OPERATION KOMOREBI

REALIZATION CALCULUS



RESEARCH LIBRARY

Sufficient Resolution

Exact reduction, hidden memory, and certified reconstruction

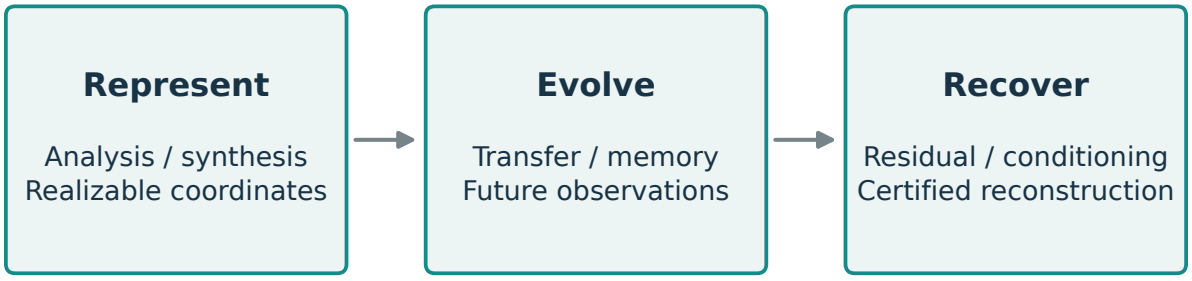
Visual guide

Komorebi P.

Standardized publication edition / 8 October 2026

REALIZATION CALCULUS

Reading scope



Every stage retains its own assumptions and error budgets.

The working route: represent, evolve, and recover with explicit contracts.

Contents

1 How to read a sufficient resolution 4

2 Exact representation begins with reconstruction 5

3 Redundancy creates a realizable coordinate space 6

4 Compatible kernels and packet transfer blocks 7

5 Compression exposes a hidden-return defect 8

6 A return can hide beyond the tested depths 9

7 Eliminating hidden dynamics produces memory 10

8 A finite certificate can establish every depth 11

9 Prepared closure and all-state closure differ 12

10 Sufficiency depends on every permitted future 13

11 A finite observation horizon can be complete 14

12 Minimality concerns the actual model image 15

13 A certificate can identify the minimal quotient 16

14 Approximate intertwining has an exact error identity 17

15 Stability controls the accumulation of defects 18

16 Representation chains retain both errors and conditioning 19

17 Continuous hidden elimination gives an actual memory equation 20

18 Jets link finite algebra to all-time closure 21

19 Continuous defects obey a transport integral 22

20 Compatibility and realizability have different kernels 23

21 The geometry separates three kinds of content 24

22 Residual size needs a conditioning constant 25

23 Enlarging realizations transfers obstruction energy 26

24 The best image can have many source preimages 27

25 Source ambiguity survives exact and noisy recovery 28

26 The obstruction has its native quotient metric 29

27 Intrinsic classes also handle inconsistent data 30

28 Positive regularization constructs a unique source 31

29 Bias, noise, and kernel ambiguity remain separate 32

30 Strong convergence is different from uniform inversion 33

31 Discrepancy selection uses the observed fitting curve 34

32 Accepted discrepancy gives a direct consistency route 35

33 Inexact solves require two residual budgets 36

34 Vanishing raw residuals can still certify the wrong source 37

35 An uncertain operator inflates the data allowance 38

36 Model residuals need transport to the true operator 39

37 Generalized acceptance generates its own source bound 40

38 Acceptance and parameter feasibility remain separate 41

39 Spanning calibration controls the whole operator 42

40 Approximate frames expose conditioning and unseen directions 43

41 The completed calibration-to-recovery chain 44

42 Evidence and reusable source 45

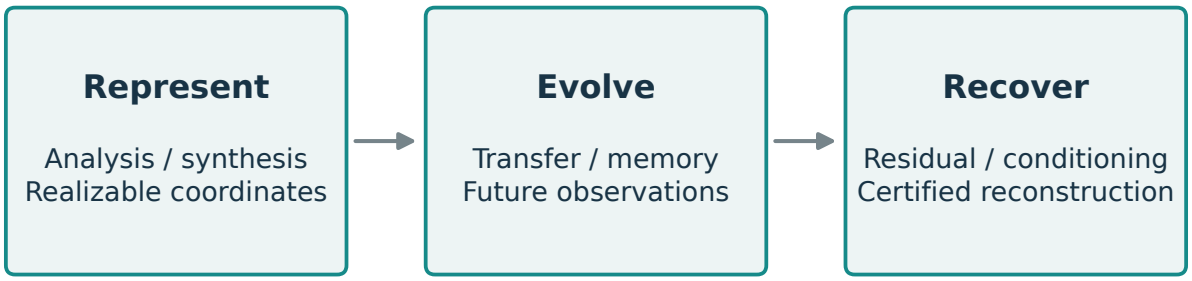
1 How to read a sufficient resolution

The governing question is operational: what must a representation retain to support a specified observation and every permitted continuation? The answer depends on the observation, the operations, the admissible models, and the metric used for error. A representation can be exact for one contract and insufficient for another.

This guide follows three connected threads. Exact coordinates and hidden-return identities explain when representation commutes with evolution. Behavioral quotients identify the distinctions future observations require. Reconstruction, regularization, discrepancy, and calibration attach quantitative certificates to recovery.

The worked examples use small rational systems. The accompanying generator checks 280 exact assertions with `fractions.Fraction`. Smooth plotted curves evaluate printed analytic formulas using floating point for illustration. Neither kind of drawing proves an infinite limit or validates an unknown physical operator.

Source anchors identify actual modules and declarations in the frozen v0.16.0 program. The traditional paper supplies consolidated theorem statements and proofs. The Lean guide distinguishes source inspection, archived compilation, and any fresh verification recorded for the final release.



Representation, evolution, and recovery share one language while retaining distinct proof obligations.

Traditional theorem groups: SR-04

Source anchors: `THEOREM_GUIDE.md;ExactResolution.lean:decode_encode;Behavioral.lean:all_future_observations_preserved`

2 Exact representation begins with reconstruction

An analysis map records an object in coordinates; a synthesis map reconstructs it. The essential hypothesis is the split identity below. Coordinates may be redundant, and neither orthogonality nor finite dimension is required in the algebraic result. The scalar setting is a semiring with additive commutative monoids.

The coordinate projector selects precisely the realizable records. Its fixed-point space is the image of analysis. A coordinate vector outside that image is a record the declared synthesis/analysis pair does not recognize as an exact representation.

For two objects, encoding a genuine linear transformation and decoding its kernel returns the original transformation. Decoding and re-encoding an arbitrary coordinate kernel instead sandwiches it between the two realizability projectors. This distinction becomes important whenever the coordinates contain unused directions.

$$a_X : X \to V_X, \quad s_X : V_X \to X, \quad s_X a_X = I_X,$$
$$E_X = a_X s_X, \quad E_X^2 = E_X, \quad E_X v = v \iff v \in \operatorname{Ran} a_X.$$
$$K(T) = a_Y T s_X, \qquad D(K) = s_Y K a_X.$$
$$D(K(T)) = T, \qquad K(D(K)) = E_Y K E_X.$$

Traditional theorem groups: SR-01

Source anchors: `THEOREM_GUIDE.md;ExactResolution.lean:projector_idempotent;ExactResolution.lean:fixed_iff_realizable;ExactResolution.lean:decode_encode;ExactResolution.lean:encode_decode`

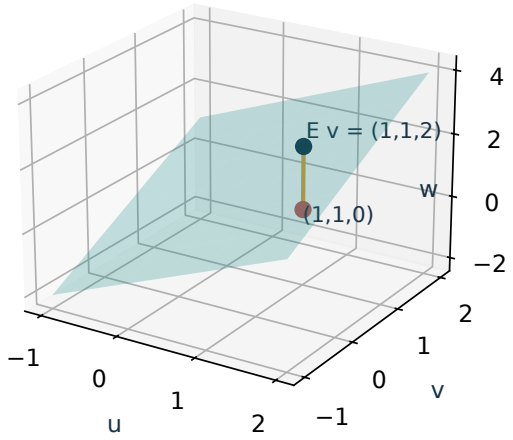
3 Redundancy creates a realizable coordinate space

Take a two-dimensional source and store three coordinates: the two entries and their sum. Synthesis reads the first two entries. These maps reconstruct every original vector, although arbitrary triples need not be realizable.

The stored record $(1,1,0)$ violates its own consistency relation. The coordinate projector returns $(1,1,2)$. The plane in the figure is the set of fixed records. The projection is oblique for the ordinary Euclidean metric; the split algebraic theorem does not assert it is an orthogonal projection.

Redundancy is useful when different operations or measurements have natural coordinate views. It also imposes a compatibility obligation: operations on all coordinate triples must respect the realizable plane before they can be interpreted as transformations of the original object.

$$a(x,y) = (x,y,x+y), \qquad s(u,v,w) = (u,v),$$
$$E(u,v,w) = (u,v,u+v).$$



The plane $w = u + v$ consists of realizable records. The highlighted correction uses the declared oblique projector.

Traditional theorem groups: SR-01

Source anchors: `THEOREM_GUIDE.md;ExactResolution.lean:fixed_iff_realizable;ExactResolution.lean:compatible_iff_existsUnique`

4 Compatible kernels and packet transfer blocks

A coordinate kernel represents one unique source transformation exactly when it satisfies the projector sandwich identity. The compatibility theorem proves both existence and uniqueness. Its content is a correspondence between actual maps and compatible kernels, rather than a convention that every coordinate array is automatically meaningful.

When coordinate spaces have finite direct-sum decompositions, a packet is one component $P_i V_X$. A transfer block $K_{ji} = Q_j K P_i$ is the action from source packet i to target packet j. Each block is itself a map, rather than necessarily a scalar weight. The full block array reconstructs K. For redundant coordinates, compatibility is a condition on that assembled kernel; individual coordinate blocks need not preserve the realizable subspace by themselves.

Encoded transformations compose by ordinary kernel multiplication. This follows from the split reconstruction identity in the intermediate space. Complete packet coordinates can therefore turn operator composition into transfer-block composition while retaining the original operator.

For redundant coordinates, even the full coordinate identity may fail compatibility. In the three-coordinate example, $EIE = E \neq I$. The appropriate identity on realizable records is E. Unused coordinate directions cannot be treated as additional mathematical states merely because the coordinate container includes them.

$$E_Y K E_X = K \iff \exists! T : \; K = a_Y T s_X,$$
$$K(ST) = K(S)K(T).$$

Traditional theorem groups: SR-01

Source anchors: `THEOREM_GUIDE.md;ExactResolution.lean:compatible_iff_existsUnique;ExactResolution.lean:kernelEquiv;ExactResolution.lean:encode_comp`

5 Compression exposes a hidden-return defect

Compression has a different contract from complete coordinate encoding. Let C retain a coarse state and let R choose a lift with $CR = I$. The retained projector is $P = RC$ and the discarded projector is $Q = I - P$. The commutative-ring hypothesis permits these differences.

The compressed transform $\Phi(T) = CTR$ need not preserve products. Its exact defect is the path that enters the discarded space during the first operation and returns during the second. A zero defect for one pair or one depth only settles that pair or depth.

Complete coarsening that groups full packet contents remains exact. Discarding a component or substituting a statistic is an additional map, with an additional kernel. The hidden-return theorem applies to the specified compression and lift; a smaller number of addresses alone does not establish information loss.

$$\Delta(S, T) = \Phi(ST) - \Phi(S)\Phi(T) = CSQTR,$$
$$\Phi(T^{m+n}) - \Phi(T^m)\Phi(T^n) = CT^mQT^nR.$$

Traditional theorem groups: SR-02

Source anchors: `THEOREM_GUIDE.md;Compression.lean:defect_factorization;Compression.lean:power_defect;Compression.lean:defect_coherence`

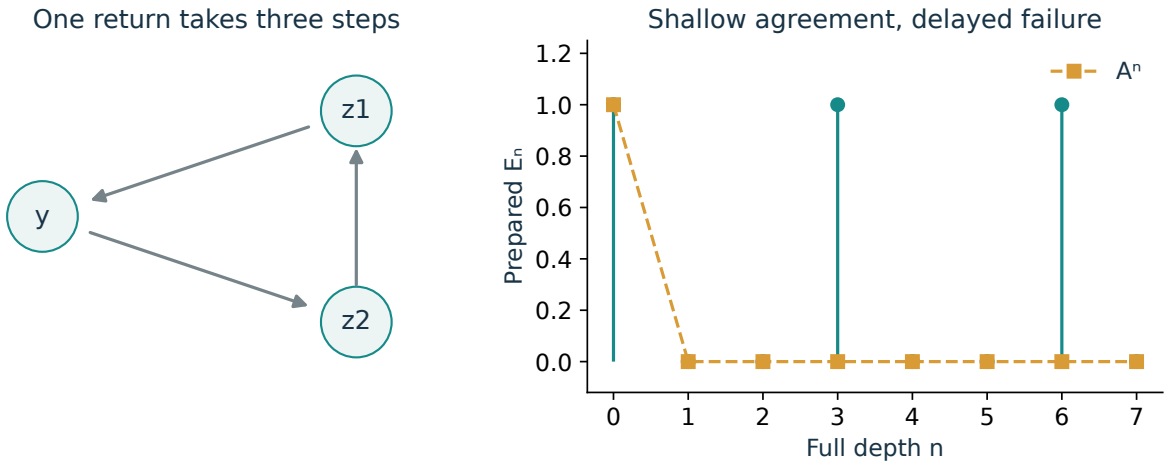
6 A return can hide beyond the tested depths

Our first worked system rotates three rational coordinates. The visible coordinate is y, while (z_1, z_2) are hidden. Preparation places the visible input in $(y, 0, 0)$. The one-step compressed operator is zero.

The first two positive depths agree with powers of that zero operator. At depth three the hidden path returns the original visible input. This is an exact counterexample to extrapolating shallow closure indefinitely. The discrepancy is an identity map, not a small numerical mismatch.

The figure follows the same input around its three-step cycle. At depth zero both prepared and coarse evolutions are identities. This convention prevents the zero-power case from being silently omitted.

$$L(y, z_1, z_2) = (z_1, z_2, y), \quad C(y, z_1, z_2) = y, \quad R(y) = (y, 0, 0),$$
$$E_0 = I, \quad E_1 = 0, \quad E_2 = 0, \quad E_3 = I, \qquad A = CLR = 0.$$



The prepared operator returns after three full steps, although the coarse operator remains zero at every positive depth.

Source anchors: `BLOCK_02_THEOREMS.md;Examples.lean:correct_through_depth_two;Examples.lean:fails_at_depth_three;MemoryExamples.lean:delayed_memory_one`

7 Eliminating hidden dynamics produces memory

Write a split linear step as $T(y,z) = (Ay + Bz, Gy + Dz)$. The maps G and B cross the visible/hidden boundary; D evolves hidden content. The source constructs the full operator and derives its trajectories, rather than assuming the recurrence as a separate conclusion.

Iterating the hidden recurrence yields an initial-state term and a finite history convolution. Substitution into the visible recurrence gives an exact memory equation. Each history coefficient $M_k = BD^kG$ is a visible-to-hidden-to-visible return with k internal hidden transitions.

Both equations include $n = 0$, where the sum is empty. The initial hidden state cannot be dropped unless preparation sets it to zero or another stated condition removes its effect. These are algebraic module identities over a commutative ring.

$$z_n = D^n z_0 + \sum_{j=0}^{n-1} D^{n-1-j} G y_j,$$

$$y_{n+1} = Ay_n + BD^n z_0 + \sum_{j=0}^{n-1} M_{n-1-j} y_j, \qquad M_k = BD^kG.$$

Traditional theorem groups: SR-05

Source anchors: `BLOCK_02_THEOREMS.md;DiscreteMemory.lean:hidden_elimination;DiscreteMemory.lean:discrete_memory_equation;DiscreteMemory.lean:forced_recurrence_formula`

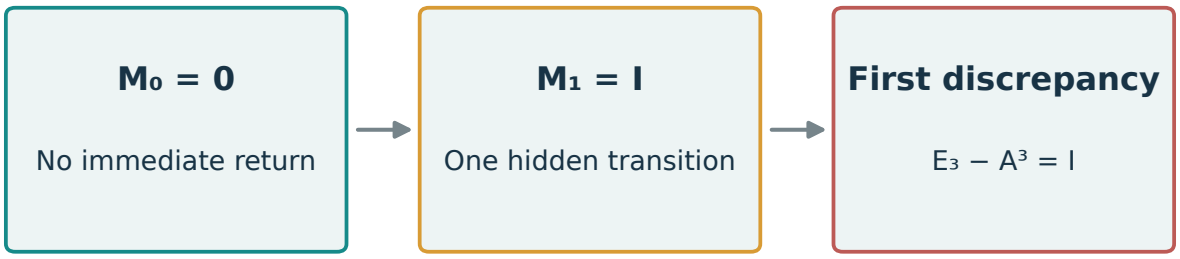
8 A finite certificate can establish every depth

For prepared states, exact visible closure at every natural depth is equivalent to vanishing of every feedback moment. It is also equivalent to the prepared semigroup law. Necessity is proved by successively isolating the feedback coefficients; it does not require the entry or return maps separately to vanish.

Over a field, a supplied monic annihilating polynomial for D reduces the all-depth moment obligation to its first degree-many moments. If the hidden space has finite dimension d, Cayley–Hamilton supplies that polynomial. The visible space may remain infinite-dimensional.

The certificate tests equality of whole linear maps. A finite list of sampled states does not satisfy that hypothesis without an independently justified spanning or whole-map argument. The zero-dimensional hidden case is included.

$$\big[\forall n,\ E_n = A^n\big] \iff \big[\forall k,\ BD^kG = 0\big],$$
$$\dim H = d \implies \big[\forall n,\ E_n = A^n\big] \iff \big[M_0 = \cdots = M_{d-1} = 0\big],$$
$$\big[\forall n,\ E_n = A^n\big] \iff \big[E_n = A^n \text{ for } 0 \le n \le d+1\big].$$



Moment k is first visible at full depth k + 2.

A return moment indexed by k first contributes at full depth $k+2$. In the worked system $d = 2$ and $M_1 = I$.

Traditional theorem groups: SR-06, SR-07

Source anchors: `BLOCK_02_THEOREMS.md;MemoryClosure.lean:prepared_powers_iff_memory_zero;FiniteMemory.lean:prepared_powers_iff_finite_horizon;MemoryClosure.lean:first_return_discrepancy`

9 Prepared closure and all-state closure differ

Prepared closure describes trajectories whose hidden initial state is zero. All-state closure describes every initial hidden state. The latter holds exactly when the return map B vanishes. Prepared closure only requires every composite BD^kG to vanish.

For example, take $A = I$, $D = I$, $B(z_1, z_2) = z_1$, and $G(y) = (0, y)$. Both boundary maps are nonzero, but their active directions are disconnected: visible inputs enter z_2, while B reads only z_1. Every prepared feedback moment vanishes. An initial $z_1 \neq 0$ immediately affects the visible trajectory.

An arbitrary split compression has a canonical hidden space $\ker C$. Its decomposition $x \mapsto (Cx, Qx)$ and reconstruction $(y, z) \mapsto Ry + z$ are actual inverse maps. The corresponding block system is conjugate to the fine dynamics, so these memory criteria apply using the actual discarded dimension.

$$\left[\forall n, y_0, z_0,\ y_n = A^n y_0\right] \iff B = 0,$$
$$x \longmapsto (Cx, Qx), \qquad (y, z) \longmapsto Ry + z.$$

Traditional theorem groups: SR-06, SR-08

Source anchors: `BLOCK_02_THEOREMS.md;MemoryClosure.lean:all_state_closure_iff_B_zero;SplitMemory.lean:block_synthesis_analysis;SplitMemory.lean:exact_powers_iff_canonical_finite_horizon`

10 Sufficiency depends on every permitted future

Choose an observation O and allowed operations T_i. A state difference is behaviorally invisible when no finite permitted word can reveal it. The empty word includes the present observation. Chronological words apply the listed first operation first, so their operator products run in reverse written order.

The invisible submodule is invariant under every allowed operation and is the largest invariant submodule contained in $\ker O$. Consequently, the quotient carries actual induced operations and an actual observation. It preserves every original future observation.

Changing the observation or enlarging the operation alphabet can shrink the invisible submodule and require a larger state representation. The theorem assumes the operators are specified. It does not infer them from finite training data or identify every nonlinear encoding.

$$N = \bigcap_{w} \ker(OT_w), \qquad T_{i::w} = T_w T_i,$$

$$[x] = [y] \text{ in } X/N \iff OT_w x = OT_w y \text{ for every finite } w.$$

Traditional theorem groups: SR-04

Source anchors: `THEOREM_GUIDE.md;Behavioral.lean:invisible_invariant;Behavioral.lean:invariant_submodule_le_invisible;Behavioral.lean:quotient_eq_iff_same_future`

11 A finite observation horizon can be complete

For finite-dimensional X over a field, let N_r retain the differences invisible through word length r. These nested subspaces satisfy the recursion below. A single plateau forces invariance and certifies the all-future invisible space.

The a priori horizon is $k = \dim \ker O$, including $k = 0$. Distinct behavioral states therefore have a distinguishing word of length at most k. The operations need not commute. An infinite alphabet is allowed in the theorem, although its word family is not then a finite enumeration.

With a finite alphabet, the complete observation table has coordinates $OT_w x$ for words up to that horizon. Its state space is the realizable range of the table map, rather than the product of every arbitrary table. The source constructs dynamics on this range that reproduce observations at every later depth.

$$N_0 = \ker O, \quad N_{r+1} = \ker O \cap \bigcap_i T_i^{-1} N_r,$$

$$N_{r+1} = N_r \Longrightarrow N_r = N, \qquad N_k = N, \quad k = \dim \ker O.$$

$$H_k x = (OT_w x)_{|w| \le k}, \qquad X/N \cong \operatorname{Ran} H_k.$$

Traditional theorem groups: SR-09, SR-11

Source anchors: `BLOCK_03_THEOREMS.md;FiniteHorizon.lean:visibleKernel_stabilizes;FiniteHorizon.lean:same_future_iff_bounded;FiniteObservableModel.lean:coordinate_all_future_fidelity`

12 Minimality concerns the actual model image

An exact linear competitor $q : X \to Z$ must intertwine the specified operations and factor the observation. Its kernel is contained in N: it cannot erase a difference that some permitted future reveals. The induced comparison from $\operatorname{Ran} q$ to X/N is surjective.

This yields a minimum dimension among exact linear models of all original states and all specified continuations. Unused coordinates in Z do not count as necessary state. The final comparison with $\dim Z$ requires finite-dimensional Z; the model-image theorem does not assume q is onto Z.

Minimum dimension is attained exactly when $\ker q = N$. In that case the source constructs an actual linear equivalence from the model image to the behavioral quotient. This is a precise model class, not an unrestricted minimum over nonlinear encodings, positive realizations, or probabilistic states.

$$qT_i = U_i q, \qquad Vq = O,$$
$$\dim(X/N) \le \dim(\operatorname{Ran} q) \le \dim Z,$$
$$\dim(\operatorname{Ran} q) = \dim(X/N) \iff \ker q = N.$$

Traditional theorem groups: SR-10

Source anchors: `BLOCK_03_THEOREMS.md;BehavioralMinimality.lean:finrank_state_le_image;BehavioralMinimality.lean:image_dimension_minimal_iff;BehavioralMinimality.lean:minimalImageEquiv`

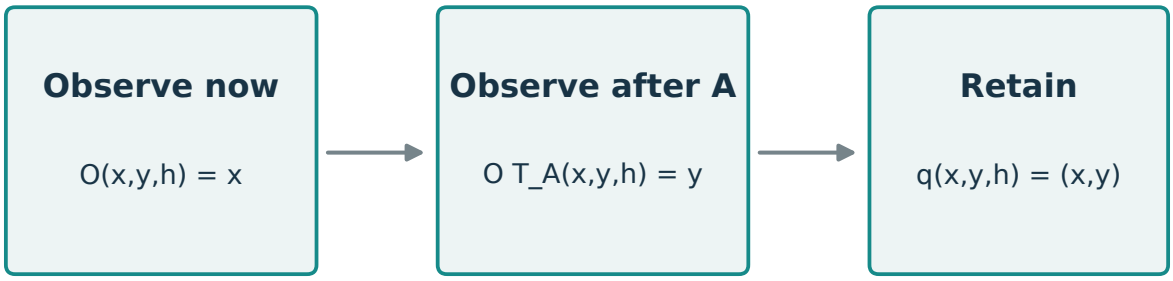
13 A certificate can identify the minimal quotient

A finite certificate combines three kinds of information. The model coordinates are generated by finitely many original observations, the operations intertwine, and the readout factors. Generation prevents the model from retaining a distinction invisible to every original observation. Intertwining and readout prevent it from discarding any future-visible distinction.

Together these identities establish $\ker q = N$. This kernel result does not require finite-dimensional X or a predeclared lower bound on the certificate depth. A certificate found early already establishes minimality.

The worked noncommuting system observes x now and y after operation A. The hidden coordinate h remains irrelevant under every word. The two-coordinate model is exact and has the smallest dimension, while the current scalar observation alone is insufficient. The rational implementation checks certificates; it is not itself formally verified or extracted from Lean.

$$q = GH_r, \qquad qT_i = U_i q, \qquad Vq = O \quad \Longrightarrow \quad \ker q = N.$$
$$T_A(x,y,h) = (y,0,h), \quad T_B(x,y,h) = (0,x,h), \quad O(x,y,h) = x.$$



Two depth-one observation coordinates determine the exact minimal behavioral model; the operations remain noncommuting.

Traditional theorem groups: SR-12

Source anchors: `BLOCK_03_THEOREMS.md;FiniteCertificate.lean:certificate_kernel_eq;ObservableExamples.lean:exactCertificate;ObservableExamples.lean:operations_do_not_commute`

14 Approximate intertwining has an exact error identity

For a candidate representation q, define the one-step defect $\delta_i = qT_i - U_i q$ and the word defect $E_w = qT_w - U_w q$. Their transport law is exact before any norm estimate is applied. Each injected one-step defect is propagated by the preceding true operations and subsequent model operations.

Whole-map norms then give a computable envelope when valid operator bounds are available. Readout error has a separate term. Its omission would confuse state-representation fidelity with accuracy of the actual requested output.

The algebraic identity uses commutative-ring modules. The analytic estimates use continuous linear maps over a nontrivially normed field. Finite dimension, completeness, a finite alphabet, and surjectivity of q are not assumed. Sampling finitely many states alone cannot justify the whole-map defect allowance.

$$E_{i::w} = E_w T_i + U_w \delta_i, \qquad E_{u++v} = E_v T_u + U_v E_u,$$
$$\|OT_w - VU_w q\| \le \eta A_w + \|V\| D_w, \quad \eta \ge \|O - Vq\|.$$

Traditional theorem groups: SR-13, SR-14

Source anchors: `BLOCK_04_THEOREMS.md;ApproximateAlgebra.lean:wordDefect_cons;ApproximateAlgebra.lean:wordDefect_eq_sum;ApproximateBounds.lean:norm_outputError_le`

15 Stability controls the accumulation of defects

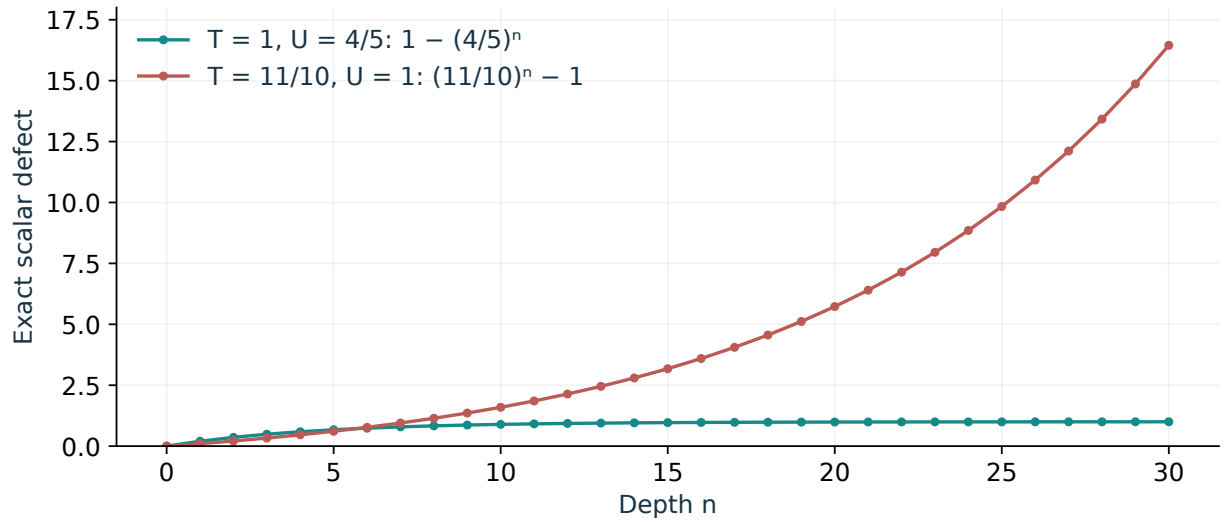
With uniform bounds $\|T_i\| \le a$, $\|U_i\| \le b$, and $\|\delta_i\| \le \varepsilon$, the two-rate response below bounds every word of length n. Equal rates use the denominator-free sum and do not divide by zero.

If the true operations are nonexpansive and the reduced operations contract with $0 \le b < 1$, the word defect is uniformly bounded by $\varepsilon/(1-b)$. If both sides are merely nonexpansive, the general estimate is $n\varepsilon$. A common operator-norm bound is essential; individual spectral-radius estimates do not supply it.

The scalar examples in the figure attain the exact response. A small one-step defect with $T = 1 + \varepsilon$ can grow without bound. This violates the true nonexpansive hypothesis and establishes the boundary of the stable result.

$$H_n(a,b) = \sum_{j=0}^{n-1} a^j b^{n-1-j}, \qquad \|E_w\| \le \varepsilon H_n(a,b),$$

$$(a-b)H_n = a^n - b^n, \qquad H_n(a,a) = na^{n-1} \quad (n \ge 1).$$



Exact rational samples contrast a stable scalar pair with an expanding true operator. The displayed curves do not infer stability from data.

Traditional theorem groups: SR-15, SR-16

Source anchors: `BLOCK_04_THEOREMS.md;ScalarBudgets.lean:response_mul_sub;StableApproximation.lean:norm_wordError_contractive;ApproximateExamples.lean:unstable_error_unbounded`

16 Representation chains retain both errors and conditioning

Composing two approximate representations propagates both defects. The first is multiplied by the second representation map; the second is multiplied by the first representation map. The corresponding readout mismatch decomposes in the same way.

A split re-encoding $S : Z \to W$, $J : W \to Z$ with $JS = I$ changes coordinate norms and may enlarge operator bounds by $\|S\|\|J\|$. Nevertheless, the actual represented output remains unchanged. A poorly conditioned coordinate choice can worsen a norm certificate without changing the mathematical output.

For arbitrary initialization z, the coarse-state mismatch has its own propagated factor. In a contractive model its influence decays geometrically. This term is separate from the defect generated by the dynamics and the mismatch of readout maps.

$$E_w(pq; T, S) = pE_w(q; T, U) + E_w(p; U, S)q,$$
$$\|E_w(pq)\| \le \|p\|\varepsilon_1 + \varepsilon_2\|q\|,$$
$$\|qT_w x - U_w z\| \le D_w\|x\| + B_w\|qx - z\|.$$

Traditional theorem groups: SR-14

Source anchors: `BLOCK_04_THEOREMS.md;ApproximateComposition.lean:wordError_two_stage;ApproximateComposition.lean:reencode_outputError;ApproximateBounds.lean:norm_state_error_le`

17 Continuous hidden elimination gives an actual memory equation

For bounded real generators on Banach spaces, the flow is the actual operator exponential. The block generator $L(y,z) = (Ay + Bz, Gy + Dz)$ yields differentiated visible and hidden trajectories. Variation of constants eliminates the hidden state.

The resulting memory kernel is $M(t) = Be^{tD}G$. The source proves the continuity and interval integrability needed for these formulas. Oriented intervals permit every real time in the identities, although later positive-time norm estimates restrict to $t \geq 0$.

The theorem applies to trajectories constructed from the full bounded generator. It does not construct a hidden trajectory from every arbitrary proposed Volterra solution, and it does not extend to unbounded generators with domain obligations.

$$z(t) = e^{tD}z_0 + \int_0^t e^{(t-s)D}Gy(s)\,ds,$$

$$y\prime(t) = Ay(t) + Be^{tD}z_0 + \int_0^t M(t-s)y(s)\,ds.$$

Traditional theorem groups: SR-17, SR-19

Source anchors: `BLOCK_05_THEOREMS.md;VariationOfConstants.lean:variation_of_constants;ContinuousMemory.lean:hidden_elimination;ContinuousMemory.lean:continuous_memory_equation`

18 Jets link finite algebra to all-time closure

The analytic bridge identifies mapped exponential equality for every real time with equality of all generator powers. Its forward direction uses actual derivatives; its reverse direction transports the norm-convergent exponential series through bounded maps.

For a hidden space of dimension d, all-time prepared closure is equivalent to the first d return moments vanishing. It is also equivalent to the prepared derivatives matching A^n through order $d+1$. These test entire operators, rather than a finite selection of times or sampled trajectories.

In the delayed example, $D^2 = 0$ and $M(t) = tI$. The first return appears in the third derivative. Its Taylor coefficient is $I/3!$. Derivatives and Taylor coefficients therefore carry different factorial conventions.

$$Ce^{tL}R = e^{tA} \text{ for all } t \iff BD^kG = 0 \text{ for } 0 \le k < d,$$
$$E^{(n)}(0) = CL^nR, \qquad M^{(n)}(0) = BD^nG,$$
$$M(t) = tI, \qquad E^{(3)}(0) - A^3 = I.$$

Traditional theorem groups: SR-18, SR-20

Source anchors: `BLOCK_05_THEOREMS.md;ExponentialTransfer.lean:mapped_flow_eq_iff_powers_eq;ContinuousClosure.lean:prepared_exact_iff_finite_memory;ContinuousJets.lean:first_return_derivative`

19 Continuous defects obey a transport integral

For a representation q and bounded generators L, A, the generator defect is $\delta = qL - Aq$. The flow defect $E(t)$ is its exact Duhamel transport integral. Vanishing generator defect is equivalent to vanishing flow defect at every real time.

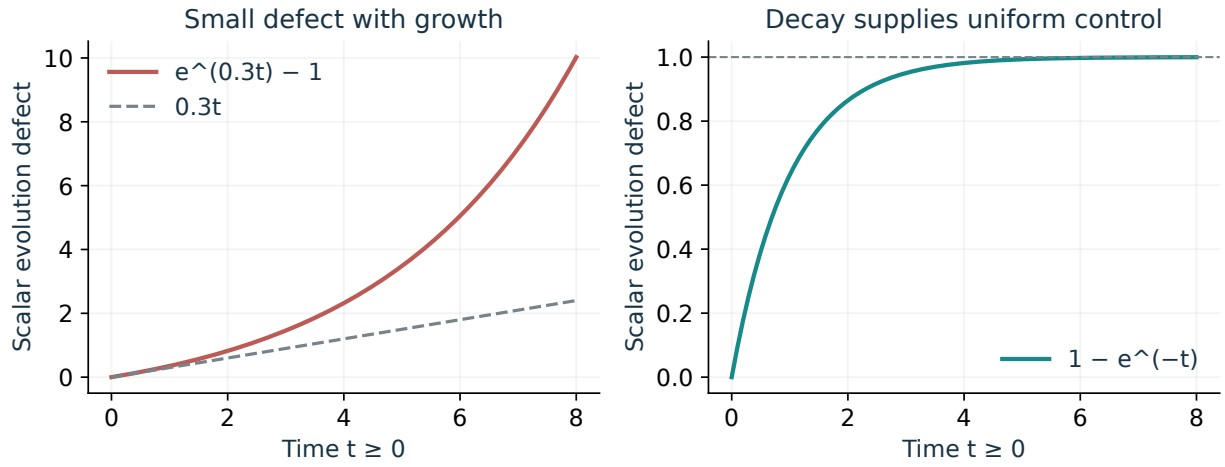
Assumed flow envelopes turn that identity into a whole-operator estimate. With a bounded true flow and a reduced flow decaying at rate $r > 0$, the response is uniformly bounded by $1/r$. Readout and initialization errors retain their own factors.

The plots evaluate scalar formulas. Their role is to show why growth assumptions matter. A small generator defect alone permits $e^{\varepsilon t} - 1$ growth. Neither eigenvalues alone nor a short sampled trajectory supply the common flow envelopes used by the theorem.

$$E(t) = qe^{tL} - e^{tA}q = \int_0^t e^{(t-s)A}\delta e^{sL}\, ds,$$

$$\|E(t)\| \le M_L M_A \varepsilon R(a,b,t), \quad R(a,b,t) = \int_0^t e^{b(t-s)}e^{as}\, ds,$$

$$R(0,-r,t) = \frac{1-e^{-rt}}{r} \le \frac{1}{r} \quad (t \ge 0).$$



Analytic scalar formulas illustrate growing and stable regimes; floating-point plotting provides no additional proof of their limits.

Traditional theorem groups: SR-21, SR-22, SR-23, SR-24

Source anchors: `BLOCK_06_THEOREMS.md;ContinuousDefect.lean:duhamel;ContinuousApproximation.lean:norm_evolution_le_response;ContinuousStability.lean:norm_evolution_uniform;ContinuousApproxExamples.lean:scalar_response_sharp`

20 Compatibility and realizability have different kernels

A linear presentation has a realization map $A : G \to V$ and a compatibility map $d : V \to W$, with $dA = 0$. Thus every realization is compatible. The reverse implication is an additional exactness condition.

The obstruction is an actual quotient of the compatible space by the realized subspace. A compatible record has zero class precisely when an original global source realizes it. The theorem constructs the quotient class and proves this equivalence; it does not declare every quotient zero.

A presentation map needs commuting realization and compatibility squares. It induces a map on obstruction spaces. Forward preservation of a zero obstruction follows. Reflection requires injectivity of the induced obstruction map, and cannot be inferred from forward transport alone.

$$R = \operatorname{Ran} A \subseteq K = \ker d, \qquad \mathcal{H} = K/R,$$
$$[v] = 0 \text{ in } \mathcal{H} \iff v \in \operatorname{Ran} A \quad (v \in K).$$

Traditional theorem groups: SR-03

Source anchors: `THEOREM_GUIDE.md;Globalization.lean:obstructionClass_zero_iff;Globalization.lean:all_obstructions_zero_iff_exact;ObstructionTransport.lean:reflects_zero_obstruction`

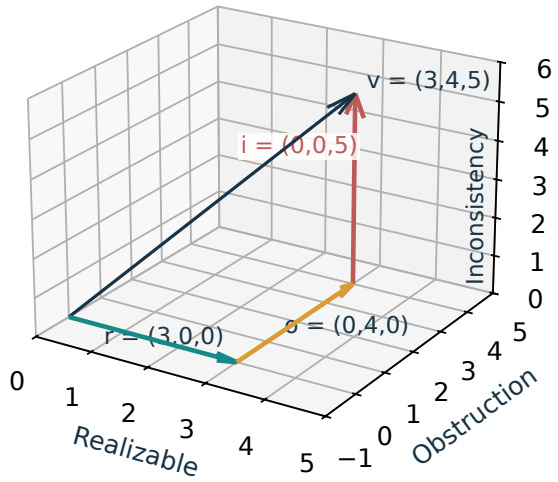
21 The geometry separates three kinds of content

Let $R \subseteq K$ admit orthogonal projections in a real inner-product space. The three components below are pairwise orthogonal: realizable content, compatible obstruction, and inconsistency. Closed subspaces of a Hilbert space and finite-dimensional subspaces provide the standard projection hypotheses.

For our rational example, $A(g_1, g_2) = (g_1 + g_2, 0, 0)$ and $d(x, y, z) = 2z$. The realized space is the x axis, compatibility is the xy plane, and the remaining axis measures inconsistency.

The datum $(3, 4, 5)$ has obstruction energy 16 and inconsistency energy 25. Its nearest realizable image is $(3, 0, 0)$, with squared error 41. A zero compatibility residual would remove the last term only. It would not force the compatible obstruction to vanish.

$$r = P_R v, \qquad o = (P_K - P_R)v, \qquad i = (I - P_K)v,$$
$$v = r + o + i, \qquad \operatorname{dist}(v, R)^2 = \|o\|^2 + \|i\|^2.$$



The three orthogonal contributions for $v = (3, 4, 5)$ preserve the exact error split $41 = 16 + 25$.

Traditional theorem groups: SR-25, SR-26

Source anchors: `BLOCK_07_THEOREMS.md;OrthogonalResolution.lean:components_sum;OrthogonalResolution.lean:infDist_sq_decomposition;HilbertGlobalization.lean:obstructionEquivHarmonic`

22 Residual size needs a conditioning constant

A compatibility residual measures inconsistency after the compatibility operator acts. To convert it into geometric distance, one needs a positive off-kernel coercivity constant for that operator. The obstruction term remains separate.

In the worked presentation $d(x, y, z) = 2z$, the constant $\sigma = 2$ gives equality in the conditioned squared-distance estimate. Finite dimension guarantees existence of some positive constant, but does not compute an optimal singular value or a family-wide lower bound.

The scalar family $d_s(x) = sx$ with $s > 0$ can have an arbitrarily small residual at a unit-distance point. The needed inverse sensitivity grows as $1/s$. Therefore a residual-only certificate cannot be uniform across changing poorly conditioned presentations. Source conditioning of A is a different obligation.

$$\sigma\|x\| \le \|dx\| \quad (x \in K^{\perp}), \quad \sigma > 0,$$

$$\operatorname{dist}(v, R)^2 \le \|o(v)\|^2 + \left(\frac{\|dv\|}{\sigma}\right)^2,$$

$$41 = 16 + \frac{100}{4}.$$

Traditional theorem groups: SR-27, SR-28

Source anchors: `BLOCK_07_THEOREMS.md;HilbertConditioning.lean:conditioned_distance_sq;HilbertConditioning.lean:exists_coercivity_finite;GlobalizationExamples.lean:no_uniform_conditioning_over_scalar_family`

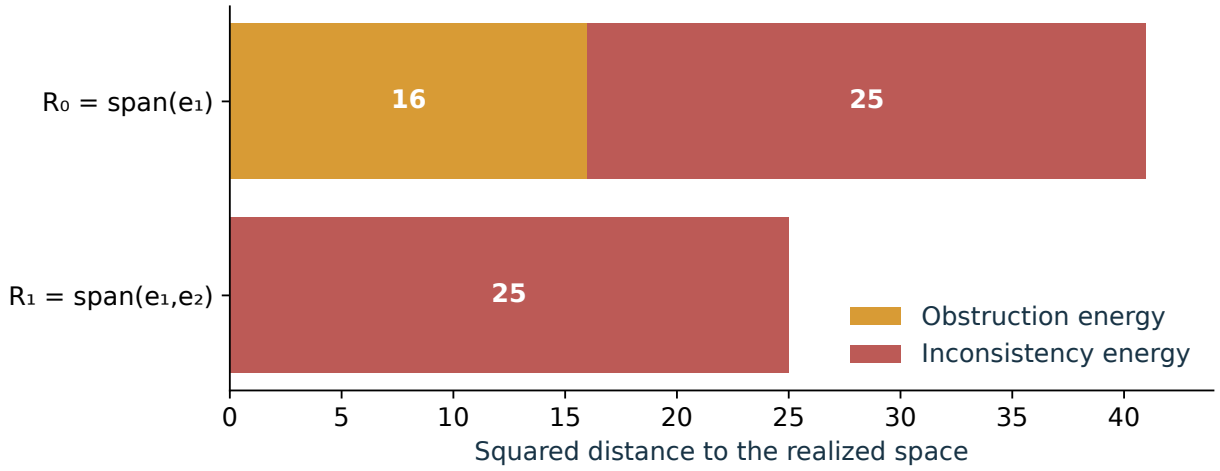
23 Enlarging realizations transfers obstruction energy

Hold the compatibility space and ambient metric fixed while enlarging $R_0 \subseteq R_1 \subseteq K$. The newly realizable component is the difference of the two orthogonal projections. The squared reduction in distance is exactly its energy.

In the worked example, enlarging the realized space from the x axis to the entire compatible plane makes $(0, 4, 0)$ realizable. The obstruction energy 16 disappears, while the inconsistency energy 25 remains. The improvement has a precise source.

The later quotient refinement map is a surjective contraction from K/R_0 to K/R_1. Its norm loss is exactly the newly realizable energy of the original minimum representative. This is a theorem for a fixed metric and compatibility space; an arbitrary change of packetization, tests, or metric has its own obligations.

$$\operatorname{gain}(v) = (P_{R_1} - P_{R_0})v,$$
$$\operatorname{dist}(v, R_0)^2 = \operatorname{dist}(v, R_1)^2 + \|\operatorname{gain}(v)\|^2.$$



Enlargement converts obstruction into realizable content without altering the inconsistency term.

Traditional theorem groups: SR-28, SR-36

Source anchors: `BLOCK_07_THEOREMS.md;GlobalizationRefinement.lean:distance_refinement_sq;MetricRefinement.lean:refinementMap_norm_le;MetricRefinement.lean:quotient_refinement_energy`

24 The best image can have many source preimages

For a bounded real map between Hilbert spaces with actual closed range, restriction of A to $\ker(A)^{\perp}$ is bijective onto its range. A continuous inverse and the range projection construct the Moore–Penrose recovery operator B. Its existence is proved from the stated range hypothesis.

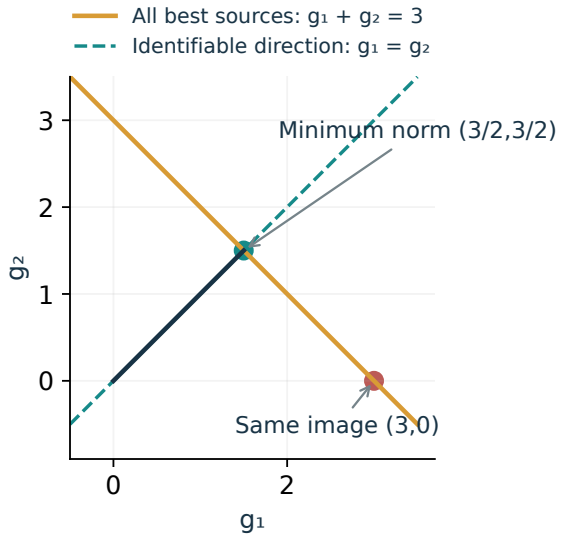
The selected source is the unique minimum-norm least-squares source. Every other least-squares source differs by a kernel vector. These distinctions concern the global source, whereas the earlier best-image theorem concerned only the realized datum.

For $A(g_1, g_2) = (g_1 + g_2, 0, 0)$ and $v = (3, 4, 5)$, the selected source is $(3/2, 3/2)$ with squared norm $9/2$. The source $(3, 0)$ has the same best image and squared norm 9. Its extra energy lies in the true kernel.

$$AB = P_{\operatorname{Ran} A}, \qquad BA = P_{\ker(A)^{\perp}},$$

$$g \text{ is least squares} \iff g = Bv + k \text{ for some } k \in \ker A,$$

$$\|Bv + k\|^2 = \|Bv\|^2 + \|k\|^2.$$



The affine family of best sources intersects the identifiable direction at the unique minimum-norm source.

Traditional theorem groups: SR-29, SR-30

Source anchors: `BLOCK_08_THEOREMS.md;MinimumNormReconstruction.lean:image_comp_recover;MinimumNormOptimality.lean:all_least_squares_sources;MinimumNormOptimality.lean:minimum_norm_unique`

25 Source ambiguity survives exact and noisy recovery

Even exact data Ag recover only the identifiable part of g. Its kernel component is invisible to the forward map. Noise propagation and source ambiguity appear as orthogonal terms in the exact source-error identity.

A positive source-conditioning constant α bounds A below on $\ker(A)^{\perp}$, giving an inverse norm at most $1/\alpha$. This constant belongs to A. Compatibility conditioning σ controls a different map and cannot supply it.

An arbitrary computed candidate needs both a fit tolerance and a kernel-leakage tolerance to certify proximity to Bv. A small fit residual alone permits a large added kernel vector. For Hilbert spaces, closed range is equivalent to positive off-kernel coercivity and to existence of bounded recovery of the identifiable source.

$$g - B(Ag) = P_{\ker A} g,$$
$$\|B(Ag + e) - g\|^2 = \|Be\|^2 + \|P_{\ker A} g\|^2,$$
$$\|\widetilde{g} - Bv\|^2 \le \left(\frac{\|A\widetilde{g} - P_R v\|}{\alpha} \right)^2 + \|P_{\ker A} \widetilde{g}\|^2.$$

Traditional theorem groups: SR-30, SR-31, SR-32

Source anchors: `BLOCK_08_THEOREMS.md;MinimumNormOptimality.lean:noisy_source_error_sq;SourceConditioning.lean:candidate_source_bound_sq;ClosedRangeReconstruction.lean:closed_range_iff_coercivity`

26 The obstruction has its native quotient metric

The obstruction quotient uses Mathlib's standard quotient seminorm. Under the projection hypotheses, the existing linear equivalence to $K \cap R^{\perp}$ is proved isometric. The harmonic vector is the unique compatible representative of minimum norm.

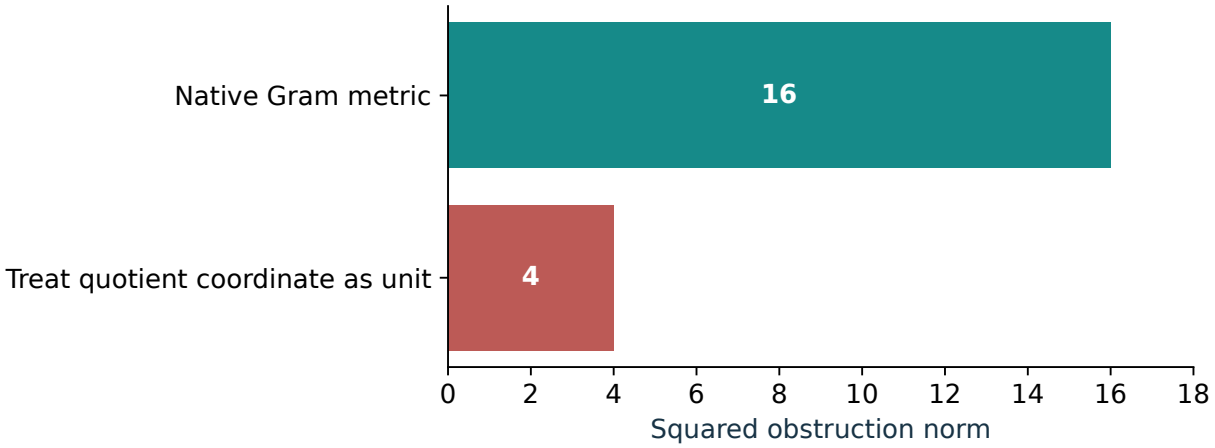
Metrics must travel with coordinate changes. With compatible basis $b_1 = (1,0,0)$ and $b_2 = (1,2,0)$, the Gram matrix below encodes the ambient inner product. The harmonic representative $(0,4,0)$ has basis coefficients $(-2,2)$ and squared Gram norm 16. After quotienting out b_1, treating the remaining quotient coordinate 2 as unit length would instead give the incorrect value 4.

Without closedness, zero quotient seminorm detects membership in the closure of the realized subspace, whereas zero quotient class detects actual membership. A nonclosed quotient can therefore contain a nonzero class of zero seminorm. The guide keeps these statements distinct.

$$\Gamma = \begin{pmatrix} 1 & 1 \\ 1 & 5 \end{pmatrix}, \qquad (-2,2)\Gamma(-2,2)^{\mathsf{T}} = 16,$$

$$\|[v]\| = \operatorname{dist}(v,R) \quad (v \in K),$$

$$\|[v]\| = 0 \iff v \in \overline{R}, \qquad [v] = 0 \iff v \in R.$$



The native quotient metric retains the actual Gram geometry; a unit-coordinate shortcut changes the answer.

Traditional theorem groups: SR-33, SR-34, SR-36

Source anchors: `BLOCK_09_THEOREMS.md;MetricObstruction.lean:obstructionIsometry;MetricRepresentatives.lean:unique_minimum_representative;QuotientMetricBoundaries.lean:zero_norm_iff_closure`

27 Intrinsic classes also handle inconsistent data

An arbitrary datum first projects to compatibility and then acquires an actual obstruction class. This produces a bounded map $c(v) = [P_K v]$ with norm at most one. Its harmonic representative is the existing obstruction component.

The distance split is intrinsic: quotient-class norm plus inconsistency norm. Zero projected obstruction by itself does not establish realizability. The compatibility residual must vanish as well.

A bounded presentation map induces the same algebraic obstruction map as a continuous linear map, with a corresponding norm bound. A reverse presentation map whose quotient composite is the identity supplies a reverse estimate. That map and its identity are explicit hypotheses, rather than consequences of a forward bound.

$$c(v) = [P_K v], \qquad \operatorname{dist}(v, R)^2 = \|c(v)\|^2 + \|i(v)\|^2,$$
$$v \in R \iff c(v) = 0 \text{ and } dv = 0,$$
$$\|\overline{f} z\| \le C\|z\|, \qquad \|z\| \le M\|\overline{f} z\| \text{ when } \overline{g}\,\overline{f} = I.$$

Traditional theorem groups: SR-34, SR-35

Source anchors: `BLOCK_09_THEOREMS.md;IntrinsicGlobalization.lean:intrinsic_distance_sq;IntrinsicGlobalization.lean:realizable_iff_class_and_residual_zero;MetricObstructionTransport.lean:two_sided_bound`

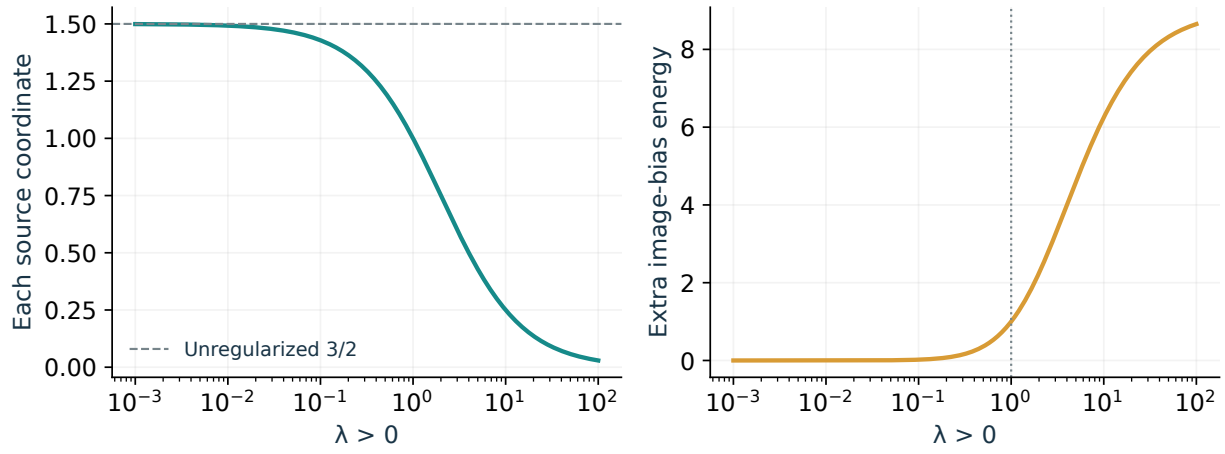
28 Positive regularization constructs a unique source

For a bounded real operator between complete inner-product spaces and $\lambda > 0$, the normal operator is coercive. Lax–Milgram constructs its continuous inverse. The resulting Tikhonov source is an actual map, not a variable assumed to satisfy the desired minimum.

The objective increment formula proves uniqueness of the global minimizer and its equivalence with the normal equation. Neither closed range nor finite dimension is needed. Positive λ is indispensable; the zero-parameter objective can have many minimizers.

In the worked reconstruction chain, the regularized source is $(3/(2+\lambda), 3/(2+\lambda))$. At $\lambda = 1$ it becomes $(1, 1)$. Its fitting error is $42 = 16 + 25 + 1$ and its objective is 44. The extra unit is deliberate image bias, alongside the unchanged obstruction and inconsistency.

$$N_\lambda = A^*A + \lambda I, \qquad B_\lambda = N_\lambda^{-1} A^*,$$
$$F_\lambda(v, x) = \|Ax - v\|^2 + \lambda \|x\|^2,$$
$$F_\lambda(v, x) = F_\lambda(v, x_\lambda) + \|A(x - x_\lambda)\|^2 + \lambda \|x - x_\lambda\|^2.$$



Positive regularization trades source size against additional fitting bias; the irreducible geometric errors remain.

Traditional theorem groups: SR-37, SR-38, SR-40

Source anchors: `BLOCK_10_THEOREMS.md;TikhonovConstruction.lean:normal_resolvent;TikhonovOptimality.lean:objective_increment;TikhonovGlobalization.lean:regularized_intrinsic_error`

29 Bias, noise, and kernel ambiguity remain separate

The exact noisy-error identity contains regularization bias and propagated data error. These two contributions need not be orthogonal. Their sum lies in the identifiable subspace and is orthogonal to the true source's kernel component.

The fixed-parameter noise norm is bounded by $1/(2\sqrt{\lambda})$ and is sharp for scalar examples. The normal-residual correction of an arbitrary computed source is bounded by its residual divided by λ. These are two different scales.

A source-condition rate requires an explicit representation $g = A^*Aw$. Strong convergence later removes that condition by density, but a quantitative uniform rate is not inferred for every identifiable source. Changing the operator at a fixed positive parameter has a proved perturbation estimate; it is not uniform as the parameter tends to zero.

$$B_\lambda(Ag+e)-g = -\lambda N_\lambda^{-1}g + B_\lambda e,$$
$$\|B_\lambda(Ag+e)-g\|^2 = \|-\lambda N_\lambda^{-1}p + B_\lambda e\|^2 + \|k\|^2,$$
$$p = P_{\ker(A)^\perp}g, \quad k = P_{\ker A}g, \qquad \|B_\lambda\| \le \frac{1}{2\sqrt{\lambda}}.$$

Traditional theorem groups: SR-38, SR-39, SR-40

Source anchors: `BLOCK_10_THEOREMS.md;TikhonovBias.lean:noisy_source_error_squared;TikhonovBounds.lean:regularized_opNorm_bound;TikhonovPerturbation.lean:regularized_operator_difference_bound`

30 Strong convergence is different from uniform inversion

For a fixed bounded true operator, positive parameters tending to zero give strong convergence of $B_\lambda A g$ to the identifiable part of each fixed source. The proof uses the established density of $\mathrm{Ran}(A^*A)$ in $\ker(A)^\perp$. No closed-range assumption is needed for this source-data statement.

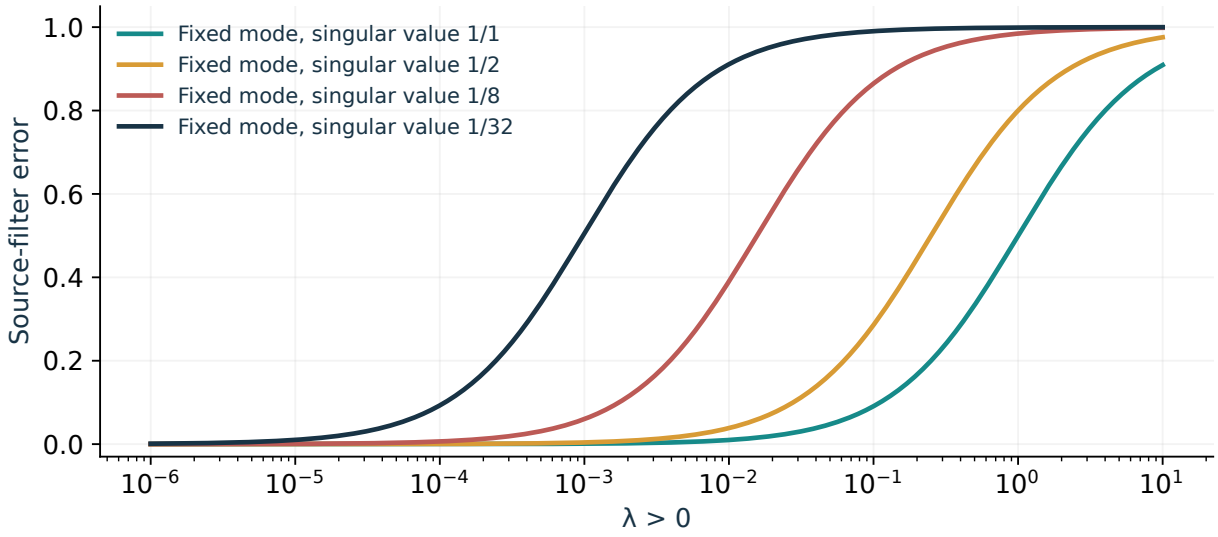
Uniform source-filter convergence holds exactly when the actual range is closed. If the range is nonclosed, the operator norm of the source-filter error equals one at every positive parameter, even while every fixed source converges strongly.

The plot shows several finite diagonal modes. Each fixed mode tends to zero as λ decreases. A finite collection has a positive smallest singular value and converges uniformly. The general nonclosed-range theorem concerns arbitrarily poorly controlled directions; this finite plot illustrates the distinction without testing that theorem.

Noisy source consistency requires $\|e_i\|/\sqrt{\lambda_i} \to 0$ for the stated schedule theorem. An inexact schedule additionally needs its actual normal-residual norm divided by λ_i to vanish. Fitted data need only $e_i \to 0$, and converge to the projection onto the closure of the range. Those weaker data-space hypotheses do not replace the source-space hypotheses.

$$B_\lambda A g \longrightarrow P_{\ker(A)^\perp} g,$$

$$\mathrm{Ran}\, A \text{ nonclosed} \Longrightarrow \|B_\lambda A - P_{\ker(A)^\perp}\| = 1 \quad (\lambda > 0).$$



Modewise filter errors $\lambda/(s^2+\lambda)$ for a finite diagonal family. The plot is illustrative, while the nonclosed-range statement is a separate Lean theorem.

Traditional theorem groups: SR-41, SR-42, SR-43, SR-44

Source anchors: `BLOCK_11_THEOREMS.md;TikhonovConvergence.lean:tendsto_source;UniformRegularization.lean:nonclosed_sourceError_norm;RegularizedDataLimit.lean:fitted_data_tendsto;NoisyRegularization.lean:noisy_inexact_solver_tendsto`

31 Discrepancy selection uses the observed fitting curve

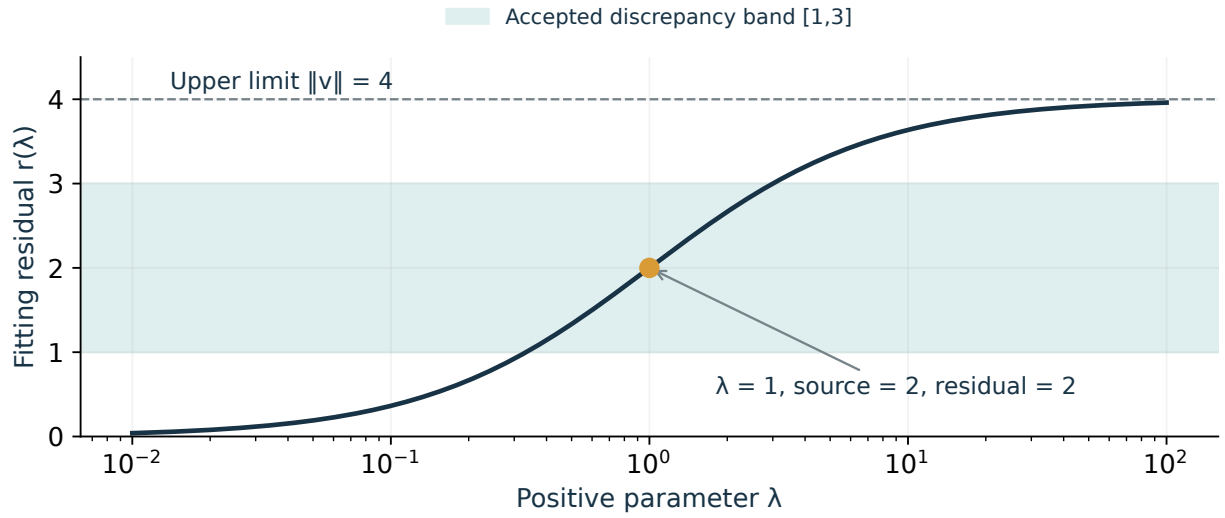
The actual residual curve $r_v(\lambda) = \|AB_\lambda v - v\|$ is continuous on positive intervals and monotone. It is strictly increasing when $A^*v \neq 0$; orthogonal data instead give a constant curve.

A unique positive parameter hits a target strictly between the distance to the actual range and $\|v\|$. The source proves bracketing, uses the intermediate value theorem, and records a noncomputable exact real root. This is distinct from the rational companion's bounded numerical search.

Given a valid positive noise allowance and $\tau > 1$, the small-data branch returns zero when $\|v\| \leq \tau\eta$. Otherwise the selector targets $\tau\eta$. A residual band may allow approximate targeting, but the source must still be the actual Tikhonov solution at the chosen parameter.

$$\operatorname{dist}(v, \operatorname{Ran} A) < \rho < \|v\| \Longrightarrow \exists! \lambda > 0 : \; r_v(\lambda) = \rho,$$

$$A = I, \quad v = 4 : \qquad r_v(\lambda) = \frac{4\lambda}{1+\lambda}.$$



For $\eta = 1$ and $\tau = 3$, $\lambda = 1$ gives source 2 and fitting residual 2 within the accepted band.

Traditional theorem groups: SR-45, SR-47, SR-48

Source anchors: `BLOCK_12_THEOREMS.md;DiscrepancySelection.lean:exists_unique_parameter_of_distance;DiscrepancyRule.lean:selectedSource_accepted;DiscrepancyBoundaries.lean:identity_parameter_formula`

32 Accepted discrepancy gives a direct consistency route

An accepted exact output has no larger norm than the identifiable true source and has a true-data residual bounded by $(\tau + 1)\eta$. These two controls yield strong source consistency as a valid noise allowance tends to zero.

The proof does not first assume selected parameters tend to zero or impose $\eta/\sqrt{\lambda} \to 0$. The acceptance rule supplies a different route to convergence. This does not license arbitrary parameter choices: the actual minimizer and discrepancy conditions are essential.

For an arbitrary true source, the limit is its projection off the true kernel and the error norm tends to the kernel norm. An explicit source condition $g = A^*w$ yields a squared rate. A common rate for every identifiable source, an inferred noise level, and verification of the Python search remain outside the contract.

$$\|x_\eta\| \le \|P_{\ker(A)^\perp} g\|, \qquad \|Ax_\eta - Ag\| \le (\tau + 1)\eta,$$
$$\eta_i \to 0 \Longrightarrow x_i \to P_{\ker(A)^\perp} g,$$
$$\|x_i - g\| \to \|P_{\ker A} g\|.$$

Traditional theorem groups: SR-46, SR-47

Source anchors: `BLOCK_12_THEOREMS.md;DiscrepancyConvergence.lean:norm_control_residual_tendsto;DiscrepancyRule.lean:selectedSource_consistent;DiscrepancyRule.lean:accepted_error_norm_limit`

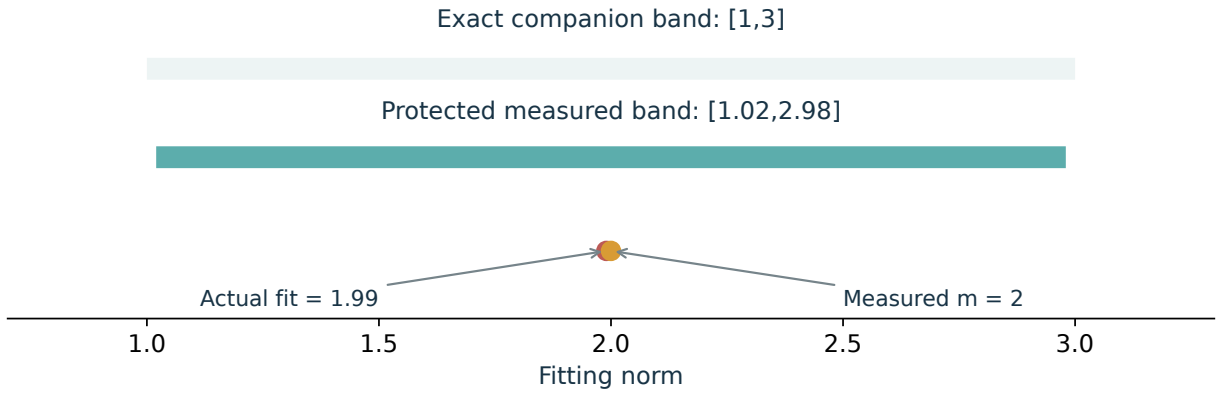
33 Inexact solves require two residual budgets

A computed source need not satisfy its normal equation. Its actual normal residual r controls the source correction at scale $1/\lambda$ and the image correction at the sharper scale $1/(2\sqrt{\lambda})$. A reported fitting norm may also have a justified measurement enclosure.

The protected band below ensures that the exact companion satisfies the original discrepancy band. The scalar margin is nonempty exactly when $2(\omega+\xi) \le (\tau-1)\eta$. This condition concerns available margin; it does not manufacture a candidate meeting all fields.

The worked source $201/100$ is genuinely different from its exact companion 2. Its normal residual is $1/50$, its actual fitting norm is $199/100$, and the reported norm 2 has enclosure $1/100$. The squared-certificate implementation checks actual inequalities without relying on floating square roots.

$$\|x - x_\lambda\| \le \frac{\|r\|}{\lambda}, \qquad \|A(x - x_\lambda)\| \le \frac{\|r\|}{2\sqrt{\lambda}},$$
$$\eta + \omega + \xi \le m \le \tau\eta - \omega - \xi.$$



The measured band reserves both solver and measurement error. Its exact companion is therefore accepted.

Traditional theorem groups: SR-49, SR-50

Source anchors: `BLOCK_13_THEOREMS.md;InexactResidual.lean:correction_bound;InexactResidual.lean:image_resolvent_bound;InexactDiscrepancy.lean:measured_certificate`

34 Vanishing raw residuals can still certify the wrong source

Near acceptance supplies an actual accepted companion within source distance κ. For families with valid noise allowance $\eta_i \to 0$ and correction allowance $\kappa_i \to 0$, the computed sources converge strongly to the identifiable true source. Numerical kernel leakage is at most κ.

The zero-operator counterexample is decisive: let $A = 0$, $\lambda = s > 0$, datum $v = s$, and computed source $x = 1$. Both its normal-residual norm and fitting norm equal s and vanish as $s \to 0$. Its source error remains one. The scaled normal residual $\|r\|/\lambda = 1$ exposes the failure.

A squared positive certificate records positivity, nonnegative error budgets, a normal-residual envelope, a fitting enclosure, and the protected interval. Its soundness theorem supplies the near-accepted witness. An arbitrary solver output without those fields is not covered merely because its raw residual is small.

$$\|x\| \le \|P_{\ker(A)^\perp} g\| + \kappa, \qquad \|P_{\ker A} x\| \le \kappa,$$

$$A = 0,\ \lambda = s,\ v = s,\ x = 1: \qquad \|r\| = s, \quad \frac{\|r\|}{\lambda} = 1.$$

Traditional theorem groups: SR-51, SR-52

Source anchors: `BLOCK_13_THEOREMS.md;InexactConsistency.lean:near_consistent;InexactToleranceCertificates.lean:certificate_sound;InexactExamples.lean:small_raw_residuals_do_not_force_convergence`

35 An uncertain operator inflates the data allowance

Let A be the fixed true operator and C the available model on the same real Hilbert spaces. A true-data allowance η, whole-operator allowance δ, and source norm cap M imply a model-data allowance $\eta + \delta M$. The source proves this transfer; it does not infer the allowances from a fit.

Operational certificates can use C, observed data, allowances, and normal-residual measurements without receiving A. Their consistency conclusion still depends on the external true bounds. For changing models, vanishing data, operator, and source-correction allowances give convergence under the stated certificate and feasibility conditions. Agreement of model and true kernels is not required; the source ambiguity always belongs to the fixed true operator.

If model error persists, a model fit can converge to the wrong source. The checked scalar example $A = 1$, $C = 2$, and true datum/source 1 produces accepted model sources tending to $1/2$. It violates the vanishing operator-error hypothesis rather than contradicting the consistency theorem.

$$\|Ag - v\| \le \eta, \quad \|C - A\| \le \delta, \quad \|g\| \le M$$
$$\Longrightarrow \quad \|Cg - v\| \le \eta + \delta M.$$

Traditional theorem groups: SR-53, SR-55, SR-56

Source anchors: `BLOCK_14_THEOREMS.md;OperatorUncertainty.lean:effective_data_noise;OperatorCertificates.lean:certified_consistent;OperatorUncertaintyExamples.lean:biased_model_wrong_limit;ChangingConsistency.lean:projected_consistent`

36 Model residuals need transport to the true operator

A numerical solve and an imperfect forward model create different errors. For a bounded operator D, define its actual normal residual $r_D = (D^*D + \lambda I)x - D^*v$. The two residuals for A and C have an exact difference identity. Their regularization terms cancel because the same parameter and source are used.

A valid whole-operator allowance $\|C - A\| \le \delta$ converts this identity into the available-operator bound below. It uses only the model fitting norm, model norm, source norm, and the justified allowance on its right-hand side.

If the computed model residual is bounded by ρ, adding ρ to that envelope bounds the true normal residual. A small model solve error cannot erase model uncertainty. The identity and its norm estimate do not require $\lambda > 0$; inverse correction and regularization certificates do.

This is an a posteriori transfer theorem for an actual computed source. It complements inflated-noise discrepancy selection. Its external operator allowance remains essential, and a raw fit on sampled vectors does not establish that whole-map allowance.

$$r_A - r_C = (A^* - C^*)(Cx - v) + A^*(A - C)x,$$
$$\|r_A - r_C\| \le \delta\big(\|Cx - v\| + (\|C\| + \delta)\|x\|\big).$$

Traditional theorem groups: SR-54

Source anchors: `BLOCK_14_THEOREMS.md;OperatorResidualTransport.lean:normal_residual_difference;OperatorResidualTransport.lean:true_normal_residual_bound`

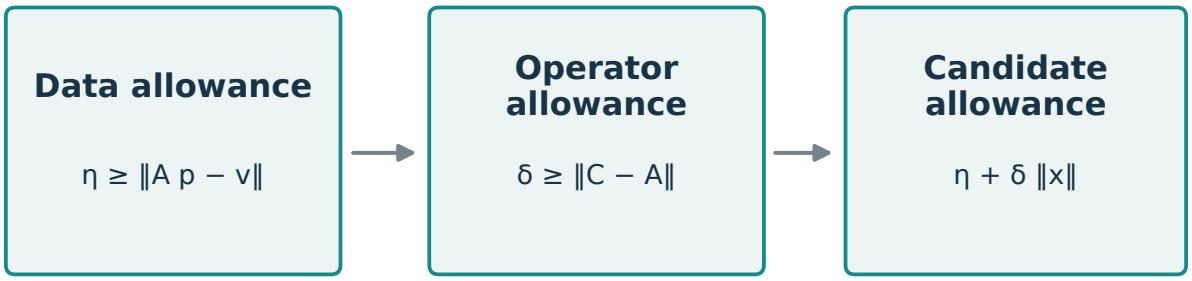
37 Generalized acceptance generates its own source bound

The generalized discrepancy predicate replaces the external source cap by a candidate-dependent allowance. Acceptance requires either its small-data zero branch or an actual positive-parameter Tikhonov source whose model residual lies in the generalized band.

A feasible comparison source p obeys $\|Cp - v\| \leq \eta + \delta\|p\|$. Comparing literal Tikhonov objectives proves that a candidate satisfying the lower generalized discrepancy bound has $\|x\| \leq \|p\|$. Otherwise the comparison source would have no larger fitting term and a strictly smaller positive penalty.

The true bounds then yield a true-operator residual estimate. With a fixed true identifiable source/operator and valid allowances tending to zero, accepted candidates converge strongly. A near-accepted computed source adds its separate source allowance. External truth of data and operator allowances remains essential.

$$\begin{aligned} \mathrm{candidateNoise}(\eta, \delta, x) &= \eta + \delta\|x\|, \\ \eta + \delta\|x\| \leq \|Cx - v\| &\leq \tau(\eta + \delta\|x\|), \\ \|x\| \leq \|p\|, \qquad \|Ax - Ap\| &\leq (\tau + 1)(\eta + \delta\|p\|). \end{aligned}$$



Acceptance uses the allowance. Truth of η and δ requires separate evidence.

The acceptance rule uses a candidate-dependent allowance; it does not validate its own true-data or true-operator assumptions.

Traditional theorem groups: SR-57, SR-58, SR-60

Source anchors: `BLOCK_15_THEOREMS.md;GeneralizedDiscrepancy.lean:regularized_norm_domination;GeneralizedDiscrepancy.lean:accepted_true_residual;GeneralizedInexact.lean:near_consistent`

38 Acceptance and parameter feasibility remain separate

Self-bounding acceptance does not imply an acceptable parameter exists for every available model. A sufficient pointwise condition is the relative model-accuracy inequality below, alongside $\eta > 0$, $\tau > 1$, and the true bounds. It lets the existing discrepancy selector find a root and proves the resulting generalized acceptance.

The exact true model $A = I$ and zero available model $C = 0$ provide a nonfeasibility boundary. For $p = v = 1$, $\eta = 1/10$, $\delta = 1$, and $\tau = 2$, the true allowances are valid. Every positive-parameter model source is zero and has residual one, exceeding its upper candidate allowance $1/5$.

The generic selection uses a noncomputable real root and choice. A finite rational checker can search under explicit caps and verify its output; that algorithm and parser are separate implementation evidence. Recovery does not follow from merely declaring a large operator-error allowance.

$$\delta\|p\| < (\tau - 1)\eta,$$

$$\delta \le \gamma\eta, \quad \gamma\|p\| < \tau - 1 \quad \Longrightarrow \quad \delta\|p\| < (\tau - 1)\eta.$$

$$C = 0: \quad x_\lambda = 0, \quad \|Cx_\lambda - v\| = 1 > \tfrac{1}{5}.$$

Traditional theorem groups: SR-59, SR-60

Source anchors: `BLOCK_15_THEOREMS.md;GeneralizedDiscrepancySelection.lean:exists_accepted;GeneralizedFeasibility.lean:relative_accuracy_of_scaled_error;GeneralizedDiscrepancyExamples.lean:zero_model_not_accepted`

39 Spanning calibration controls the whole operator

Calibration synthesis $U : K \to G$ may be redundant. A bounded analysis $R : G \to K$ with $UR = I$ ensures that the calibration directions span the source in a controlled way. Every bounded source operator factors through this frame.

A measured response operator Y has two different comparisons: justified calibration truth error relative to AU, and measured model fit relative to CU. Their sum controls the response difference, and the reconstruction norm converts it into whole-operator uncertainty.

The calibration-truth allowance is an external justified input. The theorem converts it; it does not estimate it from individual noisy samples or a physical instrument. Exact agreement identifies the whole operator only with the frame condition. The redundant scalar example $U(u,v) = u + v$, $Rx = (x/2, x/2)$ satisfies that condition.

$$UR = I, \qquad T = (TU)R,$$
$$\|Y - AU\| \le \eta_{\mathrm{cal}}, \quad \|Y - CU\| \le \mathrm{fit}$$
$$\Longrightarrow \quad \|C - A\| \le (\eta_{\mathrm{cal}} + \mathrm{fit})\|R\|.$$

Traditional theorem groups: SR-61

Source anchors: `BLOCK_16_THEOREMS.md;CalibrationFrame.lean:factor_through_frame;CalibrationFrame.lean:measured_operator_bound;CalibrationFrame.lean:exact_calibration_identifies`

40 Approximate frames expose conditioning and unseen directions

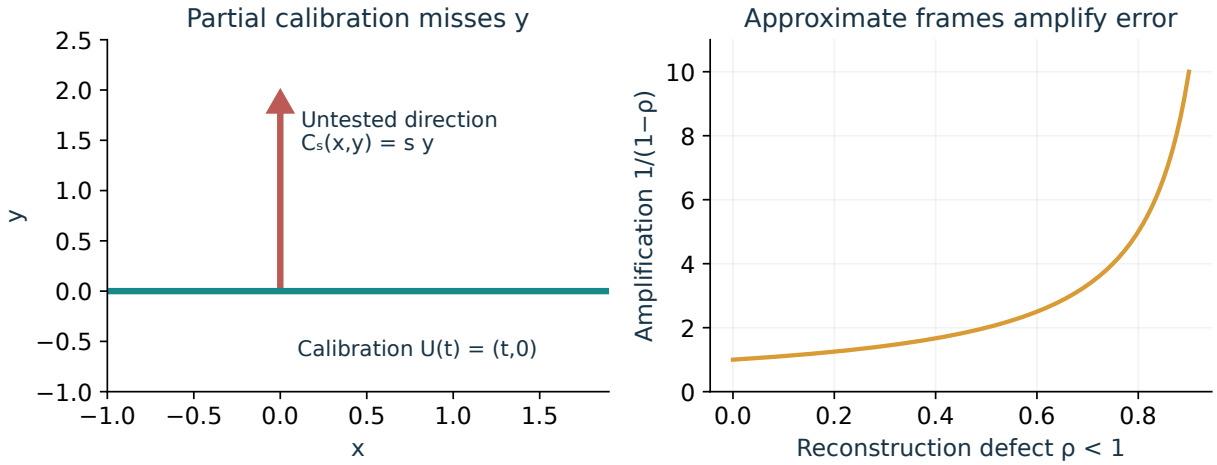
An approximate reconstruction frame satisfies $\|I - UR\| \le \rho < 1$. The exact decomposition of T below permits the residual term to be absorbed, giving a whole-operator bound with amplification $1/(1-\rho)$. At $\rho = 0$ it reduces to the split-frame estimate.

The amplification is meaningful only below one. Partial calibration $U(t) = (t, 0)$ on a two-dimensional source leaves the y direction unseen. Operators $C_s(x, y) = sy$ agree perfectly on every calibration input while their norms are unbounded as s grows.

The two panels show why a good fit needs a spanning/conditioning contract. Hierarchical split frames compose with a reconstruction-norm product bound. This permits genuine multistage calibration, while retaining sensitivity from each stage.

$$T = (TU)R + T(I - UR), \qquad \|T\| \le \frac{\|TU\|\|R\|}{1-\rho},$$

$$\|C - A\| \le \frac{(\eta_{\text{cal}} + \text{fit})\|R\|}{1-\rho}.$$



Nonspanning calibration leaves an uncontrolled direction. Approximate reconstruction frames carry explicit amplification.

Traditional theorem groups: SR-63, SR-64

Source anchors: `BLOCK_16_THEOREMS.md;ApproximateCalibration.lean:approximate_frame_bound;CalibrationExamples.lean:perfect_partial_calibration_unbounded;CalibrationHierarchy.lean:hierarchical_operator_bound`

41 The completed calibration-to-recovery chain

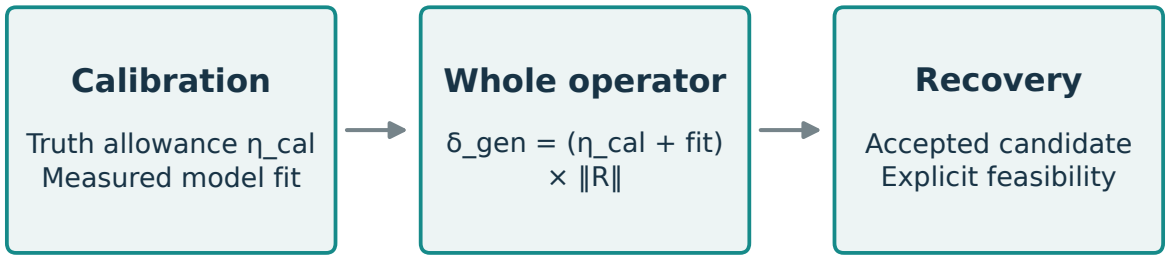
A fixed split frame supplies generated operator allowance $\delta_{\text{gen}} = (\eta_{\text{cal}} + \text{fit})\|R\|$. The frame theorem proves that this allowance bounds the whole model error. Vanishing calibration truth error and measured fit make the generated allowance vanish.

Accepted source-bound-free outputs using that allowance converge strongly to a fixed identifiable true source when the true-data allowance also vanishes. This consistency theorem fixes the real Hilbert spaces, the bounded true operator A, the source $p \in \ker(A)^{\perp}$, and the band constant $\tau \geq 0$. A feasible selector additionally requires $\eta > 0$, $\tau > 1$, and the pointwise relative-accuracy margin. The parameter construction remains noncomputable; it is not a Lean-extracted numerical search.

This closes the established contract from justified operator-level calibration data to certified recovery. Two significant bridges remain open: deriving the calibration-truth allowance from individual noisy samples, and connecting behavioral minimality directly to stable target-only prediction without full-source recovery. Neither bridge is supplied by a small residual alone.

$$\delta_{\text{gen}} = (\eta_{\text{cal}} + \text{fit})\|R\|,$$

$$\delta_{\text{gen}}\|p\| < (\tau - 1)\eta \quad \text{for the stated feasible selection route.}$$



Strong consistency: valid allowances vanish, fixed identifiable true source.

Individual noisy-sample derivation and stable task-only recovery remain open.

The chain separates justified calibration truth, model fit, whole-operator uncertainty, acceptance, and parameter feasibility.

Traditional theorem groups: SR-62

Source anchors: `BLOCK_16_THEOREMS.md;CalibrationRecovery.lean:generated_operator_bound;CalibrationRecovery.lean:calibrated_consistent;CalibrationSelection.lean:calibrated_selected_consistent`

42 Evidence and reusable source

The visual companion contains 18 original vector figures and 280 exact rational assertions about its finite worked examples. Assertion labels and figure hashes are recorded in `validation/visual_example_checks.json`. The complete Markdown preserves every displayed equation, figure caption, and source anchor. Editable TeX and the figure generator are included.

The frozen v0.16.0 source contains 998 theorem/lemma declarations in 93 production modules. These include supporting lemmas, examples, and classical specializations. They are not 998 independent discoveries. The final Lean guide and release records specify which source closure was compiled, which evidence is inherited, and whether any fresh replay succeeded.

The core contracts cover specified linear operations, bounded analytic generators, real Hilbert recovery, and justified whole-operator/noise allowances. Python parsers, elimination routines, certificate searches, and instrumentation are not formally verified by those mathematical declarations.

Packets and transfer blocks provide a working language. The completed sufficient-resolution results explain when a representation supports all declared futures, how failure returns as memory or obstruction, and what quantitative evidence permits recovery. Further nonlinear, positive, integral, arithmetic, and physical realization claims require their own proofs.

Traditional theorem groups: SR-12, SR-58, SR-62

Source anchors: `BLOCK_16_THEOREMS.md;FiniteCertificate.lean:certificate_kernel_eq;GeneralizedDiscrepancy.lean:accepted_consistent;CalibrationRecovery.lean:calibrated_consistent`