

Compact Ordered Nodal Variation*

Anonymous

Abstract—Glory first be to The Holy One, Blessed be He, from whom all good things come. We introduce Compact Ordered Nodal Variation (CONV), a deterministic method for interpolation and matched multiresolution analysis/synthesis on Cartesian grids. CONV combines high-order local reconstruction, an exact factorwise variation bound, and matched conservative multiresolution transfer. The interpolant accepts arbitrary endpoint-aligned target lattices; dyadic analysis and integer-scale polyphase evaluation are specializations of that continuous operator. Within each Cartesian factor, ringing is suppressed by projecting the high-order proposal onto a signed current fibre whose derivative-sign variation cannot exceed that of the sampled differences.

Index Terms—Interpolation, resampling, variation diminishing, Bernstein polynomial, Eikonal coordinates, structure tensor, lifting scheme, reversible transform, ringing.

I. INTRODUCTION

Regular-grid interpolation is commonly posed as reconstruction by a fixed kernel. Nearest-neighbor and multilinear rules preserve range, but attenuate high spatial frequencies; cubic convolution and windowed-sinc kernels recover more of the passband by using signed lobes, thereby admitting new extrema near strong transitions [3], [2]. Edge-directed methods instead adapt the reconstruction to local image structure [12], [13]. Although their direction estimates improve selected geometries, neither a law for transported variation nor an inverse representation follows from the direction estimate alone.

Compact Ordered Nodal Variation (CONV) is obtained by treating the observed forward differences as signed currents on a regular Cartesian grid. Here “ordered” refers to their admitted sign lineage. A fourth-order interior two-jet yields a quintic proposal, which is projected in Euclidean current coordinates onto the observed order of variation and then integrated in the Bernstein basis. Interpolation consequently involves three distinct objects: the endpoint difference, its sign lineage, and the within-cell distribution of current. Keeping these objects separate is what allows fidelity to be improved without surrendering the variation law.

Indeed, total positivity shows that an admitted one-dimensional factor cannot synthesize a derivative sign pair absent from the sampled ledger, while cardinality preserves both the observed point values and the total current of each cell. The same predictor is embedded in a dyadic transform using positive block averages, such

that cell mass is preserved and the Nyquist component is annihilated. When the residual centered moments are retained, the resulting analysis and synthesis maps are exact inverses. Fixed-order Cartesian composition then carries the same construction to channel arrays.

In two dimensions, the two Cartesian factor orders are coupled by a source-conditioned Riemannian action. The nodal action identity reduces every source-site moment to a Kronecker value; therefore, the apparent frame bank and nested chart family collapse exactly to the reverse factor order. Structure-tensor anisotropy supplies the remaining scalar admission in $[0, 1]$, and the final operator is obtained as a convex blend of the two orders without source ownership or a source-specific raster. When a global continuation is required, the same nodal coordinate also yields the weighted-harmonic construction of Section XII.

Since CONV defines a continuous source-cell profile before selecting a target grid, arbitrary $N \rightarrow M$ synthesis merely evaluates that profile at the endpoint-aligned target coordinates. Dyadic cell-average analysis and integer-phase compilation are consequently specializations of one interpolant: the former supplies a matched reversible state, whereas the latter supplies a repeated phase bank for efficient evaluation.

II. RELATION TO EXISTING RECONSTRUCTION

A. Sampling and raster reconstruction

Classical sampling theory obtains exact bandlimited reconstruction from the sinc kernel [1]. In practice, finite implementations replace sinc by a windowed kernel, of which Lanczos reconstruction is a widely used instance [2]; cubic convolution instead trades spectral reach for shorter support and polynomial accuracy [3]. Linearity and translation invariance make the frequency response of these methods transparent. However, a nontrivial finite low-pass approximation generally involves signed coefficients, and its oscillatory response is consequently unconstrained by the sign topology of the sampled differences.

Raster reconstruction developed along a complementary engineering line in which local covariance is used to orient an edge-directed interpolant [12]. EASU, for example, evaluates a single-pass radial Lanczos approximation adapted in direction and anisotropy; since the resulting signed lobes can ring, the reference construction limits the output by the nearest four-sample range [13]. These methods make two-dimensional structure and independent target-pixel evaluation primary design concerns. CONV addresses the same transfer problem at an earlier representational stage: differential current is admitted before polynomial synthesis, rather than correcting a value after it has been synthesized.

*The abbreviation CONV invites an implicit implementation question: can the formal operator be compiled into a convolutional realization? Section XII answers that question.

B. Shape-preserving spline interpolation

Shape-preserving interpolation has a substantial independent history. Monotone cubic Hermite methods constrain nodal derivatives using neighboring secants [4], [5]. Costantini constructed local co-monotone splines of arbitrary degree from Bernstein polynomials [6]. Dougherty, Edelman, and Hyman derived derivative constraints and modification algorithms for shape-preserving cubic and quintic Hermite interpolation [7]. Ulrich and Watson later gave sharp necessary-and-sufficient positivity conditions for the quartic derivative of a monotone quintic Hermite segment [8]. MQSI uses those conditions with local first- and second-derivative estimates and a B-spline realization to construct a C^2 monotone quintic in linear time [9].

Thus monotone quintic interpolation and its bound on derivative sign changes are established results. In a co-monotone construction, each knot interval receives the sign of its own secant. CONV instead retains the order of the observed secant transitions while allowing a witnessed transition to occupy the nearby current slots selected by (20); consequently, the transition is not pinned to the sampling knot. The signed Bernstein fibre (21) is therefore a convex sufficient region, rather than the full quintic monotonicity region characterized in [8], and its particular geometry yields Euclidean admission, endpoint-current conservation, and the variation-diminishing proof within one finite operator.

Linear complexity alone does not make the execution structures identical. MQSI involves derivative refinement followed by construction of a B-spline representation [9]. At a fixed scale, Section XII instead reduces CONV to FIR proposal channels, one scalar-threshold admission, and a polyphase synthesis matrix. The resulting bounded local primitive is the point at which the shape-preserving lineage meets raster reconstruction.

Shape preservation on rectangular grids also predates this work. Carlson and Fritsch gave derivative constraints and an algorithm for monotone piecewise bicubic interpolation [10]; Costantini and Manni built local differentiable co-monotone surfaces from variable-degree Bezier operators [11]. CONV's Cartesian construction is instead an ordered composition of the same one-dimensional current operator. The factorwise theorem proved below is the property of that composition; it is not a claim of first bivariate shape preservation.

C. Total positivity and product representations

Bernstein bases are nonnegative, form a partition of unity, and are totally positive. Consequently, a Bernstein polynomial has no more sign changes than its ordered coefficient sequence [18], [19]. Variation diminishing results also exist for tensor-product Bernstein operators and special multivariate domains [20], [21]. These results do not supply a general arbitrary-path variation theorem for a tensor-product surface. We therefore state and prove the property used by grid transfer: variation cannot be created within any admitted factor. Transverse features formed by

composition are products of recorded factor currents, not unrecorded factor oscillations.

D. Adaptive and invertible transfer

Edge-directed interpolation estimates orientations from the values being resampled [12], [13]. CONV instead derives its admissible current order from the full signed differential lineage of the sampled field. Its multidimensional Eikonal geometry likewise derives its Riemannian action supports and directional admissions from the source current; it accepts no external orientation as an additional input.

Harten's multiresolution algorithms already represent a hierarchy of nested cell averages by the coarsest averages and the errors made when predicting finer-grid averages from the next coarser level [22]. Cell averages plus prediction residuals are therefore established multiresolution coordinates. CONV uses that architecture with a different predictor: the fine centered moment is predicted by the same admitted ordered-current reconstruction used for interpolation, and its zero-detail synthesis obeys the variation law proved in Section IX.

Lifting schemes convert a prediction rule into an invertible transform by storing its residual [23]. Direction-adaptive lifting has also combined content-oriented image prediction with perfect reconstruction [14]. CONV applies the lifting principle directly to transported variation. Point interpolation uses the admitted Bernstein current. Multiresolution analysis uses the cell average as its invariant state and predicts the complementary centered moment from the same admitted profile. The residual moment is the fine state not represented by zero-detail synthesis.

III. INTERPOLATION REQUIREMENTS

The CONV construction follows from six operator requirements, which together determine the polynomial degree, the admission rule, and the inverse representation used below.

A. Observation consistency

For an integer refinement s , the refined line has $(N - 1)s + 1$ sites. Let D_s select indices $0, s, 2s, \dots$. Observation consistency requires

$$D_s \mathcal{T}_s = I. \quad (1)$$

Constants and affine coordinate fields must also be reproduced:

$$\mathcal{T}_s(\alpha i + \gamma)(r/s) = \alpha r/s + \gamma. \quad (2)$$

Consequently, neither post-synthesis sharpening nor range correction may move an observed sample.

B. Endpoint-compatible current coordinates

The current coordinates are chosen such that their cell mass equals the observed endpoint difference. Thus, if d_h is the fine forward difference and B_s sums a block of s consecutive fine differences, one obtains

$$B_s d_h \mathcal{T}_s = d_c, \quad (3)$$

where $d_c y = (\delta_0, \dots, \delta_{N-2})$. Since the endpoints are cardinal, (3) follows by telescoping. It also gives the representational identification used throughout the paper: internal current mass and restricted output difference are the same observed quantity.

C. Variation order

Range preservation and absence of ringing are not identical: a function may remain inside the global sample range while creating an alternating derivative pair, whereas an existing extremum may move between samples without creating such a pair. CONV therefore controls the variation topology through the requirement

$$V((\mathcal{T}y)') \leq \text{var}(\delta_0, \dots, \delta_{N-2}). \quad (4)$$

from which it can be seen that no new alternating rise-fall or fall-rise pair is admitted, although an observed turning point need not remain pinned to a grid node.

D. Directness and reversibility

Direct synthesis involves a finite local polynomial evaluation and a deterministic finite admission step; no iterative fit or selected stopping condition is present. Embedding the same predictor in the standard analysis-prediction-residual lifting construction yields an analysis map $\mathcal{L}$ whose retained detail obeys

$$\mathcal{L}^{-1} \mathcal{L} = I. \quad (5)$$

E. Canonical degree

An endpoint two-jet contains the value, first derivative, and second derivative at both ends of an interval, hence six scalar constraints. The least polynomial degree carrying all six without an additional choice is five; differentiation then yields five Bernstein coefficients, which form the finite current state. The quintic degree is consequently fixed by the compact two-jet rather than selected as a smoothness parameter.

IV. PROBLEM AND NOTATION

Let $y = (y_0, \dots, y_{N-1})$, $N \geq 5$, be samples on a unit-spaced one-dimensional grid. Multichannel data are handled componentwise in a declared basis. Define

$$\delta_i = y_{i+1} - y_i, \quad 0 \leq i < N-1. \quad (6)$$

For a finite real sequence z , let $\text{var}(z)$ denote its number of sign changes after zero entries are deleted. For a piecewise function g , let $V(g)$ be the supremum of $\text{var}(g(x_1), \dots, g(x_m))$ over ordered points within open cells

and one-sided knot limits, again omitting zeros. The Bernstein basis of degree n is

$$B_j^n(u) = \binom{n}{j} u^j (1-u)^{n-j}, \quad 0 \leq u \leq 1. \quad (7)$$

For scale vector $\mathbf{s}$, the fixed Cartesian order is part of the operator:

$$\mathcal{T}_{\mathbf{s}} = \mathcal{T}_{s_d}^{(d)} \circ \dots \circ \mathcal{T}_{s_1}^{(1)}. \quad (8)$$

The reverse order is used by the corresponding analysis transform.

A. Continuous query coordinates

CONV defines a continuous piecewise-polynomial interpolant before a target lattice is chosen. For a target line of length M , endpoints are aligned by

$$x_r = \frac{r(N-1)}{M-1}, \quad 0 \leq r < M. \quad (9)$$

Integer refinement is obtained when $M = (N-1)s + 1$, in which case (1) applies exactly. Using the same coordinate map on each Cartesian axis makes arbitrary $N \rightarrow M$ evaluation the primary interpolation map; the dyadic transform of Section IX and the integer-scale phase bank of Section XII are, respectively, its conservative analysis/synthesis and compiled-evaluation specializations.

B. Operator factorization

The one-dimensional method factors as

$$y \xrightarrow{\mathcal{J}} a \xrightarrow{\mathcal{E}} \sigma \xrightarrow{\Pi_{\mathcal{C}}(\sigma, \delta)} c \xrightarrow{\mathcal{I}} Q \xrightarrow{\mathcal{B}} \mathcal{T}y. \quad (10)$$

Here $\mathcal{J}$ forms the compact two-jet proposal, after which $\mathcal{E}$ orders the observed sign lineages and $\Pi_{\mathcal{C}}$ performs Euclidean nearest-point admission in current coordinates. Finally, $\mathcal{I}$ integrates the admitted current and $\mathcal{B}$ evaluates the Bernstein curve. The equations below specify each arrow, such that value observations, differential topology, and polynomial evaluation remain separate throughout the construction.

V. COMPACT TWO-JET

The raw local model is obtained without fitted parameters. At an interior node $2 \leq i \leq N-3$, the two-jet is

$$m_i = \frac{y_{i-2} - 8y_{i-1} + 8y_{i+1} - y_{i+2}}{12}, \quad (11)$$

$$q_i = \frac{-y_{i+2} + 16y_{i+1} - 30y_i + 16y_{i-1} - y_{i-2}}{12}. \quad (12)$$

Second-order one-sided closure is used at the outer two nodes:

$$m_0 = \frac{-3y_0 + 4y_1 - y_2}{2}, \quad m_1 = \frac{y_2 - y_0}{2}, \quad (13)$$

$$q_0 = 2y_0 - 5y_1 + 4y_2 - y_3, \quad q_1 = y_0 - 2y_1 + y_2, \quad (14)$$

with the reflected formulas used at $N-2$ and $N-1$. Since (11)–(14) are exact on affine data, the same equations both close the boundary and preserve first-degree reproduction.

Using the two endpoint jets on cell i yields the six raw quintic Bezier controls

$$\begin{aligned} P_{i0} &= y_i, & P_{i1} &= y_i + \frac{m_i}{5}, \\ P_{i2} &= y_i + \frac{2m_i}{5} + \frac{q_i}{20}, & P_{i3} &= y_{i+1} - \frac{2m_{i+1}}{5} + \frac{q_{i+1}}{20}, \\ P_{i4} &= y_{i+1} - \frac{m_{i+1}}{5}, & P_{i5} &= y_{i+1}. \end{aligned} \quad (15)$$

The corresponding raw variation current is obtained by differencing adjacent controls,

$$a_{ij} = P_{i,j+1} - P_{ij}, \quad 0 \leq j < 5. \quad (16)$$

which, by telescoping, satisfies $\sum_{j=0}^4 a_{ij} = \delta_i$ identically.

VI. ORDERED LINEAGE AND CONVEX CURRENT ADMISSION

A. Lineage ledger

The lineage component carries the order of the sign transitions observed in the sampled differences. Each nonzero change of secant sign introduces one boundary in the flattened current coordinate, and these locations are admitted in increasing order,

$$0 = b_0 < b_1 < \dots < b_L < J, \quad (17)$$

such that the next location minimizes the squared sign-conflict objective subject to (17). A later lineage can therefore neither cross nor precede an earlier one.

Each nonzero coarse current carries its sign. A zero secant inherits the preceding nonzero sign; a leading zero run inherits the first following nonzero sign. If all secants vanish, the ledger is zero. Put $I = N - 1$, $J = 5I$, and flatten five current slots per cell. Let $\kappa_1 < \dots < \kappa_L$ be the coarse transition knots and set $b_0 = 0$. For transition ℓ , define

$$\mathcal{B}_\ell = \{\max(b_{\ell-1} + 1, 5\kappa_\ell - 4), \dots, \min(J - 1, 5\kappa_\ell + 4)\}, \quad (18)$$

$$\mathcal{N}_\ell = \{\max(0, 5\kappa_\ell - 5), \dots, \min(J - 1, 5\kappa_\ell + 4)\}. \quad (19)$$

For candidate b , let $\sigma_k(b)$ equal the coarse sign left of the knot for $k < b$, and the sign right of it otherwise. The admitted boundary is

$$b_\ell = \min \arg \min_{b \in \mathcal{B}_\ell} \sum_{k \in \mathcal{N}_\ell} [\min(0, \sigma_k(b) a_k)]^2. \quad (20)$$

The outer minimum fixes ties at the smallest index. Applying (18)–(20) from left to right consequently produces a unique sign ledger $\sigma_{ij} \in \{-1, 0, 1\}$. Thus the data supply precisely the discrete current signs that interpolation must transport, without selecting either a kernel family or a continuous tuning parameter.

B. Signed mass fibre

Once the ledger is fixed, the admissible current on each cell belongs to the closed convex fibre

$$\mathcal{C}_i = \{c \in \mathbb{R}^5 : \mathbf{1}^\top c = \delta_i, \quad \sigma_{ij} c_j \geq 0 \text{ for all } j\}. \quad (21)$$

and the admitted current is obtained by Euclidean projection,

$$c_i = \arg \min_{c \in \mathcal{C}_i} \|c - a_i\|_2^2. \quad (22)$$

Consequently, (22) is an orthogonal projection in the canonical Euclidean coordinates of the five current increments, rather than a componentwise clamp; whenever the raw two-jet already belongs to the signed mass fibre, it passes unchanged. Nonemptiness follows because the ledger contains the sign of δ_i , and strict convexity then gives a unique solution. Since the dimension is five, this solution is obtained directly by a scalar threshold. Indeed, a zero ledger implies constant source data and therefore $a_i = c_i = 0$. Otherwise $\sigma_{ij} \in \{-1, 1\}$, and, for $\lambda \in \mathbb{R}$, put

$$c_{ij}(\lambda) = \begin{cases} \max\{0, a_{ij} - \lambda\}, & \sigma_{ij} = +1, \\ \min\{0, a_{ij} - \lambda\}, & \sigma_{ij} = -1. \end{cases} \quad (23)$$

The equality multiplier is consequently a solution of

$$F_i(\lambda) := \sum_{j=0}^4 c_{ij}(\lambda) = \delta_i. \quad (24)$$

The function F_i is continuous, nonincreasing, and piecewise affine, with breakpoints only at $a_{i0}, \dots, a_{i4}$. On any interval between successive breakpoints, its active index set is

$$A(\lambda) = \{j : \sigma_{ij}(a_{ij} - \lambda) > 0\}, \quad (25)$$

and the affine mass equation yields the candidate root

$$\lambda_A = \frac{\sum_{j \in A} a_{ij} - \delta_i}{|A|}. \quad (26)$$

Sorting the five breakpoints exposes at most six such intervals. The unique admitted current is therefore obtained by testing (26) against them and applying the five clips in (23).

Proposition 1 (Exact scalar admission). *Equations (23)–(26) return the Euclidean projection (22). If (24) has more than one multiplier, every such multiplier gives the same current.*

Proof. The Karush–Kuhn–Tucker conditions for (22) give (23) for the equality multiplier λ . Feasibility is exactly (24). On every interval between successive breakpoints the active set is fixed, and solving the affine mass equation gives (26). Nonemptiness of $\mathcal{C}_i$ supplies a feasible sign for δ_i . If both ledger signs occur, then $F_i(\lambda)$ tends to $+\infty$ and $-\infty$ at the two ends of the real line. If only the positive sign occurs, $\delta_i \geq 0$ and the range of F_i is $[0, \infty)$; if only the negative sign occurs, $\delta_i \leq 0$ and its range is $(-\infty, 0]$. Continuity therefore gives a root in every case. The objective is strictly convex in c_i , so its current is unique even in a zero-mass case where a multiplier interval may remain. $\square$

Lemma 1 (Fixed-ledger contraction). *For a fixed ledger and cell mass, Euclidean admission is firmly nonexpansive:*

$$\|\Pi_{\mathcal{C}_i} a - \Pi_{\mathcal{C}_i} b\|_2^2 \leq (\Pi_{\mathcal{C}_i} a - \Pi_{\mathcal{C}_i} b)^\top (a - b). \quad (27)$$

Consequently $\|\Pi_{\mathcal{C}_i} a - \Pi_{\mathcal{C}_i} b\|_2 \leq \|a - b\|_2$.

Proof. Since $\mathcal{C}_i$ is nonempty, closed, and convex, the standard variational inequality for orthogonal projection gives (27); applying Cauchy–Schwarz then yields the contraction bound. $\square$

C. Synthesis

Synthesis now involves a single integration of the admitted current. Starting from $Q_{i0} = y_i$, put

$$Q_{i,j+1} = Q_{ij} + c_{ij}, \quad 0 \leq j < 5. \quad (28)$$

The endpoint constraint in (21) then gives $Q_{i5} = y_{i+1}$, and CONV cell synthesis is obtained as

$$(\mathcal{T}y)(i+u) = \sum_{j=0}^5 Q_{ij} B_j^5(u), \quad 0 \leq u \leq 1. \quad (29)$$

Equivalently, the entire representation is the anchored current state

$$\mathcal{A}(y) = (y_0, \{c_{ij}\}_{i,j}), \quad y_i = y_0 + \sum_{k < i} \sum_{j=0}^4 c_{kj}. \quad (30)$$

VII. MULTIDIMENSIONAL CONV CONSTRUCTION

A. Ordered factor composition

Let $Y^{(0)} = Y$. For axis order $1, \dots, d$, define

$$Y^{(k)} = \mathcal{T}_{s_k}^{(k)} Y^{(k-1)}, \quad \mathcal{T}_C Y = Y^{(d)}. \quad (31)$$

At stage k , a fresh two-jet and lineage ledger are constructed on every line parallel to axis k , so that all factor currents remain represented rather than being collapsed to a principal orientation. Since admission is nonlinear, the factors need not commute; the order declared in (31) is therefore part of the operator and is carried along by a permutation of the coordinate axes. The resulting $\mathcal{T}_C Y$ is the Cartesian factor proposal consumed by the source-conditioned construction below.

B. Axis-permutation equivariance

The same one-dimensional law is used on every factor. If P_π permutes the array axes and the declared factor order is permuted with them, then

$$\mathcal{T}_{\pi(s)} P_\pi = P_\pi \mathcal{T}_s. \quad (32)$$

Thus (32) gives axis-permutation equivariance when the operator order is carried with the axes. It does not assert rotational isotropy: because nonlinear factors need not commute, fixed x -then- y and y -then- x compositions may differ. In two dimensions, synthesis uses horizontal (axis 1) followed by vertical (axis 0), and analysis uses the reverse order.

C. Source-conditioned Eikonal factor transport

Applying the one-dimensional construction in two dimensions yields two Cartesian compositions. Let $\mathcal{T}_{xy}$ apply the horizontal factor and then the vertical factor, and let $\mathcal{T}_{yx}$ apply them in the reverse order. With $\mathcal{T}_C = \mathcal{T}_{xy}$ as the declared Cartesian proposal, CONV couples the two orders through one source-conditioned Riemannian action field, without constructing a separate raster for every source.

Let the source rectangle be $\Omega = [0, W-1] \times [0, H-1]$, with lattice sites x_p . For channel c , define

$$g_p^{(c)} = (D_x Y_p^{(c)}, D_y Y_p^{(c)})^T, \quad G_p = \sum_{c=1}^C g_p^{(c)} g_p^{(c)T}, \quad \omega_p = \text{tr } G_p. \quad (33)$$

Using the same two-jet stencils and boundary closures already present in the Cartesian factors for D_x, D_y , define the dimensionless nodal quadratic form

$$Q_p = \begin{cases} I + G_p/\omega_p, & \omega_p > 0, \\ I, & \omega_p = 0, \end{cases} \quad M_p = \frac{Q_p}{\det(Q_p)^{1/2}}. \quad (34)$$

The result is SPD with $\det M_p = 1$; scalar gradient magnitude has been removed, leaving only the measured directional shape. For x in a source cell, let $b_p(x)$ be its four bilinear barycentric weights and put

$$\widetilde{M}(x) = \sum_{p \in \text{cell}(x)} b_p(x) M_p, \quad M(x) = \frac{\widetilde{M}(x)}{\det(\widetilde{M}(x))^{1/2}}. \quad (35)$$

Since positive barycentric combinations preserve positive definiteness, (35) gives one continuous determinant-one SPD field on Ω .

For each source site $s_a = x_a$, $1 \leq a \leq A = HW$, the action is the intrinsic Riemannian distance

$$D_a(x) = \inf_{\substack{\gamma(0)=s_a, \gamma(1)=x \\ \gamma([0,1]) \subseteq \Omega}} \int_0^1 \sqrt{\dot{\gamma}(t)^T M(\gamma(t)) \dot{\gamma}(t)} dt. \quad (36)$$

Proposition 2 (Well-posed constrained action). *There exist constants $0 < \mu_- \leq \mu_+ < \infty$ such that*

$$\mu_- I \preceq M(x) \preceq \mu_+ I, \quad x \in \Omega. \quad (37)$$

Consequently, for the length metric d_M induced by (36),

$$\sqrt{\mu_-} \|x - z\|_2 \leq d_M(x, z) \leq \sqrt{\mu_+} \|x - z\|_2. \quad (38)$$

The constrained infimum is attained for every $x, z \in \Omega$. Hence $D_a(x) = d_M(s_a, x)$ is finite and continuous, vanishes only at s_a , and is single-valued even where several minimizing curves have the same length.

Proof. Continuity and positive definiteness of M on compact Ω give (37). The lower inequality in (38) follows by bounding the metric length of every admissible curve below by its Euclidean length; the upper inequality follows from the straight segment in the convex rectangle. Thus d_M is bi-Lipschitz equivalent to the Euclidean metric, so (Ω, d_M) is a compact length space. A compact length space is geodesic, and therefore the infimum is attained. $\square$

Equivalently, the same action is obtained as the nonnegative viscosity solution of

$$\nabla D_a(x)^T M(x)^{-1} \nabla D_a(x) = 1, \quad D_a(s_a) = 0, \quad (39)$$

where paths are constrained to the closed source rectangle. In practice, anisotropic fast marching provides a direct Cartesian evaluation of (36) [16]. Possible nonuniqueness of a minimizing curve introduces no branch choice in CONV, because (40) depends only on its minimum length.

Using these distances, the normalized inverse-square actions are

$$w_a(x) = \frac{D_a(x)^{-2}}{\sum_b D_b(x)^{-2}}, \quad w_a(x) \geq 0, \quad \sum_a w_a(x) = 1. \quad (40)$$

At a source site, continuous extension of the quotient assigns unit mass to its zero-action source site and zero mass to all others. Thus (40) is Shepard inverse-distance weighting [17] in the source-derived Riemannian distance; normalization removes the distance units and yields a direct all-source partition without selecting an owner.

Lemma 2 (Kronecker action reduction). *For every pair of source sites,*

$$w_a(x_p) = \delta_{ap}. \quad (41)$$

Consequently, for every source-indexed quantity F_p ,

$$\sum_p w_a(x_p) F_p = F_a. \quad (42)$$

In particular, the action-weighted support inertia about the source site is identically zero:

$$\sum_p w_a(x_p) (x_p - s_a)(x_p - s_a)^T = 0. \quad (43)$$

Proof. Since the Riemannian distance vanishes at its own source site and is strictly positive between distinct lattice sites, the zero-action extension of (40) gives (41). Substitution then yields (42) and (43). $\square$

The lemma removes an apparent geometric degree of freedom: any frame inferred from (43) lies necessarily in its repeated-eigenvalue branch. With the declared Cartesian tie $t = e_x$, $n = e_y$, the nonzero nested route is therefore

$$\mathcal{C}[\{\mathcal{C}[Y_{:,i}](v)\}_i](u) = \mathcal{T}_{yx} Y(u, v), \quad (44)$$

independently of the source index. Consequently, a source frame bank and its nested chart construction are algebraically redundant under the action definitions, and (44) need be evaluated only once.

What remains is to determine how strongly the reverse order is admitted. The current tensor already supplies the required directional energy, and its normalized vertical component gives

$$\eta_a = \begin{cases} \frac{(G_a)_{yy}}{\text{tr } G_a}, & \text{tr } G_a > 0, \\ \frac{1}{2}, & \text{tr } G_a = 0. \end{cases} \quad (45)$$

Such normalized tensor energies are standard directional measures [15]. Since $(G_a)_{yy} \geq 0$ and $(G_a)_{yy} \leq \text{tr } G_a$,

equation (45) gives $0 \leq \eta_a \leq 1$ directly, without an eigendecomposition, threshold, or fitted gain. The zero-current value is one half, while exchanging the Cartesian axes sends η_a to $1 - \eta_a$, exactly as it exchanges the two factor orders. The scalar reverse-order admission is consequently

$$\beta(x) = \sum_a w_a(x) \eta_a, \quad 0 \leq \beta(x) \leq 1. \quad (46)$$

and the complete two-dimensional CONV reconstruction is

$$\boxed{\mathcal{T}Y(x) = [1 - \beta(x)] \mathcal{T}_{xy} Y(x) + \beta(x) \mathcal{T}_{yx} Y(x)}. \quad (47)$$

Proposition 3 (Source-conditioned factor convexity). *For every target point, the two coefficients in (47) are non-negative and sum to one. Consequently $\mathcal{T}Y(x)$ belongs to the componentwise convex hull of the two Cartesian factor-order candidates. The operator preserves constants and reproduces affine coordinate fields whenever the constituent factors do.*

Proof. From (45), $0 \leq \eta_a \leq 1$; using (40) then gives $0 \leq \beta \leq \sum_a w_a = 1$, and the two coefficients in (47) sum to one. Since both constituent orders return the same constant or affine field, their convex combination returns it as well. $\square$

Proposition 4 (Linewise blend criterion). *Let $\gamma : [a, b] \rightarrow \Omega$ be a C^1 curve and, at every regular point, put*

$$A = \mathcal{T}_{xy} Y \circ \gamma, \quad B = \mathcal{T}_{yx} Y \circ \gamma, \quad \theta = \beta \circ \gamma. \quad (48)$$

Then the final reconstruction obeys the exact identity

$$(\mathcal{T}Y \circ \gamma)' = (1 - \theta)A' + \theta B' + \theta'(B - A). \quad (49)$$

Suppose $sA' \geq 0$ and $sB' \geq 0$ for a fixed $s \in \{-1, 1\}$. The blend has $s(\mathcal{T}Y \circ \gamma)' \geq 0$ exactly when

$$s\theta'(B - A) \geq -\{(1 - \theta)sA' + \theta sB'\}. \quad (50)$$

In particular, the readily checked bound

$$|\theta'| |B - A| \leq (1 - \theta)sA' + \theta sB' \quad (51)$$

is sufficient.

Proof. Applying the product rule to (47) gives (49). Multiplication by s and rearrangement then yields (50), while (51) bounds its left-hand term from below. $\square$

Proposition 5 (Pointwise convexity is not variation diminution). *Nonnegative pointwise blending of two monotone functions with common endpoints does not by itself preserve monotonicity.*

Proof. On $[0, 1]$, take

$$A(t) = t, \quad B(t) = t + t(1 - t), \quad \theta(t) = \frac{1 + \sin(4\pi t)}{2}. \quad (52)$$

Both A and B are nondecreasing and share their endpoint values. For $T = (1 - \theta)A + \theta B$, equation (49) gives

$$\begin{aligned} T'(0) &= \frac{3}{2} > 0, & T'(\tfrac{3}{4}) &= \frac{3}{4} - \frac{3\pi}{8} < 0, \\ T'(1) &= \frac{1}{2} > 0. \end{aligned} \quad (53)$$

Continuity therefore forces at least two derivative-sign transitions. Hence the convex-hull statement alone cannot prove a sign-pair bound. $\square$

The factorwise ringing-incapability theorem applies unconditionally to each Cartesian composition in its declared order. The final spatially varying blend adds no amplitude outside the componentwise interval spanned by those two admitted values; its linewise sign topology is governed by (50), not by convexity alone.

D. Componentwise multichannel interpolation

For $Y : \mathbb{Z}^d \rightarrow \mathbb{R}^C$, the coordinate map and factor order are shared, while two-jets, ledgers, and fibre projections are evaluated for each component in a declared basis. The Cartesian factor proposals are therefore componentwise. In two dimensions, (33) sums the channel gradient outer products and supplies one shared metric and scalar admission; no channel selects a separate factor order. Component currents remain individually admitted inside that shared geometry. The factorwise guarantees hold per component and are preserved by independent affine changes of channel units; they do not assert covariance under an arbitrary basis rotation that mixes channels.

E. Compact dependence and causal state

On interior cell i , the raw current depends on the six samples from y_{i-2} through y_{i+3} . Bernstein evaluation uses only that cell's five admitted coefficients. The lineage is formed in one causal pass. Over an exact zero-secant run, the pass carries the last nonzero sign; a leading run receives the first following sign. The numerical jet and polynomial are compact. Only the sign state can traverse a plateau, preventing that plateau from creating a false pair of fronts.

F. Intercell regularity

Every cell polynomial is smooth in its open interval and the assembled interpolant is continuous. Its one-sided derivatives at an interior node are

$$\begin{aligned} (\mathcal{T}y)'(i^-) &= 5c_{i-1,4}, \\ (\mathcal{T}y)'(i^+) &= 5c_{i,0}, \end{aligned} \quad (54)$$

$$\begin{aligned} (\mathcal{T}y)''(i^-) &= 20(c_{i-1,4} - c_{i-1,3}), \\ (\mathcal{T}y)''(i^+) &= 20(c_{i,1} - c_{i,0}). \end{aligned} \quad (55)$$

Consequently, CONV is C^1 at node i exactly when $c_{i-1,4} = c_{i,0}$, and it is C^2 there when the corresponding second condition in (55) also agrees. Since admission imposes neither equality, a witnessed change of differential state is retained rather than concealed merely to enforce intercell smoothness.

VIII. ANALYTICAL PROPERTIES

Proposition 6 (Cardinality and commutation). *For every input y , $(\mathcal{T}y)(i) = y_i$. For every scalar α and constant γ ,*

$$\mathcal{T}(\alpha y + \gamma) = \alpha \mathcal{T}(y) + \gamma.$$

Proof. Since $Q_{i0} = y_i$ and $Q_{i5} = y_{i+1}$, cardinality follows immediately. Translation leaves every current unchanged, while scaling multiplies the raw current, fibre, projection, and integrated controls by α ; when $\alpha < 0$, the ledger reverses with the current. The deterministic tie rules are invariant under both operations, which proves the stated commutation. $\square$

A. Reproduction and ringing incapability

Proposition 7 (Affine reproduction). *If $y_i = \alpha i + \gamma$, then*

$$(\mathcal{T}y)(x) = \alpha x + \gamma$$

throughout the sampled interval.

Proof. The finite differences, including the one-sided closures, give $m_i = \alpha$ and $q_i = 0$, and therefore $a_{ij} = \alpha/5$. Since this current already lies in its signed fibre, projection leaves it unchanged; (29) is consequently the degree-elevated affine polynomial. $\square$

Corollary 1 (Cartesian consistency). *The product operator (31) is cardinal on the nested source lattice and reproduces every affine coordinate field*

$$f(x) = \gamma + \sum_{k=1}^d \alpha_k x_k.$$

Proof. Each factor fixes its source nodes and exactly reproduces the affine restriction of f to every line parallel to that factor. Induction over (31) gives both statements. $\square$

Definition 1. *A factor ringing pair is a pair of consecutive sign transitions of the derivative synthesized inside one Cartesian factor that has no corresponding pair in that factor's admitted current ledger.*

Theorem 1 (Factorwise ringing incapability). *CONV synthesis cannot create a factor ringing pair. More precisely, with $c^\flat$ denoting the flattened current ledger,*

$$V((\mathcal{T}y)') \leq \text{var}(c^\flat) \leq \text{var}(\delta_0, \dots, \delta_{N-2}). \quad (56)$$

Proof. Differentiating (29) gives

$$(\mathcal{T}y)'(i+u) = 5 \sum_{j=0}^4 c_{ij} B_j^4(u). \quad (57)$$

Using total positivity of the Bernstein basis bounds the sign variation of (57) by that of its ordered coefficients. At a knot, nonzero one-sided derivative signs come from the terminal current coefficients of the incident cells. Those coefficients are consecutive in $c^\flat$. The cell bounds therefore give $V((\mathcal{T}y)') \leq \text{var}(c^\flat)$. Since admission places one ordered boundary for each nonzero coarse sign transition, projection cannot introduce a sign outside the ledger and can delete a transition only by zeroing its incident coefficients. Therefore $\text{var}(c^\flat) \leq \text{var}(\delta)$, which completes (56). $\square$

Corollary 2. *If the admitted ledger is nonnegative or nonpositive on a cell, synthesis is monotone on that cell and remains within its endpoint range.*

The theorem is deliberately factorwise. A multivariate tensor-product surface may have transverse critical structure not visible on a single coordinate line. Ordered Cartesian composition transports a separate ledger at each stage. Such transverse crossings are formed from recorded factor currents; they are not factor ringing and are not excluded.

B. Consistency order

Proposition 8 (Uncontracted smooth-cell accuracy). *Restore physical mesh width h , and let $f \in C^6$. On an interior cell $2 \leq i \leq N-4$ whose raw current is already admissible,*

$$\sup_{0 \leq u \leq 1} |(\mathcal{T}f_h)(x_i + uh) - f(x_i + uh)| = O(h^5). \quad (58)$$

On the outer and near-outer cells the stated one-sided closure gives $O(h^3)$.

Proof. Taylor expansion of (11) and (12) gives

$$m_i = hf'(x_i) + O(h^5), \quad q_i = h^2 f''(x_i) + O(h^6).$$

Exact endpoint two-jets determine the quintic Hermite interpolant with $O(h^6)$ remainder. The first-derivative proposal therefore limits the combined interior order to $O(h^5)$. At the boundary, (13) has scaled error $O(h^3)$, while (14) has scaled error $O(h^4)$. If the raw current is admissible, projection is the identity and these orders pass unchanged to synthesis. $\square$

Away from a true critical point, a sufficiently resolved smooth monotone profile has constant secant sign and its raw current eventually lies in the same signed fibre. At a transition, CONV retains the global absolute bound below rather than assuming that admission is inactive.

C. Absolute stability and error

Let

$$|y|_{1,h} = \max_i \frac{|y_{i+1} - y_i|}{h}.$$

Set

$$\Delta_k = \frac{y_{k+1} - y_k}{|y|_{1,h}h}.$$

Substitution in (11)–(16) gives $a_i = A_t \Delta |y|_{1,h}h$. Here t denotes an outer, near-outer, or interior stencil. The exact squared row-sum certificates are

$$\sum_j \left(\sum_k |(A_t)_{jk}| \right)^2 \in \left\{ \frac{49}{100}, \frac{81}{100}, \frac{37}{45} \right\}. \quad (59)$$

The right boundary gives the reflected matrices. Since $\|\Delta\|_\infty \leq 1$, the largest value in (59) is $37/45 = 185/225$, and therefore

$$\|a_i\|_2 \leq \frac{\sqrt{185}}{15} |y|_{1,h}h. \quad (60)$$

The fibre contains a vector with one entry equal to δ_i and all others zero. Metric projection and the triangle inequality therefore give

$$\|c_i\|_2 \leq \left(1 + \frac{2\sqrt{185}}{15} \right) |y|_{1,h}h. \quad (61)$$

Because Bernstein synthesis is a convex combination of integrated controls,

$$|(\mathcal{T}y)(i+u) - y_i| \leq \beta |y|_{1,h}h, \quad \beta = \sqrt{5} \left(1 + \frac{2\sqrt{185}}{15} \right). \quad (62)$$

Proposition 9 (Amplitude excursion). *Let $m = \min_i y_i$, $M = \max_i y_i$, and $R = M - m$. Define*

$$\text{exc}(z; [m, M]) = \max \left\{ \left(\sup_x z - M \right)_+, \left(m - \inf_x z \right)_+ \right\}. \quad (63)$$

For one factor,

$$\text{exc}(\mathcal{T}y; [m, M]) \leq \beta R. \quad (64)$$

On every cell whose ledger has one sign, the sharper value is zero relative to that cell's endpoint range. Hence amplitude excursion can occur only in a cell carrying an admitted sign transition. After d ordered Cartesian factors,

$$\text{exc}(\mathcal{T}_s Y; [m, M]) \leq \frac{(1+2\beta)^d - 1}{2} R. \quad (65)$$

These bounds are independent of the number of source or target samples.

Proof. For one factor, $|y|_{1,h}h = \max_i |\delta_i| \leq R$, so (62) places every synthesized value in $[m - \beta R, M + \beta R]$, proving (64). A one-sign ledger makes (57) one-signed by nonnegativity of the Bernstein basis, so the cell is monotone between its observed endpoints. For the Cartesian bound, put $c = (m + M)/2$. Translation equivariance and (66) give $\|\mathcal{T}_s(Y - c)\|_\infty \leq (1+2\beta)^d R/2$. Subtracting the original half-range $R/2$ yields (65). $\square$

Proposition 10 (BIBO stability). *On one factor,*

$$\|\mathcal{T}y\|_\infty \leq (1+2\beta) \|y\|_\infty. \quad (66)$$

After d Cartesian factors the bound is $(1+2\beta)^d \|Y\|_\infty$.

Proof. $|\delta_i| \leq 2 \|y\|_\infty$. Apply (62) with $h = 1$, then compose the one-factor inequality. $\square$

Theorem 2 (Grid-size-independent bound). *Let f be L -Lipschitz and let its samples be transferred through d ordered Cartesian factors of mesh width at most h . Then*

$$|\hat{f}(x) - f(x)| \leq E_d L h, \quad E_d = (1+\beta) \frac{(1+2\beta)^d - 1}{2\beta}, \quad (67)$$

and

$$\text{MSE}(\hat{f}, f) \leq E_d^2 L^2 h^2. \quad (68)$$

The constants do not depend on lattice cardinality or the number of queried points.

Proof. From (62), one obtains the one-factor amplification. Thus, if $E_k L h$ bounds the error after k factors, the next factor obeys

$$E_{k+1} = (1+2\beta)E_k + 1 + \beta, \quad E_0 = 0.$$

Solving this recurrence yields (67); squaring and averaging then yields (68). $\square$

Within a chamber having fixed lineage boundaries and active clip sets, the operator is affine and hence Lipschitz. Across chamber boundaries, (61)–(68) provide the global absolute bound without assuming differentiability of the discrete selection map.

IX. RESTRICTION AND MATCHED MULTIREOLUTION GEOMETRY

A. Point-current restriction

Point interpolation and raster analysis involve different restriction functionals. Let d_h be the forward difference on an interpolated line, and let B_s sum each block of s consecutive fine differences. Using cardinality gives

$$\begin{aligned} (B_s d_h \mathcal{T}_s y)_i &= \sum_{r=0}^{s-1} [(\mathcal{T}_s y)_{is+r+1} - (\mathcal{T}_s y)_{is+r}] \\ &= y_{i+1} - y_i. \end{aligned} \quad (69)$$

Consequently,

$$D_s \mathcal{T}_s = I, \quad B_s d_h \mathcal{T}_s = d_c, \quad (70)$$

where D_s selects source sites and d_c is the source forward difference. Thus (70) identifies the conserved mass of the interpolating current; site selection itself is not used as an antialiased raster downscale.

B. Conservative scale coordinates

The centered-moment construction below uses the fixed three-point Gauss–Legendre nodes and weights

$$\xi = (-\sqrt{3/5}, 0, \sqrt{3/5}), \quad \mathbf{w} = (5/9, 8/9, 5/9). \quad (71)$$

At a multiresolution interface, raster entries are interpreted as cell averages. Accordingly, for an even fine line $x = (x_0, \dots, x_{2N-1})$, define the parent average and centered child moment by

$$z_i = \frac{x_{2i} + x_{2i+1}}{2}, \quad (72)$$

$$\mu_i^* = \frac{x_{2i+1} - x_{2i}}{2}. \quad (73)$$

The resulting analysis low-pass is positive and normalized, with transfer function

$$H(\omega) = \frac{1 + e^{-i\omega}}{2} = e^{-i\omega/2} \cos(\omega/2), \quad (74)$$

which yields $H(\pi) = 0$ exactly.

The centered-moment proposal is obtained from the already defined CONV profile, rather than from a second interpolator. Using the Gauss–Legendre nodes and weights in (71), for $1 \leq i < N - 1$, put

$$\begin{aligned} \bar{z}_i^- &= \frac{1}{2} \sum_{k=0}^2 w_k(\mathcal{T}z)(i - \tfrac{1}{4} + \tfrac{1}{4}\xi_k), \\ \bar{z}_i^+ &= \frac{1}{2} \sum_{k=0}^2 w_k(\mathcal{T}z)(i + \tfrac{1}{4} + \tfrac{1}{4}\xi_k), \end{aligned} \quad (75)$$

$$p_i = \frac{\bar{z}_i^+ - \bar{z}_i^-}{2}. \quad (76)$$

Since $\mathcal{T}z$ is quintic on each half-cell, three-point Gauss–Legendre quadrature is exact. At the two outer cells, the fixed closure is

$$p_0 = m_0/4, \quad p_{N-1} = m_{N-1}/4, \quad (77)$$

using the one-sided jet in (13); this proposal is exact for affine cell averages, for which $p_i = (z_{i+1} - z_i)/4$.

C. Admissible moment polytope

Put $\Delta_i = z_{i+1} - z_i$, fill zero signs by the ledger rule of Section V, and denote the resulting interval signs by s_i . When every Δ_i vanishes, set $\mathcal{P}(z, t) = \{0\}$ and $\hat{\mu} = 0$. Otherwise the filled signs are nonzero, and each cell is assigned a sign t_i , with $t_0 = s_0$, $t_{N-1} = s_{N-2}$, and, in the interior,

$$t_i = \begin{cases} s_{i-1}, & s_{i-1} = s_i, \\ \arg \max_{s \in \{s_{i-1}, s_i\}} s p_i, & s_{i-1} \neq s_i, \end{cases} \quad (78)$$

with the left sign selected on a tie. The second line therefore locates an existing turning point in one of its two adjacent child cells.

When μ is synthesized with zero detail, the fine currents are

$$g_{2i} = 2\mu_i, \quad g_{2i+1} = \Delta_i - \mu_i - \mu_{i+1}. \quad (79)$$

Consequently, for fixed t , the complete admissible set is

$$\mathcal{P}(z, t) = \{\mu : t_i \mu_i \geq 0, \quad s_i(\Delta_i - \mu_i - \mu_{i+1}) \geq 0\}. \quad (80)$$

Lemma 3 (Compact moment fibre). *For every nonconstant z and every assignment t produced by (78), $\mathcal{P}(z, t)$ is a nonempty compact convex polytope.*

Proof. The set is a finite intersection of closed half-spaces and contains zero, so only an unbounded ray remains to be excluded. Let h belong to its recession cone and put $v_i = t_i h_i \geq 0$. The homogeneous face inequalities are

$$(s_i t_i) v_i + (s_i t_{i+1}) v_{i+1} \leq 0. \quad (81)$$

Suppose $v_k > 0$ for the least such k . If $k = 0$, then $t_k = s_k$. If $k = N - 1$, the preceding inequality contradicts $v_k > 0$ because $t_{N-1} = s_{N-2}$. If $0 < k < N - 1$, the preceding inequality and $v_{k-1} = 0$ either contradict $v_k > 0$, or force $t_k = -s_{k-1} = s_k$. Thus every surviving case has $k < N - 1$ and $t_k = s_k$. Equation (81) on face k then forces $t_{k+1} = -s_k$ and $v_{k+1} \geq v_k$. By (78), this is possible only when $s_{k+1} = -s_k$ and $t_{k+1} = s_{k+1}$. Repetition propagates a positive nondecreasing sequence to the right endpoint, where $t_{N-1} = s_{N-2}$ contradicts the required opposite sign. Hence the recession cone is $\{0\}$, and the closed polyhedron is bounded. $\square$

On an increasing run, the constraints reduce to

$$\mu_i \geq 0, \quad \mu_i + \mu_{i+1} \leq \Delta_i. \quad (82)$$

The finite admission contains no limiter constant. First form

$$p_i^+ = t_i \max(t_i p_i, 0). \quad (83)$$

Let $\rho(i)$ denote the maximal sign lineage assigned to cell i , and let $\rho(\ell)$ denote the lineage of face ℓ . Its proposed consumption is

$$C_\ell = \sum_{\substack{j \in \{\ell, \ell+1\} \\ \rho(j) = \rho(\ell)}} |p_j^+|. \quad (84)$$

and, for each lineage, retain the scale

$$\alpha_\rho = \min \left(1, \min_{\substack{\ell: \rho(\ell) = \rho \\ C_\ell > 0}} \frac{|\Delta_\ell|}{C_\ell} \right). \quad (85)$$

where the inner minimum is $+\infty$ when the lineage has no positive consumption. Oppositely oriented moments at a transition relax rather than consume the face mass and therefore do not enter C_ℓ . The result is the largest independent lineage-ray scale satisfying every incident capacity bound, and it is denoted by

$$\hat{\mu}_i = \alpha_{\rho(i)} p_i^+ = (\mathcal{A}_z p)_i. \quad (86)$$

Lemma 4 (Admitted-moment bound). *The admission (86) satisfies*

$$\|\hat{\mu}(z)\|_\infty \leq \max_\ell |\Delta_\ell| \leq 2 \|z\|_\infty. \quad (87)$$

Proof. Every cell i has an incident face ℓ in its assigned lineage, as forced at the endpoints and by the definition of t_i in the interior. Consequently $C_\ell \geq |p_i^+|$. If $C_\ell = 0$, then $\hat{\mu}_i = 0$. Otherwise (85) gives

$$|\hat{\mu}_i| = \alpha_{\rho(i)} |p_i^+| \leq |\Delta_\ell|.$$

Finally, $|z_{\ell+1} - z_\ell| \leq 2 \|z\|_\infty$, which proves the second inequality. $\square$

D. Exact conservative lifting pair

The detail coordinate is obtained as the child moment not represented by the transported coarse variation,

$$r_i = \mu_i^* - \hat{\mu}_i(z). \quad (88)$$

so analysis is $\mathcal{L}x = (z, r)$, while synthesis is

$$x_{2i} = z_i - \hat{\mu}_i(z) - r_i, \quad (89)$$

$$x_{2i+1} = z_i + \hat{\mu}_i(z) + r_i. \quad (90)$$

Theorem 3 (Conservative CONV lifting). *For every fine line of length $2N$, $N \geq 5$, in exact arithmetic,*

$$\mathcal{L}^{-1} \mathcal{L}x = x, \quad R \mathcal{L}^{-1}(z, r) = z, \quad (91)$$

where R is the positive restriction (72). If $r = 0$, the synthesized current has no more nonzero sign changes than Dz .

Proof. Substituting (88) into (89)–(90) returns $z_i \mp \mu_i^* = x_{2i}, x_{2i+1}$. Their average is z_i , proving both identities. By (80), the nonzero fine-current signs are a subsequence of

$$t_0, s_0, t_1, s_1, \dots, s_{N-2}, t_{N-1}.$$

Within a sign run the entries agree. At a transition, t_i equals one of its two adjacent signs; that triple therefore

contains one transition, not two. Deleting zero currents cannot increase the count. $\square$

The transform consequently distinguishes exact recovery from discarded-detail reconstruction. Retaining r recovers every fine cell, whereas setting $r = 0$ retains cell mass and transported variation while omitting the unpredicted centered moment. For two detail states on the same coarse raster, one obtains

$$\|\mathcal{L}^{-1}(z, r) - \mathcal{L}^{-1}(z, r')\|_\infty = \|r - r'\|_\infty. \quad (92)$$

Proposition 11 (Multilevel max-norm stability). *The positive restriction satisfies $\|R\|_\infty \leq 1$. A J -level synthesis obeys*

$$\|x\|_\infty \leq 3^J \|z^{(J)}\|_\infty + \sum_{j=1}^J 3^{j-1} \|r^{(j)}\|_\infty. \quad (93)$$

The constant depends on the declared depth, not the line length.

Proof. Since the two restriction weights are positive and sum to one, $\|R\|_\infty \leq 1$. At one synthesis level, using (89)–(90) together with (87) gives

$$\|x^{(j-1)}\|_\infty \leq 3 \|x^{(j)}\|_\infty + \|r^{(j)}\|_\infty.$$

Iterating this recurrence from $x^{(J)} = z^{(J)}$ yields (93); induction on the one-level theorem likewise shows that ordered sign variation is nonincreasing at every zero-detail level. $\square$

E. Two-dimensional moment geometry

For a fine 2×2 block, introduce one invariant mean and three centered moments,

$$\begin{aligned} z &= \frac{1}{4}(x_{00} + x_{01} + x_{10} + x_{11}), \\ \mu_x &= \frac{1}{4}(-x_{00} + x_{01} - x_{10} + x_{11}), \\ \mu_y &= \frac{1}{4}(-x_{00} - x_{01} + x_{10} + x_{11}), \\ \mu_{xy} &= \frac{1}{4}(x_{00} - x_{01} - x_{10} + x_{11}). \end{aligned} \quad (94)$$

from which the inverse is obtained as

$$\begin{aligned} x_{00} &= z - \mu_x - \mu_y + \mu_{xy}, & x_{01} &= z + \mu_x - \mu_y - \mu_{xy}, \\ x_{10} &= z - \mu_x + \mu_y - \mu_{xy}, & x_{11} &= z + \mu_x + \mu_y + \mu_{xy}. \end{aligned} \quad (95)$$

Thus the four child cells are replaced without redundancy. The horizontal and vertical proposals are obtained by applying (75) along their corresponding factors, and applying the same half-cell functional to the horizontal proposal along the vertical factor yields the mixed proposal.

Inside a block, the horizontal face currents are

$$2(\mu_x - \mu_{xy}), \quad 2(\mu_x + \mu_{xy}), \quad (96)$$

and the vertical face currents are

$$2(\mu_y - \mu_{xy}), \quad 2(\mu_y + \mu_{xy}). \quad (97)$$

For horizontally adjacent blocks A, B , put $\Delta_x = z_B - z_A$. Their upper and lower cross-face currents are

$$\begin{aligned} h_{AB}^- &= \Delta_x - \mu_{x,A} - \mu_{x,B} + \mu_{y,A} - \mu_{y,B} + \mu_{xy,A} + \mu_{xy,B}, \\ h_{AB}^+ &= \Delta_x - \mu_{x,A} - \mu_{x,B} - \mu_{y,A} + \mu_{y,B} - \mu_{xy,A} - \mu_{xy,B}. \end{aligned} \quad (98)$$

For vertically adjacent blocks A, C , put $\Delta_y = z_C - z_A$. Their left and right currents are

$$\begin{aligned} v_{AC}^- &= \Delta_y + \mu_{x,A} - \mu_{x,C} - \mu_{y,A} - \mu_{y,C} + \mu_{xy,A} + \mu_{xy,C}, \\ v_{AC}^+ &= \Delta_y - \mu_{x,A} + \mu_{x,C} - \mu_{y,A} - \mu_{y,C} - \mu_{xy,A} - \mu_{xy,C}. \end{aligned} \quad (99)$$

Let t_x, t_y be block signs assigned by (78) on each coarse row and column, and let s_x, s_y be the corresponding coarse face signs. The face-current admissible set is

$$\begin{aligned} t_x(\mu_x \pm \mu_{xy}) &\geq 0, & s_x h^\pm &\geq 0, \\ t_y(\mu_y \pm \mu_{xy}) &\geq 0, & s_y v^\pm &\geq 0. \end{aligned} \quad (100)$$

This admissible set is a finite intersection of affine half-spaces containing the zero moment field. Admission first enforces the internal cone

$$t_x \mu_x \geq |\mu_{xy}|, \quad t_y \mu_y \geq |\mu_{xy}|. \quad (101)$$

For a cross face written $g = \Delta + r_A + r_B$, only opposed proposal mass consumes the witnessed capacity. Define

$$\begin{aligned} w_A &= \max(-sr_A, 0), & w_B &= \max(-sr_B, 0), \\ \beta &= \min\left(1, \frac{|\Delta|}{w_A + w_B}\right), \end{aligned} \quad (102)$$

with $\beta = 1$ when the denominator is zero. Both incident block scales are bounded by β , and each block retains the minimum bound from all of its faces. Since helpful components are not required for feasibility, a later reduction on another face cannot violate an admitted inequality. Consequently, (101)–(102) yield a finite local box inside (100).

Every pair of fine rows has a horizontal sign sequence interleaving $t_x, s_x, t_x, s_x, \dots$; every pair of fine columns has the analogous vertical sequence. The one-dimensional sign-count proof applies to each. The two-dimensional restriction response is

$$H(\omega_x, \omega_y) = e^{-i(\omega_x + \omega_y)/2} \cos(\omega_x/2) \cos(\omega_y/2), \quad (103)$$

which annihilates both axial Nyquist lines and their checkerboard intersection. Exact reconstruction is then obtained by restoring the three residual moment fields in (94).

F. Multilevel state

Repeated analysis produces the state

$$\mathcal{L}^{(J)} x = \left(z^{(J)}, r^{(J)}, r^{(J-1)}, \dots, r^{(1)} \right). \quad (104)$$

In one dimension, $r^{(j)}$ is a centered-moment field; in two dimensions, it contains the horizontal, vertical, and mixed residual moments. Analysis continues only while every resulting coarse axis has at least five cells, and synthesis

restores the stored levels in reverse order. Induction on (91) then gives

$$(\mathcal{L}^{(J)})^{-1} \mathcal{L}^{(J)} = I. \quad (105)$$

For a one-dimensional fine line of length $2N$, the N parent averages and N residual moments are necessary and sufficient coordinates. For a two-dimensional $2N \times 2M$ raster, the NM means and three NM residual moment fields likewise account for every scalar degree of freedom.

X. DIRECT EVALUATION AND ENGINEERING PROPERTIES

A. Line analysis

For each component of a source line, direct evaluation involves the following fixed sequence.

- 1) Compute m_i, q_i from (11)–(14).
- 2) Form P_{ij} and the raw current a_{ij} .
- 3) Fill zero secant signs and locate the ordered boundaries (18)–(20).
- 4) Sort the five current breakpoints, locate the scalar interval, and apply (23)–(26).
- 5) Integrate c_i once and evaluate each requested point by (29) or de Casteljau recursion.

The boundary action compares at most nine candidate locations over at most ten current slots, while the projection dimension remains five. A line with N samples and C components therefore has $O(NC)$ analysis cost, $O(NC)$ state, and constant cost per target evaluation. Consequently, the compact form contains neither an image-sized matrix nor a global eigendecomposition or convergence loop.

For comparison, the unprojected two-jet uses the same derivative stencils and six quintic controls, but omits the ledger and scalar-threshold admission. A direct Lanczos-3 evaluation uses at most six taps per Cartesian factor, whereas EASU uses twelve two-dimensional fetches per target together with its direction and anisotropic-weight calculations. Separable PCHIP first constructs coupled nodal derivatives and then evaluates one cubic per factor. CONV's additional work is therefore concentrated in the ordered-boundary scan, one fixed five-key sort, at most six scalar interval tests, and five clips per cell and component. Cartesian implementations of CONV, the raw quintic, Lanczos, and PCHIP require a factor barrier; EASU instead evaluates a direct two-dimensional target rule. Section XI reports reference-CPU wall times in order to expose the scale of this implementation cost.

B. Parallel structure

During two-jet formation, ledger construction, projection, and evaluation, all lines of one Cartesian factor are independent and may be processed in parallel. A barrier is required only between factors, since the next factor analyzes the values produced by the preceding one. When several target lattices query the same source, the arrays $\{Q_{ij}\}$ or $\{c_{ij}\}$ may be cached.

C. Boundary evaluation

Equations (13) and (14) give the complete boundary rule, so padded samples are not required. At a nested source node, the stored observation may be returned directly. At a cell end, assigning the final control as $Q_{i5} = y_{i+1}$, while forming the interior controls from cumulative admitted current, preserves cardinality independently of the summation order used for those interior controls.

D. Exact arithmetic and finite precision

The identities in Sections VI–IX are algebraic. Direct floating-point evaluation, however, incurs ordinary rounding error, and an approximate numerical cycle does not prove (91). Bitwise reversible coding therefore requires a reversible number representation: for integer data, the lifting residual may be stored in an integer lifting realization; for floating data, an error-free sum/difference expansion may be stored with the detail when bit identity is required. Without that additional representation, the theorem remains an exact-arithmetic statement, and ordinary floating-point addition and subtraction are not claimed to be bitwise identical.

XI. SYNTHETIC EVALUATION

A. Declared analytic design

The evaluation uses a deterministic census of closed-form fields, without a natural image, interpolated surrogate, trained parameter, or selected favorable case. For a target field f , the source nodes form the exact even–even sublattice, and error is evaluated on the remaining target nodes after a four-target-node boundary crop:

$$\text{MSE}_H(\hat{f}, f) = \frac{1}{|H|} \sum_{x \in H} |\hat{f}(x) - f(x)|^2. \quad (106)$$

Thus the reference in (106) is the analytic field itself rather than another reconstruction method.

The one-dimensional census contains 56 profiles at source lengths 9, 17, 33, 65, 129: 32 sinusoidal carriers from four frequencies and eight phases, eight translated smooth transitions, eight Gaussian bumps, and eight quadratic chirps. The two-dimensional census contains 28 fields at source sides 9, 17, 33: smooth transitions and windowed carriers at six off-axis orientations and two subpixel phases, followed by two radial transitions and two crossing-wave phases. The discontinuity census contains a step and a box at 16 phases for each of five source lengths, giving 160 topology cases.

The one-dimensional census evaluates formal CONV and its raw two-jet quintic before Π_C . The two-dimensional census evaluates the bounded-work CONV* realization of Section XII and audits it against formal Eikonal CONV separately. Comparators are normalized endpoint-aligned Lanczos-3, endpoint-aligned bilinear interpolation, and the EASU construction specified below [2], [13]. Every method uses its stated fixed rule. Since the declared cases span many orders of magnitude, MSE is aggregated by geometric mean, while paired win–loss counts retain the individual case outcomes.

B. Generator specification

All analytic coordinates lie in $[0, 1]$. For $k = 0, \dots, 7$, let $\phi_k = k\pi/4$. The one-dimensional families are

$$\begin{aligned} f_{\text{car}}(x) &= 0.5 + 0.38 \sin(2\pi\nu x + \phi_k), \\ f_{\text{edge}}(x) &= 0.5 + 0.45 \tanh((x - c)/w), \\ f_{\text{bump}}(x) &= 0.08 + 0.84 \exp(-((x - c)/\sigma)^2), \\ f_{\text{chirp}}(x) &= 0.5 + 0.34 \sin(2\pi(0.5x + 2.5x^2) + \phi_k). \end{aligned} \quad (107)$$

Here $\nu \in \{0.75, 1.5, 2.25, 3\}$, $c \in \{0.44, 0.48, 0.52, 0.56\}$, $w \in \{0.025, 0.05\}$, and $\sigma \in \{0.045, 0.09\}$. The carrier uses all (ν, ϕ_k) pairs; edge and bump use all stated (c, w) and (c, σ) pairs; the chirp uses all ϕ_k .

For the two-dimensional fields, put

$$\begin{aligned} p_\theta &= x \cos \theta + y \sin \theta, \\ n_\theta &= (x - \tfrac{1}{2}) \cos \theta + (y - \tfrac{1}{2}) \sin \theta, \\ e(x, y) &= \sin^2(\pi x) \sin^2(\pi y), \\ r &= \sqrt{(x - 0.47)^2 + (y - 0.53)^2}. \end{aligned} \quad (108)$$

With $h = (N - 1)^{-1}$, the four families are

$$\begin{aligned} F_{\text{edge}} &= 0.5 + 0.43 \tanh((n_\theta - \delta)/0.035), \\ F_{\text{car}} &= 0.5 + 0.34e \sin(2\pi(3.25)p_\theta + \phi), \\ F_{\text{curved}} &= 0.5 + 0.42 \tanh((r - 0.27 - \delta)/0.032), \\ F_{\text{cross}} &= 0.5 + 0.19\{g_1 + g_2\}, \\ g_1 &= \sin(2\pi(2.75x + 0.9y) + \phi), \\ g_2 &= \sin(2\pi(-0.9x + 2.75y) - \phi). \end{aligned} \quad (109)$$

Here $\theta \in \{7.5, 22.5, 37.5, 52.5, 67.5, 82.5\}^\circ$. Edge fields use $\delta = \pm h/4$, carriers use $\phi \in \{0, \pi/2\}$, curved fields use $\delta = \pm h/4$, and crossing fields use $\phi \in \{0, \pi/2\}$.

For a source line of length N , topology phase $s_k = (k + 1/2)/16$ gives

$$\begin{aligned} y_j^{\text{step}} &= \mathbf{1}\{j \geq N/2 - 1/2 + s_k\}, \\ y_j^{\text{box}} &= \mathbf{1}\{|j - (N/2 - 1/2 + s_k)| \leq N/6\}. \end{aligned} \quad (110)$$

Here $k = 0, \dots, 15$. Equations (107)–(110), the listed grids, and the held-out rule in (106) determine the entire population without an external corpus.

C. Comparator definitions

Let $\text{sinc}(t) = \sin(\pi t)/(\pi t)$ and $L_3(t) = \text{sinc}(t) \text{sinc}(t/3)$ for $|t| < 3$, zero otherwise. Resampling n nodes to m nodes uses

$$\begin{aligned} x_r &= \frac{r(n-1)}{m-1}, \\ \rho &= \min\left\{1, \frac{m-1}{n-1}\right\}, \\ w_{rj} &= \frac{L_3(\rho(x_r - j))}{\sum_q L_3(\rho(x_r - q))}. \end{aligned} \quad (111)$$

Only in-domain j are summed; hence the boundary rule is support truncation followed by renormalization. This gives the endpoint-aligned Lanczos arm for the point-sample interpolation census. Bilinear uses the same endpoint coordinates with support restricted to the two neighboring

nodes, whereas the cell-centered conservative transform ablation is stated separately in (113)–(114).

The EASU arm is a binary64 CPU transcription of the FSR 1.0 EASU construction. Scalar fields are replicated to RGB. It uses the twelve EASU fetch locations, exact binary64 reciprocal and square root rather than the shader approximations, and no RCAS pass. Target (r, s) is evaluated at source coordinate $(r/2, s/2)$; this is the complete coordinate-constant specialization for the nested endpoint-aligned grid. Out-of-domain fetches are clamped to the nearest source pixel. The raw two-jet arm uses the same stencils, quintic controls, Cartesian order, and evaluation rule as CONV but sets $c_i = a_i$, omitting the ledger and Π_C . PCHIP appears only in the timing audit; SciPy’s `PchipInterpolator` is applied horizontally and then vertically on the same endpoint-aligned nested coordinates, without padding.

D. Resolution and paired fidelity

Table I compares formal one-dimensional CONV and two-dimensional CONV* with Lanczos-3, such that a ratio below one favors the corresponding CONV form. At nine nodes neither population is resolved, and Lanczos-3 has the lower aggregate error. From 17 nodes onward, however, CONV wins the majority of one-dimensional cases; at 129 nodes, it wins all 56. In two dimensions, CONV* has the lower geometric-mean error at 17 and 33 nodes. The paired count at 33 nodes is even because Lanczos retains a small advantage on steep transition fields, whereas CONV* has larger advantages on the resolved carrier and crossing fields. The aggregate ratio therefore records the magnitude of this distinction without discarding its paired count.

TABLE I
PAIRED POINT-FIDELITY RESULTS VERSUS LANCZOS-3

Source nodes	1-D point		2-D point	
	W-L	MSE ratio	W-L	MSE ratio
9	28-28	1.247	9-19	1.277
17	36-20	0.0929	19-9	0.800
33	44-12	0.00751	14-14	0.126
65	52-4	0.000481	–	–
129	56-0	0.0000275	–	–

The EASU arm establishes a different engineering point. At the 9×9 source size it has the lowest two-dimensional geometric-mean MSE, 2.38×10^{-3} . Its nested-source maximum error is 1.81×10^{-1} , whereas CONV*, Lanczos-3, and bilinear are cardinal to their stated floating-point precision. At 33×33 , CONV* gives 4.31×10^{-8} geometric-mean MSE, compared with 3.42×10^{-7} for Lanczos-3, 3.23×10^{-5} for bilinear, and 4.67×10^{-5} for EASU. These measurements separate interpolation fidelity from a reconstruction rule that is permitted to alter source nodes.

E. Ordered-current ablation

The raw two-jet arm isolates the admission map by retaining the compact quintic proposal and removing only Π_C . From Table II, it can be seen that admission does not obtain its topology result by replacing the high-order proposal: aggregate MSE is slightly lower on every reported

grid and becomes indistinguishable from the raw proposal in the resolved regime.

TABLE II
RAW TWO-JET VERSUS ADMITTED CONV GEOMETRIC-MEAN MSE;
THE 2-D COLUMNS USE CONV*

Source nodes	1-D		2-D	
	Raw	CONV	Raw	CONV
9	2.95×10^{-4}	2.50×10^{-4}	4.02×10^{-3}	3.87×10^{-3}
17	7.02×10^{-7}	6.88×10^{-7}	2.63×10^{-5}	2.48×10^{-5}
33	5.51×10^{-10}	5.49×10^{-10}	4.31×10^{-8}	4.31×10^{-8}

The distinction is decisive on the discontinuity census. The raw quintic has $S > 0$ in all 160 step/box cases and reaches a surplus of eight sign changes; CONV has $S = 0$ in all 160. Thus the compact jet supplies the resolved-field accuracy, while ordered-current admission removes the unwitnessed variation and, on this census, does not increase aggregate MSE.

F. Topology and excursion

For the discontinuity census, the measured topology quantity is

$$S = \max\{0, \text{var}(d_h \hat{f}) - \text{var}(d_c f)\}. \quad (112)$$

CONV has $S = 0$ in all 160 cases, agreeing with (56); bilinear also has zero surplus. The raw quintic and Lanczos-3 each have $S > 0$ in all 160 cases and reach eight additional sign changes. This isolates the effect of admission from that of the compact jet and distinguishes CONV from a sharp signed-lobe kernel. CONV’s maximum amplitude excursion on these step and box cases is 0.1008, compared with 0.1194 for Lanczos-3 and zero for bilinear. The zero-surplus result therefore does not conceal the separately measured motion of an existing extremum.

The conservative zero-detail transform ablation uses the dyadic state of Section IX. Fine arrays have 16, 32, or 64 cells per axis. The one-dimensional census contains one Nyquist sequence, eight phases of a 0.92π carrier, and eight phases of a quadratic chirp at each size. The two-dimensional census contains axial and checkerboard Nyquist arrays, eight half-plane interfaces over four angles and two offsets, and four chirp phases. Half-plane pixels are their analytic polygonal cell coverage; no supersampling constructs the reference.

Each method uses its own analysis and synthesis. CONV-C denotes (72) followed by zero-residual synthesis (89)–(90) and its two-dimensional moment form. Box-bilinear uses the same positive analysis and pixel-center linear synthesis. For Lanczos-3, if $\rho = M/N$, the target center is

$$u_j = \frac{j + 1/2}{\rho} - \frac{1}{2}, \quad \kappa = \min(1, \rho), \quad (113)$$

and its normalized in-domain weight is

$$w_{jn} \propto \text{sinc}(\kappa(u_j - n)) \text{sinc}\left(\frac{\kappa(u_j - n)}{3}\right), \quad (114)$$

$$|\kappa(u_j - n)| < 3.$$




Fig. 1. Resolution census against direct analytic target values. Each point is the geometric mean over the complete declared population at that source size. The coarse regime exposes underresolution; the resolved regime exposes the approximation order of each fixed operator.




Fig. 2. One-dimensional discontinuity topology. The plot counts every declared step or box phase that creates a derivative-sign pair absent from the samples.

The same definition is used in analysis and synthesis; finite boundaries are truncated and renormalized. MSE is evaluated against the original fine cell array after a four-cell boundary crop.

TABLE III
CONSERVATIVE TWO-DIMENSIONAL ZERO-DETAIL ABLATION

Fine side	CONV-C	L3	Box-bilin.	Edge ratio
16	2.78×10^{-2}	3.16×10^{-2}	4.57×10^{-2}	0.845
32	1.45×10^{-2}	1.64×10^{-2}	2.35×10^{-2}	0.807
64	8.66×10^{-3}	9.86×10^{-3}	1.41×10^{-2}	0.793

CONV-C has the lowest aggregate MSE at all three sizes. Its mean diagonal interface MSE is 15.5, 19.3, and 20.7 percent below Lanczos-3 as size increases. The chirp comparison is phase dependent: CONV-C has the lower four-phase mean at side 16, while Lanczos-3 is lower by 0.50 and 1.14 percent at sides 32 and 64. Thus the aggregate result does not conceal the remaining high-frequency case.

The conservative restriction maps axial Nyquist and checkerboard inputs to zero exactly, agreeing with (74) and (103). CONV-C has zero horizontal and vertical sign surplus in all 42 two-dimensional cases. Six retained-detail audits—one line and one three-channel raster at each size—have maximum binary64 round-trip residual 8.88×10^{-16} ; the exact identity is proved in Section IX.

The preceding census isolates the reversible conservative state. The technical fields in Fig. 4 instead use the current principal raster operator in a matched $64^2 \rightarrow 32^2 \rightarrow 64^2$ cycle. CONV* applies basin restriction followed by its present Q_1 -compiled synthesis. Lanczos-3 uses its endpoint-aligned, scale-widened kernel in both directions, while bilinear uses endpoint-aligned linear interpolation in both directions. Thus no row inherits the CONV-C zero-detail synthesis shown in Fig. 3.




Fig. 3. Conservative zero-detail transform ablation. Left: geometric-mean MSE over all 14 two-dimensional cases at each size. Right: mean MSE over the eight polygonal diagonal interfaces. Each curve uses its method’s declared analysis and synthesis.




Fig. 4. Fine-grid truth and current matched resize cycles. For the selected diagonal interface, MSE is 1.86×10^{-3} for CONV*, 1.37×10^{-3} for Lanczos-3, and 2.64×10^{-3} for bilinear; the corresponding range excursions are 3.91×10^{-2} , 9.63×10^{-2} , and zero. On the quadratic chirp, MSE is 5.37×10^{-2} , 5.20×10^{-2} , and 6.00×10^{-2} in the same order.

G. Technical synthetic fields

Figure 5 contains one declared case from each two-dimensional family at source side 17. The images use a dense 513×513 endpoint-aligned query lattice; its comparison grid retains the scored target side 33 and reports held-out MSE, maximum excursion, and nested-source error. Lanczos is slightly better on the selected edge and radial transition; CONV* is better on the carrier and crossing field and has the smaller maximum excursion while remaining cardinal. The tabulated raw row records the projection’s local effect on the same fields.

H. Eikonal compilation conformance

CONV* evaluates the two formal Cartesian factor orders and replaces the all-source Eikonal coefficient β by the

bilinear extension $\beta^* = I_h \beta$ of its cardinal nodal trace. Subtraction from the formal operator gives the exact relation

$$\mathcal{T}^* Y - \mathcal{T} Y = (\beta^* - \beta)(\mathcal{T}_{yx} Y - \mathcal{T}_{xy} Y). \quad (115)$$

Therefore the fast form is exact wherever the factor orders commute; away from that set, its complete discrepancy is precisely the product in (115).

Table IV evaluates that discrepancy on 16 fixed edge, carrier, curved-interface, and crossing fields. Every row has the same 65×65 target lattice; only source density and scale change. RMS is formed over all target values and cases. The maximum column is a census maximum, not a confidence estimate. Across the three rows, replacing the formal action changes geometric-mean truth MSE by less than 0.01 percent.




Fig. 5. Technical synthetic fields and their local comparison grid. Images use a dense endpoint-aligned query lattice and one common $[0, 1]$ display scale; the table reports unquantized binary64 values at target side 33 against direct analytic evaluation.

TABLE IV
CONV* CONFORMANCE TO FORMAL EIKONAL CONV AT TARGET
SIDE 65

Source	Scale	RMS	Median L_∞	Max. L_∞	Formal ms	CONV* ms
9	8	1.75×10^{-4}	1.50×10^{-3}	4.63×10^{-3}	11.0	0.402
17	4	5.67×10^{-5}	6.05×10^{-4}	2.43×10^{-3}	43.6	0.460
33	2	5.95×10^{-6}	1.56×10^{-5}	8.50×10^{-4}	236	0.659

I. Aggregate Green admission audit

The optional Green coordinate (121) was evaluated on the same 16 fixed fields using the source-derived determinant-one metric, $\rho = 1$, no-flux boundaries, a monotone anisotropic Selling stencil, and the cardinal weighted-harmonic form (123) [16]. Table V compares the result with the local coordinate $\beta^* = I_h \beta$. Although the coordinate changes materially, the exact commutator

identity (126) suppresses most of that difference in the reconstructed field.

TABLE V
CARDINAL GREEN ADMISSION VERSUS LOCAL Q_1 COMPILATION

Source	Local MSE	Green MSE	Ratio	β RMS	Output RMS
9^2	2.3888571×10^{-3}	2.3890219×10^{-3}	1.000069	4.76×10^{-2}	7.78×10^{-5}
17^2	1.8625653×10^{-4}	1.8626100×10^{-4}	1.000024	3.89×10^{-2}	8.50×10^{-6}

Because its discrete diagonal is finite, the raw finite-grid ratio V/U has maximum source error 0.542 and 0.628 on the two lattices. Imposing the nodal traces makes that error exactly zero. Consequently, the global coordinate supplies a closed metric-harmonic option and a formal comparator; however, this census yields no fidelity reason to displace the local Q_1 compilation.

J. Native CPU cost

Table VI reports three-channel float32 wall time on an M4 Mac mini. Native CONV*, Lanczos-3, bilinear, and EASU use ARM NEON and the same ten-worker pool; PCHIP is SciPy’s compiled implementation. Each entry is the median of 21 calls after three warmups, without affinity. The largest median absolute deviation in the table is 0.065 ms. Fidelity and timing use the same formulas, but the EASU timing arm uses the native shader translation rather than the binary64 audit port.

TABLE VI
NATIVE RGB POINT-SYNTHESIS TIME TO 512² (MS)

Method	from 64 ²	from 256 ²
CONV*	2.985	9.725
Lanczos-3	2.045	3.147
Bilinear	0.282	0.834
EASU	1.715	1.704
PCHIP	3.750	15.131

The native result places the cost rather than hiding it: at twofold synthesis CONV* is 3.09 times the Lanczos-3 time and 5.71 times the EASU time. At eightfold synthesis it is 1.46 times the Lanczos-3 time and 1.74 times the EASU time. These measurements characterize the present CPU realization; they are not used in any analytical claim.

Table VII is an implementation audit, not a substitute for the preceding exact-arithmetic proofs. It uses 80 deterministic Gaussian three-channel lines over five lengths. The current-mass, cardinal, and lifting residuals remain at binary64 rounding scale, and no admitted current has sign variation exceeding its observed secant sequence.

TABLE VII
BINARY64 CONFORMANCE AUDIT OVER 80 DETERMINISTIC TRIALS

Quantity	Maximum observed value
Cell-current mass residual	8.88×10^{-16}
Nested cardinal error	1.33×10^{-15}
Retained-detail lifting error	4.44×10^{-16}
Sign-variation surplus	0

The executable generator, complete per-case records, and figure construction form the supplementary material accompanying the paper. Aggregate curves and tables are generated from those records without case removal or post hoc weighting; the analytic population itself is fully specified above.

XII. CAN CONV BE MADE EFFICIENT USING CONVOLUTION?

Because CONV specifies the two-jet, lineage ledger, signed-fibre projection, Bernstein synthesis, and Eikonal factor blend independently of an execution backend, the same operator admits a bounded-work FIR realization. Denoted by CONV*, this realization eliminates the one-dimensional jet algebra into fixed convolution banks, batches the exact projection, compiles repeated phases into a polyphase output bank, and replaces the all-source action sum by the local nodal compilation derived below. The asterisk therefore denotes a computational refinement

of the same admitted fibre, rather than another interpolation model. In two dimensions, the FIR bank replaces each one-dimensional occurrence of $\mathcal{C}$ in $\mathcal{T}_{xy}$ and $\mathcal{T}_{yx}$; at a fixed endpoint-aligned scale, the same polyphase bank is used in opposite factor orders and the resulting proposals are blended pointwise, without forming a source-specific frame or chart bank.

A. Local compilation of Eikonal admission

Let η_p be the nodal order coordinate in (45). A local compilation is obtained by replacing the all-source action interpolant β with its bilinear nodal extension,

$$\beta^*(x) = \sum_{p \in \text{cell}(x)} b_p(x) \eta_p, \quad \mathcal{T}^*Y = (1 - \beta^*)\mathcal{T}_{xy}Y + \beta^*\mathcal{T}_{yx}Y. \quad (116)$$

Proposition 12 (Cardinal compilation identity). *Let I_h denote cellwise bilinear Q_1 interpolation on the source lattice. Then*

$$\boxed{\beta^* = I_h \beta}. \quad (117)$$

Consequently, β^ is the unique continuous cellwise bilinear extension of the formal Eikonal admission’s nodal trace.*

Proof. The Kronecker action identity (41) gives $\beta(x_p) = \eta_p$ at every source node. Applying I_h therefore yields $(I_h \beta)(x) = \sum_{p \in \text{cell}(x)} b_p(x) \beta(x_p)$, which is exactly (116). $\square$

Thus the compilation is not an independently chosen approximation rule, but the canonical Q_1 interpolant of the formal admission data; it is parameter-free, bounded in $[0, 1]$, and permutation covariant. Subtracting (47) from (116) gives the exact implementation defect

$$\mathcal{T}^*Y - \mathcal{T}Y = (\beta^* - \beta)(\mathcal{T}_{yx}Y - \mathcal{T}_{xy}Y), \quad (118)$$

and therefore, componentwise,

$$|\mathcal{T}^*Y - \mathcal{T}Y| \leq |\beta^* - \beta| |\mathcal{T}_{yx}Y - \mathcal{T}_{xy}Y|. \quad (119)$$

Therefore the local compilation is exact wherever the factor orders commute, including constant and affine coordinate fields. Moreover, (118) isolates the sole two-dimensional approximation introduced by CONV*, since the current construction within each factor remains unchanged.

For direct evaluation, η_p requires only the squared first-jet channels already produced by the two Cartesian front ends:

$$\eta_p = \frac{\sum_c (D_y Y_p^{(c)})^2}{\sum_c [(D_x Y_p^{(c)})^2 + (D_y Y_p^{(c)})^2]}, \quad (120)$$

with the value one half when the denominator vanishes. Consequently, neither an eigendecomposition nor the cross term of the current tensor is required to evaluate the order coordinate. At a target point, the four nodal values in (116) may be combined with the final convex blend in the same output write. This evaluates $I_h \eta$ exactly while avoiding a separate target-sized coordinate raster and its two Cartesian interpolation passes.

B. Aggregate Green admission

A global elliptic continuation of the same nodal coordinate is obtained as follows. Put $K = M^{-1}$, let ρ be the source density per unit Riemannian area, and define

$$\begin{aligned} (\rho - \Delta_M)U &= \sum_a \delta_{s_a}, \\ (\rho - \Delta_M)V &= \sum_a \eta_a \delta_{s_a}, \\ \beta_G &= V/U. \end{aligned} \quad (121)$$

subject to the no-flux boundary condition $n^T K \nabla U = n^T K \nabla V = 0$. In the unit source-lattice coordinates used throughout the paper, $\rho = 1$. Under $x \mapsto cx$, both ρ and Δ_M scale by c^{-2} , so the screening length is fixed by the sampling density rather than by a fitted bandwidth. Moreover, the two right-hand sides use one SPD operator and one factorization, without forming an individual distance field for each source.

Proposition 13 (Green transport coordinate). *The continuum coordinate in (121) is cardinal by continuous extension, satisfies*

$$\min_a \eta_a \leq \beta_G(x) \leq \max_a \eta_a, \quad (122)$$

and, away from the source sites, obeys

$$\operatorname{div}(U^2 M^{-1} \nabla \beta_G) = 0. \quad (123)$$

Equivalently, it minimizes

$$\mathcal{E}_U[\vartheta] = \frac{1}{2} \int_{\Omega} U^2 \nabla \vartheta^T M^{-1} \nabla \vartheta \, d\operatorname{vol}_M \quad (124)$$

subject to the nodal traces $\vartheta(s_a) = \eta_a$.

Proof. For $\rho > 0$, the shifted Neumann form is coercive and its Green kernel is strictly positive. Consequently, V/U is a positive normalized Green-kernel blend of the η_a , which proves (122) [24]. In two dimensions, the diagonal singularity is

$$G_\rho(x, s_a) = -\frac{1}{2\pi} \log d_M(x, s_a) + O(1). \quad (125)$$

and therefore the self term dominates at s_a , giving $\lim_{x \rightarrow s_a} \beta_G(x) = \eta_a$. Away from the sources, substituting $V = U\beta_G$ into the two homogeneous resolvent equations cancels the common mass term; the product rule then yields (123), whose weak form is the Euler equation of (124). $\square$

This construction continues the same nodal order coordinate, but it does not fuse the inverse-square Eikonal admission exactly. Indeed, (125) is logarithmic, whereas (40) uses d_M^{-2} ; hence no second-order local elliptic Green function in two dimensions can make the two kernels identical. Since a finite-stencil diagonal also removes the continuum singularity, the raw discrete ratio V/U is not exactly cardinal. Exact discrete cardinality is obtained by solving (123) with the nodal traces imposed.

Both the global Green admission and the local compilation therefore have the trace η_p . The latter is the unique

cellwise Q_1 continuation satisfying locality, linearity in the four nodal values, cardinality, and separate affine exactness, whereas the former is the global minimizer of (124). Their difference in the reconstructed field is isolated exactly by

$$\mathcal{T}_{\beta_G} Y - \mathcal{T}_{\beta^*} Y = (\beta_G - \beta^*)(\mathcal{T}_{yx} Y - \mathcal{T}_{xy} Y). \quad (126)$$

Consequently, Green admission supplies a global metric-harmonic coordinate when that property is required, while Q_1 remains the compact CONV* realization.

One Cartesian factor has the dataflow

$$y \xrightarrow{\text{FIR}} (a, \delta) \xrightarrow{\text{scan}} \sigma \xrightarrow{\text{tensorized projection}} c \xrightarrow{\text{polyphase}} \hat{y}. \quad (127)$$

Since the factors retain CONV's horizontal-vertical barrier, the second factor analyzes the values synthesized by the first. No single separable two-dimensional convolution therefore represents the complete nonlinear composition.

C. Compact FIR current bank

On an interior cell, let $v_i = (y_{i+r})_{r=-2}^3$. Eliminating the intermediate jets and Bezier controls yields

$$a_i = K v_i, \quad K = \frac{1}{240} \begin{bmatrix} 4 & -32 & 0 & 32 & -4 & 0 \\ 3 & -16 & -30 & 48 & -5 & 0 \\ -7 & 39 & -130 & 130 & -39 & 7 \\ 0 & 5 & -48 & 30 & 16 & -3 \\ 0 & 4 & -32 & 0 & 32 & -4 \end{bmatrix}. \quad (128)$$

Hence the complete interior proposal is a fixed bank of five six-tap FIR filters. An equivalent realization first computes the two five-tap jet channels and the secant, and then forms the five currents by pointwise algebra. The latter requires fewer arithmetic operations on a conventional CPU, whereas the fused form maps directly to a multi-output convolution primitive. At the outer and near-outer cells, the stated second-order closure is retained rather than inferred from (128) or supplied by zero padding.

D. The irreducible sign scan

Starting from the sampled signs $s_i = \operatorname{sign}(\delta_i)$, CONV* fills a zero from the last witnessed nonzero sign, using the first future witness only for an initial zero run. Applying the same local least-disagreement rule to the resulting transition sites gives a segmented prefix scan followed by a short causal transition scan.

Proposition 14 (No bounded convolutional ledger). *No fixed finite-receptive-field operator can reproduce the zero-sign filling rule for every line length.*

Proof. Fix a receptive-field radius R and a zero secant at a site whose previous R secants are also zero. Construct two lines with identical secants on that complete receptive field, but place a positive witness immediately before it in one line and a negative witness there in the other. The formal ledger assigns opposite signs at the selected zero site. A bounded local operator receives identical inputs and must return the same value, a contradiction. $\square$

Therefore the scan must be retained rather than approximated by a bounded convolution. It involves linear work with one sign of persistent state per component, while all lines of a Cartesian factor remain independent.

E. Vectorized scalar-threshold admission

CONV* batches the exact scalar projection (23)–(26) by using a fixed sorting network for the five breakpoints in each cell. The six resulting affine intervals are tested in parallel, after which five signed clips return the admitted current. Chunking consequently bounds temporary storage independently of image size, and neither a combinatorial search nor an iterative numerical solve is involved.

F. Polyphase synthesis

Let $B_j^5(u)$ denote the degree-five Bernstein basis. Eliminating the integrated controls gives

$$\hat{y}(i+u) = y_i + \sum_{k=0}^4 g_k(u) c_{ik}, \quad g_k(u) = \sum_{j=k+1}^5 B_j^5(u). \quad (129)$$

For integer refinement r , only the phases $u = p/r$, $p = 1, \dots, r-1$, are required. Their g_k values form a fixed polyphase matrix. At twofold refinement,

$$\hat{y}(i + \tfrac{1}{2}) = y_i + \tfrac{1}{32} [31, 26, 16, 6, 1] c_i. \quad (130)$$

Thus neither the Bezier controls nor their cumulative current sums need be stored. Nested source nodes are copied directly for point interpolation, whereas conservative multiresolution compiles the half-cell functional (75), block averages, and centered moment channels of Section IX; it does not replace them by strided anchor selection. Accordingly, (130) is a compiled twofold phase rather than the domain of CONV. For an arbitrary target length M , the continuous operator evaluates (129) at the coordinates (9); whenever the target uses a repeated finite phase set, the same expression yields a fixed polyphase bank.

G. Equivalence audit

The binary64 audit uses 72 deterministic Gaussian three-channel lines over lengths 5, 9, 17, 33, 65, 129. The maximum difference from the formal current is 9.44×10^{-16} ; the maximum synthesized-value difference is 8.88×10^{-16} . Cardinal error is zero, current-mass residual is at most 1.78×10^{-15} , and no sign-variation surplus is observed.

A second audit in Section XI compares the local action compilation with the formal all-source Eikonal operator. It reports the defect in (118), rather than conflating that approximation with the binary64-equivalent FIR factor compilation.

H. FIR refinement

Collecting the preceding reductions, CONV* comprises compact FIR analysis, exact segmented lineage, vectorized scalar-threshold admission, polyphase synthesis, and the cardinal Q_1 action coordinate (117). Its one-dimensional factor is algebraically identical to CONV; in two dimensions, the only departure is the explicitly bounded defect (118).

I. Can CONV* Compete with EASU?

EASU produces each target pixel in a single two-dimensional pass over a fixed neighborhood [13]. CONV*, however, retains a barrier between its Cartesian factors, since the second factor analyzes the nonlinear output of the first. No throughput ordering follows from the formulas alone; nevertheless, every nonlinear cell operation has now been reduced to the fixed scalar-threshold and scan primitives specified above.

1) *Parallel lineage*: Sign inheritance is obtained by seeded prefix propagation, and the boundary objective likewise has a scan form. For one transition window $[p, q]$, let s^- and s^+ be its left and right coarse signs and define

$$P^-(b) = \sum_{k=p}^{b-1} [\min(0, s^- a_k)]^2, \quad S^+(b) = \sum_{k=b}^q [\min(0, s^+ a_k)]^2. \quad (131)$$

Using these partial sums, every candidate cost in (20) becomes

$$C(b) = P^-(b) + S^+(b), \quad (132)$$

so that one prefix sum and one suffix sum replace repeated evaluation of the overlapping windows.

The dependence between successive transition boundaries is finite as well. Indeed, if $\kappa_\ell - \kappa_{\ell-1} \geq 2$, then

$$b_{\ell-1} + 1 \leq 5\kappa_\ell - 5 < 5\kappa_\ell - 4, \quad (133)$$

and the preceding boundary cannot restrict the candidate set (18); dependence remains only inside a run of adjacent coarse transitions. Writing $r_\ell = b_\ell - 5\kappa_\ell \in \{-4, \dots, 4\}$, each transition in such a run defines a deterministic map G_ℓ from the preceding offset to the next, including the stated first-minimum tie rule. Since function composition is associative, a parallel prefix composition of these nine-state maps yields the exact greedy result.

2) *Factor schedule*: The remaining work of one factor is fixed local analysis, the lineage scans, the scalar solve (24), and the polyphase dot product (129). Node jets may be formed once and shared by their two incident cells. Fixed-scale phases are constants, and components in the declared channel basis may be evaluated together. A tiled implementation may write the synthesized factor through an on-chip transpose so that the next factor again reads contiguous lines; it may not remove the intervening factor barrier without changing the operator.

These reductions preserve the boundary closures, ordered tie rules, channel basis, and Cartesian factor order of CONV. They reduce the remaining nonlinear state to scans and a five-value piecewise-affine root. Whether that arithmetic advantage offsets the intermediate factor traffic is a property of a concrete implementation and target architecture. No performance claim is made here; equations (23)–(133) specify the exact reduced form for such an implementation.

XIII. INVERSION IN CONVENTIONAL RESAMPLING

Windowed-sinc, cubic, and edge-directed raster interfaces are ordinarily specified as synthesis maps from one grid to another.¹ Their multiresolution analysis and synthesis filters are separately constructed, and the identity $D_s U_s = I$ on coarse inputs does not make the discarded fine array recoverable. CONV adds a native analysis state: conserved cell averages, admitted centered moments, and residual moment fields have explicit analysis and synthesis maps. The comparison is representational rather than a claim that invertible multiscale methods are new; wavelet lifting already establishes the general analysis–prediction–detail principle [23], and Harten’s construction already represents nested cell averages by coarse averages and fine-grid prediction errors [22]. Here those established coordinates are joined to compact quintic transport and a factorwise ringing-incapability theorem.

XIV. CONCLUSION

Compact Ordered Nodal Variation formulates interpolation and matched multiresolution transfer through a conservative signed-current representation. A high-order two-jet supplies the current proposal, ordered admission restricts its derivative-sign lineage before Bernstein synthesis, and a source-derived Riemannian coordinate couples the two Cartesian factor orders. The same representation yields arbitrary-scale basin reduction, exact residual-moment lifting, and the FIR/polyphase realization CONV*.

The analysis establishes cardinality, affine reproduction, cell-current conservation, factorwise variation diminution, fixed-ledger nonexpansiveness, sample-count-independent error bounds, well-posed Riemannian action, exact retained-detail inversion, and Nyquist annihilation by conservative restriction. The deterministic synthetic results confirm the predicted sign topology and rapid resolved-field convergence; matched-cycle comparisons show competitive fidelity with substantially smaller range excursion than Lanczos on the selected diagonal interface, while native timings identify the remaining cost of nonlinear admission and the Cartesian factor barrier.

Appendix A develops the joint two-dimensional current admission required for direct affine and projective warps, extending the factorwise construction without intermediate rasterization.

APPENDIX A DIRECT JOINT CONV WARPS

The continuous CONV representation admits a direct extension from lattice resize to rotation, skew, anisotropic affine transfer, and perspective warp. The necessary change is not a decomposition of the coordinate map into one-dimensional resampling passes. It is a joint admission

¹Glory first be to The Holy One, Blessed be He (הוא ברוך הקדוש), from whom all good things come. To Him we attribute the ability that has provided for all good to come and be demonstrated anywhere herein, and declare ourselves unworthy servants. If in any way we have erred, or our work is insufficient, we ask mercy and kindness, for our effort is good and earnest.

of the two-dimensional current before that map is evaluated. The construction below consequently retains one inverse map, one continuous source atlas, and one positive raster functional. No intermediate raster occurs.

A. One inverse map and its induced geometry

Let $\Omega \subset \mathbb{R}^2$ denote the normalized source rectangle and let $\hat{\Omega} \subset \mathbb{R}^2$ denote the target rectangle. A nonsingular inverse warp

$$F : \hat{\Omega} \longrightarrow \Omega, \quad q \longmapsto p = F(q), \quad (134)$$

is declared once. Given the admitted source atlas U_Y , point synthesis is

$$(\mathcal{W}_F Y)(q) = U_Y(F(q)). \quad (135)$$

Thus a shear factorization of an affine map is neither required nor used. Indeed, repeated admitted interpolation would generally define a different nonlinear operator, because every intermediate raster changes the current state from which the next admission is formed.

When the source current geometry is expressed by the SPD field $M(p)$, the same inverse map induces the target quadratic form

$$G_F(q) = J_F(q)^T M(F(q)) J_F(q), \quad J_F = DF, \quad (136)$$

and its determinant-one shape is

$$\overline{M}_F(q) = \frac{G_F(q)}{\det(G_F(q))^{1/2}}. \quad (137)$$

Equation (136) combines sampling scale, anisotropy, and the warp before a square root or directional support is formed. Consequently it does not assume that Cartesian spacing commutes with the eigenframe of M . The scalar $\det(G_F)^{1/4}$ carries linear scale, while (137) carries directional shape.

B. Compiled bivariate proposal

Index the source cells by C , and write $(u, v) \in [0, 1]^2$ for local cell coordinates. The joint source potential on cell C is represented in the tensor Bernstein basis as

$$U_C(u, v) = \sum_{i=0}^5 \sum_{j=0}^5 P_{C,ji} B_i^5(u) B_j^5(v). \quad (138)$$

The proposal controls are not fitted by an optimizer. Let

$$A_{ri} = B_i^5(r/5), \quad 0 \leq r, i \leq 5, \quad (139)$$

and let Z_C contain the values of the complete two-order CONV proposal at the 6×6 fifth-step tensor lattice of C . Since the Bernstein collocation matrix A is nonsingular, the unique bidegree-(5, 5) proposal is

$$\tilde{P}_C = A^{-1} Z_C A^{-T}. \quad (140)$$

All entries of every Z_C are obtained from one global fifth-step evaluation. Adjacent cells therefore have the same six samples on their common edge; uniqueness in (139) then gives the same complete Bernstein edge. Equivalently, the patches are stored in one shared global control lattice. The

resulting proposal is C^0 , and its current is the gradient of one continuous piecewise-polynomial potential.

Let B denote the degree-elevated Q_1 interpolant of the source samples on that shared lattice. It is cardinal, continuous, and affine exact. Admission will move B toward $\tilde{P}$, while the source vertices remain fixed.

C. Support-range admission

For each source cell C , let $\mathcal{S}_C$ be the complete Cartesian two-jet sample support, clipped at the outer boundary. If shared control entity e is incident to the cell set $\mathcal{I}(e)$, define, for channel c ,

$$I_e^{(c)} = \bigcap_{C \in \mathcal{I}(e)} \left[\min_{p \in \mathcal{S}_C} Y_p^{(c)}, \max_{p \in \mathcal{S}_C} Y_p^{(c)} \right]. \quad (141)$$

This intersection is nonempty: the corresponding Q_1 control is a convex combination of common incident vertices and belongs to every interval in (141). The proposal is first replaced by

$$\hat{P}_e^{(c)} = \Pi_{I_e^{(c)}} \tilde{P}_e^{(c)}. \quad (142)$$

Unlike a post-reconstruction clamp, (142) constrains the continuous potential. Subsequent current admission remains on the segment from $B_e^{(c)}$ to $\hat{P}_e^{(c)}$, so every final control remains in $I_e^{(c)}$.

Theorem 4 (Continuous support-range admission). *For every channel and source cell,*

$$\min_{p \in \mathcal{S}_C} Y_p^{(c)} \leq U_C^{(c)}(u, v) \leq \max_{p \in \mathcal{S}_C} Y_p^{(c)}, \quad (u, v) \in [0, 1]^2. \quad (143)$$

Consequently, point synthesis and every positive warped-pixel average have no range excursion beyond the union of the source supports they traverse.

Proof. For cell C , every one of its 36 shared controls belongs to the interval on the right of (143). The tensor Bernstein basis is nonnegative and forms a partition of unity, hence U_C is a convex combination of those controls. Point evaluation gives the first result, and integration against a nonnegative normalized measure gives the second. $\square$

D. Support-current cone

The one-dimensional theorem controls current along a declared factor. A rotated or projective query instead requires a joint source current law. For cell C , collect the four corner gradients of every Q_1 cell contained in $\mathcal{S}_C$:

$$\mathcal{G}_C = \{\nabla Q_D(z) : D \subset \mathcal{S}_C, z \in \text{vert}(D)\}. \quad (144)$$

Zero vectors are removed. If the remaining directions generate a proper closed planar cone, let

$$K_C = \text{cone } \mathcal{G}_C; \quad (145)$$

if no nonzero current is observed, or if the observed currents positively span the plane, set $K_C = \mathbb{R}^2$. Thus a clean interface gives a narrow cone, curvature gives a wedge, and

crossing currents remove a directional restriction rather than forcing a false owner or normal.

Every proper planar cone occurring here has a finite half-space description

$$K_C = \{z \in \mathbb{R}^2 : \ell_{Cr}^T z \geq 0, 1 \leq r \leq R_C\}, \quad R_C \leq 3. \quad (146)$$

A wedge uses its two inward boundary normals. A ray additionally requires a forward inequality, while a line uses two opposing transverse inequalities. The whole-plane case has $R_C = 0$.

For each r , the directional derivative

$$L_{Cr}(u, v) = \ell_{Cr}^T \nabla U_C(u, v) \quad (147)$$

is degree elevated to the bidegree-(5, 5) Bernstein basis. Denote its 36 linear control functionals by $a_{Crk}^T P$. The finite admission conditions are

$$a_{Crk}^T P \geq 0, \quad 1 \leq r \leq R_C, \quad 1 \leq k \leq 36. \quad (148)$$

Nonnegativity of these Bernstein coefficients is a sufficient certificate for $L_{Cr} \geq 0$ on the complete cell.

The degree-elevated Q_1 baseline is a constructive feasible point. Its gradient is a convex combination of the four corner gradients of its source cell; those gradients occur in $\mathcal{G}_C$. Hence

$$\nabla B_C(u, v) \in K_C \quad (149)$$

everywhere, and every margin in (148) is nonnegative.

E. Finite topological-entity admission

Assign each nonvertex control to exactly one shared-edge or cell-interior entity g . With the range-admitted proposal from (142), write

$$\hat{P} = B + \sum_g d_g, \quad (150)$$

where the entity corrections have disjoint control ownership. For a global constraint $a_k^T P \geq 0$, put

$$m_k = a_k^T B, \quad H_k = \sum_g \max(0, -a_k^T d_g). \quad (151)$$

If $\hat{P}$ already satisfies all constraints, it is accepted unchanged. Otherwise each harmful entity is assigned the simultaneous capacity

$$\alpha_g = \min \left\{ 1, \min_{k: a_k^T d_g < 0, H_k > 0} \frac{m_k}{H_k} \right\}, \quad (152)$$

with an empty inner minimum interpreted as one. The first admitted state is

$$P^{(0)} = B + \sum_g \alpha_g d_g. \quad (153)$$

Finally, let $R = \hat{P} - P^{(0)}$ and take the maximal common ray

$$\tau = \min \left\{ 1, \min_{k: a_k^T R < 0} \frac{a_k^T P^{(0)}}{-a_k^T R} \right\}, \quad P = P^{(0)} + \tau R. \quad (154)$$

Again, an empty inner minimum is one. Equations (152)–(154) contain neither an iteration nor an active-set search.

Theorem 5 (Finite joint current admission). *The control lattice P obtained from (152)–(154) satisfies every constraint in (148). Consequently,*

$$\nabla U_C(u, v) \in K_C, \quad (u, v) \in [0, 1]^2, \quad (155)$$

for every source cell. The construction is finite, order independent, cardinal, C^0 , and affine exact.

Proof. Fix constraint k . Contributions with $a_k^T d_g \geq 0$ can be discarded in a lower bound. For every harmful entity, $\alpha_g \leq m_k/H_k$, and therefore

$$a_k^T P^{(0)} \geq m_k - \frac{m_k}{H_k} \sum_g \max(0, -a_k^T d_g) = 0. \quad (156)$$

Equation (154) is the exact first intersection of the remaining affine ray with the finite half-space boundary, hence $a_k^T P \geq 0$ as well. Applied to all derivative coefficients, Bernstein positivity gives $\ell_{Cr}^T \nabla U_C \geq 0$, which is equivalent to (155) through (146).

Source vertices are not assigned to a movable entity. Shared-edge controls are single variables, so cardinality and C^0 continuity are retained. For affine data, the proposal equals the degree-elevated affine baseline and already satisfies the cone and range constraints; it is consequently passed unchanged. $\square$

The theorem also explains why the construction is not an edge detector. It does not select one current from $\mathcal{G}_C$. Instead it admits the complete high-order proposal whenever all of its derivative coefficients lie in the cone generated by the currents actually present in the declared support.

F. Transport through affine and projective maps

Let $\gamma : [a, b] \rightarrow \widehat{\Omega}$ be a C^1 target curve. At every regular point contained in one source cell, the chain rule gives

$$\frac{d}{ds} U_Y(F(\gamma(s))) = [J_F(\gamma(s))\gamma'(s)]^T \nabla U_Y(F(\gamma(s))). \quad (157)$$

Let

$$K_C^* = \{t \in \mathbb{R}^2 : t^T z \geq 0 \text{ for all } z \in K_C\} \quad (158)$$

be the dual current direction cone.

Corollary 3 (Warped-current transport). *Suppose $F(\gamma(s))$ crosses only finitely many source-cell boundaries and*

$$J_F(\gamma(s))\gamma'(s) \in K_{C(s)}^* \quad (159)$$

for almost every s , where $C(s)$ contains $F(\gamma(s))$. Then $U_Y \circ F \circ \gamma$ is nondecreasing on $[a, b]$.

Proof. Equations (155), (158), and (157) make the derivative nonnegative in every open patch segment. The atlas is continuous and piecewise C^1 ; therefore its composition

with the curve is absolutely continuous after splitting at the finite knot set. The fundamental theorem gives

$$U_Y(F(\gamma(b))) - U_Y(F(\gamma(a))) = \int_a^b \frac{d}{ds} U_Y(F(\gamma(s))) ds \geq 0. \quad (160)$$

The knot points have measure zero and require no derivative convention. $\square$

Affine maps transport the admitted directions by one constant Jacobian. Homographies transport them by the spatially varying Jacobian of the same declared map. In either case, no target-space owner, orientation, or second interpolation operation is introduced.

G. Conservative warped-pixel analysis

Point synthesis is appropriate for enlargement. For reduction, let $\{P_j\}$ be the endpoint-aligned target Voronoi basins. The matched raster functional is

$$(\mathcal{B}_F Y)_j = \frac{1}{|P_j|} \int_{P_j} U_Y(F(q)) dq. \quad (161)$$

Since the basins partition the target domain,

$$\sum_j |P_j| (\mathcal{B}_F Y)_j = \int_{\cup_j P_j} U_Y(F(q)) dq. \quad (162)$$

Thus (161) is conservative relative to the admitted continuous profile. It is also a positive average and inherits (143).

For affine F , a target basin maps to a source parallelogram. Splitting at the integer source-knot lines produces convex polygons, each contained in one patch. Positive fan triangulation introduces no overlap. On every triangle, U_C has total degree at most ten. Under a Duffy map the area Jacobian raises one separate degree to at most eleven; consequently, the tensor three-point rule is insufficient, whereas tensor Gauss–Legendre order six is exact. Division by the mapped parallelogram area then gives (161) exactly.

For a homography, source-coordinate integration yields

$$(\mathcal{B}_F Y)_j = \frac{1}{|P_j|} \int_{F(P_j)} U_Y(p) |\det DF^{-1}(p)| dp. \quad (163)$$

With homogeneous coordinate normalizations absorbed into the nonsingular matrix H , the inverse-map density is

$$|\det DF^{-1}(p)| = \frac{|\det H^{-1}|}{|\ell(p)|^3}, \quad (164)$$

where ℓ is the affine denominator of H^{-1} . On a clipped triangle choose a nonzero center c of the denominator interval and write $\ell = c(1+z)$, with $|z| \leq r < 1$. Then

$$(1+z)^{-3} = \sum_{k=0}^K (-1)^k \binom{k+2}{2} z^k + R_K(z). \quad (165)$$

The tail admits the closed bound

$$|R_K(z)| \leq T_{K+1}(r), \quad (166)$$

$$T_n(r) = \frac{r^n}{2} \left[\frac{r(1+r)}{(1-r)^3} + \frac{(2n+3)r}{(1-r)^2} + \frac{(n+1)(n+2)}{1-r} \right].$$

The finite series multiplied by U_C is polynomial and is integrated exactly by tensor Gauss order

$$q_K = \left\lceil \frac{K+12}{2} \right\rceil. \quad (167)$$

If $U_C = P_0 + (U_C - P_0)$, the constant P_0 is integrated exactly before (165) is applied. Hence constants have zero certified residual. For triangle T , the remaining error is bounded by

$$E_T \leq |T| \frac{|\det H^{-1}|}{|c|^3} T_{K+1}(r) \|U_C - P_0\|_{L^\infty(T)}. \quad (168)$$

Bernstein controls and first differences provide the required value bound. Four-way triangle subdivision is invoked only when the sum of (168) exceeds the declared pixel budget. The exact integral (163) remains the operator; the finite implementation returns its value together with a rigorous enclosure.

H. Multichannel form

For $Y : \Omega \cap \mathbb{Z}^2 \rightarrow \mathbb{R}^C$, the control topology, support sets, and inverse map are shared. Range and current admission are applied to each component in the declared channel basis, yielding the product set

$$\prod_{c=1}^C \left\{ U^{(c)} : \nabla U_C^{(c)} \in K_C^{(c)}, U_C^{(c)} \in I_C^{(c)} \right\}. \quad (169)$$

The preceding proofs therefore hold componentwise. As with the main CONV operator, this statement is invariant under independent affine changes of channel units; it does not assert invariance under an arbitrary basis change that mixes channels.

I. Exact-truth evaluation

The evidence uses only procedural sources whose point values and warped-pixel means are independent of every tested interpolator. The sharp corpus contains a half-plane with normal angle 37° , a width-0.065 bar at 27° , two width-0.055 crossing bars at 31° and 121° , and a 0.62×0.34 box at -19° , all centered at $(1/2, 1/2)$. The four target-to-source maps are

$$\begin{aligned} F_{\text{rot}} : (\theta, s_x, s_y) &= (21^\circ, 0.72, 0.72), \\ F_{\text{skew}} : (s_x, s_y, h_x) &= (0.70, 0.75, 0.42), \\ F_{\text{aniso}} : (\theta, s_x, s_y) &= (-14^\circ, 0.50, 0.84), \\ F_{\text{proj}} : (g_x, g_y, \theta, s_x, s_y) &= (0.32, -0.18, 9^\circ, 0.67, 0.72), \end{aligned} \quad (170)$$

where affine maps use $F(q) = \frac{1}{2} + A(q - \frac{1}{2})$, with $A = R_\theta \begin{pmatrix} 1 & h_x \\ h_y & 1 \end{pmatrix} \text{diag}(s_x, s_y)$, and the projective map uses

$$F(q) = \frac{1}{2} + \frac{A(q - \frac{1}{2})}{1 + (g_x, g_y)^T (q - \frac{1}{2})}. \quad (171)$$

Exact polygon membership supplies point truth. Since a source half-plane remains a target half-plane under a pole-free homography, polygon clipping in target coordinates supplies exact basin coverage.

The smooth census uses the real Fourier fields

$$f(x) = b + \sum_{k=1}^K a_k \cos(2\pi \xi_k^T x + \varphi_k). \quad (172)$$

For an affine map, the exact mean over a target rectangle with center m and side lengths h_x, h_y is obtained by multiplying each transformed mode by

$$\text{sinc}((A^T \xi_k)_x h_x) \text{sinc}((A^T \xi_k)_y h_y). \quad (173)$$

The scalar probes use the two mode sets

$$\begin{aligned} \mathcal{F}_1 &= \{((5, 2), .22, .13), ((9, 3), .09, 1.17), \\ &\quad ((2, -4), .07, -.41)\}, \\ \mathcal{F}_2 &= \{((6, 1), .17, .2), ((-1, 6), .17, -.6), \\ &\quad ((4, 4), .09, 1.1)\}, \end{aligned} \quad (174)$$

with $b = 0.5$. A three-channel probe uses wavevectors $(4, 1), (1, 5), (5, -3)$, phases $0, 0.7, -0.4$, and amplitude rows

$$\begin{bmatrix} .22 & .05 & .08 \\ .04 & .20 & .07 \\ .06 & .04 & .18 \end{bmatrix}, \quad b = (.5, .5, .5). \quad (175)$$

Table VIII reports the 16 sharp cases on a 17×17 source and 13×13 target. CONV has the lowest MSE in seven point cases and seven warped-pixel cases. More importantly for the proved admission, its maximum range excursion remains at floating roundoff, whereas the factor-only ablation and Lanczos admit substantial excursion.

TABLE VIII
SHARP AFFINE/PROJECTIVE EXACT-TRUTH CENSUS

Method	Point MSE	Pixel MSE	Point exc.	Pixel exc.
CONV	.021002	.007377	2.22×10^{-16}	3.18×10^{-14}
Factor ablation	.021137	.007705	.2287	.1433
Lanczos-3	.022928	.008741	.2817	.1989
Bilinear	.021513	.007576	0	0

The intrinsic topology census evaluates normal profiles rather than using MSE as a surrogate for ringing. For a profile $z_0, \dots, z_n$ whose exact orientation is increasing, the reported wrong-way current is

$$V_-(z) = \sum_{i=0}^{n-1} \max(0, z_i - z_{i+1}). \quad (176)$$

Straight interfaces use six source angles, three maps, 513 samples per profile, and source sides 17 and 33. CONV has no normal sign changes, $V_- \leq 1.34 \times 10^{-14}$, and excursion below 4.45×10^{-16} . Lanczos-3 produces four to nine sign




Fig. 6. Direct joint CONV warp evidence. Top: one homography applied to sources with exact projective basin truth; each displayed certificate is the computed enclosure of the projective polynomial integral. Bottom: topology and exact-truth fidelity over the declared corpus. Values plotted at the numerical floor are below 2×10^{-16} . The factor-only result is an ablation of joint current admission, not a second CONV definition.

changes, mean $V_- = 0.276\text{--}0.277$, and mean excursion 0.166.

For curved interfaces, circles and ellipses are tested at 12 phases under three maps, giving 72 profiles per source size. Table IX shows that the admitted cone eliminates wrong-way normal current at sides 13, 17, and 25. At side 9, the range theorem remains exact, while the six-node support spans competing conic normals and the cone correctly widens; the remaining wrong-way current

therefore measures support ambiguity rather than failure of the finite inequalities.

TABLE IX
CURVED-INTERFACE INTRINSIC NORMAL CENSUS FOR CONV

Source side	Mean V_-	Max. V_-	Max. excursion
9	1.94×10^{-3}	2.80×10^{-2}	2.22×10^{-16}
13	1.67×10^{-15}	4.77×10^{-15}	4.44×10^{-16}
17	2.23×10^{-15}	6.44×10^{-15}	4.44×10^{-16}
25	3.31×10^{-15}	5.44×10^{-15}	4.44×10^{-16}

On the nine band-admitted Fourier cases, mean point MSE is 1.48764×10^{-4} for CONV, 1.06820×10^{-4} for the factor-only ablation, 4.57424×10^{-5} for Lanczos-3, and 1.22893×10^{-3} for bilinear. The corresponding warped-pixel MSEs are 6.27668×10^{-5} , 4.09234×10^{-5} , 1.88691×10^{-5} , and 7.63065×10^{-4} . Thus the complete support maximum principle has a measurable smooth-field fidelity cost; the experiment does not support universal MSE dominance.

Finally, affine conservation equates the weighted target sum and an independent exact integral of the admitted atlas to 3×10^{-12} . On a nonconstant projective plane, independent 40-point tensor integration differs from the certified result by 1.18×10^{-11} , within its 5.84×10^{-10} enclosure. The largest certificate returned by the sharp projective corpus is 1.52×10^{-11} . These computations verify the exact polynomial and error-enclosure implementations; the cardinality, integrability, range, current-cone, and transport statements follow from the preceding formulas rather than from the measurements.

The direct warp therefore extends CONV without changing its governing idea: the high-order proposal is admitted in the representation that carries the current, after which point evaluation and positive basin analysis are two functionals of the same continuous object. The principal remaining boundary is source resolution itself. When one compact two-jet support contains genuinely competing interface normals, the observed cone contains both and cannot certify a narrower direction without adding information absent from the samples.

AI GENERATED WORK DISCLAIMER

Substantial portions of this work were generated with the assistance of OpenAI Codex general reasoning models.

REFERENCES

- [1] C. E. Shannon, "Communication in the presence of noise," *Proc. IRE*, vol. 37, no. 1, pp. 10–21, Jan. 1949.
- [2] C. Lanczos, *Applied Analysis*. Englewood Cliffs, NJ, USA: Prentice-Hall, 1956.
- [3] R. G. Keys, "Cubic convolution interpolation for digital image processing," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 29, no. 6, pp. 1153–1160, Dec. 1981.
- [4] F. N. Fritsch and R. E. Carlson, "Monotone piecewise cubic interpolation," *SIAM J. Numer. Anal.*, vol. 17, no. 2, pp. 238–246, Apr. 1980.
- [5] J. M. Hyman, "Accurate monotonicity preserving cubic interpolation," *SIAM J. Sci. Stat. Comput.*, vol. 4, no. 4, pp. 645–654, Dec. 1983.
- [6] P. Costantini, "Co-monotone interpolating splines of arbitrary degree—a local approach," *SIAM J. Sci. Stat. Comput.*, vol. 8, no. 6, pp. 1026–1034, Nov. 1987, doi: 10.1137/0908083.
- [7] R. Dougherty, A. Edelman, and J. M. Hyman, "Nonnegativity-, monotonicity-, or convexity-preserving cubic and quintic Hermite interpolation," *Math. Comp.*, vol. 52, no. 186, pp. 471–494, Apr. 1989, doi: 10.1090/S0025-5718-1989-0962209-1.
- [8] G. Ulrich and L. T. Watson, "Positivity conditions for quartic polynomials," *SIAM J. Sci. Comput.*, vol. 15, no. 3, pp. 528–544, May 1994, doi: 10.1137/0915035.
- [9] T. Lux, L. T. Watson, T. Chang, and W. Thacker, "Algorithm 1031: MQSI—monotone quintic spline interpolation," *ACM Trans. Math. Softw.*, vol. 49, no. 1, Art. no. 6, pp. 1–17, Mar. 2023, doi: 10.1145/3570157.
- [10] R. E. Carlson and F. N. Fritsch, "An algorithm for monotone piecewise bicubic interpolation," *SIAM J. Numer. Anal.*, vol. 26, no. 1, pp. 230–238, Feb. 1989, doi: 10.1137/0726013.
- [11] P. Costantini and C. Manni, "A local scheme for bivariate co-monotone interpolation," *Comput. Aided Geom. Des.*, vol. 8, no. 5, pp. 371–391, Nov. 1991, doi: 10.1016/0167-8396(91)90011-Y.
- [12] X. Li and M. T. Orchard, "New edge-directed interpolation," *IEEE Trans. Image Process.*, vol. 10, no. 10, pp. 1521–1527, Oct. 2001.
- [13] T. Lottes, "AMD FidelityFX Super Resolution 1.0," Advanced Micro Devices, Santa Clara, CA, USA, 2021.
- [14] W. Ding, F. Wu, X. Wu, S. Li, and H. Li, "Adaptive directional lifting-based wavelet transform for image coding," *IEEE Trans. Image Process.*, vol. 16, no. 2, pp. 416–427, Feb. 2007, doi: 10.1109/TIP.2006.888341.
- [15] J. Bigun and G. H. Granlund, "Optimal orientation detection of linear symmetry," in *Proc. 1st IEEE Int. Conf. Comput. Vis.*, London, U.K., 1987, pp. 433–438.
- [16] J.-M. Mirebeau, "Anisotropic fast-marching on Cartesian grids using lattice basis reduction," *SIAM J. Numer. Anal.*, vol. 52, no. 4, pp. 1573–1599, 2014, doi: 10.1137/120861667.
- [17] D. Shepard, "A two-dimensional interpolation function for irregularly-spaced data," in *Proc. 23rd ACM Natl. Conf.*, New York, NY, USA, 1968, pp. 517–524, doi: 10.1145/800186.810616.
- [18] S. Karlin, *Total Positivity*, vol. 1. Stanford, CA, USA: Stanford Univ. Press, 1968.
- [19] A. Pinkus, *Totally Positive Matrices*. Cambridge, U.K.: Cambridge Univ. Press, 2010.
- [20] A. S. Cavaretta, Jr. and A. Sharma, "Variation diminishing properties and convexity for the tensor product Bernstein operator," in *Functional Analysis and Operator Theory*, Lecture Notes in Mathematics, vol. 1511. Berlin, Germany: Springer, 1992, pp. 18–32, doi: 10.1007/BFb0093794.
- [21] T. N. T. Goodman, "Variation diminishing properties of Bernstein polynomials on triangles," *J. Approx. Theory*, vol. 50, no. 2, pp. 111–126, 1987, doi: 10.1016/0021-9045(87)90002-5.
- [22] A. Harten, "Multiresolution algorithms for the numerical solution of hyperbolic conservation laws," *Commun. Pure Appl. Math.*, vol. 48, no. 12, pp. 1305–1342, Dec. 1995, doi: 10.1002/cpa.3160481201.
- [23] W. Sweldens, "The lifting scheme: A construction of second generation wavelets," *SIAM J. Math. Anal.*, vol. 29, no. 2, pp. 511–546, 1998, doi: 10.1137/S0036141095289051.
- [24] D. Gilbarg and N. S. Trudinger, *Elliptic Partial Differential Equations of Second Order*. Berlin, Germany: Springer, 2001.