(y) = y$ and $V_L = 0$, would give the uniform lower bound $c = 1$ since $K_{0,n_L} \succeq 0$. This counterexample explains why finite-dimensional output alone is not the hypothesis.

7 Constructed derivative repairs and tangent observations

The bounded-rank obstruction does not exhaust the available repairs. For the actual componentwise forward gradient $\mathcal{G}A = (D_k A_i)_{x,i,k}$, finite periodic summation and cancellation of the mixed differences give

$$\|\mathcal{G}A\|^2 = \langle A, K_{0,n} A\rangle.$$

This is an identity on the whole cochain space, with unordered plaquettes counted once; it is not an assumed comparison on a supplied band. Let W be the diagonal multiplication by real weights $w_{x,i,k}$, and construct $V_w = \mathcal{G}^* W \mathcal{G}$. If $w_{x,i,k} \le M$ and $M \ge 0$, then

$$\langle A, V_w A\rangle = \sum_{x,i,k} w_{x,i,k} \|(D_k A_i)(x)\|^2 \le M \langle A, K_{0,n} A\rangle. \tag{6}$$

A lower bound on the weights is unnecessary for this upper estimate. Consequently the constructed Fourier band satisfies

$$q_w(A) := \langle A, (K_{0,n} + V_w)A\rangle \le (1 + M) b_R(n) \|A\|^2 \qquad (A \in S_R).$$

These operator identities and inequalities are formally proved in the separately pinned source reported below.

Theorem 7.1 (Finite-dimensional tangent obstruction)**.** *Let E be a finite-dimensional real inner-product space, F a finite-dimensional real normed space, and $S \subseteq E$ a linear subspace with $\dim S > \dim F$. Let $q : E \to \mathbb{R}$ satisfy $q(x) \le b\|x\|^2$ on S. Suppose $C : E \to F$ is Fréchet differentiable at zero, $C(0) = 0$, and for some $\delta > 0$ and finite $K \in \mathbb{R}$,*

$$\Phi(y) \le K\|y\|^2 \qquad (\|y\| < \delta). \tag{7}$$

If $c > b$, the full-space estimate $c\|x\|^2 \le q(x) + \Phi(C(x))$ is impossible.

Proof. Put $L = DC(0)$. Rank-nullity applied to the actual restriction $L|_S$ produces $v \in S \cap \ker L$, $v \neq 0$. Differentiability gives $\|C(tv)\|/\|tv\| \to 0$ as nonzero real $t \to 0$. For small such t, (7) applies and the proposed bound implies

$$c \le b + K \frac{\|C(tv)\|^2}{\|tv\|^2}.$$

Taking the limit gives $c \le b$, a contradiction. The kernel vector and the remainder limit are derived, not introduced as hypotheses of the theorem. □

The three focal Lean endpoints include this theorem, its remainder lemma and the actual physical specialization with threshold $(1 + M)b_R(n)$. The proof does not require an explicit assumption $\Phi(0) = 0$. A local-ball version is also valid by the same shrinking-ray argument, while the original formal endpoint above is stated on the full space. The later local kernel and average-weight endpoint in Section 9 has its own separately reproduced proof.

This result trades hypotheses. The global linearity of C is replaced by differentiability at zero, while an arbitrary zero-vanishing penalty is replaced by genuine quadratic upper control. Merely asking C and Φ to be C^1 with zero derivatives at zero is insufficient: for $1 < \alpha < 2$,

$$C(x) = \|x\|^\alpha, \qquad \Phi(y) = |y|^{2/\alpha} \quad \Longrightarrow \quad \Phi(C(x)) = \|x\|^2.$$

Both derivatives at zero vanish. Off zero, $\|DC(x)\| = \alpha\|x\|^{\alpha-1}$ and $|\Phi'(y)| = (2/\alpha)|y|^{2/\alpha-1}$ tend to zero, so both maps are indeed C^1 there and at zero. But $\Phi(y)/|y|^2 = |y|^{2/\alpha-2}$ is unbounded above near zero. On the actual lattice, setting $w = 0$ gives a uniform lower bound with $c = 1$ because $K_{0,n} \succeq 0$. This is an exact counterexample, not a numerical test.

8 Exact nonlinear regular sectors

The next theorem is an analytic consequence of rank-nullity and the classical implicit function theorem. The actual curve construction and local-energy obstruction now have separately pinned formal proofs, identified in the evidence section. The pinned formal library supplies an implicit parametrization with derivative the inclusion of the derivative kernel [6]; we construct and use that parametrization explicitly below rather than assuming the curve needed for the conclusion.

Theorem 8.1 (Regular-sector tangent obstruction)**.** *Let E be a finite-dimensional real inner-product space and F, G finite-dimensional real normed spaces. Let $T : E \to E$ be linear and set $q(x) = \langle x, Tx\rangle$. Suppose a linear subspace S satisfies $q(v) \le b\|v\|^2$ on S and $\dim S > \dim F + \dim G$. Let C be differentiable at zero with $C(0) = 0$ and let Φ satisfy (7). Let $P : E \to G$ be C^1 on a neighborhood of zero with $P(0) = 0$ and $A = DP(0)$ surjective. For every $c > b$ and every $r > 0$, there is no estimate*

$$c\|x\|^2 \le q(x) + \Phi(C(x)) \qquad (P(x) = 0,\ \|x\| < r). \tag{8}$$

Proof. Put $L = DC(0)$. Apply rank-nullity to $(L, A)|_S : S \to F \times G$ to obtain $v \in S$, $v \neq 0$, with $Lv = Av = 0$. Choose a linear complement W of $H = \ker A$. Surjectivity implies $A|_W : W \to G$ is bijective. Apply the implicit function theorem to $(u, w) \mapsto P(u + w)$ at $(0, 0) \in H \times W$. It produces a C^1 map ψ near zero in H with

$$\psi(0) = 0, \quad P(u + \psi(u)) = 0, \quad D\psi(0) = -(A|_W)^{-1}A|_H = 0.$$

Thus $\gamma(t) = tv + \psi(tv)$ lies in the exact sector for small real t, $\gamma(0) = 0$, and $\gamma(t) = tv + e(t)$ with $\|e(t)\| = o(|t|)$. In particular $\|\gamma(t)\|/|t| \to \|v\| > 0$, so the nonzero points of this curve eventually lie in the stipulated ball.

The curve need not remain in S. This is why the quadratic form assumption is used: direct expansion gives

$$q(\gamma(t)) = t^2 q(v) + t\langle v, Te(t)\rangle + t\langle e(t), Tv\rangle + \langle e(t), Te(t)\rangle.$$

Since T is bounded in finite dimension, division by t^2 gives $q(\gamma(t))/t^2 \to q(v)$. Differentiability of C and $Lv = 0$ give $C(\gamma(t)) = Le(t) + o(|t|) = o(|t|)$. Apply the assumed bound along this

exact curve and use (7). Division by t^2 and the limit yield $c\|v\|^2 \le q(v) \le b\|v\|^2$, contrary to $c > b$. □

Unlike the preceding full-space result, an arbitrary energy with an upper bound only on S would not justify this proof: the constructed exact-sector curve generally leaves S. The displayed expansion supplies the missing comparison for the actual quadratic energy.

A concrete curved example also shows why an exact vector in S is not the right object. Write $E = \mathbb{R}^4$ with coordinates (u_1, u_2, y, z), put $P(u, y, z) = z - \|u\|^2$, $C(u, y, z) = y$, $\Phi(y) = y^2$, $S = \{z = 0\}$ and $q(u, y, z) = z^2$. Then $DP(0)$ is surjective and $\dim S = 3 > 1 + 1$. But $S \cap \ker C \cap P^{-1}(0) = \{0\}$: the linear band contains no nonzero exact-sector unseen vector. Nevertheless $\gamma(t) = (t, 0, 0, t^2)$ lies in the exact sector and has repaired Rayleigh quotient $t^2/(1 + t^2) \to 0$. The implicit-curve argument is genuinely needed.

Nor can quadratic q in Theorem 8.1 be replaced merely by a C^1 energy stationary at zero with the same upper bound on S. With the same E, C, Φ, S, instead put

$$P(u, y, z) = z - \|u\|^{4/3}, \qquad e(u, y, z) = |z|^{3/2}.$$

Both maps are C^1, $DP(0)(u, y, z) = z$ is surjective, $De(0) = 0$, and $e = 0$ on S, so $b = 0$. On the exact sector,

$$e(x) + \Phi(C(x)) = \|u\|^2 + y^2, \qquad \|x\|^2 = \|u\|^2 + y^2 + \|u\|^{8/3}.$$

For $\|x\| < 1$, $\|u\|^{8/3} \le \|u\|^2$, so the local sector bound with $c = 1/2 > b$ holds. Differentiability of the norm powers at zero follows from their exponents being greater than one; their gradients off zero tend to zero. This exact counterexample refutes a proposed C^1-energy weakening, while the actual quadratic Hodge/repair energy retains the required expansion.

The regularity condition cannot simply be dropped. Take $E = \mathbb{R}^3$, $F = G = \mathbb{R}$, $C = 0$, $\Phi = 0$, $T = 0$, $S = E$, and $P(x) = \|x\|^2$. Then P is smooth, $DP(0) = 0$, and the strict dimension budget holds, but $P^{-1}(0) = \{0\}$ and the estimate with $c = 1$ is vacuous on the sector. This refutes the proposal to replace submersion by bounded derivative rank. It does not refute Theorem 8.1, because $DP(0)$ is not surjective. Critical constraint maps require a separate analysis.

Proposition 8.2 (Uniform physical consequence)**.** *Let $n_L \to \infty$ with fixed d, N_c, and let F, G have fixed finite dimensions. For each L, allow its own C_L, P_L, Φ_L satisfying the hypotheses of Theorem 8.1. Let q_L be the actual flat Hodge form plus the constructed weighted derivative repair with weights $w_L \le M_L$, $M_L \ge 0$, and $M_L/n_L^2 \to 0$. Even on the exact sectors $P_L = 0$ and arbitrarily small balls of positive radius r_L, no positive uniform coercivity constant exists.*

Proof. Choose a fixed integer R with $2(2R+1)^d > \dim F + \dim G$. For all large L, $2R < n_L$ and the constructed band obeys $q_L \le (1 + M_L) b_R(n_L) \|\cdot\|^2$. Its coefficient tends to zero. For any proposed $c > 0$, choose one such L with this coefficient below c and apply Theorem 8.1 at that fixed lattice. The curve is shrunk only after choosing L: no uniform implicit-function radius, inverse-derivative norm, penalty constant or r_L is required. □

These are finite unit-spacing obstructions for regular nonlinear sectors, not identifications of a physical gauge-fixed domain or a quotient norm.

9 Directional mean weights and local nonlinear observations

Pointwise bounds on a derivative repair are stronger than the following actual spatial-mean condition. Fix $d \geq 2$, $N_c \geq 2$, and distinct directions i, j. For nonnegative componentwise weights $w_n(k,(x,l))$, put

$$Q_n(A) = \langle A, K_{0,n}A\rangle + \sum_{x,k,l} w_n(k,(x,l))\|A_l(x+e_k) - A_l(x)\|^2, \qquad b_{ij}(n) = n^{-d}\sum_x w_n(i,(x,j)).$$

The repair is the actual G^*W_nG of the componentwise forward gradient. In particular $Q_n \geq 0$; signed weights are not included in this section. Let $D_n = \ker\delta$ be the actual flat divergence-free space.

Theorem 9.1 (Local nonlinear average-weight obstruction)**.** *For each $n \geq 1$, let F_n be a finite-dimensional real normed space, $C_n : E_n \to F_n$ be Fréchet differentiable at zero with $C_n(0) = 0$, and let $\Phi_n : F_n \to \mathbb{R}$ satisfy*

$$\Phi_n(y) \leq K_n\|y\|^2 \quad (\|y\| < \delta_n), \qquad \delta_n > 0.$$

Suppose one fixed $c > 0$ and positive radii r_n satisfy

$$c\|A\|^2 \leq Q_n(A) + \Phi_n(C_n(A)) \quad (A \in D_n,\ \|A\| < r_n). \tag{9}$$

For every integer $0 \leq R < n$,

$$c(2R - \dim F_n) \leq 2(1 + b_{ij}(n))\sum_{s=1}^{R}|e^{2\pi\mathrm{i}s/n} - 1|^2 \leq \frac{8\pi^2(1+b_{ij}(n))R^3}{n^2}. \tag{10}$$

Consequently, if $(1 + b_{ij}(n))/n^2 \to 0$, then $\dim F_n \to \infty$. No uniform bound on K_n, δ_n^{-1} or r_n^{-1} is required.

Proof. Write $L_n = DC_n(0)$. If $v \in D_n \cap \ker L_n$, differentiability gives $C_n(tv) = o(|t|)$. For all sufficiently small nonzero t, both the local coercivity and penalty bound apply. Divide by t^2 and let $t \to 0$ to get

$$c\|v\|^2 \leq Q_n(v). \tag{11}$$

This argument is performed separately at each n; it does not impose uniform neighborhoods or derivative remainder rates.

Use the two internal coordinates u, v of Section 2 to regard a complex number as a real internal vector. For $1 \leq s \leq R$ and $a \in \{1, \mathrm{i}\}$, define $\psi_{s,a}$ by taking only component j nonzero and setting it to $n^{-d/2}$ times the internal vector corresponding to $ae^{2\pi\mathrm{i}sx_i/n}$. The finite-character orthogonality identity and the real orthogonality of the two quadratures give an orthonormal family of $2R$ fields. This also covers even n; the two quadratures use two internal coordinates, not two allegedly independent scalar sine and cosine functions. Since $i \neq j$, these fields are independent of x_j, so $\delta\psi_{s,a} = 0$. Their span $S_{R,n}$ is thus an actual $2R$-dimensional subspace of D_n.

Every nonzero forward difference of $\psi_{s,a}$ is in direction i and has the same squared norm at every site, namely $n^{-d}|e^{2\pi\mathrm{i}s/n} - 1|^2$. The exact Hodge and repair energies are therefore

$$Q_n(\psi_{s,a}) = (1 + b_{ij}(n))|e^{2\pi\mathrm{i}s/n} - 1|^2.$$

In particular the trace of the compression of Q_n to $S_{R,n}$ is the middle expression in (10); no invariance of $S_{R,n}$ under the weighted operator is asserted or needed.

For completeness, the positive trace/rank argument is as follows. Set $U = S_{R,n} \cap \ker L_n$. Rank-nullity gives $\dim U \geq 2R - \dim F_n$. Extend an orthonormal basis of U to one of $S_{R,n}$. Equation (11) bounds the terms on U below by c; positivity of Q_n makes every remaining diagonal term nonnegative. Hence $c(2R - \dim F_n) \leq \operatorname{tr}_{S_{R,n}} Q_n$. Finally $|e^{2\pi \mathrm{i} s/n} - 1| \leq 2\pi s/n \leq 2\pi R/n$, which proves the last inequality in (10).

Fix R. The right side tends to zero under the stated mean condition. Eventually it is less than c. If $\dim F_n < 2R$, integrality would make the left side at least c. Thus $\dim F_n \geq 2R$ eventually. Arbitrarily large fixed R prove the claimed divergence. □

The condition concerns one directional spatial average, not the maximum of all weights. For example, a weight of size n^{d+1} at one site in component (i, j) and zero elsewhere has $b_{ij}(n) = n$, while its maximum is not subquadratic. Theorem 9.1 still excludes bounded output dimension under its other hypotheses. This example is a direct calculation, not an additional formal endpoint. The nonnegative-weight condition is essential to the trace argument as used here: the complement of the derivative kernel must contribute nonnegative energy. This is not a universal extension of older statements allowing signed weights with a pointwise absolute cap.

9.1 A constructed quadratic-weight boundary example

The exponent 2 cannot be replaced by an unrestricted allowance of quadratic-size mean weights. Let $\mu_j(A) = n^{-d} \sum_x A_j(x)$, with output in the same fixed space $F = (\mathbb{R}^{N_c^2 - 1})^d$ for every n. A finite path comparison on the actual periodic lattice gives

$$\|A - \operatorname{lift} \mu(A)\|^2 \leq dn^2 \|GA\|^2, \qquad \|\operatorname{lift} \mu(A)\|^2 = n^d \|\mu(A)\|^2.$$

One way to obtain the first inequality is to write each difference of site values as a path of at most dn forward edges, apply finite Cauchy–Schwarz, and average the pairs of sites; the displayed constant is deliberately nonoptimal. The exact modular-torus reindexing and its transport to the physical cochain norm are also proved in the companion formal construction. Using $\|x + y\|^2 \leq 2\|x\|^2 + 2\|y\|^2$, choose actual weights $w_n(k, (x, l)) = 2dn^2$ and the quadratic penalty $\Phi_n(y) = 2n^d \|y\|^2$. Then

$$\|A\|^2 \leq \langle A, G^* W_n G A \rangle + \Phi_n(\mu(A)) \leq Q_n(A) + \Phi_n(\mu(A))$$

on the entire space, hence also on D_n. The mean observation is an actual continuous linear map with finite-dimensional output type independent of n, not a postulated observation certificate. Here $b_{ij}(n) = 2dn^2$ and the subquadratic hypothesis fails. The example allows $K_n = 2n^d$ to vary, exactly as the obstruction does. It proves neither an optimal boundary constant nor positive coercivity for an interacting gauge operator. The path/mean mechanism is classical; see the primary exposition in [8].

9.2 Why the penalty assumption remains

The C^1 counterexample after Theorem 7.1 also applies to the actual nonnegative Q_n and to D_n. A second useful counterexample retains a polynomial observation and a globally Lipschitz penalty:

$$C(A) = \|A\|^2 \in \mathbb{R}, \qquad \Phi(y) = c|y|, \qquad DC(0) = 0, \qquad \Phi(C(A)) = c\|A\|^2.$$

Thus output dimension one suffices even when $Q_n = 0$. No finite local quadratic upper constant for Φ exists, since $c|y|/y^2 = c/|y|$. Likewise, $C(A) = \|A\|^{3/2}$ and $\Phi(y) = c|y|^{4/3}$ are both C^1 with zero derivatives at zero, yet their composition is exactly quadratic. These are complete analytic counterexamples, not new Lean claims.

There is a distinct analytic regularity tradeoff, not a hypothesis removal for merely differentiable observations. If, on each fixed derivative-kernel direction, $C(tv) = t^2 H_v + o(t^2)$, and $\Phi(0) = 0$, $D\Phi(0) = 0$, then $\Phi(C(tv)) = o(t^2)$. Indeed $\|C(tv)\|/t^2$ is bounded and $|\Phi(y)| \le \epsilon\|y\|$ near zero, including at zero; multiply the two estimates and let $\epsilon \downarrow 0$. A C^2 observation and a nonnegative differentiable penalty vanishing at zero suffice for this step, because a differentiable local minimum has zero derivative. No actual observation's C^2 regularity is proved here, and no formal composition of this alternate route is claimed.

The strengthened mean-weight result is an output-rank divergence theorem. Its fixed-R argument does not assert a positive density of observations per lattice volume. Fixed $c > 0$, nonnegative weights and the stated penalty condition remain explicit. In particular no interacting or continuum Yang–Mills hypothesis, ultraviolet activity estimate, or physical spectral gap is supplied by this section.

10 Why dimension is not a sufficiency criterion

The following analytic statement is included to prevent the reverse inference from a necessary rank budget. It is not a new Lean-verified endpoint here. The proof below is the classical singular-Schur range condition [3, 4], expressed as boundedness below on each measurement fiber; we claim no new general matrix criterion. Let E be a finite-dimensional real inner-product space, $H = H^*$ and $B = H - cI$. Fix linear $C : E \to F$, and decompose $E = N \oplus W$ with $N = \ker C$, $W = N^\perp$. In this decomposition write

$$B = \begin{pmatrix} A & D \\ D^* & G \end{pmatrix}.$$

Proposition 10.1. *A finite real-valued function Φ with $\Phi(0) = 0$ can make $\langle x, Bx\rangle + \Phi(Cx) \ge 0$ for every x if and only if $A \succeq 0$ and $\ker A \subseteq \ker D^*$. In that event a quadratic penalty suffices.*

Proof. Taking $x \in N$ proves $A \succeq 0$. For $z \in \ker A$ and fixed $w \in W$, the value at $tz + w$ contains $2t\langle z, Dw\rangle$, while $C(tz + w) = Cw$ is fixed. Boundedness below for all real t forces this coefficient to vanish.

Conversely, $\ker A \subseteq \ker D^*$ implies $D(W) \subseteq \operatorname{ran} A$. Let $A^\dagger$ be the inverse on $\operatorname{ran} A$ and zero on $\ker A$. Completing the square gives

$$\langle n + w, B(n + w)\rangle = \langle n + A^\dagger Dw, A(n + A^\dagger Dw)\rangle + \langle w, (G - D^* A^\dagger D)w\rangle.$$

The restriction $C|_W$ is injective; extend its inverse on $C(W)$ to a linear map $T : F \to W$. The quadratic function $\Phi(y) = \langle Ty, (D^* A^\dagger D - G)Ty\rangle$ cancels the second term and vanishes at zero. If F has a chosen Euclidean norm, a sufficiently large nonnegative multiple of $\|y\|^2$ also suffices, by finite-dimensional norm bounds. □

For a concrete counterexample let $H = \operatorname{diag}(0, 2)$, $c = 1$ and $C(x, y) = x + y$. On $\ker C$, $\langle x, (H - I)x\rangle = 0$, and C detects the negative eigendirection of $H - I$. Nevertheless, on $(t, 1 - t)$ the repaired form is $1 - 2t + \Phi(1)$, unbounded below for every finite $\Phi(1)$. Kernel nonnegativity and low-mode detection alone do not suffice; cross-coupling must be checked.

11 Operator scaling and the actual consumer

In the repository, the flat fixed-volume covariance endpoint uses $K_0 + aQ^*Q$ with the physical counting ℓ^2 norm and the same unscaled flat Hodge operator. More precisely, the declared precision shell includes $-\Sigma$. Specializing every perturbation and its summable majorant to zero and taking $a > 0$ gives the available certified coercivity lower bound $\min(1, a)/C_P$, with $C_P > 0$ from the full-periodic block-Poincare estimate. This is not a claim that this value is optimal. At fixed d, N_c, N' the block codomain is fixed and finite. Inspection therefore identifies the unscaled flat base route with the setting of Theorem 6.1, with $\Phi_L(y) = a_L\|y\|^2$ and no fine repair. Even allowing arbitrary scalar a_L or rescaling Q_L does not yield a positive uniform full-space constant. This is a source-level specialization in this text; no new Lean adapter to that consumer is claimed. At the new source anchor below, the module `PhysicalGaugeFixedPrecision` defines `flatGaugeFixedPrecisionCLM`; its explicit-Poincare coercivity theorem is recorded by its full name and specialization in the linked declaration map.

The inspected endpoint is explicitly fixed-volume. Its positive rate may depend on L, N' and C_P, so there is no contradiction. Nor does the obstruction exclude all possible uniform Combes–Thomas arguments. Its locality radius and matrix-entry bounds also vary with L; coercivity alone would not settle the rate. An interacting Hessian, a growing measurement budget, a different sector norm or a physically rescaled operator must each be audited separately.

In particular, writing $\ell = L + 1$ so that $n_L = \ell N'$, replacing K_0 by $\ell^2 K_0$ changes the band threshold exactly to $4\pi^2 dR^2/(N')^2$. The proof above no longer drives it below every fixed $c > 0$. This observation is an escape from this proof, not an established coercivity theorem for the programme's physical operator. No ultraviolet activity estimate, continuum limit, Wilson-Hessian identification or Yang–Mills mass gap is obtained.

12 Formal evidence and publication context

The source-tree anchor for the new physical-band and uniform-rank results is

`e31b428917aa767b4a1aad097db5cf92fb9cd23e`.

The original factorized-kernel and repair development in Section 3 instead has the distinct tested source anchor

`7b25bfb39357406cf61880e3ac0428d7d1ba3449`.

Its exact modules are not present in the new main tree merely because a reproduction supplement is present there. The two full preserved focal runs at that older SHA agree in compiled object, oracle log and source manifest; their six named reports use only the three standard axioms. The reproduced factorized endpoint is dimension four. The general-dimensional witness calculation, trigonometric closed form, spectral multiplicity, liminf and witness-line statements are analytic here; they are not all silently included in that formalization claim. The external declaration map keeps both anchors. The public source repository is THE-ERIKSSON-PROGRAMME on GitHub. The corrected map is `declaration-map.json` and the preserved evidence is `reproducibility/`. Exact source SHAs, not moving branch heads, identify the formal claims. The revised map and manuscript have their own revision SHA in the accompanying package.

The geometric and Fourier modules are ordinary proof dependencies, not assumptions about a newly supplied band. Each development stage has complete same-source fresh-clone Colab evidence. For the uniform endpoints the tested source is

`17ca395d4e9de992b1c575454a112da9e6f97e86`.

Its two runs agree in compiled object, raw oracle log and source-manifest hashes. The standalone budget lemma uses `propext, Quot.sound`; the other four uniform-module reports also use `Classical.choice`. No extra axiom is reported. Actual-log acceptance and three real-field rejection tests pass under normal and optimized Python. These tests concern evidence handling, not mathematical independence.

The additional whole-space gradient/Hodge identity and constructed weighted repair bounds use tested source

`dfd49d0b0730a4d36f58a4cbb7a17d6a41fb60bc`.

The full-space tangent endpoints in Section 7 use tested source

`81cf07b2e0d30543682c4073a4a2d82950d638b7`.

Each has two complete same-source fresh Colab Linux reconstructions with matching compiled object, raw oracle and LF-source-manifest hashes. For the tangent source, all three named endpoints use only `propext`, `Classical.choice` and `Quot.sound`. Its oracles and pinned Colab reproduction drivers are separate from the manuscript render. Failed source candidates and their full logs are retained in the programme evidence archive. Later, separately reproduced sources formalize the actual implicit-function curve construction at

`ef81f028ce6e7e8aaca7e9229eb00aa683752a41`.

The local quadratic-energy regular-sector obstruction uses

`9bb90cfd77ec5cea5c4447667fe8baf8c21429a1`.

The displayed sector proofs remain readable analytic arguments; their uniform consequence is not silently included in those four-endpoint reports. No theorem about a nonlinear sector is inferred from a full-space no-go alone.

Short labels and an external declaration map replace long wrapped Lean names in this manuscript. The companion source map records exact names, source SHAs, and evidence directories. Reproduction commands remain in the existing copyable notebooks rather than being split into misleading line fragments. The fixed-map Schur criterion, its counterexample calculation and consumer specialization above are analytic or source-inspection arguments; the whole manuscript is not described as machine checked.

The actual mean-weight finite/asymptotic theorem of Section 9 uses tested source

`ac60c0df885d8866329a839e2b4e73c90595d806`.

Its two own fresh Colab Linux clones 076/076b pass six exact standard-axiom reports and all nine stages. Compiled module SHA-256 is

`28fcc004795a1d8ab74ca87b5e67cdf90f557551bc02b1bb912da38324c26e50`.

Raw oracle and LF-source-manifest hashes also match; full records and the strict mutation-tested validator are preserved in the mean-weight reproduction package. The proof composes the actual character strip, its exact trace, positivity, local derivative-kernel bound and multiplier estimate; these are derived lemmas, not desired conclusions supplied as hypotheses. Nonnegative weights, fixed positive c and the quadratic penalty upper bound remain. The constructed quadratic-weight mean observation instead uses tested source

7f37d6437a8f69a1d3fddf3dc0bcfad74c6bf9b6.

It has two same-source fresh-clone reproductions 067/067b with six exact standard-axiom reports and matching compiled, raw-oracle and LF-source hashes. The formal output is an actual fixed finite-dimensional type; the displayed numeric dimension and nonoptimal path explanation are not additional oracle endpoints. The penalty examples and regularity tradeoff are analytic only. Classical quadratic-form, Fourier and trace mechanisms, as discussed in [7], establish no new general principle or priority here.

Model assessments and review context

This candidate was prepared by the Codex agent in the author's continuing programme chat, with access to earlier objections, source audits and supplied reproduction records. The served backend model identifier and reasoning setting are not exposed, and no technical isolation or scientific independence is claimed for this preparation.

The public v2 of this record has canonical PDF SHA-256

b37f2fe86b7a8a526ad02f8b28491c300a9ff72d5653536bc5532e8c9468ed8f.

A subagent explicitly authorized by the author and configured as `gpt-6-astra` assessed that exact eight-page artifact after first preserving the earlier candidate's objections. Actual backend identity and reasoning setting were not exposed. The workflow used the same account, provider and programme context; it was not blind or organizationally independent. The complete earlier and corrected-artifact reports remain in the v2 package. The corrected-artifact report regarded O1–O5 as resolved and retained O6: no independent complete Lean rebuild or dependency-closure audit. No numerical score or human peer review was supplied.

This v4 candidate adds the mean-weight theorem and quadratic boundary construction, with later separately anchored formal sector results and analytic penalty counterexamples. No previous report, score or approval is an assessment of these new bytes, and no new external model has reviewed this candidate. The preparing agent's displayed mathematical arguments, source inspections and eventual visual layout checks are not independent review. The formal results above are separately anchored; the whole manuscript is not machine checked.

Lluis Eriksson is both author and AIRR operator. The current public v3 is a Working paper, not admitted to the AIRR accepted collection, released on 31 October 2026 under CC BY 4.0 with CC0 metadata. Its exact PDF SHA-256 is

c9de923a66038ed66c3f68ecf8256b3f738a3ccf27e41f77e470bf7dc3efb5b9.

Public v1 and v2 remain preserved. This v4 candidate has not been deposited, accepted or published. No previous assessment, public license permission or exact-artifact deposit grant is represented as covering the new PDF bytes. No external model has assessed this candidate. The full prior reports and objections remain version-specific; no score or approval is inherited. No new human peer review or independent organizational review is claimed.

References

[1] L. Eriksson, *Exact-Kernel Fourier Families Obstruct Every Penalty Factored Through a Block Map in a Flat Lattice Gauge Form*, ai.viXra 2608.0015v1 (2026), AIRR

ARR-2026-7EPM3V3DP796BBRR v1.
https://airr.science/papers/ARR-2026-7EPM3V3DP796BBRR/versions/v1/. Historical import, not peer reviewed or assessed by AIRR; checked 30 October 2026.

[2] J. Dodziuk, *Finite-difference approach to the Hodge theory of harmonic forms*, American Journal of Mathematics **98** (1976), 79–104. https://doi.org/10.2307/2373615. Institutional source: https://academicworks.cuny.edu/gc_pubs/180/.

[3] A. Albert, *Conditions for Positive and Nonnegative Definiteness in Terms of Pseudoinverses*, SIAM Journal on Applied Mathematics **17** (1969), 434–440. https://doi.org/10.1137/0117041.

[4] D. Carlson, E. Haynsworth and T. Markham, *A Generalization of the Schur Complement by Means of the Moore–Penrose Inverse*, SIAM Journal on Applied Mathematics **26** (1974), 169–175. https://doi.org/10.1137/0126013.

[5] J. Dimock, *Covariant Axial Gauge*, arXiv:1410.2798v2 (2014). https://arxiv.org/abs/1410.2798v2.

[6] Y. Kudryashov and Mathlib contributors, *Implicit function theorem*, Mathlib source at commit 07642720480157414db592fa85b626dafb71355b. Pinned primary formal source, particularly HasStrictFDerivAt.to_implicitFunction and its exact level-set identity; inspected 30 October 2026. This attributes the classical parametrization; the separately pinned sector formalizations are identified in the evidence section.

[7] D. A. Spielman, *Graphs, Vectors, and Matrices*, Bulletin of the American Mathematical Society **54** (2017), 45–61. Author-hosted primary exposition: https://cs-www.cs.yale.edu/homes/spielman/PAPERS/bull1557.pdf.

[8] D. A. Levin and Y. Peres, with contributions by E. L. Wilmer, *Markov Chains and Mixing Times*, second edition, American Mathematical Society (2017), Chapter 13. Author-hosted text: https://pages.uoregon.edu/dlevin/MARKOV/markovmixing.pdf.