

Frequency interactions and curvature in Fourier density models

Scott Schubr

October 21, 2026

Abstract

We classify the intrinsic sectional-curvature sign at the uniform density for every three-frequency aligned-cosine affine density model on the unit circle. All sectional curvatures are strictly negative except for the spectrum $\{m, 2m, 3m\}$, which has the sharp attained range $\pi^2 m^2 [-25 - \sqrt{1649}, -25 + \sqrt{1649}]$. A finite Fourier Gauss formula identifies the normal frequency interactions responsible for this exception. In arbitrary finite dimension, successive frequency ratios exceeding $1 + \sqrt{2}$ force a diagonal Fourier curvature operator and give an explicit attained negative curvature range; the constant is optimal for this universal diagonality assertion. Two four-frequency examples show that the triple criterion does not extend by inspecting subtriples and that a positive curvature-operator direction need not represent positive sectional curvature. As a computational application, we certify corrected heat-kernel Sinkhorn metrics and their parameter derivatives through order two on a fixed two-parameter box at five bandwidths. The smallest bandwidth has a uniform relative metric-error improvement greater than 870.37. We prove bandwidth-uniform magnitude bounds for the cancellation-preserving interval iteration and provide exact and interval reproducibility code. The curvature theorems concern intrinsic finite-dimensional model geometry; the finite-band certificates do not assert a new all-bandwidth convergence rate or a complexity bound.

1 Introduction

Parametric density models inherit a transport metric from the Wasserstein space. Even when a one-dimensional transport problem admits a quantile representation, a restricted statistical model can be curved. The metric and general submanifold viewpoint are established [3, 4, 1]. Here we ask a more explicit question: which finite sets of density frequencies allow positive intrinsic sectional curvature at their common uniform member?

For a finite set S of distinct positive integers we study

$$f_\theta(x) = 1 + \sum_{j \in S} \theta_j \cos(2\pi jx), \quad x \in \mathbb{T} = \mathbb{R}/\mathbb{Z}, \quad \sum_{j \in S} |\theta_j| < 1.$$

The parameters enter affinely in the density, rather than through an affine deformation of the spatial variable. The model is equipped with the pullback Wasserstein/Otto metric. All explicit frequency classifications below are evaluated at $\theta = 0$ and use aligned cosine phases.

Our first result is a complete three-frequency classification: the harmonic ratio $1 : 2 : 3$ is the sole exception to strict negative sectional curvature. We derive it by an exact finite-frequency Gauss calculation, exhausting the possible off-diagonal resonances and their projection exceptions. We then give an arbitrary-dimensional sufficient criterion and sharpen its spacing consequence to the optimal universal Fourier-diagonality threshold $1 + \sqrt{2}$. This optimality does not concern the sign of sectional curvature itself.

The four-dimensional examples delimit the method. A spectrum without a harmonic triple can have positive sectional curvature, while another has strictly negative sectional curvature despite a positive direction of its curvature operator. The latter uses the classical Finsler–Thorpe four-form modification [5]; the explicit transport-density realization and its rational certificate are given here.

The computational part links these exact model metrics to finite-bandwidth heat-kernel Sinkhorn geometry. The continuum leading-bias and positive fourth-order correction are results of the predecessor paper [9], building on [8]; they are not reproved as new convergence theorems here. We execute the corrected tensor on an unchanged density box at five bandwidths, retain cancellations in interval arithmetic, and certify uniform matrix-error comparisons and metric ordering. The exact geometry supplies analytic benchmarks, while the interval certificates measure the finite-band approximation on a nonconstant parameter patch.

Related work and scope. Chen–Li give the primitive-over-density metric for one-dimensional statistical models [3]; Hamm et al. distinguish induced model geometry from ambient transport geometry [4]. Lott develops the general Wasserstein connection and curvature calculus [1, 2]. The recent circle calculation [6] concerns ambient curvature, which must not be identified with the intrinsic curvature computed here. Flatness classifications for other statistical families, such as location–scale–shape models [7], concern different parametrizations; our classification is arithmetic in cosine frequencies. Neither quantile pullback nor the Gauss equation is a new method of this paper. Our geometric contribution is the explicit sign classification, sharp frequency-spacing theorem and exact limitations in these affine density models. Verified Fourier continuation and interval derivatives also have substantial predecessors [11, 12]; our numerical contribution is the specialized cancellation estimate and the executed continuum metric certificates. No empirical-density, global-geodesic, training-performance or runtime result is claimed.

2 The affine Fourier metric and its Gauss formula

Let $\mathbb{T} = \mathbb{R}/\mathbb{Z}$ have normalized Lebesgue measure, put $k = 2\pi$, and let S be a finite set of distinct positive integers. We consider

$$f_\theta(x) = 1 + \sum_{j \in S} \theta_j \cos(kjx), \quad \sum_{j \in S} |\theta_j| < 1. \quad (1)$$

The tangent metric is the pullback of the Wasserstein metric: if $F_j = \cos(kjx)$ and $-(f_\theta \phi'_j)' = F_j$ with periodic potentials, then $G_{ij} = \int_{\mathbb{T}} f_\theta \phi'_i \phi'_j \, dx$. All curvature assertions below concern this finite-dimensional intrinsic metric at $\theta = 0$, unless a neighborhood is explicitly specified. They do not assert negative ambient Wasserstein curvature. The distinction between ambient and induced geometry is essential; see [1, 3, 4, 6].

Write $B_j = \sin(kjx)/(kj)$ and $H_S = \text{span}\{\sin(kjx) : j \in S\} \subset L^2(0, 1)$. The transport velocity associated with a density direction $\dot{\theta}$ at zero is $-\sum_j \dot{\theta}_j B_j$. Thus H_S is the tangent space in velocity coordinates, and its metric is the L^2 inner product. Let P_S denote its orthogonal projection.

Proposition 2.1 (Finite Fourier Gauss formula). *For $u, v \in H_S$ define $A_S(u, v) = (I - P_S)(uv)'$. For linearly independent u, v ,*

$$K_S(u, v) = \frac{\langle A_S(u, u), A_S(v, v) \rangle - \|A_S(u, v)\|_2^2}{\|u\|_2^2 \|v\|_2^2 - \langle u, v \rangle^2}. \quad (2)$$

In the orthonormal basis $e_j = \sqrt{2} \sin(kjx)$,

$$\begin{aligned} A_S(e_i, e_i) &= 2ki \sin(2kix) \mathbf{1}_{2i \notin S}, \\ A_S(e_i, e_j) &= k \left[(i+j) \sin(k(i+j)x) \mathbf{1}_{i+j \notin S} - |i-j| \sin(k|i-j|x) \mathbf{1}_{|i-j| \notin S} \right], \quad i \neq j. \end{aligned} \quad (3)$$

In particular,

$$K_S(e_i, e_j) = -\frac{k^2}{2} \left[(i+j)^2 \mathbf{1}_{i+j \notin S} + (i-j)^2 \mathbf{1}_{|i-j| \notin S} \right]. \quad (4)$$

Proof. Parity gives $\int_{\mathbb{T}} B_j / f_\theta = 0$, so the periodic potential derivative is $\phi'_j = -B_j / f_\theta$ and $G_{ij} = \int B_i B_j / f_\theta$. Define $C_\theta(x) = x + \sum_j \theta_j B_j(x)$ on $[0, 1]$ and its inverse X_θ . Positivity of $f_\theta = C'_\theta$ gives a smooth Hilbert-space immersion into $L^2(0, 1)$. Implicit differentiation yields $X_i = -B_i(X) / f_\theta(X)$; changing variables $t = C_\theta(x)$ gives $\langle X_i, X_j \rangle_{L^2(dt)} = \int B_i B_j / f_\theta = G_{ij}$. At zero, $X_i = -B_i$ and $X_{ij} = (B_i B_j)'$. In velocity coordinates the normal second derivative is therefore $(I - P_S)(uv)'$. The Gauss equation for an immersion into a flat Hilbert space gives (2). Differentiating $2 \sin(kix) \sin(kjx) = \cos(k(i-j)x) - \cos(k(i+j)x)$ gives (3). The pure normal forms have different doubled frequencies, hence zero mutual inner product when $i \neq j$. Orthogonality and $\int \sin^2(k\ell x) = 1/2$ give (4). No global identification of circle transport distances with fixed-origin quantile distances is used. $\square$

We use the following precise operator convention. Divide the normal forms in (3) by k and denote them by A_{ij} . On the orthonormal wedge basis indexed by $i < j$, define

$$M_{ij,\ell r} = \langle A_{i\ell}, A_{jr} \rangle - \langle A_{ir}, A_{j\ell} \rangle. \quad (5)$$

Then the curvature numerator on $u \wedge v$ is $k^2 \langle M(u \wedge v), u \wedge v \rangle$. This fixes all sign and scale conventions. In dimension three every bivector is decomposable; this fails in higher dimensions.

3 Complete three-frequency classification

Theorem 3.1. *Let $S = \{a, b, c\}$ with positive integers $a < b < c$. At the uniform density, every sectional curvature of (1) is strictly negative unless $S = \{m, 2m, 3m\}$ for some integer $m \geq 1$. In that exceptional case the exact attained range is*

$$\pi^2 m^2 [-25 - \sqrt{1649}, -25 + \sqrt{1649}], \quad (6)$$

so both signs occur. In particular, among aligned three-cosine families, $1 : 2 : 3$ is the unique frequency ratio permitting a nonnegative sectional curvature at zero.

Proof. In wedge order (ab, ac, bc) , put $D_{ij} = -M_{ij,ij}$. Formula (4) gives

$$D_{ij} = \frac{1}{2} [(i+j)^2 \mathbf{1}_{i+j \notin S} + (i-j)^2 \mathbf{1}_{|i-j| \notin S}].$$

Both outputs of (a, b) can belong to S only when $b - a = a$ and $a + b = c$. All other pairs have positive D_{ij} . Thus a zero diagonal entry occurs only for $S = \{a, 2a, 3a\}$.

For completeness we exhaust the off-diagonal possibilities. A shared output between $A_{ii}, A_{j\ell}$, or between $A_{ij}, A_{i\ell}$, requires an equality among $2i, j + \ell, |j - \ell|, i + j, |i - j|, i + \ell$, and $|i - \ell|$. Solving these equalities for three distinct ordered indices gives exactly

$$c = 2b - a, \quad c = 2a + b, \quad c = a + 2b.$$

Projection can delete such a coefficient but cannot add another possibility. The third equality cannot overlap either of the first two; the first two overlap exactly at $(a, 3a, 5a)$. Without any of these relations M is diagonal and strictly negative. The remaining cases follow directly from (3) as follows.

If $c = a + 2b$, only $M_{ab,bc}$ can be nonzero off the diagonal, and it equals $t = (b - a)(3b + a)/2$. Here $D_{bc} = b^2 + c^2 > 5b^2$ and $D_{ac} = a^2 + c^2$. Also $D_{ab} = a^2 + b^2$ unless $b = 2a$, in which case $D_{ab} = 9a^2/2$. Consequently $D_{ab} \geq b^2$ and $0 < t < 3b^2/2$, so $D_{ab}D_{bc} > t^2$.

If $c = 2a + b$ and $b \neq 3a$, only $M_{ab,ac}$ can be nonzero. For $b \neq 2a$ its value and diagonal magnitudes are

$$t = (b - a)(b + 3a)/2, \quad D_{ab} = a^2 + b^2, \quad D_{ac} = a^2 + c^2.$$

Since $0 < t \leq (a+b)^2/2 \leq D_{ab} < D_{ac}$, the coupled block is negative definite. If $b = 2a$, the corresponding values are $D_{ab} = t = 9a^2/2$, $D_{ac} = 17a^2$, $D_{bc} = 18a^2$, giving the same conclusion.

If $c = 2b - a$ and $b \notin \{2a, 3a\}$, only $M_{ab,bc}$ can be nonzero, with

$$t = (b-a)^2/2 - 2b^2, \quad D_{ab} = a^2 + b^2, \quad D_{bc} = b^2 + c^2.$$

The exact identity

$$D_{ab}D_{bc} - t^2 = \frac{1}{4}(b-a)^2[3(b-a)^2 + 8b^2] > 0$$

proves negative definiteness of this block. The remaining D_{ac} is positive; its precise value is not needed.

At the overlap $(a, 3a, 5a)$,

$$M/a^2 = \begin{pmatrix} -10 & 6 & -16 \\ 6 & -26 & 0 \\ -16 & 0 & -34 \end{pmatrix}.$$

Its leading principal minors are $-10, 224, -960$, so Sylvester's criterion gives negative definiteness.

Finally, for $(a, b, c) = (m, 2m, 3m)$,

$$M/m^2 = \begin{pmatrix} 0 & 0 & -8 \\ 0 & -8 & 0 \\ -8 & 0 & -25/2 \end{pmatrix}. \quad (7)$$

Its eigenvalues are -8 and $(-25 \pm \sqrt{1649})/4$, times m^2 . Every unit bivector in dimension three represents a plane, so the Rayleigh range, multiplied by k^2 , is exactly (6). This also proves attainment of both endpoints. $\square$

For an explicit positive plane in the exceptional family, take $u = \sin(kmx) + \sin(3kmx)$ and $v = \sin(2kmx)$. Their Gram determinant is $1/2$. Formula (2) has numerator $7k^2m^2/8$, hence curvature $7\pi^2m^2$. The two endpoints of (6) grow quadratically with m with opposite signs. Thus no frequency-independent upper or lower intrinsic curvature bound holds across these changing families, even at their common uniform density. This does not contradict estimates imposing spatial regularity or conditioning bounds.

4 Higher-dimensional sufficient conditions and sharp spacing

Proposition 4.1 (Unique normal outputs). *Call an unordered pair $\{i, j\} \subset S$, allowing $i = j$, a representation of a positive normal output $\ell \notin S$ if $\ell = i + j$ or $\ell = |i - j|$. If every normal output has at most one representing pair, then M is diagonal and negative semidefinite. It is negative definite if each distinct pair has at least one output outside S .*

Proof. For each normal sine frequency, its coefficient matrix in A_{ij} has either one diagonal entry or one symmetric off-diagonal pair. A single diagonal entry contributes zero to the Gauss form. A symmetric off-diagonal pair contributes a negative multiple of $(u_i v_j - u_j v_i)^2$. Summing over orthogonal normal frequencies proves the assertion and gives diagonal entries $-D_{ij}$. $\square$

Theorem 4.2 (Sharp universal spacing threshold). *Let $S = \{s_1 < \dots < s_d\}$, $d \geq 2$, and suppose $s_{r+1}/s_r > Q = 1 + \sqrt{2}$ for every r . Then M is diagonal, with entries $-(s_i^2 + s_j^2)$, and the exact attained sectional-curvature range is*

$$[-4\pi^2(s_{d-1}^2 + s_d^2), -4\pi^2(s_1^2 + s_2^2)]. \quad (8)$$

The constant Q is optimal for this universal diagonality conclusion. No optimality for negative sectional curvature is asserted.

Proof. Because successive ratios exceed $Q > 2$, no sum or nonzero difference of two spectrum elements, including a doubled element, belongs to S . Otherwise the largest term in a signed three-occurrence relation would exceed the sum of the other two. Thus both outputs of every distinct pair survive projection and $D_{ij} = i^2 + j^2$.

A nonzero off-diagonal entry of (5) requires a common output of two normal forms, hence a signed relation among four frequency occurrences. If equal frequencies cancel with opposite signs, the remaining two must also cancel, giving the same unordered wedge pair and a diagonal entry. A genuinely off-diagonal entry can therefore be considered after excluding such cancellations. Alternation excludes any pattern with three or four equal occurrences. Let D be the largest frequency and C the largest smaller distinct frequency. If D occurs twice with the same sign, $2D$ exceeds the sum of the other two terms. If D occurs once, the sum of the other three magnitudes is at most one of

$$C(1 + Q^{-1} + Q^{-2}), \quad C(2 + Q^{-1}), \quad C(1 + 2Q^{-1}),$$

according as all are distinct, C is repeated, or a smaller frequency is repeated. Each is at most QC , since $Q = 2 + Q^{-1}$. But $D > QC$, a contradiction. All off-diagonal entries vanish. For a unit decomposable bivector, squared wedge coordinates sum to one, so its curvature is a convex combination of the diagonal values $-k^2(s_i^2 + s_j^2)$. Coordinate planes attain the two endpoints. The Grassmannian of unoriented two-planes is connected, so every intermediate value occurs (the range is a singleton for $d = 2$).

For sharpness, fix $q < Q$ and choose a rational $r = b/a < Q$ sufficiently close to Q that $r > \max\{q, 1\}$ and $2 + 1/r > q$. Put $c = a + 2b$. Both successive ratios exceed q , while the three-frequency formula gives $M_{ab,bc} = (b - a)(3b + a)/2 > 0$. Thus no smaller constant forces universal diagonality. These triples are nevertheless strictly negatively curved by Theorem 3.1, demonstrating the limitation of the sharpness assertion. $\square$

5 Why triple tests and operator signs do not suffice

Adding modes changes P_S in the Gauss formula. Hence the curvature of a three-mode coordinate submanifold is not the curvature of its tangent plane inside a larger affine density manifold.

Proposition 5.1. *The spectrum $S = \{1, 3, 4, 5\}$ contains no $1 : 2 : 3$ triple but has a positive sectional curvature. Conversely, for $S = \{1, 4, 6, 7\}$ the curvature operator has a positive eigenvalue while every sectional curvature is strictly negative.*

Proof. For the first spectrum, use its orthonormal sine basis in increasing frequency order and take $u = e_1 + e_5/2$, $v = e_3$. Direct substitution into (3) gives the curvature numerator divided by k^2 equal to $11/2$, and the Gram determinant is $5/4$. Thus $K/k^2 = 22/5 > 0$.

For the second spectrum, in wedge order $(12, 13, 14, 23, 24, 34)$, where the labels refer to positions in the ordered spectrum, formula (5) gives

$$M = \begin{pmatrix} -17 & 21/2 & 0 & 0 & -55/2 & 0 \\ 21/2 & -25/2 & 0 & 0 & 0 & 0 \\ 0 & 0 & -32 & 0 & 0 & 0 \\ 0 & 0 & 0 & -52 & 0 & 0 \\ -55/2 & 0 & 0 & 0 & -65 & 0 \\ 0 & 0 & 0 & 0 & 0 & -169/2 \end{pmatrix}.$$

Let H be the symmetric matrix whose quadratic form is $w_{12}w_{34} - w_{13}w_{24} + w_{14}w_{23}$. The leading elimination pivots of $-(M - 24H)$, in the displayed order, are exactly

$$17, \quad 409/68, \quad 32, \quad 95/2, \quad 13401/818, \quad 80505/2978.$$

All are positive, proving negative definiteness of $M - 24H$. Every decomposable bivector satisfies the Plücker relation $\langle Hw, w \rangle = 0$, so M is strictly negative on all nonzero decomposable bivectors. On the other hand, the principal minor of $-M$ on positions $(1, 2, 5)$ is $-22455/8$. Thus M is not negative semidefinite and has a positive eigenvalue. The use of a four-form shift exploits the familiar distinction between operator and sectional positivity; compare [5]. Here the matrices give a fully explicit instance in an affine Fourier density family. $\square$

Corollary 5.2 (Persistence near uniform density). *For every nonexceptional three-frequency spectrum, strict negative sectional curvature persists on some open parameter neighborhood of zero. For the exceptional spectrum $\{m, 2m, 3m\}$, both positive and negative planes persist on a common neighborhood. Strict negativity likewise persists for the classes in Proposition 4.1 when its strictness hypothesis holds, and in Theorem 4.2.*

Proof. The metric and curvature depend smoothly on θ while $f_\theta > 0$. At zero, strict negativity is uniform on the compact Grassmannian of tangent two-planes. Continuity therefore gives a common neighborhood preserving it. In the mixed-sign case, fix one strictly positive and one strictly negative plane and continue these planes using their fixed coefficient vectors smoothly; continuity preserves their signs on a common neighborhood. $\square$

This is a non-effective consequence: it supplies no radius, certified parameter box, or finite-bandwidth sign claim. The classification is restricted to aligned cosine phases and three frequencies; Proposition 5.1 rules out extending it by checking all triples.

6 Certified approximation by corrected Sinkhorn metrics

The preceding results concern the limiting transport metric. We now give a finite-bandwidth application on a fixed, nonconstant two-mode parameter patch. This patch is not the three-mode mixed-curvature model. The continuum correction and its fourth-order smooth-input accuracy are inherited from [9]. The results here are effective finite comparisons, a cancellation-preserving realization and its interval magnitude estimate.

6.1 Population tensor and prescribed family

Put $H_h = \exp(h\partial_x^2/2)$, $h > 0$. For a positive probability density f , let $wH_h w = f$, set $v = H_h w$, and define

$$P_h \phi = v^{-1} H_h (w \phi).$$

This is reversible in $L^2(f dx)$ with constant ground state. Inverses at the singular root below are taken on the density-mean-zero complement, or in any equivalent constant gauge when only invariant pairings are used. For $F_i = \partial_i f$ and $Y_i = F_i / f$, let

$$X_{zi} = h(I - zP_h)^{-1} Y_i, \quad z \in \{1, -1, i, -i\}, \quad T_i = \frac{1}{2}(X_{1i} - X_{-1i}),$$

and put

$$C_i = \frac{1}{6}(3X_{1i} - 5X_{-1i} + X_{ii} + X_{-ii}), \quad V = \frac{f''}{4f} - \frac{(f')^2}{8f^2}, \quad K = fV''.$$

The original and corrected pullback metrics, in the normalization of [9], are respectively

$$U_{h,ij} = \int_{\mathbb{T}} F_i T_j dx, \tag{9}$$

$$C_{h,ij} = \int_{\mathbb{T}} F_i C_j dx - \frac{h^2}{6} \int_{\mathbb{T}} K T_i' T_j' dx. \tag{10}$$

The rational base in (10) is compact on the tangent complement; its scalar spectral filter is

$$\frac{h}{6} \left[\frac{3}{1-p} - \frac{5}{1+p} + \frac{2}{1+p^2} \right].$$

The four-root form is holomorphic in parameter jets; taking a real part before differentiating would not be the same analytic calculation. Spatial derivatives in (10) eliminate any constant gauge.

We use the same family and box at every bandwidth:

$$f_\theta = 1 + \theta_1 \cos(2\pi x) + \theta_2 \cos(4\pi x), \quad B = [.019, .021] \times [.00475, .00525]. \quad (11)$$

Here $F_i = \cos(2\pi i x)$, $i = 1, 2$. Parity gives the limiting metric

$$G_{ij}(\theta) = \int_{\mathbb{T}} \frac{B_i B_j}{f_\theta} dx, \quad B_i = \frac{\sin(2\pi i x)}{2\pi i}. \quad (12)$$

This formula is evaluated independently by an exact Fourier expansion of f^{-1} , rather than by the finite-band root solver.

6.2 Executed certificates and finite comparisons

For each $h \in \{.01, .005, .0025, .00125, .000625\}$, a full-box calculation and a center calculation enclose the three symmetric entries of each of U_h, C_h, G , with all six parameter multiindices of total order at most two. Thus the ten calculations contain 540 scalar enclosures. Center calculations retain nonzero formal derivative jets even though the degree-zero parameter intervals are points.

For a symmetric tensor A , define its invariant relative distortion by

$$\mathcal{E}_G(A; \theta) = \sup_{\xi \neq 0} \frac{|\xi^T (A - G) \xi|}{\xi^T G \xi}. \quad (13)$$

The coordinate spectral norm is denoted $\|\cdot\|_2$. Table 1 gives conservative lower bounds on the ratios of original to corrected errors, valid at every $\theta \in B$. The first norm uses the declared amplitude coordinates; the

Table 1: Uniform finite-band improvement factors. All displayed ratios are strict certified lower bounds.

h	$\ U_h - G\ _2 / \ C_h - G\ _2$	$\mathcal{E}_G(U_h) / \mathcal{E}_G(C_h)$
.01	> 4.28	> 4.31718
.005	> 14.51	> 14.61797
.0025	> 55.43	> 55.82108
.00125	> 218.66	> 220.48633
.000625	> 846.45	> 870.37762

quantity (13) does not depend on a change of parameter coordinates. Its enclosing constants are nevertheless computed in these coordinates.

The same-point error enclosures also prove the strict quadratic-form ordering

$$U_h < C_h < G \quad \text{on } B, \quad h = .01, .005, .0025, .00125. \quad (14)$$

For $h = .000625$ they prove $U_h < C_h$; the upper inequality $C_h < G$ is unresolved by these intervals. Failure of that certification is not a counterexample. At the first four bandwidths every nonzero tangent-vector squared length is increased toward the limiting value without overshoot.

At the smallest bandwidth we certify

$$\mathcal{E}_G(C_h; \theta) < 4.661841 \cdot 10^{-7} \quad (\theta \in B).$$

For any bound $0 < \epsilon < 1$ on this distortion, $(1 - \epsilon)G \leq C_h \leq (1 + \epsilon)G$. Integrating square roots along the same permitted curves gives length factors $\sqrt{1 - \epsilon}$ and $\sqrt{1 + \epsilon}$. Taking infima gives the same comparison for path distances restricted to this box. This does not compare unrestricted ambient Wasserstein distances.

All five corrected matrices have positive Gram margins and strictly negative intrinsic Gaussian curvature throughout B . At $h = .000625$, the eigenvalue lower bound exceeds .00309432 and curvature is enclosed in $[-181.893, -175.947]$. These bounds do not establish an improved curvature-error rate relative to G .

Table 2: Corrected whole-box finite-construction C^2 allowances. These bound numerical construction versus the exact finite- h tensor, not its bias versus G or the width of the box range.

h	allowance upper bound
.01	$5.568 \cdot 10^{-16}$
.005	$4.448 \cdot 10^{-15}$
.0025	$1.41952 \cdot 10^{-14}$
.00125	$2.52352 \cdot 10^{-14}$
.000625	$3.11072 \cdot 10^{-14}$

Each allowance in Table 2, and its center counterpart, is less than the corresponding exact h^4 . At the smallest bandwidth the center allowance is at most $1.63264 \cdot 10^{-14}$, whereas $h^4 = 1.52587890625 \cdot 10^{-13}$.

For illustration of finite-band behavior, successive corrected $(2, 2)$ -entry absolute-error halving ratios, uniformly on B , are enclosed by

$$[12.8481, 12.8515], \quad [15.0673, 15.0737], \quad [15.7419, 15.7711], \quad [15.7584, 16.1222].$$

The center C^2 maximum-entry error-norm halving ratios are approximately 7.323, 11.032, 14.061, 15.433: derivatives reach the asymptotic regime later than metric values. Center ratios are not uniform-box C^2 rates. Five discrete bandwidths do not prove a new asymptotic theorem or a statement for every intermediate bandwidth; the asymptotic theory is [9].

6.3 Retaining the shared heat cancellation

Writing $w = 1 + u$ and $e_n = \exp(-2\pi^2 h n^2)$, the scaled root numerator a_z solves

$$(H_h w) a_z - z w H_h a_z = h F_i, \quad X_{zi} = a_z / w.$$

At $z = 1$ we omit the zero equation and set the numerator's zero coefficient to zero. The omitted row vanishes by heat symmetry; its kernel changes X only by a constant and does not affect (9)–(10). The preconditioned update is

$$a_{\text{new}, n} = \frac{h F_{i, n}}{1 - z e_n} - \sum_{k+m=n} u_k a_m r_z(k, m), \quad r_z(k, m) = \frac{e_k - z e_m}{1 - z e_{k+m}}. \quad (15)$$

The full convolution is formed before Fourier projection, so omitted tails through degree $N + M$ remain available for certification. Combining the scalar multiplier before multiplying the interval jet $u_k a_m$ avoids bounding two copies of the same product independently.

For $z = 1$, $n = k + m \neq 0$,

$$|r_1(k, m)| \leq \max \left(1, \frac{|k^2 - m^2|}{(k + m)^2} \right) \leq 1 + 2|k|. \quad (16)$$

Indeed, for $A \geq 0, B > 0$, $(1 - e^{-tA})/(1 - e^{-tB}) \leq \max(1, A/B)$; factor the smaller heat exponential from the numerator. The last inequality uses $|k - m| \leq |k + m| + 2|k|$ and $|k + m| \geq 1$. For $z = -1, i, -i$, $|r_z| \leq 2$. When $k = 0$ use the exact scalar 1, and when $z = 1, k^2 = m^2$ use exact zero after excluding $n = 0$. The remaining singular-root multiplier can be evaluated by

$$r_1 = e_m \frac{\expm1[-2\pi^2 h(k^2 - m^2)]}{-\expm1[-2\pi^2 h n^2]}.$$

Stable exponential evaluation and avoidance of interval dependency are established tools [10, 12]; the specialization is the effective heat-root cancellation and its use in this certificate.

Theorem 6.1 (Uniform interval magnitudes). *Fix $\rho > 0$, a finite parameter-jet order, and multiplicative weights $\tau^\alpha > 0$. In exact-real natural interval arithmetic use complex rectangles with magnitude $M_1(I + iJ) = M(I) + M(J)$, where $M(I) = \max(|\inf I|, |\sup I|)$. Let $\|\cdot\|_A$ be the sum of coefficient magnitudes weighted by $e^{\rho|n|}\tau^\alpha$ and $\|\cdot\|_B$ the same norm with additional Fourier factor $1 + 2|n|$. Assume the zero-mean density perturbation jet g satisfies $\|g\|_B \leq \delta \leq 1/8$. Set*

$$r = \frac{1 - \sqrt{1 - 4\delta}}{2}, \quad \kappa = 2\sqrt{2}, \quad \Gamma = \max \left\{ \frac{h_0}{1 - e^{-2\pi^2 h_0}}, \sqrt{2}h_0 \right\}.$$

For $0 < h \leq h_0$, projected normalization and precombined scaled-root iterations started at zero obey, for every cutoff and iteration count,

$$\|u_{\text{iter}}\|_B \leq r, \quad \|a_{z,\text{iter}}\|_A \leq \frac{\Gamma \|F_i\|_A}{1 - \kappa r}, \quad \kappa r \leq \sqrt{2} - 1 < 1,$$

provided the forcing's zero Fourier coefficient vanishes. The singular root always excludes its zero mode.

Proof. The interval norms are submultiplicative: interval magnitudes, $e^{\rho|n|}$, $1 + 2|n|$ and the multiplicative jet weights have that property, and projection only deletes terms. The normalization update is

$$u_{\text{new}} = \Pi_N(I + H_h)^{-1}(g - uH_h u).$$

Both scalar heat multipliers and $(I + H_h)^{-1}$ are contractions, so $\|u_{\text{new}}\|_B \leq \delta + \|u\|_B^2 \leq \delta + r^2 = r$. For the root iteration, (16) bounds the singular convolution by $\|u\|_B \|a\|_A$. At the other roots, the complex-rectangle magnitude is at most $\sqrt{2}$ times the scalar complex modulus, giving the common bound $2\sqrt{2}\|u\|_A \|a\|_A \leq \kappa r \|a\|_A$. The source satisfies

$$\frac{h}{1 - e_n} \leq \frac{h}{1 - e_1} \leq \frac{h_0}{1 - e^{-2\pi^2 h_0}} \quad (n \neq 0),$$

since $x/(1 - e^{-x})$ increases. Stable-root source magnitudes are at most $\sqrt{2}h_0$. Thus $\|a_{\text{new}}\|_A \leq \Gamma \|F_i\|_A + \kappa r \|a\|_A$. Induction and a geometric sum finish the proof. $\square$

This theorem bounds interval magnitudes, not convergence of their widths. It assumes exact scalar heat/root coefficients. Finite-precision multiplier enclosures and rounding must be checked separately, as in the executed certificates. It does not guarantee fixed-digit accuracy for arbitrarily small h , or running-time or bit complexity.

A Continuum certificate and comparison contracts

We describe the bounds used by the supplied interval producer so that the numerical claims have an explicit continuum interpretation. In the continuum budgets of this appendix, A_ρ and B_ρ use the usual complex modulus for exact coefficients. Rectangle magnitudes provide computable upper bounds for these norms; they are not substituted for that modulus in the analytic stable-root operator constant 2. This convention is distinct from the exact-real natural-interval magnitude theorem above, whose common rectangle constant is $2\sqrt{2}$. For (11) use Fourier strip $\rho = .2$, jet weight $\tau = .00025$, and jets $a_\alpha = \partial^\alpha a / \alpha!$ of total order at most two. The input majorant is

$$\delta \geq .021(3e^{-2}) + .00525(5e^{-4}) + \tau(3e^{-2} + 5e^{-4}) < 1/8.$$

An outward upper bound is used in the actual calculation. With $\lambda = 2r < 1$, normalization truncation after L iterations obeys

$$\epsilon_w \leq r\lambda^L + \frac{\max_\ell \eta_\ell}{1 - \lambda}, \quad \eta_\ell = \|(I - \Pi_N)\Phi(u_\ell)\|_B, \quad \Phi(u) = (I + H_h)^{-1}(g - uH_h u). \quad (17)$$

Every omitted tail is a finite polynomial through $2N$. Natural interval evaluation encloses the finite iterations, including rounding; the analytic allowance bounds their difference from the true fixed point.

For a root at the final normalization polynomial u_t , let $c_z = \kappa_z \|u_t\|_B < 1$, $\kappa_1 = 1$, $\kappa_z = 2$ otherwise, and $b_z = \|hD_z^{-1}F_i\|_A$, $D_z = \text{diag}(1 - ze_n)$. A root budget after J projected iterations is

$$\epsilon_{\text{root}} \leq \frac{c_z^J b_z + \max_\ell \zeta_\ell}{1 - c_z},$$

where ζ_ℓ encloses the full omitted update tail through $M + N$. Put $a_w = \|u_t\|_B + \epsilon_w$ and require $c_{\text{true}} = \kappa_z a_w < 1$. Transfer to the true normalization gives

$$\epsilon_a \leq \epsilon_{\text{root}} + \frac{\kappa_z \epsilon_w (\|a_t\|_A + \epsilon_{\text{root}})}{1 - c_{\text{true}}}, \quad \epsilon_X \leq \frac{\epsilon_a}{1 - a_w} + \frac{\|a_t\|_A \epsilon_w}{(1 - a_w)(1 - \|u_t\|_B)}.$$

All denominators are checked strictly positive with outward arithmetic. These bounds compare exact polynomial iterates to continuum roots; complex-rectangle interval magnitudes in the finite implementation are handled by its coefficient norms, with the scalar operator constants applied to the exact mathematical perturbation problem.

Use the strip $\rho' = .1$ and $D = 2\pi/[e(\rho - \rho')]$. An A_ρ error ϵ_X gives spatial derivative error $D\epsilon_X$ in $A_{\rho'}$. For the differential coefficient we evaluate the exact identity

$$K = \frac{f''''}{4} - \frac{3f'f'''}{4f} - \frac{(f'')^2}{2f} + \frac{7(f')^2 f''}{4f^2} - \frac{3(f')^4}{4f^3}.$$

Inverse powers have geometric majorants because $\|f - 1\|_{A_{\rho'}} < 1$. If d_i bounds the derivative error of T_i and t_i bounds its trial derivative norm, the differential product error is bounded by

$$\frac{h^2}{6} \|K\|_{A_{\rho'}} (t_i d_j + t_j d_i + d_i d_j).$$

The base pairing adds the corresponding ϵ_{C_j} , with symmetrization averaging the two entry budgets. For the Q -node periodic trapezoidal rule and any integrand I ,

$$\left| Q^{-1} \sum_{q=0}^{Q-1} I(q/Q) - \int_I I \right| \leq e^{-\rho' Q} \|I\|_{A_{\rho'}}.$$

This follows directly from Fourier aliases at nonzero multiples of Q . Each jet-norm allowance ϵ yields actual derivative errors $\alpha! \tau^{-|\alpha|} \epsilon$, hence at most $2\tau^{-2} \epsilon$ through order two. The final whole-box controls are $N = 16, L = 44, M = 20, J = 48, Q = 512$; the two larger-band center calculations use $Q = 640$ and the algebraically equivalent original root expression. All other center calculations use the precombined expression with $Q = 512$. The intervals are evaluated with 30 decimal digits and serialized outward to 25 decimal digits.

A.1 Independent limiting integration and shared-parameter errors

Writing $g = f - 1$, the limiting derivatives follow from

$$\partial^\alpha G_{ij} = (-1)^{|\alpha|} |\alpha|! \int \frac{B_i B_j F_1^{\alpha_1} F_2^{\alpha_2}}{f^{|\alpha|+1}} dx.$$

The negative-binomial expansion of $(1 + g)^{-(m+1)}$, $m = |\alpha|$, is integrated by exact Fourier coefficient extraction. If $\delta_0 = \sup_B (|\theta_1| + |\theta_2|)$, truncating after degree L has tail at most

$$\frac{m!}{4\pi^2 i j} \binom{m+L+1}{m} \frac{\delta_0^{L+1}}{(1 - \delta_0)^{m+1}}.$$

The supplied benchmark uses $L = 12$ and exact decimal parameter endpoints.

For an entry $E(\theta) = A_h(\theta) - G(\theta)$ on $B = c + [-a_1, a_1] \times [-a_2, a_2]$, use the same parameter in the difference. The center error interval can be enlarged by

$$a_1 \sup |E_1(c)| + a_2 \sup |E_2(c)| + \frac{1}{2}(a_1^2 M_{11} + 2a_1 a_2 M_{12} + a_2^2 M_{22}), \quad (18)$$

where $M_{ij} \geq \sup_B |E_{ij}|$. Intersect this with the direct box difference. Gradient ranges use their center intervals enlarged by $a_1 M_{i1} + a_2 M_{i2}$. This is ordinary Taylor's theorem with integral remainder; no third derivatives are assumed. Curvature comparisons do not follow by applying this second-order Taylor argument to curvature, which would require additional metric derivatives.

Given entrywise absolute bounds b_{ij} , $\|E\|_2 \leq \max(b_{11} + b_{12}, b_{12} + b_{22})$. An absolute diagonal lower bound gives an operator-norm lower bound. The invariant bounds are sharper: choose a rational $\epsilon > 0$ so that both $\epsilon G + E$ and $\epsilon G - E$ are positive definite for every allowed interval entry, using positive diagonal and determinant lower bounds. This proves $\mathcal{E}_G(A_h) \leq \epsilon$. A uniform lower bound for the original relative error is

$$\max_{i=1,2} \frac{\inf_B |(U_h - G)_{ii}|}{\sup_B G_{ii}}.$$

The exact rational ratios give Table 1; floating display values are not used in the positivity tests. The strict sandwich (14) follows from Sylvester or Gershgorin bounds on the same-point signed error intervals.

B A joint residual-Gram bound for the rational base

The following standard inverse-trial identity provides a useful specialized certificate. It is not an executed improvement of the numerical budgets above and does not cover the differential sandwich.

Proposition B.1. *Let Q be a bounded selfadjoint operator on a Hilbert space with $0 \leq Q \leq (1 - \gamma)I$, $\gamma > 0$, and let Y have finitely many forcing columns in this space. Set*

$$A_1 = I - Q, \quad A_- = I + Q, \quad A_2 = I + Q^2, \quad B_0 = \frac{h}{6}(3A_1^{-1} - 5A_-^{-1} + 2A_2^{-1}).$$

For conforming trial-column maps X_j put $Z_j = Y - A_j X_j$, $W_j = Y^* X_j + X_j^* Y - X_j^* A_j X_j$, and $B_t = \frac{h}{6}(3W_1 - 5W_- + 2W_2)$. Then

$$-\frac{5h}{6}Z_-^* Z_- \leq Y^* B_0 Y - B_t \leq \frac{h}{2\gamma}Z_1^* Z_1 + \frac{h}{3}Z_2^* Z_2.$$

Proof. Expand $Z_j^* A_j^{-1} Z_j$ to obtain $Y^* A_j^{-1} Y = W_j + Z_j^* A_j^{-1} Z_j$. Now $0 \leq A_1^{-1} \leq \gamma^{-1}I$ and $0 \leq A_-^{-1}, A_2^{-1} \leq I$. Apply these inequalities with the signed coefficients 3, -5 , 2. $\square$

The full residual Gram retains cross-correlations between tangent columns. Residuals must be actual continuum Hilbert-space residuals on the true ground-state complement; nodal residuals alone are insufficient. Parameter derivatives require differentiated equations and their full forcings. The differential term in (10) still requires its own spatial-derivative error estimates. The stored computations do not retain the trial columns needed to claim a new executed certificate from Proposition B.1.

C Reproducibility and limits

The accompanying code uses exact rational Fourier arithmetic for the geometric identities and the four-dimensional certificates, with a separate coordinate metric-jet implementation. The numerical data include the ten full-jet certificates, an independently integrated limiting benchmark and exact rational comparison arithmetic. Reader instructions distinguish regenerating the interval producer from checking saved mathematical enclosures and regenerating their postprocessing.

The general theorems are proved analytically; finite enumerations are supplementary checks. The full interval solver is not asserted to have been independently rewritten. Distinct-control calculations and a separate ordinary dense solver supply complementary checks, and ordinary floating-point diagnostics are not used as interval evidence. The publication makes no necessary-and-sufficient classification for arbitrary finite spectra, no arbitrary-phase classification, no effective radius for sign persistence, and no all-bandwidth signed comparison. The exact examples and the finite-band tables are benchmarks for transport-information geometry, including contexts such as Sinkhorn natural-gradient methods [13]; they do not demonstrate an optimization or industrial performance improvement.

References

- [1] J. Lott, *Some geometric calculations on Wasserstein space*, arXiv:math/0612562v2 (2007). <https://arxiv.org/abs/math/0612562v2>.
- [2] W. Li, *Transport information geometry I: Riemannian calculus on probability simplex*, arXiv:1803.06360v2 (2020). <https://arxiv.org/abs/1803.06360v2>.
- [3] Y. Chen and W. Li, *Optimal transport natural gradient for statistical manifolds with continuous sample space*, arXiv:1805.08380v4 (2020). <https://arxiv.org/abs/1805.08380v4>.

- [4] K. Hamm, C. Moosmüller, B. Schmitzer and M. Thorpe, *Manifold learning in Wasserstein space*, arXiv:2311.08549v3 (2025). <https://arxiv.org/abs/2311.08549v3>.
- [5] R. G. Bettiol and R. A. E. Mendes, *Strongly nonnegative curvature*, *Mathematische Annalen* 368 (2017), 971–986. <https://doi.org/10.1007/s00208-016-1457-3>.
- [6] A. M. de Sá Gomes, C. S. Rodrigues and L. A. B. San Martin, *The Wasserstein space of the circle is not flat*, arXiv:2610.00625v1 (2026). <https://arxiv.org/abs/2610.00625v1>.
- [7] A. Fukushi, Y. Nakanishi-Ohno and T. Matsuda, *Flatness of location-scale-shape models under the Wasserstein metric*, *Information Geometry* (2026). <https://doi.org/10.1007/s41884-026-00206-6>.
- [8] S. Schubr, *Sharp convergence of heat-kernel Sinkhorn metrics*, *Luminary Research Archive*, edition 1 (2026). <https://luminaryoddities.com/researcharchive/sharp-convergence-of-heat-kernel-sinkhorn-metrics/v/1>.
- [9] S. Schubr, *Leading bias and positive correction of heat-kernel Sinkhorn metrics*, *Luminary Research Archive*, edition 1 (2026). <https://luminaryoddities.com/researcharchive/leading-bias-positive-correction-heat-kernel-sinkhorn-metrics/v/1>.
- [10] The FLINT contributors, *Arb: arbitrary-precision real numbers*, FLINT documentation, section “Exponentials and logarithms.” <https://flintlib.org/doc/arb.html>.
- [11] M. Gameiro, J.-P. Lessard and A. Pugliese, *Computation of smooth manifolds via rigorous multi-parameter continuation in infinite dimensions*, *Foundations of Computational Mathematics* 16 (2016), 531–575. <https://doi.org/10.1007/s10208-015-9259-7>.
- [12] S. M. Rump, *Verification methods: rigorous results using floating-point arithmetic*, *Acta Numerica* 19 (2010), 287–449. <https://doi.org/10.1017/S096249291000005X>.
- [13] Z. Shen, Z. Wang, A. Ribeiro and H. Hassani, *Sinkhorn Natural Gradient for Generative Models*, *Advances in Neural Information Processing Systems* 33 (2020). <https://proceedings.neurips.cc/paper/2020/file/122e27d57ae8ecb37f3f1da67abb33cb-Paper.pdf>.