PAPER V

Multiresolution Signal Calculus

Exact coordinates, singular packets, and stable reconstruction

Forty worked pages trace exact signal coordinates through stable filter banks, finite packet trees, singular trace packets, and conditional recovery.

Komorebi P.

Visual guide • 20 October 2026

How to use this volume

Forty worked pages and thirty-seven editable vector figures follow exact signal coordinates, a fully enumerated packet tree, and a reflection orbit, then extend the route to stable banks, singular trace packets, and conditional recovery of separated jumps.

Part I / A route through the examples

The labels V1 through V40 identify worked pages independently of PDF pagination. Each page names its corresponding statements in the monograph traditional edition. The Lean guide records the separate formal contract.

V1 Three objects, three different questions

V2 An atlas of the objects

V3 Decimation is a permutation, not a loss

V4 Refine packets on a common lattice

V5 A filter becomes a matrix of actual delays

V6 One sample is not one channel step

V7 Perfect reconstruction is an operator identity

V8 A jump can hide from the first Haar split

V9 Lifting preserves information, not energy

V10 The second level resolves the hidden change

V11 Deleting coefficients is the approximation step

V12 A mixed-depth basis must cover every branch

V13 Solve every alternative without searching every alternative

V14 How far can the costs move before the tree changes?

V15 Retain the jump and approximate the other packet

V16 Why a Fourier cutoff cannot reproduce this jump pointwise

V17 Small location errors, different norms

V18 Separate amplitude error from mask error

V19 Geometry localizes error while data may leave it ambiguous

V20 Two products, two packet laws

V21 Nonlinear neighbors cross channel boundaries

V22 One reflection circle, four very different impedances

V23 A pole is a chart boundary, not a broken group action

V24 Residue labels are not impedance averages

V25 Sensitivity grows before the pole is reached

V26 What the reader can now verify

Three reading routes

Signal coordinates: V3 to V11, then V20 and V21. **Selection and error:** V12 to V19. **Nonlinear geometry:** V22 to V25.

One model per question

The sixteen-sample signal gives exact rational values. The depth-three tree has twenty-six complete prunings. The reflection orbit has four exact quarter-turn points. These bounded examples expose the paper's mechanisms without replacing its general statements.

Read the hypotheses with the picture

Zero extension, periodic indexing, and infinite sequences describe different objects. Known jumps differ from estimated jumps. A finite optimum depends on a specified cost, and a coordinate transformation does not establish an optimal sampling rule. V26 records the evidence behind this guide.

Part II / From coordinates to recovery

The expanded traditional development supplies additional theorems beyond the inherited twenty-result formal contract. These fourteen new worked pages follow its hypotheses through exact examples and quantitative bounds.

V27 What the extended developments add

V28 Six image phases are six exact address classes

V29 A positive singular-value gap survives perturbation

V30 Nonunitary tree bounds propagate along branches

V31 Derivative jumps have their own packets

V32 The approximation operator also acts on packet error

V33 An overlap matrix is an exact error ledger

V34 A shifted step and a shifted ramp have different moduli

V35 A strict contrast gap locates every separated jump cell

V36 Repeated observations shrink the noise side of the gap

V37 Haar sees a jump only through its position in the support

V38 Two half-shifted grids remove exact alignment failure

V39 Higher differences annihilate polynomial interiors

V40 Conditional Gibbs control closes the reconstruction estimate

A connected route

The lattice example extends exact addressing. Frame bounds then quantify the stability of analysis and synthesis. Trace packets remove specified singularities, while overlap and location formulas measure errors in their estimates. Three local detectors address separated jumps under contrast hypotheses. The final page combines the recovered packet with a chosen approximation operator.

Keep the evidence attached to its claim

The twenty-result core carries its preserved Lean map. Part II carries traditional proofs and the current mathematical review, with a separate exact worked-example checker. The rerun symbolic visual suite, current arithmetic runs, saved compiler transcripts, and rendered figures each retain their own scope. Neither a figure nor a finite check expands the formal theorem map.

What the example numbers mean

The additional checker uses exact rational arithmetic. Its logarithmic sampling budget is enclosed using a positive Taylor partial sum. Curve coordinates are converted to decimal positions for drawing. Displayed formulas retain exact values. All source-drawn plates are editable TikZ, and the final comparison figure is preserved from the expanded manuscript's vector source.

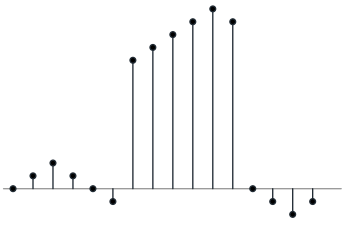
V1 / PAPER V / A WORKED VISUAL COMPANION

Three objects, three different questions

Traditional monograph: Section 1 (p. 2), Section 6 (p. 11), Section 11 (p. 19), Section 13 (p. 24).

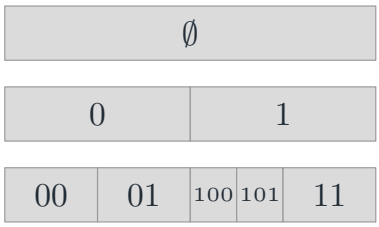
A packet calculus can reorganize a signal exactly, choose among finitely many representations, or describe a nonlinear coordinate geometry. These are not the same task. This guide keeps one concrete object in view for each, so that the reader can compute rather than only follow arrows.

SIGNAL COORDINATES



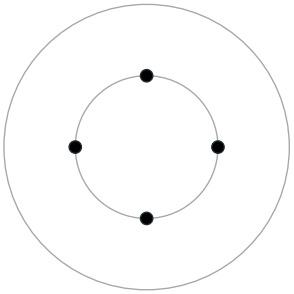
$f = s + j$ in $\mathbb{Q}^{16}$

FINITE TREE CHOICES



26 complete prunings

REFLECTION ORBIT



$\Gamma = \frac{1}{2} i^k$

A signal that can be checked by hand

Our main finite signal is $f = s + j \in \mathbb{Q}^{16}$. The background repeats

$$\tfrac{1}{4}(0, 1, 2, 1, 0, -1, -2, -1),$$

and $j[n] = 3$ at $6 \le n < 12$, with $j[n] = 0$ elsewhere. This is a discrete background model, not an assertion of continuum smoothness. Its energy is exactly 57. We use it for residue coordinates, Haar transforms, coefficient deletion, and known-packet cancellation.

A tree whose alternatives can all be seen

The full binary tree of depth three has only 26 complete prunings. We enumerate all of them, compute a unique mixed-depth optimum, and certify its stability with rational costs. The displayed objective is deliberately explicit, not claimed to be a universal signal-quality score.

A coordinate orbit with no hidden numerical poles

The four reflection points $\Gamma = i^k/2$ map to four exact impedances. A second orbit passes through infinity; homogeneous coordinates make that point an ordinary part of the construction.

> **Read this as:** Exact coordinates preserve information. An optimizer chooses coordinates. An estimator must still identify information not already supplied.

V2 / READING CONVENTIONS / KEEP THE TYPES VISIBLE

An atlas of the objects

Traditional monograph: Locator 5 (p. 34), Locator D (p. 35).

Several symbols in the paper describe different realizations of the same operator language. The following distinctions are used throughout this guide.

Object	Literal meaning	What it is not
S, δ_0	Sample shift and unit impulse	A finite truncation of every ℓ^2 signal
$P_{N,r} f$	A masked signal on the original lattice	The shorter decimated channel
$f_r[q]$	The value $f[Nq + r]$	An average of neighboring samples
D	One channel-index delay	Necessarily one original-sample delay
$E(D), R(D)$	Finite Laurent analysis and synthesis matrices	Values at one test frequency
W_α	The addressed Haar analysis operator	An edge detector by itself
$W_\alpha^* W_\alpha$	A projector on the original signal space	A scalar packet energy
Λ	A complete covering antichain of a finite tree	An arbitrary collection of nodes
$j, \hat{j}$	Known and estimated discontinuity packets	Interchangeable hypotheses
Γ, z	Reflection coefficient and normalized impedance	Coefficients that can be linearly interchanged

Three arithmetic conventions

FIR examples use bilateral sequences with finite support and *zero extension*. Haar examples use vectors of length 16, with lengths divisible by the required powers of two. The three-tap smoother and nonlinear neighbor example use *periodic* indexing. These boundary conventions are stated where they enter. None is silently substituted for another. The polynomial variable t on V20 indexes coefficients. The delay symbol D acts on channels. The Smith variable z is impedance. Keeping them separate prevents a visual resemblance from becoming an invalid identity.

> **Read this as:** A correct formula needs the correct object, product, norm, and boundary convention. A familiar-looking matrix does not supply them automatically.

V3 / EXACT ADDRESSES / THE RUNNING SIGNAL

Decimation is a permutation, not a loss

Traditional monograph: Theorem 2.1 (p. 3), Corollary 2.2 (p. 3).

For four phases, write every sample address uniquely as $n = 4q + r$ with $0 \le r < 4$. Analysis stores $f[n]$ at channel address (r, q). Synthesis puts that exact value back at $4q + r$.

FOUR PHASES: EVERY ADDRESS APPEARS ONCE

	$q=0$	$q=1$	$q=2$	$q=3$	
$r=0$	$n=0:\ 0$	$n=4:\ 0$	$n=8:\ 3$	$n=12:\ 0$	$n = 4q + r$
$r=1$	$n=1:\ \frac{1}{4}$	$n=5:\ -\frac{1}{4}$	$n=9:\ \frac{13}{4}$	$n=13:\ -\frac{1}{4}$	
$r=2$	$n=2:\ \frac{1}{2}$	$n=6:\ \frac{5}{2}$	$n=10:\ \frac{7}{2}$	$n=14:\ -\frac{1}{2}$	$-1 = 4(-1) + 3$
$r=3$	$n=3:\ \frac{1}{4}$	$n=7:\ \frac{11}{4}$	$n=11:\ \frac{13}{4}$	$n=15:\ -\frac{1}{4}$	$-5 = 4(-2) + 3$

Every cell contains an original address and its rational sample value. Shading marks the known jump packet, not an estimated feature or a frequency band.

What has changed? Only where the values are stored

The complete tuple of channels contains sixteen values, just as the original vector does. A single decimated channel is shorter, but keeping all four loses no samples. The residue projector $P_{4,r} f$ is different: it keeps the original lattice and writes zeros outside phase r.

$$\mathcal{S}_4 \mathcal{A}_4 f = f, \qquad \|f\|_2^2 = \sum_{r=0}^{3} \|f_r\|_2^2 = 57.$$

The norm on the channel tuple is the sum-of-squares Hilbert norm, not the maximum of the channel norms. This distinction matters in the infinite-dimensional statement.

Negative addresses are ordinary addresses

Euclidean division, rather than truncation toward zero, gives $-1 = 4(-1) + 3$ and $-5 = 4(-2) + 3$. The finite tests exercise both positive and negative sample addresses. The algebraic identity applies before any plotting window or finite buffer is chosen.

> **Read this as:** Downsampling one channel discards information. Splitting into all residue channels is an invertible address change.

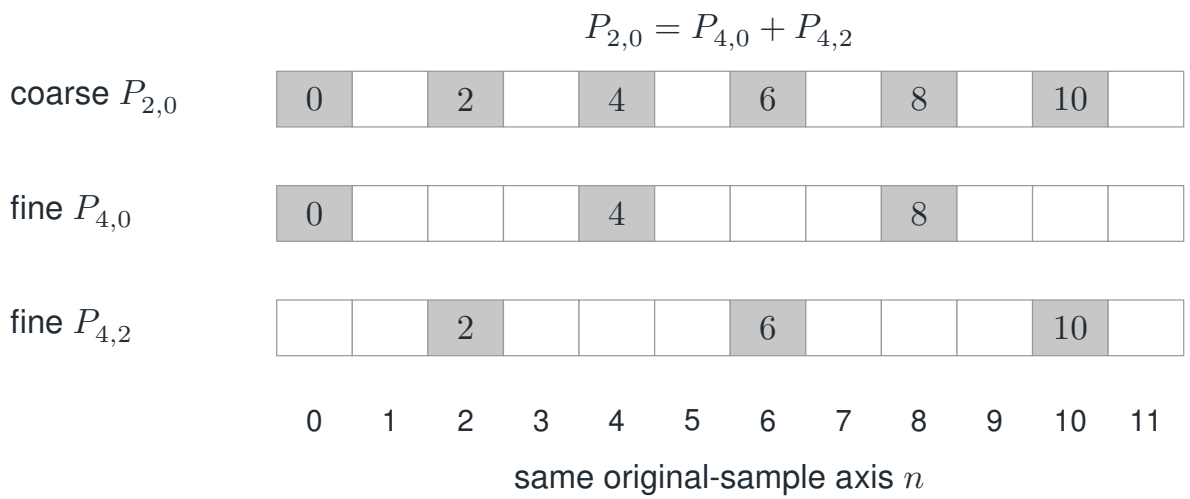
V4 / EXACT REFINEMENT / WHAT CAN BE ADDED

Refine packets on a common lattice

Traditional monograph: Theorem 2.3 (p. 5).

An even sample belongs to exactly one of residues 0 and 2 modulo four. Therefore the coarse even mask is the sum of two disjoint fine masks:

$$P_{2,0} = P_{4,0} + P_{4,2}.$$



The common horizontal axis is the original sample index. Empty cells denote masked-out addresses. There are no interpolated samples between the displayed cells.

Why the coordinate convention matters

The decimated channel $f_0[q] = f[2q]$ has its own index q. At modulus four the descendants are $f[4q]$ and $f[4q+2]$. These shorter arrays must be interleaved to recover the coarse channel. Adding them entrywise instead produces sums of different original samples. For example, at the first descendant address the values $f[0] = 0$ and $f[2] = 1/2$ are two separate values. Their sum $1/2$ cannot tell us both of them. Exact refinement is a disjoint partition, not an averaging rule or a new measurement.

The general formula

When N divides M,

$$P_{N,r} = \sum_{\substack{0 \le a < M \\ a \equiv r \pmod{N}}} P_{M,a}.$$

The divisibility hypothesis ensures that every fine residue class lies entirely inside one coarse class. Without it, this particular descendant relation is not the right refinement rule.

> **Read this as:** Finer coordinates can preserve the object exactly. Compressing their values to averages is an additional operation, with an additional information-loss question.

V5 / FIR TRANSFER / A FOUR-TAP EXAMPLE

A filter becomes a matrix of actual delays

Traditional monograph: Theorem 3.1 (p. 5).

Take the causal filter $h = (1, 2, -1, 1)$, with taps at $0, 1, 2, 3$, acting by $y[n] = \sum_q h[q] f[n-q]$. Resolve the input into even and odd channels. Let D denote one delay in the *channel* index.

$$h = (1, 2, -1, 1), \qquad D : \; u[n] \mapsto u[n-1]$$

	input 0	input 1
output 0	$1 - D$	$D^2 + 2D$
output 1	$D + 2$	$1 - D$

$$y_0 = (1 - D)x_0 + (2D + D^2)x_1$$

Reading one matrix entry

The odd-input contribution to the even output is $(2D+D^2)f_1$. The tap at original delay one reaches the preceding odd channel sample; the tap at delay three reaches the odd channel two positions back. These are different delays despite coming from the same input phase. The two channel equations are

$$y_0[n] = f_0[n] - f_0[n-1] + 2f_1[n-1] + f_1[n-2],$$
$$y_1[n] = 2f_0[n] + f_0[n-1] + f_1[n] - f_1[n-1].$$

Interleaving gives exactly the direct convolution, including its boundary samples under zero extension.

The general address calculation

For N phases, the block from input s to output k has kernel $h_{ks}[m] = h[Nm + k - s]$. This comes from the exact identity $Nn + k - q = N(n - m) + s$. It is a finite reindexing of taps, so no numerical transform of the input is needed. Evaluation of a bilateral transform is a separate step with the summability hypotheses stated in the manuscript.

> **Read this as:** The matrix does not approximate the filter. Its entries record which input phase is used, how far back to read, and what coefficient to multiply.

V6 / DELAY UNITS / AN IMPULSE MAKES THE DIFFERENCE VISIBLE

One sample is not one channel step

Traditional monograph: Theorem 4.1 (p. 7).

A one-sample delay swaps the two phases and wraps one of them into the preceding channel index. Its matrix is $J(D)$, not DI_2.

one original sample

0	D
1	0

$$J(D)^2 =$$

two original samples

D	0
0	D

$$\delta_0 \mapsto \delta_1$$

$$\delta_0 \mapsto \delta_2$$

$$y_0[n] = x_1[n-1], \quad y_1[n] = x_0[n]$$

channel delay 1 at $N = 2$ is original-sample delay 2

Trace one impulse

The input δ_0 has a single nonzero value in even channel address zero. Under $J(D)$ it appears in odd channel address zero: this is the original sample δ_1. Under DI_2 it stays even but moves to channel address one: this is δ_2. Thus

$$J(D)^2 = DI_2, \qquad H(D) = I_2 + 2J(D) - J(D)^2 + J(D)^3$$

recovers the previous page's four-tap transfer matrix. The same delay algebra explains both the off-diagonal phase swap and the channel powers.

Signed delays need the same care

For an original-sample delay d at modulus N,

$$y_k[n] = f_{(k-d) \bmod N}\left[n + \left\lfloor \frac{k-d}{N} \right\rfloor\right].$$

This includes advances when $d < 0$. The finite verification checks the formula on signed sample addresses. A bank with product $D^q I_N$ delays the original signal by Nq samples, not q samples.

> **Read this as:** Always state the unit of a delay. A phase permutation and a channel delay together can represent one original sample.

V7 / FILTER BANKS / WHY A FREQUENCY SPOT-CHECK IS INSUFFICIENT

Perfect reconstruction is an operator identity

Traditional monograph: Theorem 4.1 (p. 7).

Consider the finite Laurent bank

$$E(D) = \begin{pmatrix} 1 & D \\ 0 & 1 \end{pmatrix}, \qquad R(D) = \begin{pmatrix} 1 & -D \\ 0 & 1 \end{pmatrix}.$$

Its synthesis cancels the exact delayed cross-channel contribution introduced by analysis.

$E(D)$

1	D
0	1

$R(D)$

1	$-D$
0	1

$$R(D)E(D) = I$$

every impulse is recovered

$$B(D) = \operatorname{diag}(D, 1): \quad B(1) = I, \quad B(D) \neq I$$

testing one frequency cannot detect a channel delay

Two inverses, all inputs

Direct polynomial multiplication gives $RE = ER = I_2$. Since D acts as a shift, these identities imply perfect reconstruction on every bilateral channel sequence. With finite kernels, composition requires only finitely many operations at each address. The converse uses impulses. Insert an impulse in one input channel. Each output channel lists every coefficient of the corresponding transfer polynomial. If all such responses vanish for $RE - I$, every coefficient is zero; hence the polynomial matrix itself is zero.

A deliberately misleading test

The matrix $B(D) = \operatorname{diag}(D, 1)$ satisfies $B(1) = I_2$. At that one transform value it appears perfect. In the time domain, however, it delays only the first channel. An impulse in that channel immediately distinguishes it from the identity. This is why a small set of successful frequency evaluations is not automatically a proof of the full Laurent identity. A separate degree bound and a sufficient interpolation argument would be needed to turn finitely many evaluations into such a certificate.

> **Read this as:** Exact reconstruction concerns every coefficient of the transfer operator, not just an attractive response curve or one evaluation point.

V8 / HAAR ENERGY / EDGE ALIGNMENT

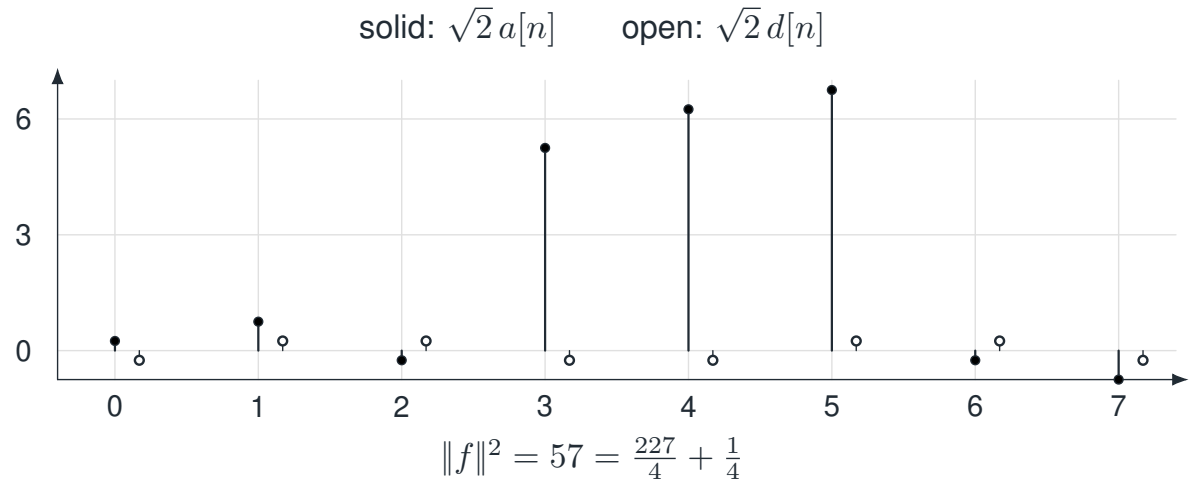
A jump can hide from the first Haar split

Traditional monograph: Corollary 4.2 (p. 8).

For the running signal, Haar analysis gives

$$a[n] = \frac{f[2n] + f[2n+1]}{\sqrt{2}}, \qquad d[n] = \frac{f[2n] - f[2n+1]}{\sqrt{2}}.$$

The plot scales both channels by $\sqrt{2}$ so that all displayed ordinates are rational.



$$\|f\|^2 = 57 = \tfrac{227}{4} + \tfrac{1}{4}$$

The exact energy calculation

The low channel has energy $227/4$, and the detail channel has energy $1/4$. Their sum is 57, exactly the input energy. The normalization is essential: unnormalized sums and differences would multiply the total energy by two.

Why is the first detail channel so small?

The jump starts at sample 6 and ends just before sample 12. Both lie *between* the pairs $(0,1), (2,3), \ldots, (14,15)$. Inside every pair the jump contributes either $(0,0)$ or $(3,3)$, so its first-level difference is zero. The small detail values come entirely from the discrete background. This does not mean the edge is absent from the transform. The next split of the low channel compares neighboring pairs, and those comparisons do cross the jump. The next pages expose that contribution explicitly.

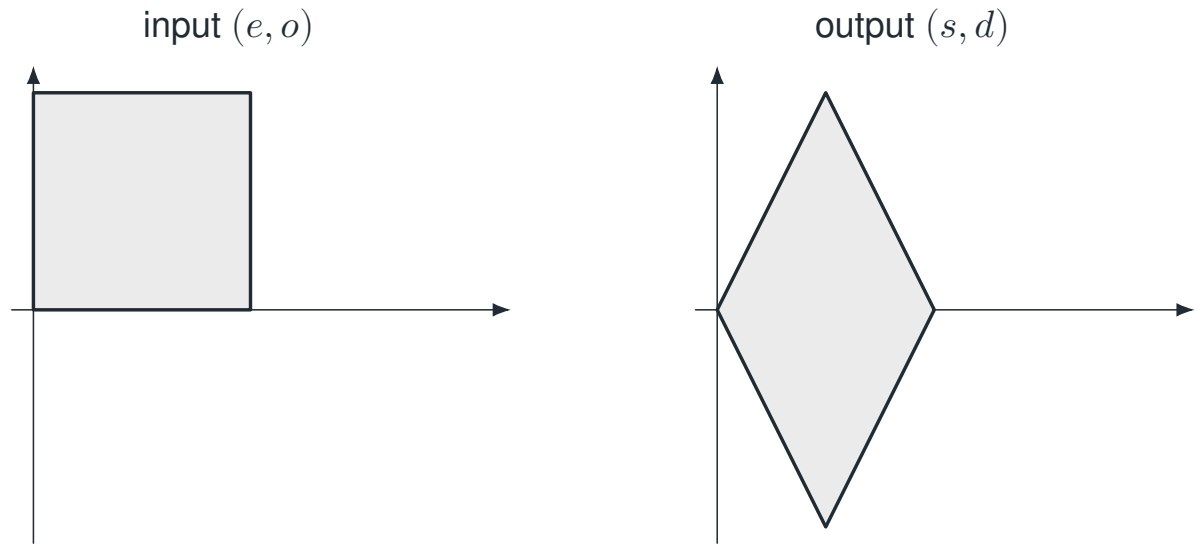
> **Read this as:** Haar energy is exact, but a particular detail level is not a universal edge detector. Edge visibility depends on alignment, scale, and the chosen packets.

V9 / PREDICT–UPDATE / INVERTIBILITY VERSUS ISOMETRY

Lifting preserves information, not energy

Traditional monograph: Theorem 4.3 (p. 10).

For an even/odd pair (e, o), form the prediction residual $d = o - e$, then update $s = e + d/2$. The inverse reads $e = s - d/2$, followed by $o = d + e$.



The square $[0, 4]^2$ and its exact image under the lifting matrix have the same area. The separate pair $(2, 6)$ illustrates that Euclidean energy need not be preserved.

$$\begin{pmatrix} s \\ d \end{pmatrix} = \begin{pmatrix} 1/2 & 1/2 \\ -1 & 1 \end{pmatrix} \begin{pmatrix} e \\ o \end{pmatrix}, \qquad \begin{pmatrix} e \\ o \end{pmatrix} = \begin{pmatrix} 1 & -1/2 \\ 1 & 1/2 \end{pmatrix} \begin{pmatrix} s \\ d \end{pmatrix}.$$

What determinant one does establish

The matrix is invertible, and in these real coordinates it preserves oriented area. The predict and update operations can be undone in reverse order. Applied to each pair, the construction gives exact finite-signal reconstruction.

What determinant one does not establish

The pair $(2, 6)$ becomes $(4, 4)$. Its squared Euclidean norm changes from 40 to 32. The lifting matrix is not unitary. Energy accounting must therefore use its actual coordinate metric or a separately normalized transform; invertibility alone cannot replace the Haar energy identity. All values in this example are rational. Integer-to-integer lifting with a rounding rule is a different algorithm. Its inverse and error conventions require their own specification and are not silently inferred from the displayed rational matrix.

> **Read this as:** Perfect reconstruction, determinant one, and norm preservation are different properties. This transform has the first two, not the third.

V10 / ACTUAL WAVELET PACKETS / FOUR RATIONAL CHANNELS

The second level resolves the hidden change

Traditional monograph: Theorem 5.1 (p. 10).

Apply Haar analysis to *both* children. A binary address records the chronological low/detail choices, with 0 for low and 1 for detail. Each depth-two channel contains four coefficients.

	coefficient positions within each packet				energy
00	$\frac{1}{2}$	$\frac{5}{2}$	$\frac{13}{2}$	$-\frac{1}{2}$	49
01	$-\frac{1}{4}$	$-\frac{11}{4}$	$-\frac{1}{4}$	$\frac{1}{4}$	$\frac{31}{4}$
10	0	0	0	0	0
11	$-\frac{1}{4}$	$\frac{1}{4}$	$-\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$

$$\sum_{\alpha\in\{0,1\}^2} W_\alpha^* W_\alpha = I_{16}$$

The edge that first-level detail missed

Packet 01 is the detail child of the first low channel. It contains $-11/4$, the large change between neighboring pair averages. Packet 10 vanishes for this particular signal because the first-level differences alternate in cancelling pairs. Neither feature is a universal identity for all inputs. The packet energies are

$$49, \qquad 31/4, \qquad 0, \qquad 1/4,$$

which again sum to 57. Two factors of $1/\sqrt{2}$ explain why the depth-two coefficients are rational halves and quarters.

Coefficients and projectors are different objects

$W_\alpha f$ is the short coefficient vector. The original-space component is $W_\alpha^* W_\alpha f$. Summing these synthesized components, not the short arrays entrywise, reconstructs f:

$$f = \sum_{|\alpha|=2} W_\alpha^* W_\alpha f.$$

The finite verification constructs the actual analysis matrices and checks their product with their adjoints. The paper's ℓ^2 statement additionally establishes boundedness and the existence of the literal Hilbert adjoints.

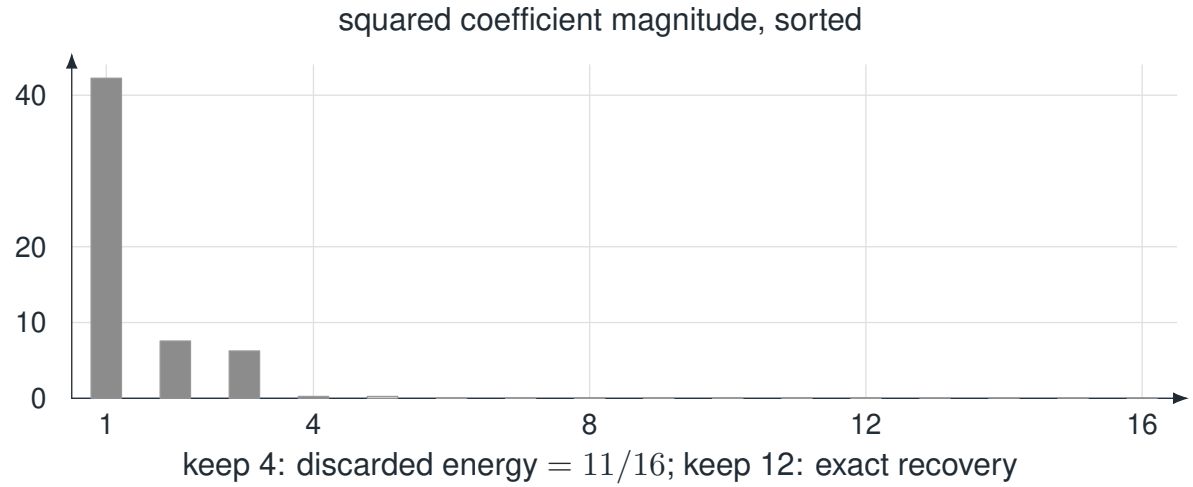
> **Read this as:** A packet address specifies an operator word. To add different packets as signals, synthesize them back into the common signal space first.

V11 / SPARSITY / AN EXACT ERROR LEDGER

Deleting coefficients is the approximation step

Traditional monograph: Section 10 (p. 18).

The complete depth-two transform is exact. Approximation begins only when we discard some of its sixteen coefficients. Order them by squared magnitude, with a deterministic address-and-position rule for ties.



Every discarded square has a meaning

In this fixed orthonormal basis, the squared reconstruction error equals the sum of the squared discarded coefficients. For our running signal:

kept coefficients k	1	2	4	8	12	16
squared error	$59/4$	$115/16$	$11/16$	$1/4$	0	0

The four zero coefficients belong to packet 10, so twelve coefficients suffice for exact reconstruction here. This is a property of the selected signal and basis, not a guarantee that all edge signals need only twelve coefficients.

Do not compare incompatible budgets

A cutoff m in a one-dimensional signed Fourier spectrum means $2m + 1$ bins, including zero. A budget of k largest coefficients is a different rule. Real Fourier synthesis may also require conjugate-pair closure. Neither count alone equals a bit-storage budget once amplitudes, phases, indices, or a selected tree must be encoded. The manuscript's Fourier-versus-Haar experiment is a finite comparison, not a proof that one basis dominates every signal class. Here we isolate the part that can be checked exactly: reconstruction error within one fixed finite basis.

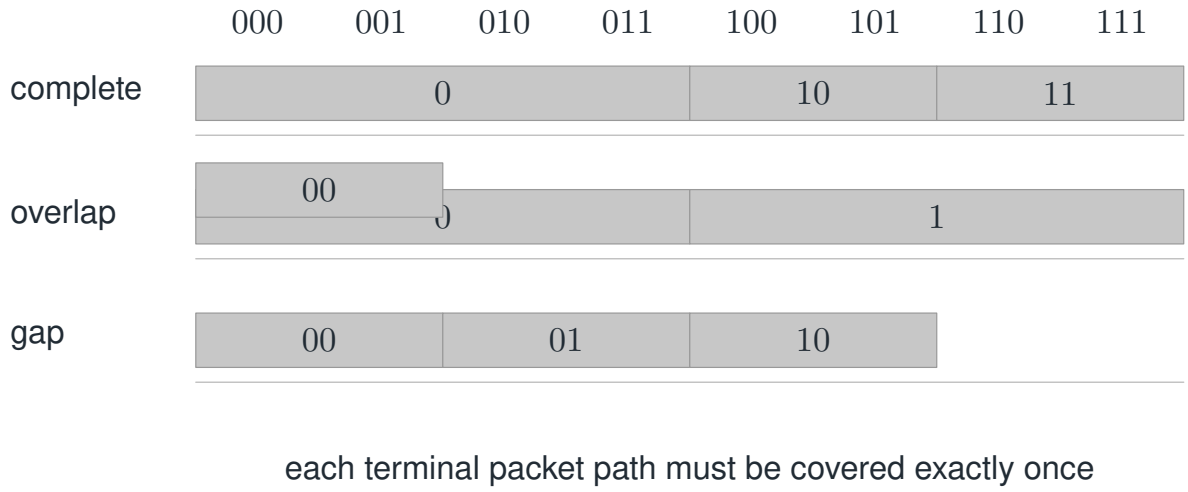
> **Read this as:** Changing coordinates creates no error. Throwing away coefficients does, and orthogonality makes that error exactly accountable.

V12 / COMPLETE PRUNINGS / NO GAPS AND NO DOUBLE COUNTING

A mixed-depth basis must cover every branch

Traditional monograph: Definition 11.1 (p. 19), Theorem 11.2 (p. 20).

A complete pruning chooses nodes of a finite refinement tree so that every root-to-terminal path meets exactly one selected node. The horizontal positions below are *terminal packet addresses*, not time intervals in the original signal.



Three candidate selections

The complete selection $\{0, 10, 11\}$ covers every terminal branch once. The selection $\{0, 00, 1\}$ covers the 00 branch twice: it is not an antichain. The selection $\{00, 01, 10\}$ misses the 11 branch entirely. For actual Haar packets, set $P_v = W_v^* W_v$. Parent-to-child conservation gives

$$P_v = P_{v0} + P_{v1}.$$

Replacing unselected parents by their children until the complete pruning is reached proves $\sum_{v \in \Lambda} P_v = I$. This is why the coefficients at different depths still form an exact coordinate system.

An exhaustive finite check

At depth three the pruning count obeys $a_0 = 1$ and $a_{d+1} = 1 + a_d^2$, giving $1, 2, 5, 26$. The supplied suite checks all 26 covering conditions and all 26 actual 16×16 projector-sum identities. These finite checks illustrate, but do not replace, the universal tree theorem.

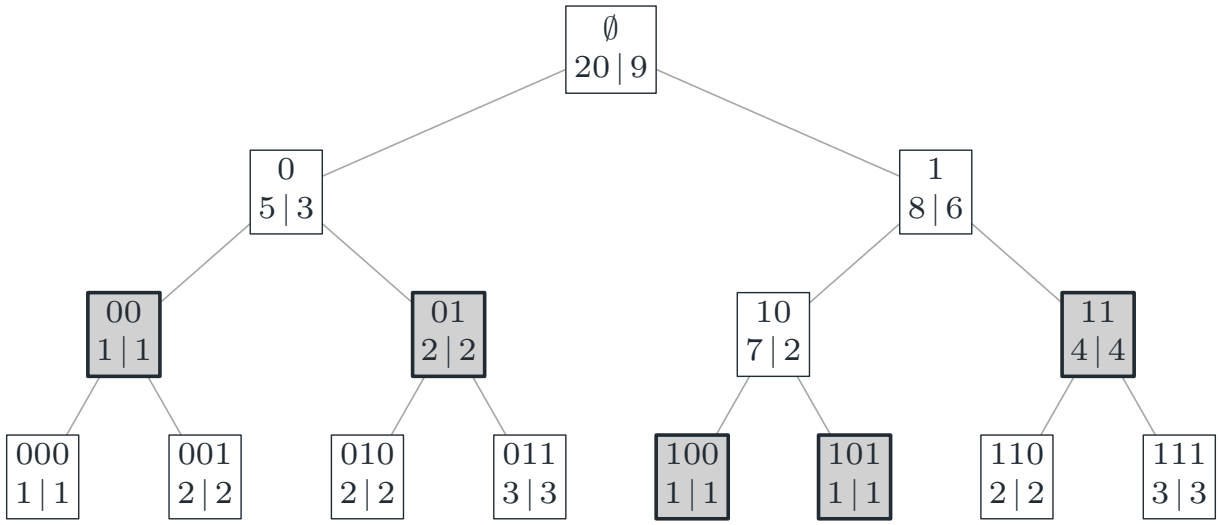
> **Read this as:** Mixed depth is allowed. Missing a terminal path or selecting an ancestor and its descendant is not.

V13 / ADDITIVE COSTS / EXACT DYNAMIC PROGRAMMING

Solve every alternative without searching every alternative

Traditional monograph: Theorem 11.3 (p. 20), Theorem 11.4 (p. 21).

Give every node the rational stop cost shown below. A node may either stop and pay its own cost, or refine and pay the already optimized costs of its children. The shaded nodes form the selected complete pruning.



node labels: stop cost c_v | optimal subtree cost $V(v)$

$$V(v) = \min\{c_v,\ V(v0) + V(v1)\}, \qquad V(v) = c_v \text{ at a terminal node.}$$

Read the calculation from the bottom upward

At 10, the children cost $1 + 1 = 2$, better than stopping for 7. At 11, stopping costs 4, better than $2 + 3 = 5$. The right subtree therefore costs 6, better than its stop cost 8. The left subtree costs $1 + 2 = 3$, better than 5. At the root, $3 + 6 = 9$ beats 20. The unique optimum is

$$\Lambda_* = \{00, 01, 100, 101, 11\}, \qquad J(\Lambda_*) = 9.$$

All 26 complete prunings are separately enumerated and give the same minimum. At equality, the rule stops, making the selector deterministic. A singleton tree needs no comparison. The costs here are supplied numbers. The theorem optimizes that stated additive objective; it does not prove those numbers encode the best denoising criterion. Coupled costs that cannot be written as a sum of selected-node costs need a different optimizer.

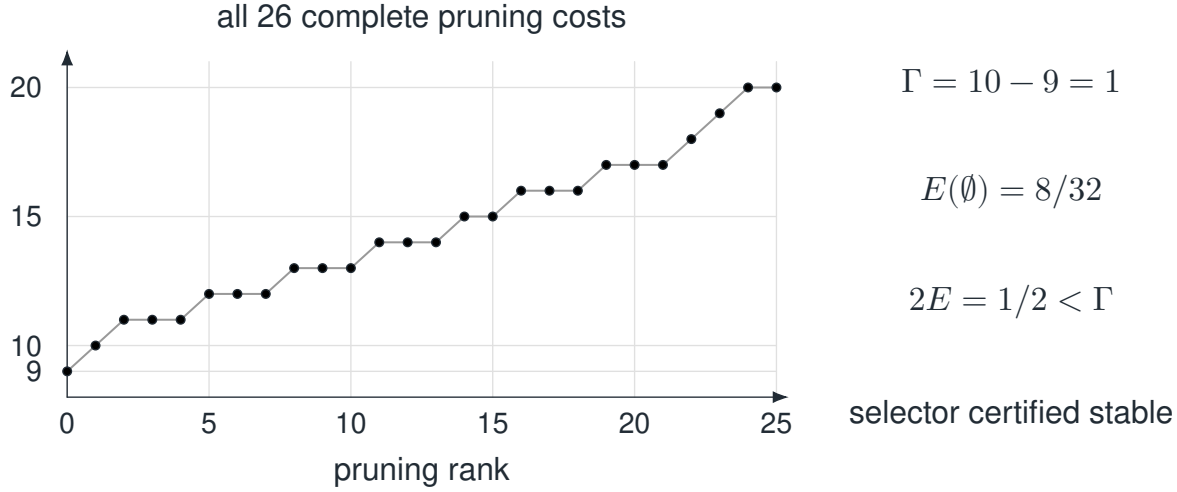
> **Read this as:** Global optimality follows from the exact stop-or-refine decomposition of a complete pruning, not from a greedy assumption.

V14 / A GLOBAL MARGIN / NOT A LOCAL HEURISTIC

How far can the costs move before the tree changes?

Traditional monograph: Theorem 11.3 (p. 20), Corollary 11.5 (p. 21).

The best pruning on the previous page costs 9; the cheapest competitor costs 10. Its exact global gap is therefore $\Gamma = 1$. The plate lists every complete-pruning cost in sorted order, giving each pruning its own horizontal rank.



Propagate an error budget as carefully as a cost

Suppose every node cost changes by at most $\varepsilon_v = 1/32$. A pruning may select as many as eight leaves, so its total cost can change by at most $8/32 = 1/4$. The recursive envelope proves this without guessing which pruning will be selected:

$$E(v) = \max\{\varepsilon_v,\ E(v0) + E(v1)\}, \qquad E(\emptyset) = \tfrac{1}{4}.$$

For any competitor Λ, the changed cost difference is bounded below by

$$J'(\Lambda) - J'(\Lambda_*) \geq \Gamma - 2E(\emptyset) = \tfrac{1}{2} > 0.$$

Thus the optimum is unchanged for *every* permitted simultaneous perturbation, not just a sampled noise realization. The envelope is over all feasible prunings, not just the five nodes in the current optimum.

Strict inequality has content

If two competing costs are 9 and 10, moving them to 9.5 and 9.5 removes uniqueness. Equality in $2E = \Gamma$ is therefore not a unique-selector certificate. For ordinary probability, the random node costs and the joint good event must be measurable. Separate nodewise probability bounds do not automatically give the probability of that joint event.

> **Read this as:** The exact selector is stable because its global gap exceeds the total worst-case movement of both competing costs.

V15 / AN EXACT IDENTITY / A CONDITIONAL INPUT

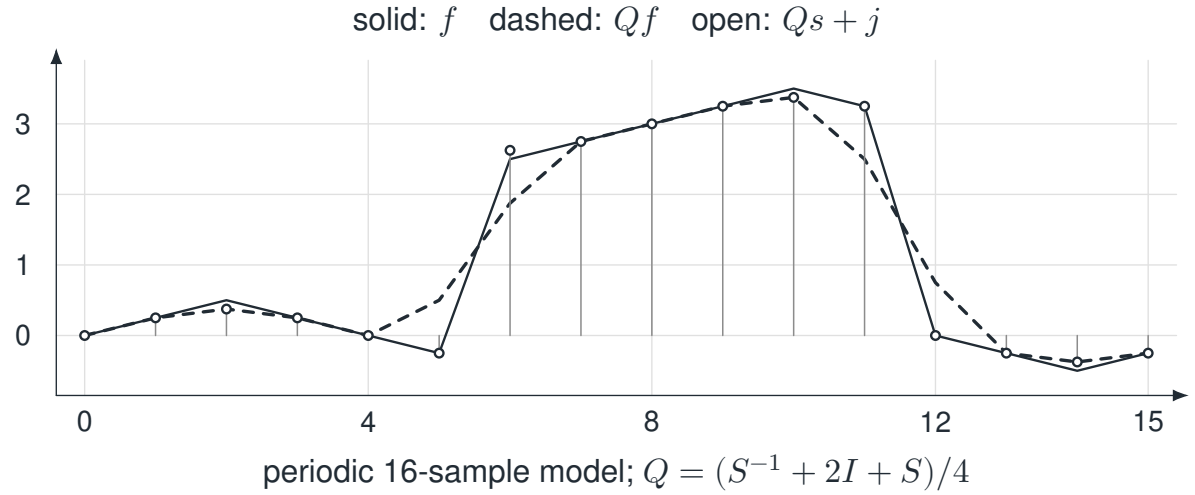
Retain the jump and approximate the other packet

Traditional monograph: Theorem 6.1 (p. 12), Corollary 6.2 (p. 12).

Return to the running signal $f = s + j$. For a concrete finite calculation, use the periodic three-tap smoother

$$(Qx)[n] = \tfrac{1}{4}x[n-1] + \tfrac{1}{2}x[n] + \tfrac{1}{4}x[n+1], \qquad n \in \mathbb{Z}/16\mathbb{Z}.$$

The plots compare smoothing the full signal with retaining its supplied jump packet. Every plotted sample is a rational number computed from the displayed formula.



The cancellation is the theorem

Define $\tilde{f} = Qs + j$. Then, sample by sample,

$$\boxed{f - \tilde{f} = s - Qs.}$$

By contrast, linearity gives

$$f - Qf = (s - Qs) + (j - Qj).$$

The second error includes a new term caused by smoothing the discontinuity. In this finite example that term is concentrated near the two jump boundaries. The identity itself requires no linearity of the approximation applied to s; linearity is used only in the comparison with Qf.

The input is not discovered by subtraction

The data here include the correct j. Replacing it with a wrong $\hat{j}$ retains the error $j - \hat{j}$, even when s is approximated perfectly. Multiple amplitudes and overlapping shapes are allowed if their full sum is supplied, but the identity does not estimate that sum from observations. Our finite periodic smoother is an illustration, not a new continuum approximation theorem. Its chosen boundary convention is part of the experiment.

> **Read this as:** Known-edge reconstruction removes the jump from the approximation problem. It does not remove the need to identify the jump.

V16 / A CONTINUOUS PULSE / EXACT ENDPOINT VALUES

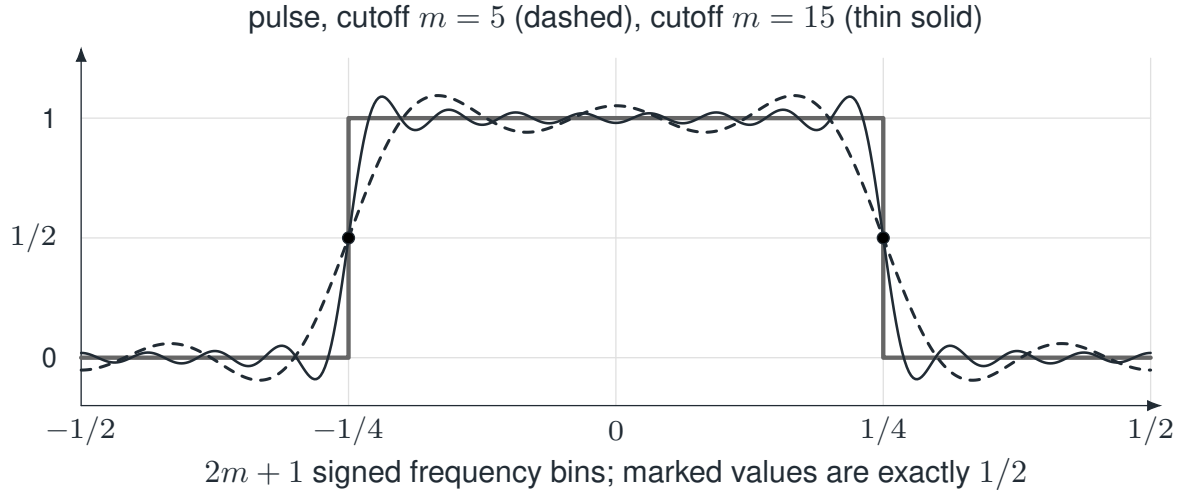
Why a Fourier cutoff cannot reproduce this jump pointwise

Traditional monograph: Section 6 (p. 11), Section 8 (p. 13).

Use a separate period-one example on $[-1/2, 1/2)$:

$$j(t) = \mathbf{1}_{\{|t|<1/4\}}, \qquad S_m j(t) = \tfrac{1}{2} + \frac{2}{\pi}\sum_{n=1}^{m} \frac{\sin(\pi n/2)}{n}\cos(2\pi n t).$$

This is the exact finite Fourier formula. The drawn curves sample it numerically; the endpoint calculation below is symbolic.



One value is fixed at every cutoff

At $t = \pm 1/4$, every summand vanishes because $\sin(\pi n/2)\cos(\pi n/2) = 0$ for integer n. Hence

$$S_m j(\pm 1/4) = \tfrac{1}{2} \qquad \text{for every } m \geq 0.$$

A larger cutoff resolves a narrower transition, but it does not make the finite trigonometric polynomial discontinuous. The two displayed cutoffs retain 11 and 31 signed frequency bins, respectively, not five and fifteen coefficients.

An exact known-packet comparison

Take $s(t) = \frac{1}{4}\cos(2\pi t) + \frac{1}{10}\cos(6\pi t)$. For $m \geq 3$ this smooth packet is already in the retained Fourier span. Thus $S_m s + j = s + j$ exactly in this special model, whereas $S_m(s + j)$ still replaces the pulse by $S_m j$. The values assigned at the two endpoints matter for pointwise error, but do not affect an L^p norm for finite p. The illustrated ringing is not a root enclosure or an asymptotic error certificate.

> **Read this as:** Increasing spectral resolution and retaining a known discontinuity are different operations, with different pointwise consequences.

V17 / MASK GEOMETRY / TRANSLATION BY ONE SIXTEENTH

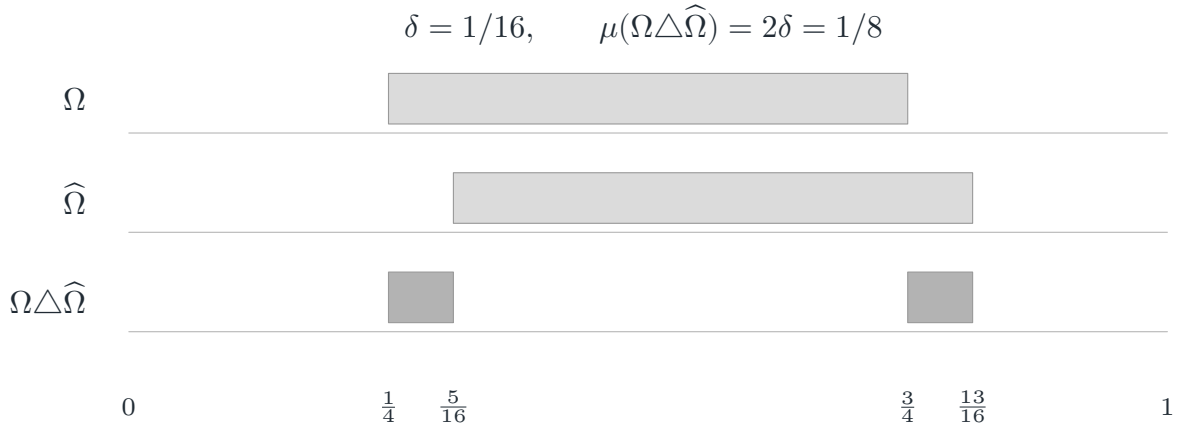
Small location errors, different norms

Traditional monograph: Theorem 7.1 (p. 13).

In Lebesgue measure on the line, let

$$\Omega = [\tfrac{1}{4}, \tfrac{3}{4}], \qquad \widehat{\Omega} = [\tfrac{5}{16}, \tfrac{13}{16}], \qquad a = \hat{a} = 3.$$

The two intervals have the same length. Their disagreement consists of one missed strip and one extra strip, each of width $1/16$.



$$\|j - \hat{j}\|_1 = 3/8, \qquad \|j - \hat{j}\|_2^2 = 9/8, \qquad \|j - \hat{j}\|_\infty = 3$$

$$|\Omega \triangle \widehat{\Omega}| = \tfrac{1}{8}, \qquad j - \hat{j} = 3(\mathbf{1}_\Omega - \mathbf{1}_{\widehat{\Omega}}).$$

error measure	L^1 norm	squared L^2 norm	L^∞ norm
exact value	$3/8$	$9/8$	3

Area shrinks while the missed jump keeps its height

For a translation by δ with $0 < |\delta| < 1/2$, the symmetric difference has measure $2|\delta|$. At finite p,

$$\|j - \hat{j}\|_p = 3(2|\delta|)^{1/p}.$$

It tends to zero as the displacement tends to zero. The essential supremum stays 3 for every nonzero displacement, because the error reaches height 3 on a set of positive measure. At exactly zero displacement all norms vanish. A nonempty disagreement set of measure zero would be different: it can affect the pointwise supremum while leaving the essential norm unchanged. This is why the paper distinguishes the two statements.

> **Read this as:** A shrinking boundary strip can give small integrated error and still contain a full-height jump error.

V18 / COUNTING MEASURE / AN EXACT FINITE BUDGET

Separate amplitude error from mask error

Traditional monograph: Theorem 7.1 (p. 13).

On the 16 sample addresses, let $\Omega = \{6, \dots, 11\}$ and $\widehat{\Omega} = \{7, \dots, 12\}$. Use true amplitude $a = 3$ and estimate $\hat{a} = 5/2$. There are six true samples, five overlapping samples, and two disagreement samples.

$$a = 3, \quad \hat{a} = 5/2, \quad |\Omega| = 6, \quad |\Omega \triangle \widehat{\Omega}| = 2$$

	4	5	6	7	8	9	10	11	12	13	14
j	0	0	3	3	3	3	3	3	0	0	0
$\hat{j}$	0	0	0	$\frac{5}{2}$	$\frac{5}{2}$	$\frac{5}{2}$	$\frac{5}{2}$	$\frac{5}{2}$	$\frac{5}{2}$	0	0
$j - \hat{j}$	0	0	3	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	$-\frac{5}{2}$	0	0

$$\|j - \hat{j}\|_1 = 3 + 5(1/2) + 5/2 = 8$$

The residual splits exactly as

$$3\mathbf{1}_\Omega - \tfrac{5}{2}\mathbf{1}_{\widehat{\Omega}} = \tfrac{1}{2}\mathbf{1}_\Omega + \tfrac{5}{2}(\mathbf{1}_\Omega - \mathbf{1}_{\widehat{\Omega}}).$$

The missed sample has error 3, the five shared samples have error $1/2$, and the extra sample has error $-5/2$. With *counting measure*,

$$\|j - \hat{j}\|_1 = 8, \qquad \|j - \hat{j}\|_2^2 = \tfrac{33}{2}, \qquad \|j - \hat{j}\|_\infty = 3.$$

Read the theorem as a budget, not an equality claim

Here the L^1 bound is attained:

$$|3 - \tfrac{5}{2}|\,|\Omega| + \tfrac{5}{2}|\Omega \triangle \widehat{\Omega}| = \tfrac{1}{2} \cdot 6 + \tfrac{5}{2} \cdot 2 = 8.$$

The L^2 triangle bound is $\frac{1}{2}\sqrt{6} + \frac{5}{2}\sqrt{2}$, and need not be sharp. A background estimate adds another term $\|s - \hat{s}\|_p$; cancellation between terms can make the total smaller than the bound. Normalized RMSE divides the L^2 norm by $\sqrt{16}$. A boundary-band RMSE uses the band size instead. Changing that denominator changes the reported quantity.

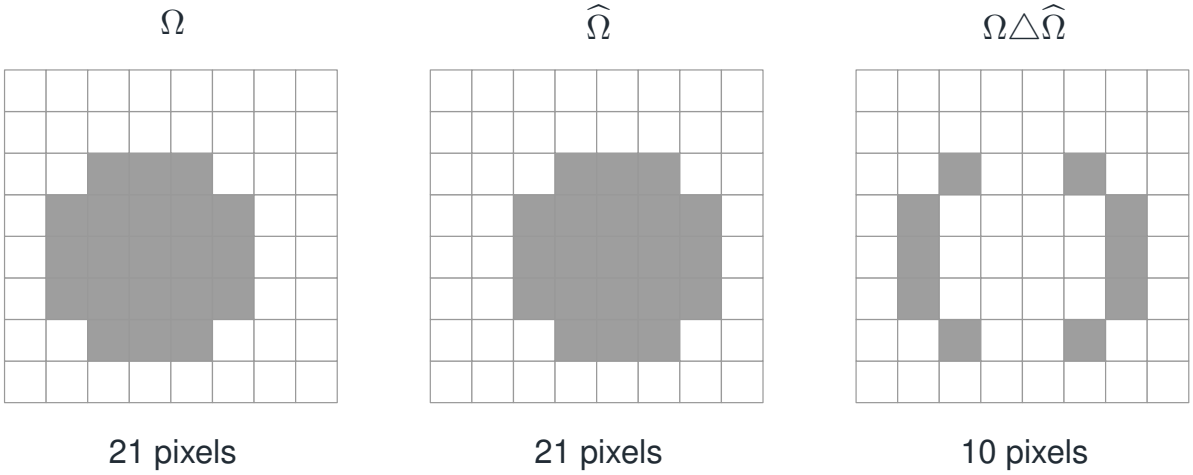
> **Read this as:** Mask error, amplitude error, and background error are separate mathematical inputs. A single reconstruction score conceals that separation.

V19 / A FINITE IMAGE / AN AMBIGUOUS OBSERVATION

Geometry localizes error while data may leave it ambiguous

Traditional monograph: Section 7 (p. 12), Section 9 (p. 15).

On the integer 8×8 grid, choose the disk-shaped mask $\Omega = \{(i,j) : (i-3)^2 + (j-3)^2 \le 5\}$ and translate it one pixel to the right. Each mask contains 21 pixels; their symmetric difference contains exactly 10.



A geometric error calculation is not a segmentation algorithm

For a common amplitude a, the squared counting-measure error is $10a^2$, entirely on the disagreement pixels. The full-image mean squared error is $10a^2/64$. Both are exact, but answer different normalization questions. Now consider the four-sample observation

$$y = (0, \tfrac{3}{2}, 3, 3).$$

It is consistent with either true signal

$$u = (0, 0, 3, 3), \qquad v = (0, 3, 3, 3)$$

under deterministic sup-norm noise at most $3/2$. Indeed $y = (u+v)/2$, and both error norms equal $3/2$. A rule receiving only y and that noise bound cannot be guaranteed to select both distinct true masks correctly.

What additional structure would have to do

A lower noise level, separation, contrast, geometry, or a specified prior can change the identification problem. The example rules out a guarantee at the *stated* noise radius and signal pair, not all recovery results. The manuscript's positive-contrast estimator is deliberately narrower than a universal segmentation method.

> **Read this as:** Exact synthesis can preserve an identified boundary. It cannot decide between two boundaries that the observation model makes indistinguishable.

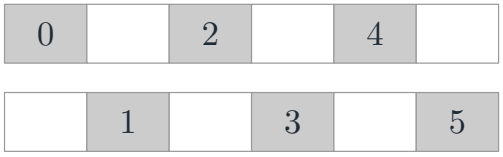
V20 / POINTWISE SIGNALS VERSUS CAUCHY COEFFICIENTS

Two products, two packet laws

Traditional monograph: Theorem 12.1 (p. 24).

The word “product” must be typed. Multiplying two signal values at the same time address is not the same operation as multiplying generating series. The distinction changes which packets can interact.

POINTWISE SIGNAL PRODUCT



$$(P_{2,0}x)\,(P_{2,1}y) = 0$$

CAUCHY COEFFICIENT PRODUCT

0	1	2
1	2	0
2	0	1

output label $(i + j) \bmod 3$

$$A = 1 + 2t + 3t^2,\ B = 2 + t^2 : \qquad P_{3,0}(AB) = 2 + 2t^3$$

Pointwise multiplication preserves the same address

For time-residue masks, $(P_{N,i}x)[n](P_{N,j}y)[n] = 0$ when $i \neq j$. Therefore

$$P_{N,k}(x \cdot y) = (P_{N,k}x) \cdot (P_{N,k}y).$$

Here the dot means multiplication at the same sample, not convolution.

Cauchy multiplication adds degrees

For formal series $A(t), B(t)$, the coefficient of t^n is a sum over $a+b = n$. Consequently the output residue is $i+j \pmod N$:

$$P_{N,k}(AB) = \sum_{i+j \equiv k\ (N)} (P_{N,i}A)(P_{N,j}B).$$

For $N = 3$, take $A = 1 + 2t + 3t^2$ and $B = 2 + t^2$. Direct multiplication gives

$$AB = 2 + 4t + 7t^2 + 2t^3 + 3t^4, \qquad P_{3,0}(AB) = 2 + 2t^3.$$

The $2t^3$ term comes from residues $1+2$, not $0+0$. At power zero, the empty product contributes 1 only in residue zero. These are formal coefficient identities and require no convergence assumption.

> **Read this as:** Before applying a packet multiplication rule, specify whether multiplication acts on sample values or on series coefficients.

V21 / FULL INPUT TUPLES / EXACT CYCLIC INDEXING

Nonlinear neighbors cross channel boundaries

Traditional monograph: Section 12 (p. 24).

For the periodic running signal define

$$F(x)[n] = x[n] + \tfrac{1}{2}x[n]x[n-1].$$

This is nonlinear and not a convolution operator. Nevertheless its two-channel representation is exact, provided the neighbor comes from the correct channel and index.

$$F(x)[n] = x[n] + \tfrac{1}{2}x[n]x[n-1]$$

$$F_0[n] = e[n] + \tfrac{1}{2}e[n]o[n-1] \qquad F_1[n] = o[n] + \tfrac{1}{2}o[n]e[n]$$

$$e[1] = 1/2, \quad o[0] = 1/4 \qquad o[1] = 1/4, \quad e[1] = 1/2$$

$$F(x)[2] = 9/16 \qquad F(x)[3] = 5/16$$

phase wrap changes the neighbor address; it does not remove the mixed term

Put $e[q] = x[2q]$ and $o[q] = x[2q+1]$, with channel indices taken modulo eight. Then

$$F(x)[2q] = e[q] + \tfrac{1}{2}e[q]o[q-1],$$
$$F(x)[2q+1] = o[q] + \tfrac{1}{2}o[q]e[q].$$

The even output needs the *previous* odd sample; the odd output needs the current even sample. This is the same carry in the address system that appeared in the one-sample delay matrix.

Check two actual outputs

The running values are $x[1] = 1/4$, $x[2] = 1/2$, and $x[3] = 1/4$, so

$$F(x)[2] = \tfrac{1}{2} + \tfrac{1}{2}\tfrac{1}{2}\tfrac{1}{4} = \tfrac{9}{16}, \qquad F(x)[3] = \tfrac{1}{4} + \tfrac{1}{2}\tfrac{1}{4}\tfrac{1}{2} = \tfrac{5}{16}.$$

Applying the same formula independently within each decimated array would use the wrong neighbors. Keeping the full channel tuple avoids that loss of information. The supplied suite compares every output sample in the two implementations using exact rational arithmetic. No convergence, contraction, or approximation property follows merely from writing these exact coordinates. Those require hypotheses on the nonlinear map and the problem being solved.

> **Read this as:** Exact packet coordinates remain available for nonlinear maps. They do not turn cross-channel interactions into independent channel problems.

V22 / THE CAYLEY CHART / EXACT QUARTER TURNS

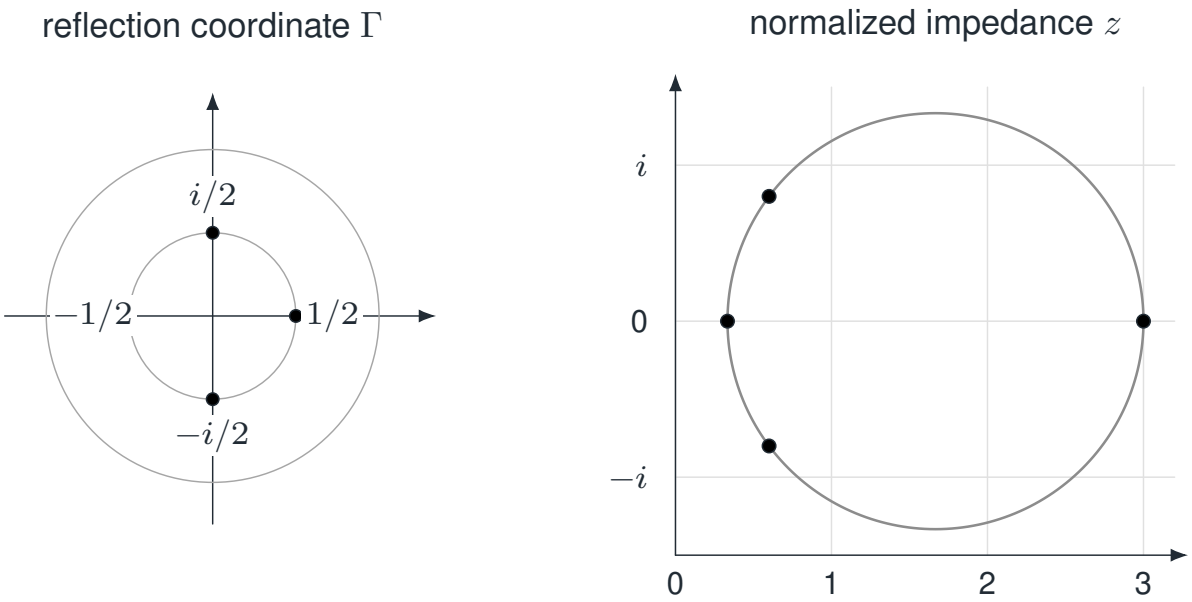
One reflection circle, four very different impedances

Traditional monograph: Theorem 13.2 (p. 26).

Use normalized impedance z and reflection coefficient Γ:

$$\Gamma = \frac{z-1}{z+1}, \qquad z = \frac{1+\Gamma}{1-\Gamma}.$$

The circle $|\Gamma| = 1/2$ is a constant-mismatch orbit. Its four quarter-turn points map to the exact impedances shown below.



$$3, \quad \tfrac{3}{5} + \tfrac{4}{5}i, \quad \tfrac{1}{3}, \quad \tfrac{3}{5} - \tfrac{4}{5}i$$

Γ	$1/2$	$i/2$	$-1/2$	$-i/2$
z	3	$3/5 + 4i/5$	$1/3$	$3/5 - 4i/5$

The coordinates preserve structure, not Euclidean distances

A quarter turn is a rotation in the reflection disk. After pulling back, it is a Möbius map, not a Euclidean rotation of the impedance plane. Equal reflection magnitudes need not give equal impedance magnitudes. The opposite points $1/2$ and $-1/2$ become 3 and $1/3$. The continuous impedance locus is the circle $(\operatorname{Re} z - 5/3)^2 + (\operatorname{Im} z)^2 = 16/9$, also drawn in the plate. Finite passive impedances with $\operatorname{Re} z \geq 0$ map to the closed disk except $\Gamma = 1$. That missing point is obtained only by adjoining impedance infinity. The interior maps to the open right half-plane. The classical Smith grid supplies the geometry; the cyclic labels specify exact rotations on it. Neither a new physical law nor a new optimal sampling theorem is implied by this coordinate description.

> **Read this as:** A simple action in one chart can be nonlinear in another. Conjugacy preserves the action law, not the apparent shape of a displacement.

V23 / HOMOGENEOUS COORDINATES / NO DIVISION BY ZERO

A pole is a chart boundary, not a broken group action

Traditional monograph: Theorem 13.2 (p. 26), Corollary 13.3 (p. 26).

Represent impedance by a nonzero pair $[u : v]$, where $[cu : cv] = [u : v]$ for $c \neq 0$. A finite impedance is $[z : 1]$; infinity is $[1 : 0]$. For a nonzero reflection multiplier w, use

$$A(w) = \frac{1}{2}\begin{pmatrix} 1+w & 1-w \\ 1-w & 1+w \end{pmatrix}.$$

$$A(i) = \tfrac{1}{2}\begin{pmatrix} 1+i & 1-i \\ 1-i & 1+i \end{pmatrix}$$



$z = 0$	$z = -i$	$z = \infty$	$z = i$
$\Gamma = -1$	$\Gamma = -i$	$\Gamma = 1$	$\Gamma = i$

$A(i)^4 = I$: the fourth step returns to $[0 : 1]$

The action is defined before a ratio is taken

The determinant is w, so the matrix cannot send a nonzero pair to zero. Direct multiplication gives

$$A(w)A(w') = A(ww'), \qquad A(1) = I.$$

For $w = i$, the orbit through zero is

$$[0 : 1] \longmapsto [-i : 1] \longmapsto [1 : 0] \longmapsto [i : 1] \longmapsto [0 : 1].$$

The third point is a legitimate projective point, even though a finite impedance formula has a zero denominator there. The second iterate is $A(-1)$, which swaps the two coordinates. Thus a half turn exchanges 0 and ∞ and sends each finite nonzero z to $1/z$. The matched point $\Gamma = 0$, equivalently $z = 1$, is fixed by every phase multiplier. It is not part of the boundary orbit above. The exact checks test homogeneous pairs and their rescaling equivalence, rather than introducing an artificial complex value for $1/0$.

> **Read this as:** Compute with homogeneous pairs first. A denominator becomes a chart-selection issue rather than an exception to the cyclic action.

V24 / DYADIC–ODD FACTORIZATION / TWELVE ADDRESSES

Residue labels are not impedance averages

Traditional monograph: Proposition 14.1 (p. 29).

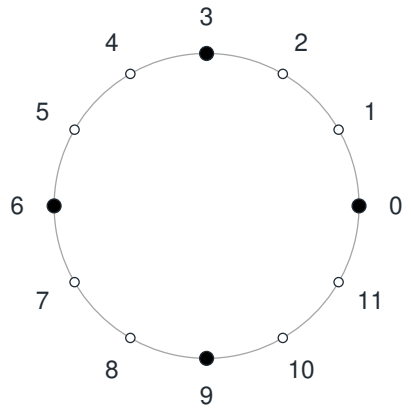
At phase order $12 = 4 \cdot 3$, the map

$$k \longmapsto (k \bmod 4,\ k \bmod 3)$$

is a bijection from twelve residue addresses to the 4×3 grid. Every cell below contains its unique original address.

	0	1	2
0	0	4	8
1	9	1	5
2	6	10	2
3	3	7	11

row $j \bmod 4$, column $j \bmod 3$



filled: the four coarse rays; open: the eight new rays

The factorization is an intersection of conditions

For integer sample addresses, including negative ones,

$$P_{12,k} = P_{4,k \bmod 4}\, P_{3,k \bmod 3}.$$

The product on the right is composition of two residue masks: both congruences must hold. Coprimality makes their intersection a single class modulo twelve. It is not a formula for multiplying or adding impedances.

Nested samples and nested projectors are different statements

The four coarse reflection rays occur at indices $0, 3, 6, 9$ of the twelve-ray grid. Thus the geometric sample set is nested. Separately, a coarse coefficient projector is the sum of its fine descendants in a common coefficient space. There is no corresponding rule saying that a coarse impedance is the average of its descendants. The inverse Cayley map is nonlinear. Even opposite reflection points have impedances 3 and $1/3$ in the preceding example; their average $5/3$ is not the impedance 1 of their averaged reflection coefficient zero.

> **Read this as:** CRT factors addresses. Refinement conserves coefficient packets. Neither operation authorizes linear averaging through a nonlinear chart.

V25 / LOCAL DERIVATIVES / FINITE INCREMENTS

Sensitivity grows before the pole is reached

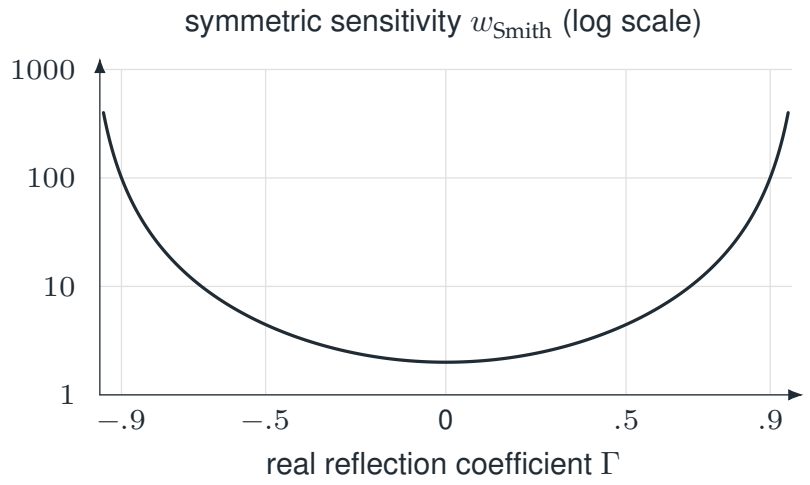
Traditional monograph: Section 14 (p. 29).

Away from $\Gamma = 1$, the inverse Cayley map satisfies

$$\frac{dz}{d\Gamma} = \frac{2}{(1-\Gamma)^2}.$$

Along the real axis, three exact derivatives already show the scale change:

$$\Gamma = \tfrac{1}{2} : \ 8, \qquad \Gamma = \tfrac{9}{10} : \ 200, \qquad \Gamma = \tfrac{99}{100} : \ 20000.$$



Γ	$\lvert dz/d\Gamma \rvert$
$1/2$	8
$9/10$	200
$99/100$	20000

$$\Delta z = 20/9 \neq 2$$

A differential is not the full finite-error formula

For a finite increment ε, with both denominators nonzero,

$$z(\Gamma+\varepsilon) - z(\Gamma) = \frac{2\varepsilon}{(1-\Gamma-\varepsilon)(1-\Gamma)}.$$

At $\Gamma = 9/10$ and $\varepsilon = 1/100$, the actual increment is $20/9$, while derivative times increment is 2. Both are calculated exactly. Their difference comes from using a differential in place of a finite increment. The derivative gives a first order approximation.

What a proposed weight can and cannot certify

Impedance is sensitive near $\Gamma = 1$, and admittance near $\Gamma = -1$. The symmetric weight

$$w(\Gamma) = |1-\Gamma|^{-2} + |1+\Gamma|^{-2}$$

reflects both chart sensitivities. Its growth motivates denser sampling near the endpoints, but does not establish an optimal grid, interpolation rule, or measurement allocation. A finite-error estimate must also control how close the perturbed point comes to either pole.

> **Read this as:** Use the exact finite difference for an error budget. A large derivative diagnoses conditioning. Choosing an optimal algorithm requires further input.

V26 / EVIDENCE GUIDE / THE FORMAL CONTRACT IS UNCHANGED

What the reader can now verify

Traditional monograph: Locator D (p. 35), Section E15 (p. 78).

Three worked structures, three distinct kinds of conclusion

Structure	What is established in the original worked layer
Finite signal	Exact signed addressing, filter transfer, Haar reconstruction and energy, sparse coefficient error, and full-tuple nonlinear evaluation.
Finite packet tree	All 26 complete prunings, actual projector sums, the unique cost-9 optimum, and a strict perturbation-margin certificate.
Reflection orbit	Exact rational-complex coordinates, a pole-safe projective action, CRT labels, and finite sensitivity calculations.

What each evidence layer means

Manuscript theorems. The preserved v1.1.0 release maps twenty theorem-like statements to 225 supporting Lean targets in 23 proof modules. Its saved compiler and axiom records remain inherited evidence. No fresh Lean run is claimed for this edition. The additional Part II developments remain outside this twenty-result map.

Current exact examples. The symbolic visual suite was rerun for this monograph. All 181 rational, algebraic, symbolic, and exhaustive finite checks passed. Seven deliberately false assertions were rejected. Numerical samples draw the continuous Fourier and sensitivity curves; they are not promoted to interval enclosures or universal error bounds.

Current executions. The fresh source-release checker passed 20 exact tests containing 305,933 assertions and passed 16 verifier tests, including positive and negative controls. The expanded-source numerical audit passed 11,052 finite checks, mixing exact and numerical calculations. The new extension examples separately passed 38 grouped exact checks and four false controls. These counts measure different operations and are not combined as theorem counts.

Reproducibility. The package preserves the original sources and their reports. Run `python3 build.py` here for an isolated PDF build. The extension finite checker is `checks/verify_extension_examples.py`; its contract and raw execution manifest are included. The original visual checker requires the listed optional dependencies. PDF assembly does not run Lean or either mathematical checker.

The remaining boundaries are not decorative disclaimers

Known-edge synthesis requires the correct packet. A finite cost optimum depends on the supplied objective. An infinite packet tree, a universal segmentation procedure, minimax recovery rates, and optimal Smith sampling require additional theorems. The core's broad edge-recovery problems are not resolved by these finite examples. V27 introduces the additional separated-jump recovery results and their explicit hypotheses. The 16-sample running signal is not the 1024-sample manuscript experiment; the 26-pruning example is not its 458,330-pruning benchmark. Their smaller sizes allow every relevant number to be inspected exactly.

> **Read this as:** Coordinates, approximation, identification, optimization, and physical interpretation are separate layers. The examples make their interfaces visible without erasing their hypotheses.

V27 / PART II / REPRESENTATION BECOMES CONDITIONAL RECOVERY

What the extended developments add

Traditional monograph: Locator D (p. 35), Section E1 (p. 39), Section E15 (p. 78).

The first twenty-six pages follow the corrected twenty-result formal core. The traditional paper's second part restores the broader mathematical developments of the earlier expanded manuscript. The next thirteen worked pages make their additional hypotheses and conclusions visible.

Four transitions, with the missing input made explicit

Transition	Additional input	Resulting control
Residues to lattices	A finite-index sublattice and one representative per coset	Exact addresses in several dimensions
Invertibility to stability	Positive lower frame bound and bounded perturbation	Quantified coefficient and reconstruction error
Known jumps to recovery	Separated jumps, sampling, a Lipschitz background, and a strict contrast gap	Exact jump cells and bounded location/amplitude error
Recovered packets to approximation	A bounded linear approximation operator and its norm	A complete two-term error estimate

The two parts have different formal evidence

The inherited Lean map covers the twenty designated results of Part I through its 225 supporting targets. The additional traditional proofs in Part II are not silently included in that map. The new pages below are mathematical explanations and exact worked specializations of those developments, with a separate finite-example checker. The original manuscript's broad edge-recovery questions remain broader than separated one-dimensional jumps. Conditions such as one jump per sample cell, an available wavelet support, and a positive detection margin do real work. In particular, the ambiguous observation on V19 still rules out an unrestricted recovery guarantee at its stated noise radius.

How to read the extension

V28 to V30 extend coordinates and stability. V31 to V34 identify what a singular packet contains and measure its error. V35 to V39 give three detection mechanisms. V40 closes the route by combining packet accuracy with a specified approximation operator.

> **Read this as:** A positive recovery theorem can coexist with an impossibility example. Their sampling models, noise bounds, and signal classes decide which conclusion applies.

V28 / LATTICE COORDINATES / A RECTANGULAR EXAMPLE

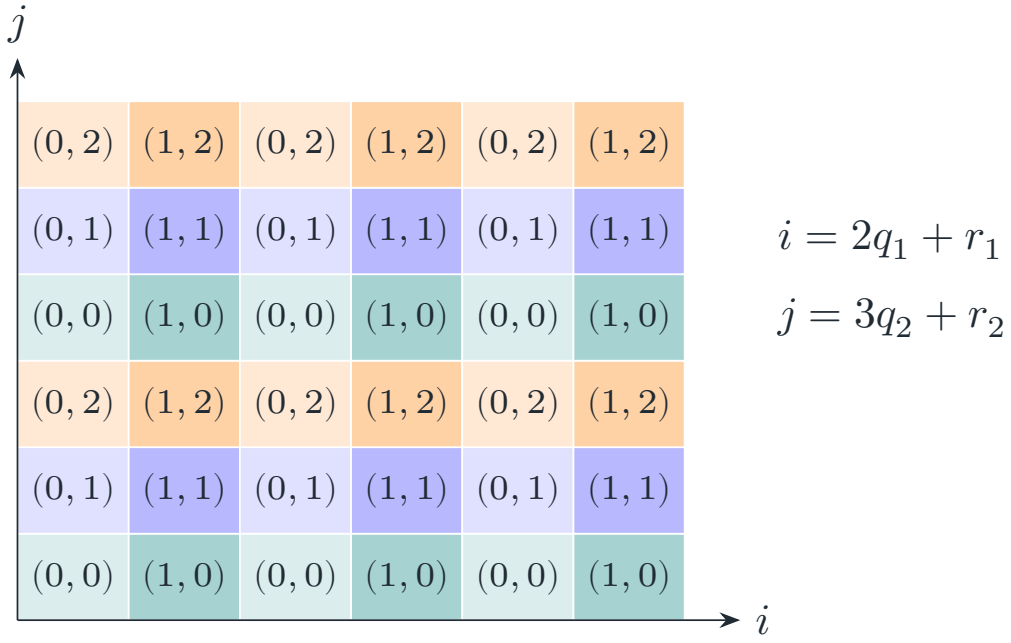
Six image phases are six exact address classes

Traditional monograph: Theorem E3.2 (p. 43), Theorem E3.5 (p. 44), Remark E2.2 (p. 43).

Take $L = 2\mathbb{Z} \times 3\mathbb{Z}$ inside $\mathbb{Z}^2$, with representatives

$$R = \{0, 1\} \times \{0, 1, 2\}.$$

Every address has a unique expression $(i, j) = (r_1, r_2) + (2q_1, 3q_2)$. The six channels are $x_r(2q_1, 3q_2) = x(r_1 + 2q_1, r_2 + 3q_2)$.



Each cell shows its residue label. Colors identify six disjoint classes; they do not average neighboring image pixels.

A negative address is still unique

The point $(-1, 7)$ has representative $(1, 1)$ and lattice displacement $(-2, 6)$, hence channel index $(-1, 2)$. Floor division gives the correct answer in both coordinates. For finite p, the channel norm uses the sum of the pth powers. At $p = \infty$, it uses their maximum. Both equal the norm on the original lattice because the cosets partition all addresses. In ℓ^∞, this coordinate identity does not imply that translates of an impulse are dense: their closure is only c_0.

Refinement partitions a class again

For $L' = 4\mathbb{Z} \times 6\mathbb{Z} \subset L$, the coarse representative $(1, 2)$ has fine representatives

$$(1, 2),\ (3, 2),\ (1, 5),\ (3, 5).$$

Their original-lattice masks sum to the coarse mask. As on V4, decimated arrays are interleaved, not added entrywise.

> **Read this as:** Dimension changes the address tuple. Unique coset decomposition still supplies exact synthesis and the correct norm identity.

V29 / FRAME STABILITY / AN EXACT TWO-CHANNEL BANK

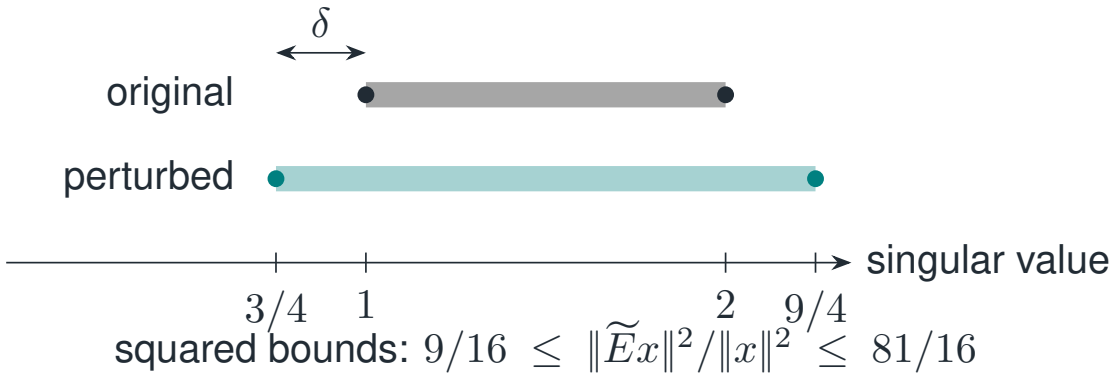
A positive singular-value gap survives perturbation

Traditional monograph: Theorem E5.2 (p. 49), Theorem E5.7 (p. 50), Corollary E5.6 (p. 50).

Consider the constant analysis bank

$$E = \begin{pmatrix} 1 & 0 \\ 0 & 2 \end{pmatrix}, \qquad E^*E = \operatorname{diag}(1, 4).$$

It has frame bounds $A = 1$, $B = 4$. For every perturbation with operator norm at most $\delta = 1/4$, the squared-norm bounds remain at least $(1 - 1/4)^2 = 9/16$ and at most $(2 + 1/4)^2 = 81/16$.



Both bounds can be attained

Choose $\Delta E = \operatorname{diag}(-1/4, 1/4)$. Then $\widetilde{E} = \operatorname{diag}(3/4, 9/4)$ has precisely those two frame bounds. Its inverse has operator norm $4/3$. A coefficient error $\nu = (3/4, 0)$ therefore produces signal error $(1, 0)$, attaining the noise bound. The strict condition $\delta < \sqrt{A}$ prevents the lower bound from vanishing. At $\delta = 1$, the perturbation $\operatorname{diag}(-1, 0)$ destroys the first channel entirely.

An inverse and an approximate inverse answer different questions

If the original synthesis $R = \operatorname{diag}(1, 1/2)$ is retained after this perturbation, then

$$R\widetilde{E} - I = \operatorname{diag}(-1/4, 1/8).$$

Its defect norm is $1/4$. Every input has relative reconstruction error at most that value, with equality on the first coordinate. For frequency-dependent banks the theorem uses an essential supremum on the unit circle, not a few sampled frequencies.

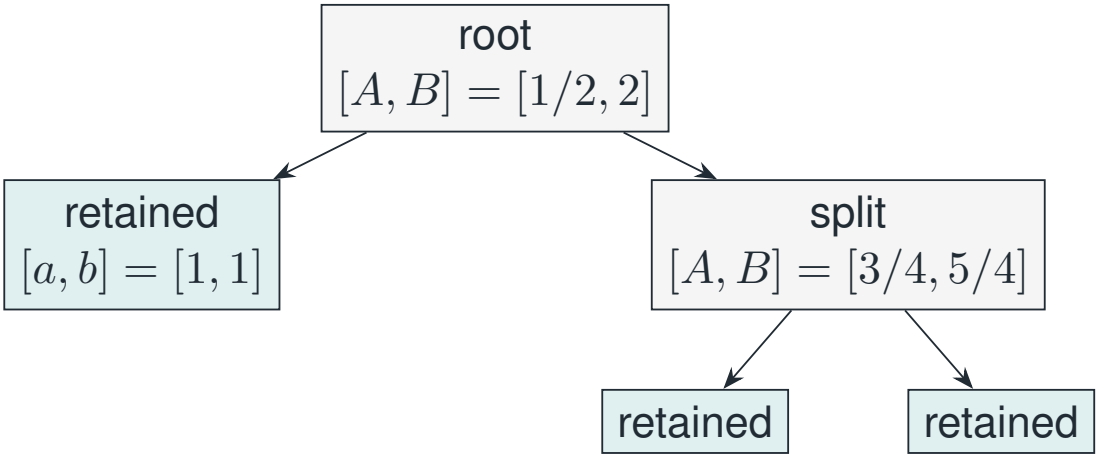
> **Read this as:** A lower singular-value gap quantifies stable recovery. A reconstruction defect quantifies the error of the actual synthesis bank being used.

V30 / LOCAL FRAME CONSTANTS / A MIXED-DEPTH CERTIFICATE

Nonunitary tree bounds propagate along branches

Traditional monograph: Proposition E6.4 (p. 53).

Suppose the root split has nonnegative squared-norm bounds $A = 1/2$ and $B = 2$. Retain the left child. Split the right child with bounds $3/4$ and $5/4$, then retain both grandchildren.



Start with one at every retained node

The left subtree has bounds $a_0 = b_0 = 1$. The right subtree has $a_1 = 3/4$ and $b_1 = 5/4$. The recursive certificate gives

$$a_{\text{root}} = \tfrac{1}{2}\min\{1, \tfrac{3}{4}\} = \tfrac{3}{8}, \qquad b_{\text{root}} = 2\max\{1, \tfrac{5}{4}\} = \tfrac{5}{2}.$$

Consequently the total leaf coefficient energy lies between $3/8$ and $5/2$ times the root energy. These are guaranteed bounds; no assertion that every admissible collection of local maps attains both constants is needed.

Why a minimum and maximum appear

After the root split, its energy may distribute unevenly between the two children. The lower estimate must survive the least favorable child lower bound. The upper estimate must survive the largest child upper bound. Replacing these by an unweighted average would presume an energy distribution that the hypotheses do not supply. When all local maps are unitary, every factor is one and the recursion recovers exact energy conservation. For general local maps, a positive lower bound alone establishes injectivity and stable inversion on the range; it does not by itself assert surjectivity onto every possible leaf tuple.

> **Read this as:** Exact tree coverage and good conditioning are separate inputs. The first chooses a coordinate route; the second bounds what that route does to norms.

V31 / TRACE DATA / A STEP AND A RAMP

Derivative jumps have their own packets

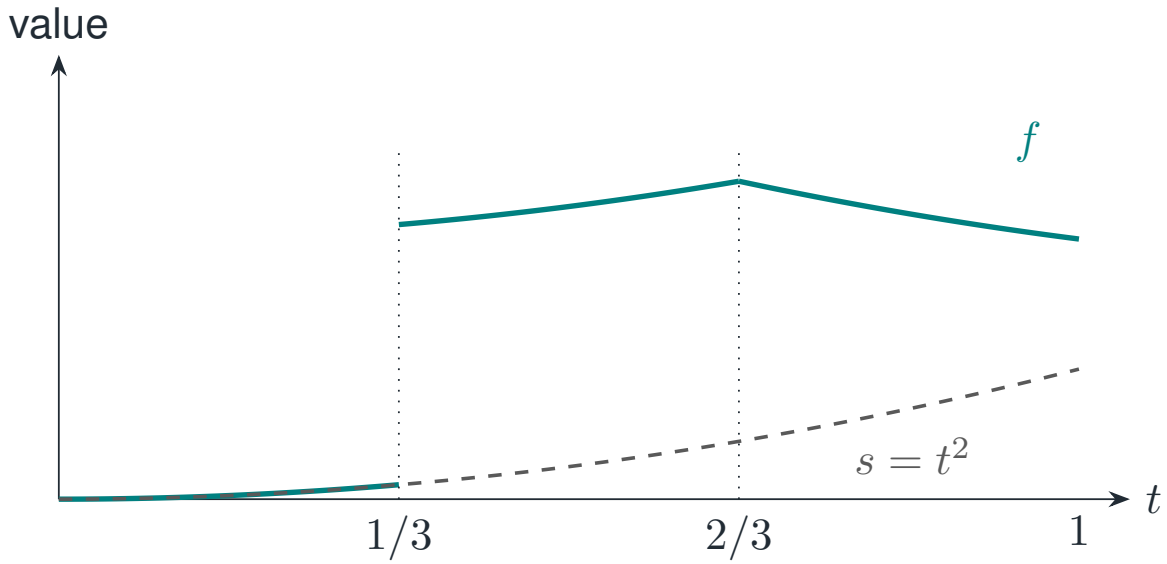
Traditional monograph: Theorem E7.2 (p. 56), Corollary E7.3 (p. 57).

On $[0,1]$, let $H_\tau = \mathbf{1}_{[\tau,1]}$ and take

$$f(t) = t^2 + 2H_{1/3}(t) - 3(t-2/3)_+.$$

The function has a value jump 2 at $1/3$ and a first-derivative jump -3 at $2/3$. Its packet and remainder are

$$j_f(t) = 2\phi_{0,1/3}(t) - 3\phi_{1,2/3}(t), \qquad s_f(t) = t^2.$$



The displayed curves are exact polynomial pieces. The solid jump is not connected across its discontinuity. The dashed curve is the global smooth remainder.

The normalization makes the trace coefficient literal

The general packet is $\phi_{m,\tau}(t) = (t-\tau)_+^m/m!$. Its derivative of order m has a unit trace jump, while its other classical trace jumps vanish. Thus its multiplier is exactly the jump being removed, with no hidden factorial. Here the remainder has matching value and first-derivative traces and belongs to $W^{2,p}(0,1)$ for every $1 \le p \le \infty$. Its second weak derivative is 2. The singular packet accounts for the distributional boundary terms that would otherwise arise at the breakpoints.

What the general theorem requires

Each piece must lie in the stated Sobolev class and possess the specified one-sided traces. Subtracting the finite trace packet then joins those pieces into a global Sobolev remainder. This is a theorem about finitely many declared breakpoints, not a representation of every possible singular measure or arbitrary image boundary.

> **Read this as:** A value jump uses a step. A slope jump uses a ramp. The derivative order determines the packet and its normalization.

V32 / ESTIMATED PACKETS / AN EXACT AMPLIFICATION EXAMPLE

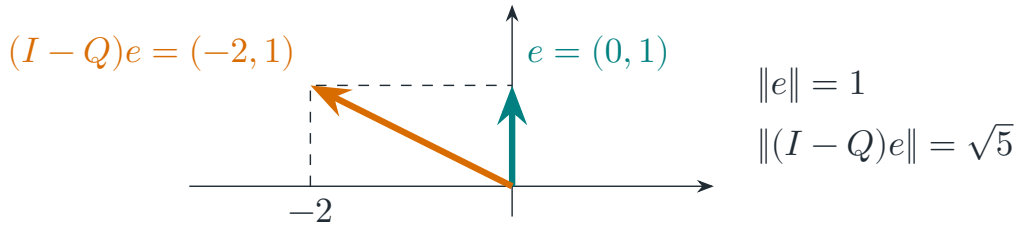
The approximation operator also acts on packet error

Traditional monograph: Theorem E7.8 (p. 58).

Let Q be a bounded *linear* operator and set $e = j - \hat{j}$. Substituting $f = s + j$ into $\hat{f} = Q(f - \hat{j}) + \hat{j}$ gives

$$f - \hat{f} = (I - Q)s + (I - Q)e.$$

The second term includes the action of Q on the packet error. Linearity is required for this split.



An oblique projection can increase the error

In $\mathbb{R}^2$, choose

$$Q = \begin{pmatrix} 1 & 2 \\ 0 & 0 \end{pmatrix}, \qquad s = (1,0), \qquad e = (0,1).$$

Then $Q^2 = Q$ and $Qs = s$, but $(I - Q)e = (-2, 1)$. The smooth error is zero while the unit packet error becomes a reconstruction error of norm $\sqrt{5}$. For a fully specified realization, take $j = e$ and $\hat{j} = 0$. Then $f = (1, 1)$, $\hat{f} = Qf = (3, 0)$, and their difference is exactly $(-2, 1)$.

Use the actual operator bound

The always-valid estimate in this setting is

$$\|f - \hat{f}\| \le \|s - Qs\| + (1 + \|Q\|)\|e\|.$$

Here $\|Q\| = \sqrt{5}$, so the bound is conservative. For an orthogonal projection, $\|I - Q\| \le 1$ gives a sharper estimate. The operator's geometry determines which bound can be used.

> **Read this as:** Recovering the packet and approximating the background are separate tasks, but the chosen approximation operator couples their final errors.

V33 / REGION VALUES / EVERY MISCLASSIFICATION HAS A WEIGHT

An overlap matrix is an exact error ledger

Traditional monograph: Theorem E8.1 (p. 59), Theorem E8.4 (p. 61).

Let two piecewise-constant signals have region values $a = (0, 2, 5)$ and $b = (0, 2, 5)$ on a counting-measure domain of twelve samples. Suppose their overlap counts are

$$M = \begin{pmatrix} 3 & 1 & 0 \\ 0 & 2 & 1 \\ 1 & 0 & 4 \end{pmatrix}.$$

The row totals $(4, 3, 5)$ are the true region sizes. The column totals are also $(4, 3, 5)$, although three samples are assigned to the wrong values.

	$b_0 = 0$	$b_1 = 2$	$b_2 = 5$
$a_0 = 0$	3	1	0
$a_1 = 2$	0	2	1
$a_2 = 5$	1	0	4

off-diagonal squared error

$$1 \cdot 2^2 + 1 \cdot 3^2 + 1 \cdot 5^2$$

$$= 38$$

Equal error counts need not mean equal reconstruction error

The three off-diagonal entries contribute one sample each. Their value differences have magnitudes 2, 3, and 5, so

$$\|j - \hat{j}\|_1 = 2 + 3 + 5 = 10, \qquad \|j - \hat{j}\|_2^2 = 4 + 9 + 25 = 38.$$

The counting-measure supremum error is 5. The full-domain mean squared error is $38/12 = 19/6$. A mismatch fraction of $3/12$ reports a different statistic and omits all three amplitude weights.

The formula requires partitions

For measurable partitions (E_i) and (F_j), with $M_{ij} = \mu(E_i \cap F_j)$,

$$\|j - \hat{j}\|_p^p = \sum_{i,j} |a_i - b_j|^p M_{ij}, \qquad 1 \le p < \infty.$$

Each location belongs to exactly one intersection. Overlapping indicator packets that are summed at the same location must first be converted to their actual combined value regions before this partition formula is applied.

> **Read this as:** Geometry supplies the overlap counts. Signal values supply the weights. Their combination gives the norm relevant to reconstruction.

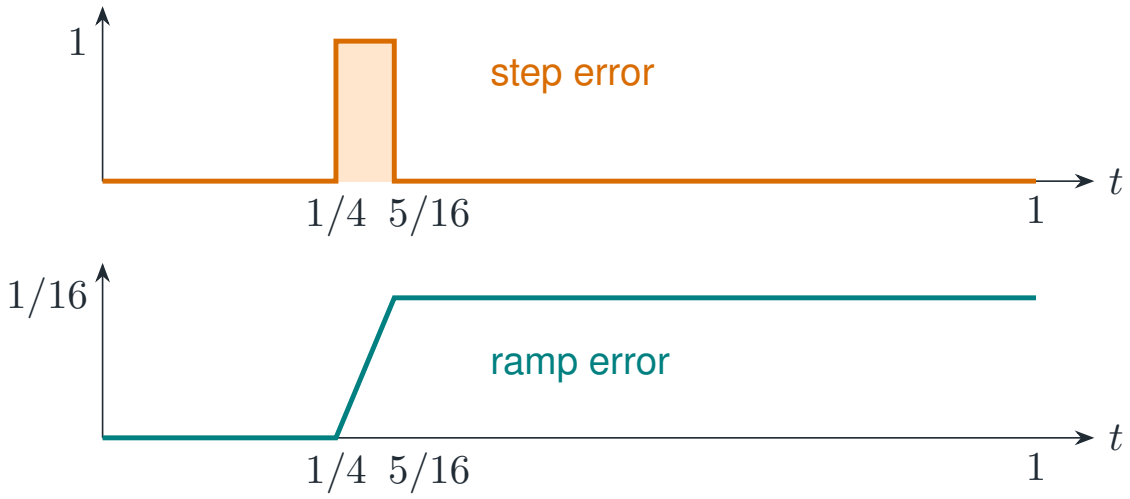
V34 / LOCATION STABILITY / A ONE-SIXTEENTH DISPLACEMENT

A shifted step and a shifted ramp have different moduli

Traditional monograph: Corollary E8.2 (p. 60), Theorem E8.3 (p. 60), Theorem E8.7 (p. 62).

Set $\tau = 1/4$, $\hat{\tau} = 5/16$, and $\delta = 1/16$. A unit step differs only on the interval between its two locations:

$$\|H_\tau - H_{\hat{\tau}}\|_2 = \sqrt{\delta} = \tfrac{1}{4}, \qquad \|H_\tau - H_{\hat{\tau}}\|_\infty = 1.$$



For a ramp, the error height shrinks too

Write $r_\tau(t) = (t - \tau)_+$. The difference grows linearly from zero to δ on $[\tau, \hat{\tau}]$ and stays at δ on $[\hat{\tau}, 1]$. Direct integration gives

$$\|r_\tau - r_{\hat{\tau}}\|_2^2 = \frac{\delta^3}{3} + (1 - \hat{\tau})\delta^2 = \frac{17}{6144} < \frac{1}{256} = \delta^2.$$

Thus the L^2 ramp error is at most δ, and its supremum error is exactly δ. Positive-order truncated powers are Lipschitz in location, with the order-dependent constants supplied by the theorem.

The geometric analogue needs a tube hypothesis

If $E \triangle F \subseteq N_\delta(\partial E)$ and the tube measure is bounded by $C\delta$ for the stated range $0 < \delta < \delta_0$, then an equal-amplitude packet has finite-p error at most $|a|(C\delta)^{1/p}$. A mere visual impression that two boundaries are close does not establish the inclusion or the tube estimate.

> **Read this as:** A displaced value jump keeps its full height on a smaller set. A displaced derivative-jump packet also decreases in height.

V35 / DETERMINISTIC RECOVERY / A NUMERICAL CERTIFICATE

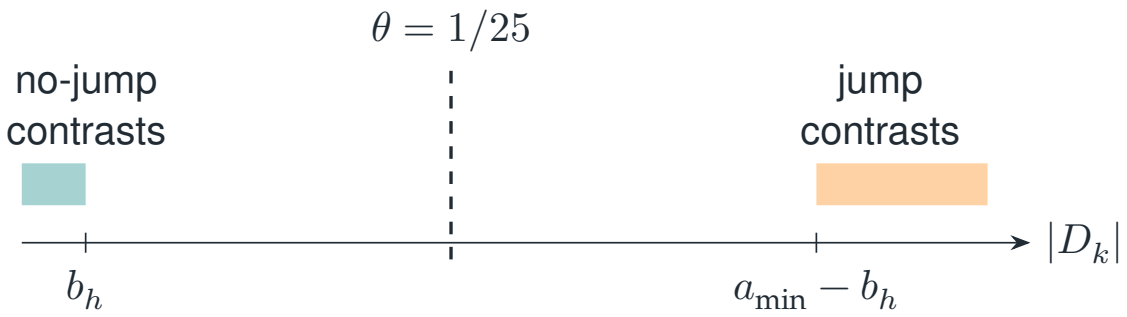
A strict contrast gap locates every separated jump cell

Traditional monograph: Lemma E9.1 (p. 63), Theorem E9.2 (p. 64), Example E9.7 (p. 65).

Assume a real signal is a Lipschitz background plus jumps. No jump occurs at a sample point, and each grid cell contains at most one jump. With spacing h, background constant M, and sample noise bounded by ε, write

$$b_h = Mh + 2\varepsilon.$$

A no-jump contrast has magnitude at most b_h. A jump contrast differs from its true amplitude by at most b_h.



Use the manuscript's exact numerical inputs

For $h = 1/256$, $M = 1/2$, $\varepsilon = 1/500$, and $a_{\min} = 2/25$,

$$b_h = \frac{381}{64000} < \frac{3}{500}, \qquad 2b_h < \frac{2}{25}.$$

The threshold $\theta = 1/25$ lies strictly between b_h and $a_{\min} - b_h$. Thresholding $|y_{k+1} - y_k|$ therefore identifies exactly the jump-containing cells. Their midpoints locate each jump within $1/512$, and each estimated amplitude differs by at most $381/64000$.

The recovered object is a cell, then an estimate

The theorem does not claim the exact real-valued breakpoint from a finite grid. It supplies the correct cell and a controlled representative of that cell. For real jumps, the amplitude error is too small to reverse the sign. Two opposite jumps in one cell could cancel in the contrast, which is why the separation assumption cannot be dropped. For Q jumps the finite-p packet bound contains both Qb_h and the location term $(\frac{h}{2}\sum_q |a_q|^p)^{1/p}$. A growing number of jumps or shrinking separation requires its own analysis.

> **Read this as:** Exact detection means exact cell identification under a strict margin. Location and amplitude remain quantitative estimates with explicit errors.

V36 / SUBGAUSSIAN RECOVERY / A CONSERVATIVE SAMPLE BUDGET

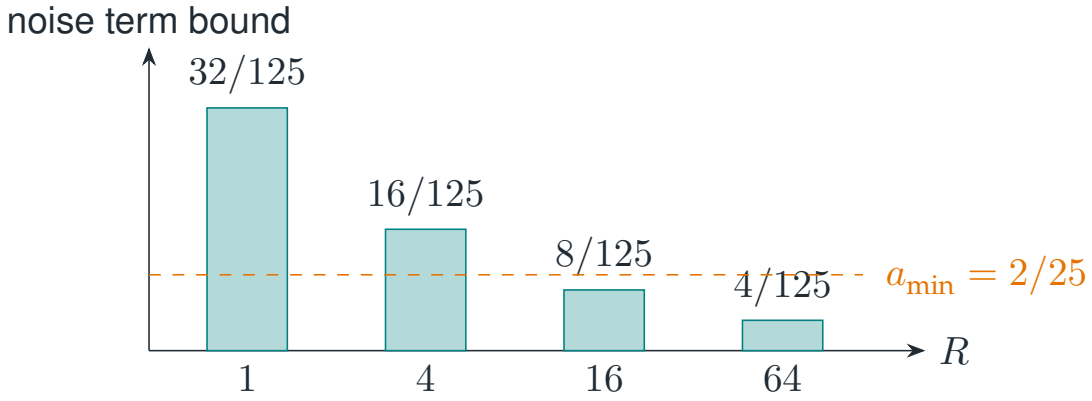
Repeated observations shrink the noise side of the gap

Traditional monograph: Theorem E9.8 (p. 67), Corollary E9.9 (p. 67), Corollary E9.11 (p. 67).

Suppose all noise observations are independent and centered, with moment-generating bound $\mathbb{E}e^{\lambda\eta} \le e^{\sigma^2\lambda^2/2}$. Averaging R observations at each of $n+1$ sample sites replaces σ by $\sigma/\sqrt{R}$. A sufficient recovery condition is

$$a_{\min} > 2Mh + \frac{4\sigma}{\sqrt{R}}\sqrt{\log(2n/\delta)}.$$

The probability of simultaneous recovery is then at least $1-\delta$, subject to the same separated-jump sampling model.



One explicit conservative choice

Take $n = 256$, $M = 1/2$, $\sigma = 1/50$, $\delta = 1/20$, and $a_{\min} = 2/25$. Since $\log(10240) < 10 < (16/5)^2$, choosing $R = 16$ makes the right-hand side strictly less than

$$\frac{1}{256} + \frac{1}{50}\frac{16}{5} = \frac{2173}{32000} < \frac{2}{25}.$$

Thus sixteen independent observations per sample site suffice for this bound. It is a sufficient budget, not a claim of optimal sample complexity. The exact checker certifies $e^{10} > 10240$ using a rational Taylor partial sum, so the displayed logarithm bound is not based on rounded numerical logarithms.

Dependence between adjacent contrasts is allowed

Neighboring differences share one sample. The proof does not require those differences to be independent: it bounds each tail and applies a union bound across all cells. Independence of the underlying sample noises supplies the stated subGaussian difference proxy. For a fixed separated signal model, a noise scale satisfying $\sigma_n\sqrt{\log n} \to 0$ gives the stated recovery consistency in probability. A fixed positive noise scale with one observation per site does not satisfy that condition.

> **Read this as:** More grid points and more repeated observations affect different terms. Finer spacing reduces background variation; averaging reduces noise.

V37 / SCALE SEPARATION / A QUANTIFIED THRESHOLD

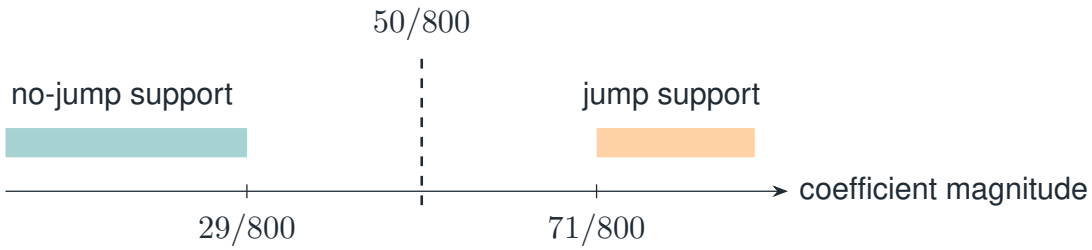
Haar sees a jump only through its position in the support

Traditional monograph: Theorem E10.1 (p. 68), Theorem E10.2 (p. 69), Corollary E10.3 (p. 69).

Use a normalized Haar wavelet on an available interval of length L on the real line, positive on its left half and negative on its right half. Here H_τ is the half-line step $\mathbf{1}_{[\tau,\infty)}$, so there is no artificial second endpoint jump. If u is the relative jump position in the support, then

$$|\langle aH_\tau, \psi\rangle| = |a|\sqrt{L}\min\{u, 1-u\}.$$

The coefficient vanishes when the jump lies exactly at a support boundary.



A worked gap with all three contributions

Take $L = 1/4$, $M = 1$, noise bound $\varepsilon = 1/100$, jump size $|a| = 1$, and relative boundary distance at least $\rho = 1/4$. The smooth and noise contributions together are bounded by

$$\frac{M}{4}L^{3/2} + \varepsilon\sqrt{L} = \frac{1}{32} + \frac{1}{200} = \frac{29}{800}.$$

The jump contributes at least $|a|\rho\sqrt{L} = 1/8$. The observed coefficient therefore has magnitude at least $71/800$. The threshold $1/16 = 50/800$ separates every such jump support from every support containing no jump.

The available support is part of the hypothesis

A fixed extension is needed if the observation domain ends inside a desired wavelet support. Also, the selected support must contain only the specified jump. An unmodeled second jump can reinforce or cancel the coefficient.

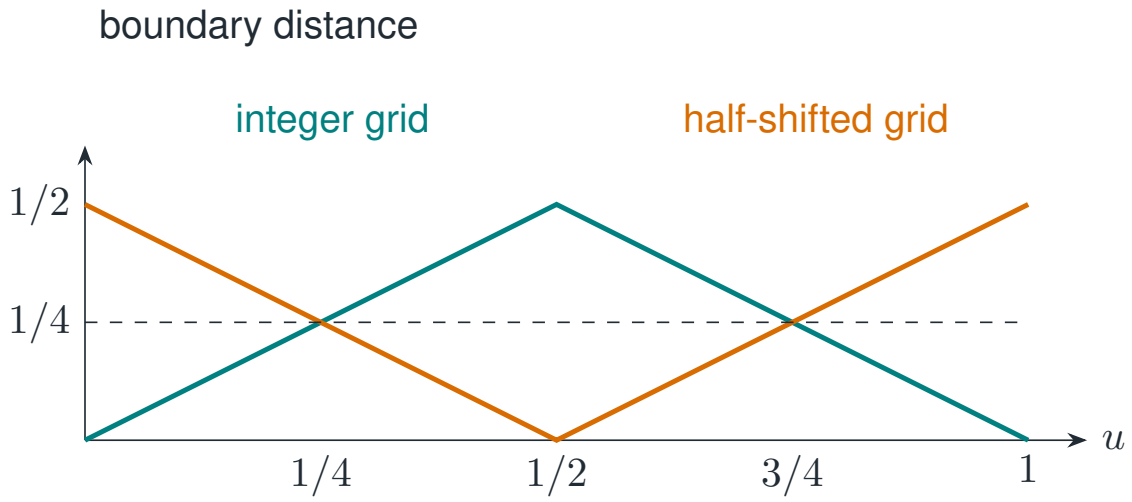
> **Read this as:** Scale determines the smooth contribution, but alignment determines the jump contribution. A detection guarantee needs both.

V38 / HAAR SUPPORT GEOMETRY / A QUARTER-CELL GUARANTEE

Two half-shifted grids remove exact alignment failure

Traditional monograph: Theorem E10.4 (p. 69), Corollary E10.5 (p. 70).

Normalize the support length to $L = 1$. The first grid has boundaries at integers; the second has boundaries at half-integers. For a point at position u modulo one, consider its distance to the nearest boundary of each grid.



At least one distance is one quarter or larger

If $1/4 \le u \le 3/4$, the original cell works. Outside that range, a shifted cell works. At $u = 1/8$, the original margin is $1/8$ but the shifted cell $[-1/2, 1/2)$ gives margin $3/8$. At $u = 1/4$, both margins equal $1/4$, showing that the uniform constant cannot be increased for these two grids. Restoring L, every jump has a support with boundary distance at least $L/4$ in one grid. If L is smaller than the minimum jump separation, that support contains no other jump.

Insert the geometric guarantee into the analytic estimate

The alignment-independent gap condition is

$$\frac{|a|}{4} > 2\varepsilon + \frac{M}{2}L.$$

For the numbers on V37, the right side is $1/50 + 1/8 = 29/200$, strictly below $1/4$. Hence one of the two grids supplies a detectable support for every jump, regardless of its position relative to the original grid. This certifies a useful support per jump. It does not by itself specify a unique rule for merging duplicate detections from two grids into exact breakpoint estimates.

> **Read this as:** A second grid repairs a geometric blind spot. The contrast inequality and separation hypotheses still supply the recovery guarantee.

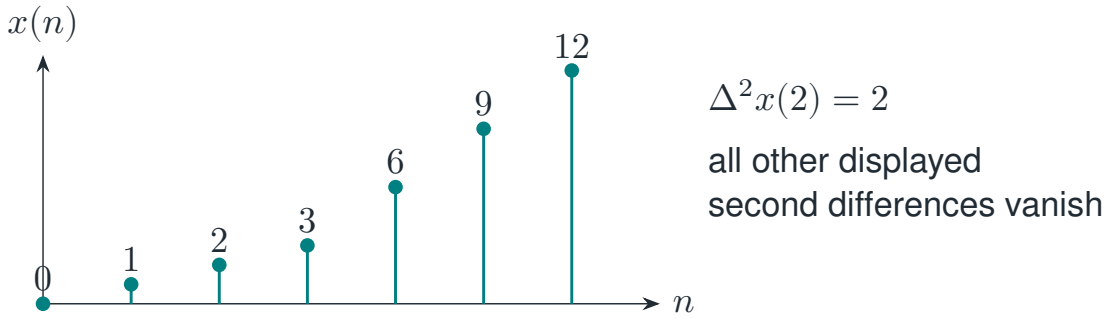
V39 / BOUNDARY STENCILS / SIGNAL RESPONSE AND NOISE GAIN

Higher differences annihilate polynomial interiors

Traditional monograph: Theorem E10.6 (p. 71), Proposition E10.7 (p. 71), Corollary E10.8 (p. 71).

Let $x(n) = n + 2(n-3)_+$ on the integer line. Its slope changes from one to three at the breakpoint 3. The second forward difference is

$$\Delta^2 x(n) = x(n+2) - 2x(n+1) + x(n).$$



An exact small stencil calculation

The values at $n = 0, \dots, 6$ are $(0, 1, 2, 3, 6, 9, 12)$. The second differences at $n = 0, \dots, 4$ are $(0, 0, 2, 0, 0)$. The only nonzero stencil straddles the slope change. For a piecewise polynomial of degree at most d, Δ^{d+1} vanishes whenever its complete stencil lies inside one polynomial piece. Candidate responses are confined to at most $d+1$ starting addresses before each breakpoint. They are boundary bands, not necessarily exact breakpoint locations.

Differencing also acts on noise

The absolute binomial coefficients of Δ^m sum to 2^m. With $\|\eta\|_\infty \le 1/16$, second-difference noise is bounded by $4/16 = 1/4$. The true response 2 remains at least $7/4$, while all responses away from the breakpoint band remain at most $1/4$. Threshold one separates them. The general robust-band statement assumes that each breakpoint has at least one sufficiently large true response. A breakpoint can fail that assumption through its amplitude or local configuration. Annihilation away from boundaries does not prove a lower bound at every boundary.

> **Read this as:** A vanishing interior response is structural. Reliable detection also requires the boundary response to exceed the amplified noise.

V40 / THE TWO ERRORS / A COMPLETE READING CONTRACT

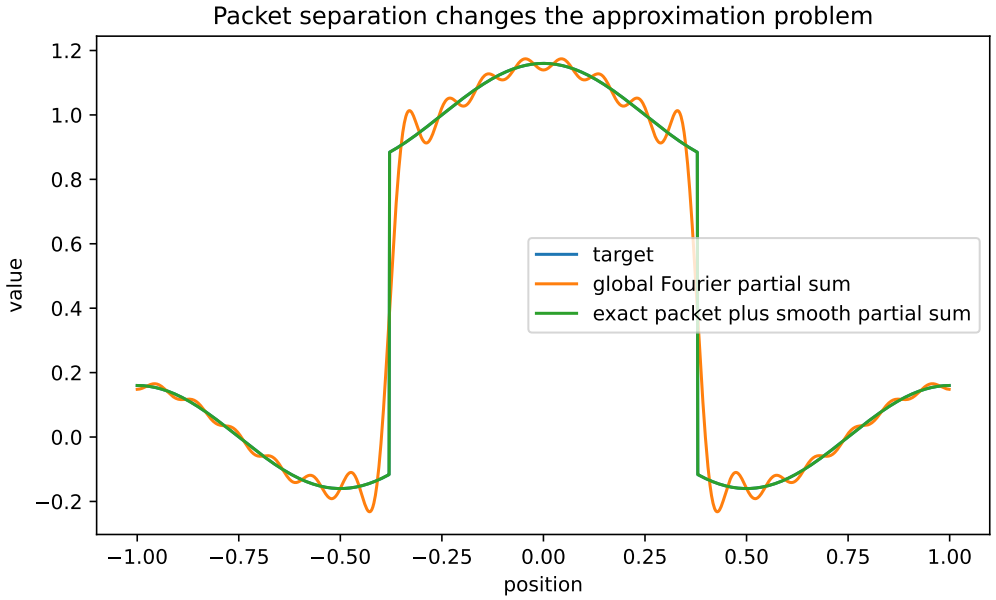
Conditional Gibbs control closes the reconstruction estimate

Traditional monograph: Theorem E11.2 (p. 72), Theorem E11.3 (p. 73).

The exact retained-packet approximation shares the same discontinuity as its target. If $f = s + j$ and $Q_m s \to s$ uniformly, then

$$\|f - (Q_m s + j)\|_\infty = \|s - Q_m s\|_\infty \longrightarrow 0.$$

The approximants themselves may be discontinuous. There is no conflict with the failure of continuous Fourier polynomials to converge uniformly to a jump.



Preserved vector figure from the expanded source. It illustrates the choice of approximation target; it is not an independent universal recovery certificate.

Estimated packets need the operator norm too

If the bounded linear operators satisfy $\sup_m \|Q_m\|_{L^p \to L^p} \le C_Q$, then

$$\|f - \hat{f}_m\|_p \le \|s - Q_m s\|_p + (1 + C_Q)\|j - \hat{j}_m\|_p.$$

For an orthogonal L^2 projection one may use $\|I - Q_m\| \le 1$, sharpening the packet multiplier to one. Such a bound must be justified in the norm actually being used. In particular, uniform boundedness of Fourier partial sums in the supremum norm is not supplied by this theorem.

What the complete guide establishes

V1 to V26 preserve the original worked layer. V27 to V40 explain the additional traditional developments, outside the inherited twenty-result Lean map. The new exact examples have their own finite arithmetic record. They do not extend the formal coverage merely by appearing beside formally mapped results. The final package keeps current computations and inherited compiler reports attached to their stated scope. The symbolic visual suite was rerun, with 181 checks passed and seven invalid controls rejected. Lean compiler evidence remains inherited. No fresh Lean run is claimed.

> **Read this as:** A complete error claim specifies the signal class, the recovered packet, the approximation operator, and the norm. Each term can then be improved and checked on its own terms.