OPERATION KOMOREBI

Komorebi Space

Worked construction and proofs

A companion to the multiresolution visual atlas

20 October 2026 | Version 4

The four surfaces in the atlas are computed graphs of four trigonometric polynomials. They are not four drawings of an unspecified manifold. Their common mathematical object is a compatible tuple of views. This note specifies that tuple, computes its detail and defect coordinates, and proves the identities used to interpret the figures.

The construction is a finite instance of the view and resolution systems in the *Komorebi Space Monograph*, especially Sections 1, 3, 6 and 10 [1]. The Fourier example was introduced for the atlas [2]. The split identity in Section 10 applies here at the level of signal spaces. These signal details are not being identified with that section's particular operator-valued dyadic residual hierarchy.

> **Central statement.** The displayed family is an element of a 289-dimensional coherent realization space inside a 404-dimensional raw family space. Its three nonzero detail bands reconstruct its finest view exactly. Its compatibility residual is zero. Altering only view 2 by $\eta = \frac{1}{5}\cos(2\pi u)$ produces the residual $(0, \eta, -\eta)$, which has an exact repair preserving the other three views.

All of these statements receive ordinary mathematical proofs below. The delivered Python additionally replays the original 8,969 exact checks and passes 3,695 added checks. The resulting 12,664 checks operate on rational Fourier coefficients, including complete bases of the relevant finite spaces. They are not a Lean compiler run and do not replace the infinite-dimensional proof in Section 8.

Reading route. Sections 1–4 build the first plate. Section 5 works through the altered family and two different repair objectives. Sections 6–7 explain transport and obstruction. Section 8 identifies the additional hypothesis needed for an infinite family to represent a Hilbert-space element. The two original plates follow the verification record.

1 Declare the views before drawing them

The parameter domain is the square torus $\mathbb{T}^2 = \mathbb{R}^2/\mathbb{Z}^2$, represented by $0 \le u, v \le 1$ with opposite edges identified. Integrals use normalized area, so $\int_{\mathbb{T}^2} 1 = 1$. Write

$$C_{a,b}(u,v) = \cos\bigl(2\pi(au+bv)\bigr), \qquad S_{a,b}(u,v) = \sin\bigl(2\pi(au+bv)\bigr).$$

Choose one representative of each nonzero pair $\{(a,b),(-a,-b)\}$ by requiring $a > 0$, or $a = 0$ and $b > 0$. This avoids duplicating cosine modes or counting a sine mode and its negative as different basis vectors.

For $q \ge 1$, let V_q be the real span of the constant function and all such sine and cosine modes with $\max(|a|,|b|) \le q$. The four views are

$$q_n = 2^n, \qquad X_n = V_{q_n}, \qquad n = 0,1,2,3.$$

These are spaces of functions, not surfaces in three-dimensional Euclidean space. A drawn sheet is the graph of *one* function belonging to one of these spaces.

Lemma 1.1 (Orthogonality and dimension)**.** *The displayed real modes form an orthogonal basis of V_q. The constant has squared norm 1. Every listed nonconstant sine or cosine has squared norm $1/2$. Consequently*

$$\dim_{\mathbb{R}} V_q = (2q+1)^2.$$

Proof. For an integer m, integration gives

$$\int_0^1 e^{2\pi i m t}\,dt = \begin{cases} 1 & m = 0, \\ 0 & m \ne 0. \end{cases}$$

Applying this in both coordinates proves orthogonality of distinct complex characters. Expressing sine and cosine in terms of the two characters at k and $-k$ gives the claimed real inner products. Our representative convention leaves no distinct basis elements with the same pair of frequencies.

There are $(2q+1)^2$ lattice frequencies in the square, including zero. The nonzero ones occur in opposite pairs. Each pair supplies two real basis vectors, so the total is $1 + 2((2q+1)^2 - 1)/2 = (2q+1)^2$. □

View	Cutoff	Dimension	What is admitted
X_0	1	9	Constant and lowest square-frequency modes
X_1	2	25	All preceding modes and the next band
X_2	4	81	All frequencies through square radius 4
X_3	8	289	All frequencies through square radius 8

1.1 The exact functions used in the atlas

Group the nine nonzero modes in the renderer into four disjoint bands:

$$b_0 = \tfrac{1}{2}C_{1,0} + \tfrac{3}{10}S_{0,1} + \tfrac{1}{5}C_{1,1}, \tag{1}$$
$$d_0 = \tfrac{1}{4}S_{2,1} + \tfrac{3}{20}C_{1,-2}, \tag{2}$$
$$d_1 = \tfrac{1}{8}C_{4,2} + \tfrac{1}{10}S_{3,-4}, \tag{3}$$
$$d_2 = \tfrac{1}{14}S_{8,3} + \tfrac{1}{16}C_{5,-7}. \tag{4}$$

The fine seed is $f = b_0 + d_0 + d_1 + d_2 \in X_3$. Its views are

$$x_0 = b_0, \quad x_1 = b_0 + d_0, \quad x_2 = b_0 + d_0 + d_1, \quad x_3 = f.$$

These are the actual coefficients in `render_komorebi.py`, not a replacement example chosen after plotting.

At $(u, v) = (0, 0)$, all sine terms vanish and all cosine terms equal one. Thus the intermediate values are

$$x_0(0,0) = \frac{7}{10}, \quad x_1(0,0) = \frac{17}{20}, \quad x_2(0,0) = \frac{39}{40}, \quad x_3(0,0) = \frac{83}{80}.$$

The increments are $3/20$, $1/8$, and $1/16$. These point evaluations illustrate the arithmetic. The proofs below compare full coefficient vectors, not a selection of sample points.

2 Restrictions, inclusions and compatibility

Let Π_q denote coefficient truncation to square cutoff q. On any trigonometric polynomial, it retains precisely the modes satisfying $\max(|a|, |b|) \le q$. Define

$$P_n = \Pi_{q_n}|_{X_{n+1}} : X_{n+1} \longrightarrow X_n, \qquad J_n : X_n \hookrightarrow X_{n+1}.$$

The inclusion J_n leaves the function unchanged and adds zero coefficients at the new frequencies. Restriction is *not* pointwise subtraction or downsampling of the rendered mesh.

Proposition 2.1 (Nested projections and sections)**.** *For any two cutoffs q, r,*

$$\Pi_q \Pi_r = \Pi_{\min(q,r)}, \qquad \Pi_q^2 = \Pi_q, \qquad P_n J_n = \mathrm{id}_{X_n}.$$

Each Π_q is an orthogonal projection for the normalized L^2 inner product.

Proof. A basis mode survives two truncations exactly when its frequency lies in both cutoff squares. Their intersection is the smaller square. This proves the first identity on every basis vector and hence on every linear combination. Idempotence is the equal-cutoff case. Every mode in X_n survives the inclusion followed by P_n, proving the section identity. The basis is orthogonal, so retaining a subset and deleting its orthogonal complement is an orthogonal projection. □

Proposition 2.2 (The plotted family is coherent)**.** *For $n = 0, 1, 2$, $P_n x_{n+1} = x_n$.*

Proof. Since $x_n = \Pi_{q_n} f$, nested projections give

$$P_n x_{n+1} = \Pi_{q_n} \Pi_{q_{n+1}} f = \Pi_{q_n} f = x_n.$$

For example, $P_1 x_2$ deletes the $(4, 2)$ and $(3, -4)$ modes in d_1 and leaves exactly x_1. □

The view category has arrows $i \to j$ for $i \le j$, while its restriction maps go from X_j to X_i. The adjacent restrictions generate the composite ones. Adjacent coherence therefore implies coherence along every arrow of this finite chain.

3 The realization space, not just the chosen realization

Put

$$E = X_0 \oplus X_1 \oplus X_2 \oplus X_3, \qquad R = X_0 \oplus X_1 \oplus X_2.$$

Here R records the three *adjacent* defects. The monograph's general arrow-indexed residual has further entries for composite arrows. Those entries are determined here by telescoping the adjacent defects, so this smaller residual space has the same zero set.

Define the linear compatibility operator

$$\delta(a_0, a_1, a_2, a_3) = (P_0 a_1 - a_0,\ P_1 a_2 - a_1,\ P_2 a_3 - a_2).$$

The finite Komorebi realization space is

$$\mathcal{K}_3 = \ker \delta = \{a \in E : P_n a_{n+1} = a_n \ (n = 0, 1, 2)\}.$$

A point is the entire compatible tuple. Each of its four components is a view of that point.

Theorem 3.1 (Finest-view parameterization)**.** *The map*

$$\Phi : X_3 \longrightarrow \mathcal{K}_3, \qquad g \longmapsto (P_0P_1P_2 g,\ P_1P_2 g,\ P_2 g,\ g)$$

is a linear isomorphism. Its inverse is evaluation at the finest view. In particular,

$$\dim E = 404, \qquad \dim \mathcal{K}_3 = 289.$$

Proof. The tuple $\Phi(g)$ satisfies every adjacent coherence equation by construction. Its finest component is g, proving injectivity. Conversely, for a coherent tuple a, the three equations imply $a_2 = P_2 a_3$, $a_1 = P_1P_2 a_3$, and $a_0 = P_0P_1P_2 a_3$. Thus $a = \Phi(a_3)$, proving surjectivity and the inverse formula. The dimension of E is $9 + 25 + 81 + 289 = 404$. □

Theorem 3.2 (Exact coherent and defect coordinates)**.** *For $r = (r_0, r_1, r_2) \in R$, define*

$$Sr = (-r_0 - P_0 r_1 - P_0P_1 r_2,\ -r_1 - P_1 r_2,\ -r_2,\ 0).$$

Then $\delta S = \mathrm{id}_R$. The operator

$$Ca = \Phi(a_3)$$

satisfies $C^2 = C$, $\operatorname{im} C = \mathcal{K}_3$, and

$$\boxed{a = Ca + S\delta a.}$$

Consequently $E \cong \mathcal{K}_3 \oplus R$, and $\operatorname{rank} \delta = 115$.

Proof. Compute the three entries of $\delta S r$. The first is

$$P_0(-r_1 - P_1 r_2) - (-r_0 - P_0 r_1 - P_0P_1 r_2) = r_0.$$

The second is $-P_1 r_2 - (-r_1 - P_1 r_2) = r_1$, and the third is $0 - (-r_2) = r_2$. Hence $\delta S = \mathrm{id}_R$, so δ is surjective.

The vector $a - S\delta a$ has zero defect and the same finest component as a, since $(Sr)_3 = 0$. The preceding theorem uniquely identifies it as Ca. Applying C to an already coherent tuple changes nothing, giving idempotence. The maps

$$a \longmapsto (Ca, \delta a), \qquad (k, r) \longmapsto k + Sr$$

are inverse: $CS = 0$, $\delta k = 0$, and $\delta S = \mathrm{id}_R$. Finally $\dim R = 9 + 25 + 81 = 115$. □

This is a concrete realization of coherent information plus removable defect. The repair C preserves the finest view. It is not generally the orthogonal projection onto $\mathcal{K}_3$ when all four views receive equal weight.

4 Retained detail and exact reconstruction

For a coherent family define

$$d_n = x_{n+1} - J_n x_n \in X_{n+1}.$$

Proposition 4.1 (The discarded band is retained as a coordinate)**.** *For each n, $P_n d_n = 0$ and $x_{n+1} = J_n x_n + d_n$. More generally,*

$$X_{n+1} = J_n X_n \oplus^{\perp} \ker P_n.$$

Proof. By coherence and the section identity,

$$P_n d_n = P_n x_{n+1} - P_n J_n x_n = x_n - x_n = 0.$$

The reconstruction equation is the definition rearranged. For arbitrary $z \in X_{n+1}$, use $z = J_n P_n z + (I - J_n P_n) z$. The second summand lies in $\ker P_n$. The two summands have disjoint frequency supports, so the sum is orthogonal and their intersection is zero. □

Substituting the reconstruction equation three times gives the fully typed identity

$$x_3 = J_2 J_1 J_0 x_0 + J_2 J_1 d_0 + J_2 d_1 + d_2.$$

Identifying each smaller function space with its inclusion into X_3 simplifies this to

$$\boxed{x_3 = x_0 + d_0 + d_1 + d_2.}$$

The corresponding coordinate decomposition is

$$\mathcal{K}_3 \cong X_0 \oplus \ker P_0 \oplus \ker P_1 \oplus \ker P_2,$$

with dimensions $9 + 16 + 56 + 208 = 289$. Thus a fixed coarse view x_0 leaves 280 real degrees of freedom at the finest level. The nine nonzero modes selected for the illustration form only one particular choice.

4.1 An exact energy calculation

Orthogonality of the disjoint bands yields

$$\|x_3\|_2^2 = \|x_0\|_2^2 + \|d_0\|_2^2 + \|d_1\|_2^2 + \|d_2\|_2^2 .$$

For example,

$$\|d_0\|_2^2 = \frac{1}{2}\left[\left(\frac{1}{4}\right)^2 + \left(\frac{3}{20}\right)^2\right] = \frac{17}{400}.$$

The complete arithmetic is

Component	Calculation	Squared norm
x_0	$\frac{1}{2}[(1/2)^2 + (3/10)^2 + (1/5)^2]$	19/100
d_0	$\frac{1}{2}[(1/4)^2 + (3/20)^2]$	17/400
d_1	$\frac{1}{2}[(1/8)^2 + (1/10)^2]$	41/3200
d_2	$\frac{1}{2}[(1/14)^2 + (1/16)^2]$	113/25088

Hence the finest energy missing from the two coarsest views is, respectively,

$$\|x_3 - x_0\|_2^2 = \frac{17}{400} + \frac{41}{3200} + \frac{113}{25088}, \qquad \|x_3 - x_1\|_2^2 = \frac{41}{3200} + \frac{113}{25088}.$$

None of these positive quantities is a compatibility defect. The family remains exactly coherent because restriction deliberately removes these modes.

5 Change one view and work out the repair

The companion plate uses

$$\eta = \frac{1}{5}C_{1,0}, \qquad y = (x_0, x_1, x_2 + \eta, x_3).$$

The perturbation belongs to X_0, so it survives every displayed restriction. Direct substitution gives

$$\begin{aligned}(\delta y)_0 &= P_0 x_1 - x_0 = 0,\\ (\delta y)_1 &= P_1(x_2 + \eta) - x_1 = \eta,\\ (\delta y)_2 &= P_2 x_3 - (x_2 + \eta) = -\eta.\end{aligned}$$

Therefore $\delta y = (0, \eta, -\eta)$. One altered view affects both adjacent comparisons, with opposite signs.

5.1 Repair while preserving views 0, 1 and 3

The correction

$$h = (0, 0, -\eta, 0)$$

is invisible on the protected views and satisfies $\delta h = (0, -\eta, \eta)$. Hence

$$y + h = x, \qquad \delta(y + h) = 0.$$

This repair is unique with these protected values: preserving view 3 already forces the entire coherent family by the finest-view theorem. Its squared correction cost in the direct-sum norm on E is

$$\|h\|_E^2 = \|\eta\|_2^2 = \frac{1}{2}\left(\frac{1}{5}\right)^2 = \frac{1}{50}.$$

The residual energy is instead $\|\delta y\|_R^2 = 2\,\|\eta\|_2^2 = 1/25$. These measure different things.

5.2 The closest repair when no view is protected

Give all four views equal weight in $\|a\|_E^2 = \sum_n \|a_n\|_2^2$. For a Fourier basis mode k, let

$$I_k = \{n : \max(|k_1|, |k_2|) \le q_n\}.$$

Coherence requires the coefficient of k to have one common value across the views in I_k. If the observed coefficients are $\alpha_{n,k}$, define

$$\bar{\alpha}_k = \frac{1}{|I_k|}\sum_{n \in I_k} \alpha_{n,k}.$$

Theorem 5.1 (Unconstrained least squares repair)**.** *Replacing each coefficient by its mean over the views in which it occurs gives the unique nearest coherent family.*

Proof. Different modes are orthogonal, so the objective separates into independent scalar problems. For one mode, its common basis norm contributes the same positive factor to each view. The identity

$$\sum_{n \in I_k} (\alpha_{n,k} - t)^2 = \sum_{n \in I_k} (\alpha_{n,k} - \bar{\alpha}_k)^2 + |I_k|(t - \bar{\alpha}_k)^2$$

proves that the unique minimizing common coefficient is $t = \bar{\alpha}_k$. Combining the finitely many minimizers yields a coherent family and the unique global minimum. □

For the displayed perturbation, the changed low-frequency mode is visible in all four views. Its mean shift is $\eta/4$. Thus

$$z = (x_0 + \eta/4,\ x_1 + \eta/4,\ x_2 + \eta/4,\ x_3 + \eta/4)$$

is the unprotected nearest coherent family. Its correction to y is

$$z - y = (\eta/4, \eta/4, -3\eta/4, \eta/4),$$

and

$$\|z - y\|_E^2 = \frac{1+1+9+1}{16}\,\|\eta\|_2^2 = \frac{3}{200} < \frac{1}{50}.$$

It is cheaper because it changes the previously protected views. The two answers solve different declared optimization problems. Neither can be substituted silently for the other.

6 An exact transfer between representations

Define the coordinate-swap transfer

$$(Tf)(u, v) = f(v, u).$$

It maps X_n to itself because swapping (a, b) to (b, a) preserves the square cutoff. The representative convention may introduce a sine sign. For example,

$$TS_{3,-4} = S_{-4,3} = -S_{4,-3}.$$

Proposition 6.1 (Naturality of the swap)**.** *$T^2 = I$, $\Pi_q T = T\Pi_q$, and therefore*

$$P_n T|_{X_{n+1}} = T|_{X_n} P_n.$$

It follows that T maps coherent families to coherent families and transports their details to the corresponding details.

Proof. Swapping twice restores the original coordinates. A mode is retained by Π_q exactly when its swapped frequency is retained, proving commutation on every mode. Linearity gives the operator identity. Thus for a coherent family,

$$P_n(Tx_{n+1}) = T(P_n x_{n+1}) = Tx_n.$$

The inclusions also commute with T, so the new detail is $Tx_{n+1} - J_n Tx_n = T(x_{n+1} - J_n x_n) = Td_n$. □

The cutoff symmetry is a hypothesis, not a stylistic choice. If a truncation retains $|a| \le 1$, $|b| \le 2$, it keeps $C_{0,2}$ but discards its swap $C_{2,0}$. That asymmetric truncation does not commute with T. The delivered verifier includes this as a negative control.

7 A residual is not yet an obstruction

For a declared linear compatibility map $\delta : E \to R$, let p_U select the values that must remain fixed. Allowed corrections form $N_U = \ker p_U$. The possible changes in residual form $\delta(N_U)$. A raw lift a of prescribed data s therefore defines

$$\mathfrak{o}_U(s) = [\delta a] \in R/\delta(N_U).$$

For coherent partial sections this is the additive extension obstruction of the monograph [1, Section 3].

Theorem 7.1 (Protected extension criterion)**.** *The class does not depend on the chosen raw lift. It vanishes exactly when some coherent family agrees with the prescribed data.*

Proof. Two raw lifts differ by an element of N_U, so their defects differ by an element of $\delta(N_U)$. This proves independence. The class vanishes precisely when $\delta a = \delta h$ for an allowed $h \in N_U$. Then $a - h$ is coherent and retains the prescribed values. Conversely, any coherent extension x gives $h = a - x \in N_U$ and $\delta h = \delta a$. □

For the altered Fourier family, protection of views $0, 1, 3$ permits changes only at view 2. Such changes have the form

$$h = (0, 0, k, 0), \qquad \delta h = (0, P_1 k, -k).$$

On the adjacent residual space the quotient can be represented explicitly by

$$\Omega(r_0, r_1, r_2) = (r_0,\ r_1 + P_1 r_2) \in X_0 \oplus X_1.$$

Indeed Ω is surjective, and $\ker \Omega = \delta(N_U)$: if $r_0 = 0$ and $r_1 = -P_1 r_2$, choose $k = -r_2$. Thus this quotient has dimension $9 + 25 = 34$, but the displayed residual has zero class:

$$\Omega(0, \eta, -\eta) = (0, 0).$$

The nonzero defect is entirely removable by an allowed change.

7.1 Why this chain does not exhibit a genuine extension obstruction

Every internally coherent partial section of this finite split tower extends. To prove this, let m be its largest prescribed index. Include the value at m into X_3 by adding zero higher coefficients, then restrict to all lower views. Coherence on the prescribed full subcategory ensures the required lower values are recovered. If no views are prescribed, the zero family is an extension.

Thus this specific chain is an excellent model of views, retained detail and repair, but not of a nonzero obstruction for internally coherent partial data. Fixing inconsistent values at two comparable levels would not supply such an example: those values would already fail to form a coherent partial section.

7.2 A small branching model with a nonzero obstruction

Use the view category $a \to w \leftarrow b$ and real fibers. Its restriction maps are

$$\rho_a(t) = t, \qquad \rho_b(t) = 2t.$$

A coherent full family has the form $(x_a, x_w, x_b) = (t, t, 2t)$. Prescribe the two endpoint views $U = \{a, b\}$ with values $(1, 3)$. This is internally coherent because there is no nonidentity arrow between the two endpoints.

A raw lift is $(1, t, 3)$ and its defect is $(t - 1, 2t - 3)$. Corrections preserving both endpoints change only t and produce defects $(h, 2h)$. The quotient map

$$\omega(r_a, r_b) = 2r_a - r_b$$

has kernel exactly $\{(h, 2h) : h \in \mathbb{R}\}$. Therefore

$$\mathfrak{o}_U(1, 3) \mathrel{\widehat{=}} 2(t - 1) - (2t - 3) = 1 \neq 0.$$

No choice of the middle value repairs the two protected observations. This branching example, unlike the altered Fourier tuple, demonstrates a genuine nonzero extension obstruction.

8 What changes for infinitely many resolutions?

Now let $X_n = V_{2^n}$ for every $n \geq 0$. Let H be the Hilbert completion of the union of these spaces under the normalized L^2 norm. Defining H this way supplies exactly the completeness and density used below, without importing a separate Fourier completeness theorem.

The algebraic realization space consists of all sequences (x_n) satisfying $P_n x_{n+1} = x_n$. This definition by itself imposes no uniform norm bound.

Theorem 8.1 (A finite energy budget is exactly the missing condition)**.** *For a coherent infinite family, the following are equivalent:*

(i) There is a unique $f \in H$ with $\Pi_{2^n} f = x_n$ for every n.
(ii) $\sup_n \|x_n\|_2 < \infty$.
(iii) $\|x_0\|_2^2 + \sum_{n\geq 0} \|d_n\|_2^2 < \infty$, where $d_n = x_{n+1} - J_n x_n$.

When these hold, $x_n \to f$ in norm and

$$\|f\|_2^2 = \|x_0\|_2^2 + \sum_{n\geq 0} \|d_n\|_2^2\,, \qquad \|f - x_n\|_2^2 = \sum_{j\geq n} \|d_j\|_2^2\,.$$

Proof. The finite orthogonal decomposition gives

$$\|x_n\|_2^2 = \|x_0\|_2^2 + \sum_{j<n} \|d_j\|_2^2\,,$$

so (ii) and (iii) are equivalent. If (i) holds, projection is contractive, giving $\|x_n\|_2 \leq \|f\|_2$ and hence (ii).

Assume (ii). For $m \geq n$, coherence says $\Pi_{2^n} x_m = x_n$, so $x_m - x_n$ is orthogonal to x_n. Consequently

$$\|x_m - x_n\|_2^2 = \|x_m\|_2^2 - \|x_n\|_2^2\,.$$

The squared norms are increasing and bounded. They converge, making (x_n) Cauchy. Completeness supplies $f \in H$. Each finite projection is a bounded operator and hence extends continuously to H. Passing to the limit in $\Pi_{2^n} x_m = x_n$ gives $\Pi_{2^n} f = x_n$.

If f and g have the same projections, $f - g$ is orthogonal to every X_n, hence to their dense union, and therefore equals zero. This proves uniqueness. Taking limits in the finite energy identities proves both remaining formulas. □

Example 8.2 (Perfect finite coherence without a Hilbert realization)**.** Set

$$w_n(u, v) = \sum_{k=1}^{2^n} \cos(2\pi k u).$$

Every adjacent restriction satisfies $P_n w_{n+1} = w_n$. Yet orthogonality gives

$$\|w_n\|_2^2 = \frac{2^n}{2} = 2^{n-1} \longrightarrow \infty.$$

Thus (w_n) is an algebraically coherent point but is not the projection family of an element of H.

The displayed nine-mode seed has a Hilbert realization: retain that same polynomial at all finer levels and set every additional detail to zero. It is only one possible continuation of the first four views. Those four views alone do not determine unknown higher-frequency coefficients. Uniqueness in the theorem concerns an entire infinite coherent family satisfying the norm condition.

9 Verification record and provenance

The following checks were actually executed for this edition:

Verification group	Passing exact checks
Original visual-atlas verifier, replayed	8,969
Additional worked-example verifier	3,695
Total	12,664

The original verifier tests projector identities, nesting, split reconstruction and coordinate-swap transport on all 289 basis vectors of X_3. The added verifier checks every basis vector of the 404-dimensional raw family space and the 115-dimensional adjacent residual space. It verifies both repairs, the right inverse, explicit obstruction coordinates, energies and point calculations. All 81 basis directions of a view-2-only correction are also checked.

Negative controls detect an omitted detail, an incorrect residual sign, confusion between pointwise difference and defect, and a cutoff that breaks coordinate-swap compatibility. The branching obstruction is evaluated at several raw lifts. Its independence from every lift is proved in Section 7.

Coefficient arithmetic uses Python's `Fraction`. Complete basis verification covers the stated linear identities over the real numbers because the operators are linear and have the tested rational matrices. Orthogonality is justified by the written proof. Trigonometric values and page geometry used for drawing remain floating-point calculations. No Lean compilation was performed. The infinite theorem is a written proof, not a consequence of a finite check count.

Reproduction

From the delivered directory, run

```
python verify_worked_example.py --output .
python render_komorebi.py --output figures --dpi 300
```

The first command writes the verification and intermediate-step JSON records. The second recreates the plates. NumPy and Matplotlib must already be installed for offline use. Rebuild this PDF with two runs of `pdflatex` on the delivered source.

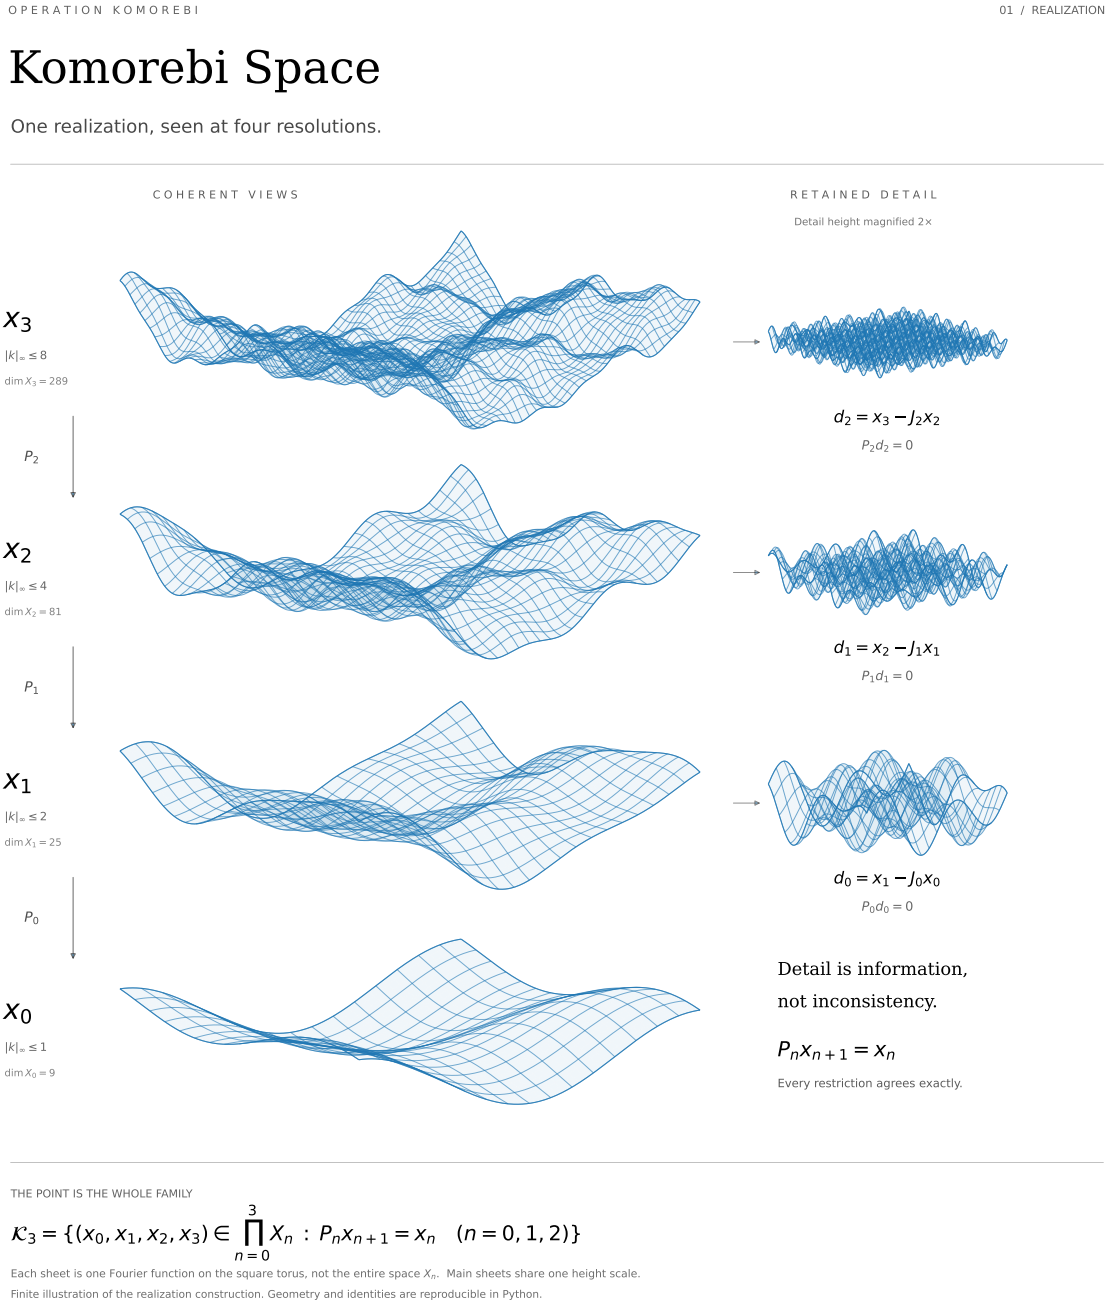
Sources

[**1**] *Komorebi Space Monograph v1.13.0*, Sections 1, 3, 6 and 10: coherent realization, additive obstruction, exact tower transport, split reconstruction and convergence.

[**2**] *Komorebi Space Visual Atlas v3.* The original renderer and its two plates. The seed, cutoffs, perturbation and plotting conventions are unchanged.

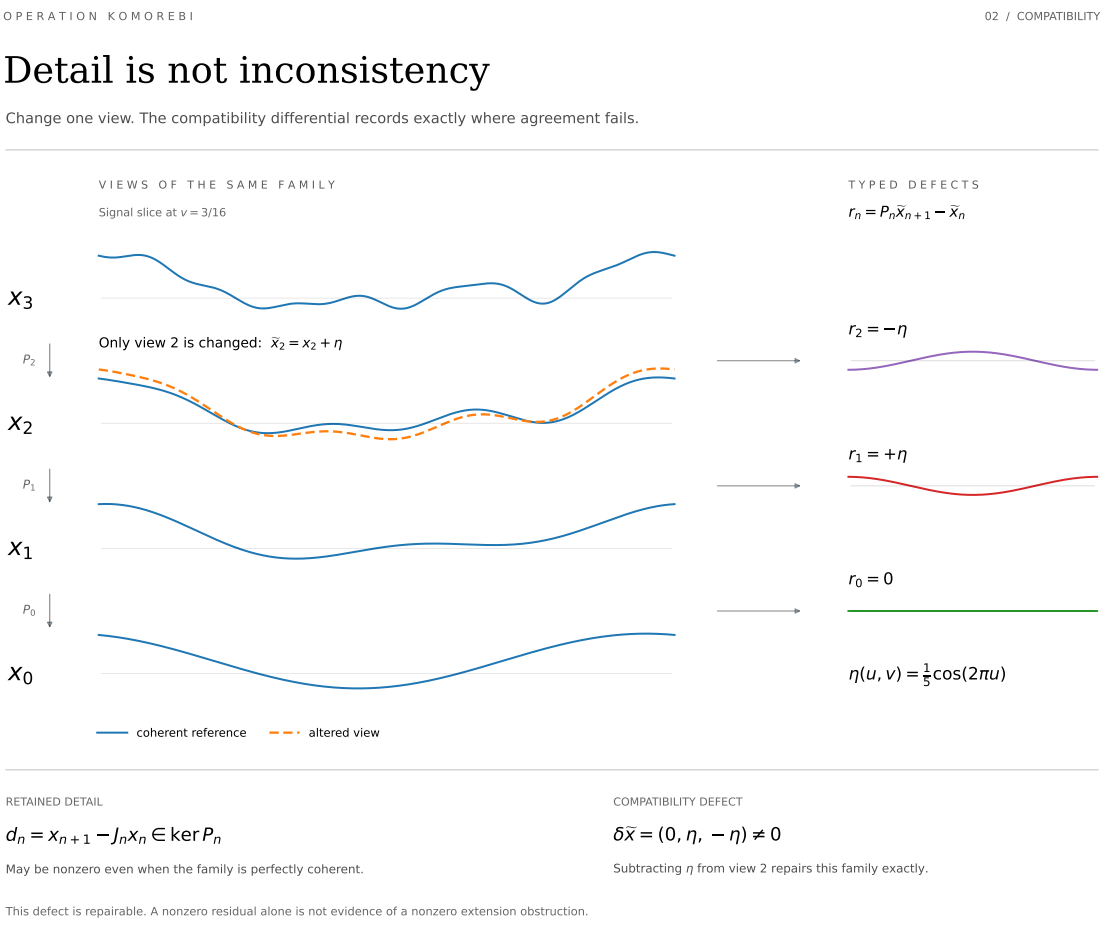
[**3**] *Worked-example audit v4.* The delivered verifier, `verification.json` and `worked_steps.json`: complete finite basis checks, intermediate coefficients and exact scalar calculations.

Plate 1 One realization at four resolutions



Main sheets are graphs of x_0, x_1, x_2, x_3 from Section 1. Small sheets show the details from Section 4 with twice the main vertical gain. Their placement separates the views for reading. It adds no mathematical coordinate to $\mathcal{K}_3$.

Plate 2 Detail, defect and exact repair



The curves are the slice $v = 3/16$ of the exact family and the altered view. The right-hand curves show the typed adjacent defects, with the same vertical gain as the signal curves. Section 5 proves their signs and the protected repair. Section 7 proves why this nonzero residual still has zero protected extension obstruction.

The distinction to retain. Detail records what a lower resolution cannot see. Defect records disagreement between resolutions. Obstruction records whether the disagreement can be removed without changing the prescribed information.