OPERATION KOMOREBI

REALIZATION CALCULUS



RESEARCH LIBRARY

Sufficient Resolution

Exact reduction, hidden memory, and certified reconstruction

Traditional paper

Komorebi P.

Standardized publication edition / 8 October 2026

REALIZATION CALCULUS

Reading scope

Reading set. This traditional paper is accompanied by a visual guide, a Lean guide, and the complete source ZIP. The ZIP includes the `.md` files, editable manuscript sources, Lean sources, Python implementations, example data, and verification records.

Scope. Exact linear representation and behavioral reduction; bounded discrete and continuous dynamics; real Hilbert reconstruction and regularization; certified discrepancy rules; and calibration-based operator allowances. Truth allowances and feasibility assumptions are stated explicitly.

Abstract

A representation is sufficient relative to an observation and a family of permitted continuations. This paper develops that principle through exact analysis and synthesis, hidden-return defects, behavioral quotients, finite certificates, quantitative approximation, and certified inverse problems. Split compression fails to preserve composition by an explicit hidden excursion; for finite-dimensional hidden state, finitely many feedback moments decide all-depth closure. Behavioral reduction retains precisely the distinctions that some permitted future observation can detect. In real Hilbert settings, compatibility, realizability, and source identifiability produce separate obstructions and conditioning requirements. Canonical reconstruction, Tikhonov regularization, discrepancy selection, inexact solves, changing operators, and redundant calibration frames are assembled into a common chain of contracts. Classical linear algebra and analysis provide the foundations. The contribution of the assembled calculus is the explicit relation between these contracts, their reusable certificates, and their sharp boundaries. The paper consolidates sixteen development blocks and maps its principal claims to existing Lean declarations. Finite rational implementations are companion evidence and are not asserted to be formally verified.

Contents

1 The problem and its mathematical contracts **3**

1.1 Conventions and domains 3

1.2 How to read the formal evidence 3

2 Exact coordinates, compression, and behavioral state **4**

3 Discrete hidden memory and finite closure **6**

4 Finite behavioral models and minimality **8**

5 Quantitative discrete sufficiency **10**

6 Continuous flow and exact hidden memory **13**

7 Continuous approximate sufficiency **15**

8 Orthogonal globalization and conditioning **17**

9 Canonical minimum-norm reconstruction **20**

10 Intrinsic metric obstruction and transport **22**

11 Tikhonov construction, bias, and perturbation **24**

12 Strong limits and the uniform boundary **27**

13 Discrepancy selection and strong consistency **29**

14 Reliable discrepancy with inexact solves **31**

15 Uncertain operators and changing identifiable models **34**

16 Candidate-dependent discrepancy without a supplied source cap **37**

17 Calibration generates an operator allowance **40**

18 What the assembled calculus establishes **42**

18.1 Exact representation and information loss 43

18.2 Relations to established mathematics 43

18.3 Open bridges and extensions 43

19 Source distribution and verification boundaries **44**

19.1 Rebuilding the reading set 44

A Coverage of the sixteen development blocks **45**

B Principal dependency paths **45**

1 The problem and its mathematical contracts

Sufficiency depends on what a representation must support. A compression that preserves today's observation can lose a distinction needed after a later operation. A coordinate enlargement can be redundant yet exact. A compatible local datum can still fail to have a global source. A small fitting residual can coexist with large source error. These are distinct phenomena with different mathematical witnesses.

The paper follows two connected lines. The first is the reduction of known linear dynamics relative to declared observations. The second is the reconstruction of a source from bounded linear data with explicit error allowances. The first yields an exact semantic state quotient. The second yields metric reconstruction and stability contracts. The present formalization does not yet identify every behaviorally minimal state with a stably recoverable target in an inverse problem. That bridge is an open extension, stated in the final section.

1.1 Conventions and domains

Operator multiplication applies the right factor first: $STx = S(Tx)$. A word $w = [i_1, \dots, i_n]$ is chronological, so $T_w = T_{i_n} \cdots T_{i_1}$; $T_\varnothing = I$. Concatenating u then v gives $T_{u{+}{+}v} = T_v T_u$. Empty sums are zero and empty products are one. All equalities of operators are equalities on the entire stated domain.

The exact redundant-coordinate theory works for modules over a semiring. Subtraction and compression defects use modules over a commutative ring. Dimension and annihilating-polynomial arguments use a field. Norm estimates use bounded linear maps on normed spaces. Continuous flow results use bounded, time-independent real generators on Banach spaces. Reconstruction and regularization use real Hilbert spaces unless a weaker projection hypothesis is printed explicitly. These changes of domain are part of the contracts.

For a subspace L admitting orthogonal projection, P_L denotes that actual projection. On a Hilbert space, closed subspaces admit it. The notation $\operatorname{dist}(v, L) = \inf_{\ell \in L} \|v - \ell\|$ always refers to the actual subspace, not a replacement chosen to make a theorem true. Closure is written $\operatorname{cl} L$.

1.2 How to read the formal evidence

The stable identifiers SR-01 through SR-64 label consolidated theorem groups. They are editorial units, with hypotheses and proofs printed here. A separate CSV/JSON concordance records their exact source declarations and source hashes. Several formal lemmas can support one printed theorem, and one source lemma can serve more than one explanatory purpose. The source inventory contains 998 theorem/lemma declarations in 93 production modules, including helper lemmas and examples. It is not a count of independent new mathematical results.

The arguments below reconstruct the mathematics of the supplied Lean sources. Classical inputs such as Cayley–Hamilton, orthogonal projection, the open mapping theorem, Lax–Milgram, and the intermediate value theorem enter through the source's Mathlib dependencies. This paper does not assert a new historical priority assessment for those inputs. The Lean guide records exactly which compilations and kernel replays are archived or freshly performed during this release. Python certificates are tested implementations of mathematical contracts; their parsers and algorithms are not Lean-verified.

2 Exact coordinates, compression, and behavioral state

The first contract supplies two maps rather than identifying a coordinate space with the object by convention. Let $a_X : X \to V_X$ and $s_X : V_X \to X$ be linear with $s_X a_X = I_X$. The ambient space V_X may have redundant coordinates. Define $E_X = a_X s_X$. For another such resolution of Y, put

$$\mathcal{K}(T) = a_Y T s_X, \qquad \mathcal{D}(K) = s_Y K a_X.$$

We call a proposed coordinate operator compatible when $E_Y K E_X = K$.

Hypotheses. Semiring modules; supplied linear analysis/synthesis pairs with s_X a_X=I and s_Y a_Y=I.

Theorem 2.1 (SR-01 Exact redundant representation)**.** *$E_X^2 = E_X$, and $E_X v = v$ exactly when $v \in \operatorname{Ran} a_X$. Moreover*

$$\mathcal{D}\mathcal{K}(T) = T, \qquad \mathcal{K}\mathcal{D}(K) = E_Y K E_X.$$

Consequently a compatible K is the encoding of a unique map $T : X \to Y$. Encoding preserves composition. Its identity on resolved coordinates is E_X.

Proof. The splitting gives $E_X^2 = a_X(s_X a_X)s_X = a_X s_X$. If $E_X v = v$, then $v = a_X(s_X v)$; conversely $E_X a_X x = a_X x$. Substituting the definitions yields

$$s_Y(a_Y T s_X)a_X = T, \qquad a_Y(s_Y K a_X)s_X = E_Y K E_X.$$

For compatible K, take $T = \mathcal{D}(K)$. If another T' encodes to K, decoding proves $T' = \mathcal{D}(K)$. Finally, an intermediate resolution of Y has $s_Y a_Y = I_Y$, so $\mathcal{K}(S)\mathcal{K}(T) = a_Z S(s_Y a_Y)T s_X = a_Z S T s_X$. Encoding I_X gives $a_X s_X = E_X$, which acts as identity on compatible kernels. No inverse on all of V_X has been assumed.

□

Formal source: `Resolution.projector_idempotent;Resolution.fixed_iff_realizable;Resolution.compatible_iff_existsUnique;Resolution.encode_comp;Resolution.encode_id`.

Example 2.2 (Redundant coordinates)**.** For $X = \mathbb{R}$ and $V_X = \mathbb{R}^2$, take $a(x) = (x, x)$ and $s(u, v) = (u + v)/2$. Then $E(u, v) = ((u + v)/2, (u + v)/2)$. The encoded scalar transformation $x \mapsto cx$ has matrix $\frac{c}{2}\left(\begin{smallmatrix} 1 & 1 \\ 1 & 1 \end{smallmatrix}\right)$. Its redundant coordinate identity is E, while the ambient identity matrix also acts on an unrealizable anti-diagonal sector. The distinction is exact.

Hypotheses. Commutative-ring modules; linear $C : X \to Z$, $R : Z \to X$ with $CR = I_Z$; arbitrary endomorphisms S, T, U of X.

Theorem 2.3 (SR-02 Hidden-return defect and coherence)**.** *Set $P = RC$, $Q = I - P$, $\Phi(T) = CTR$, and $\Delta(S, T) = \Phi(ST) - \Phi(S)\Phi(T)$. Then*

$$\Delta(S, T) = CSQTR,$$

and associativity gives

$$\Delta(S, T)\Phi(U) + \Delta(ST, U) = \Phi(S)\Delta(T, U) + \Delta(S, TU).$$

For two successive split compressions,

$$\Delta_{21}(S, T) = \Phi_2(\Delta_1(S, T)) + \Delta_2(\Phi_1(S), \Phi_1(T)).$$

Proof. Insert $I = P + Q$ between S and T: $CSTR = CS(P+Q)TR = (CSR)(CTR) + CSQTR$. This proves the factorization. Expanding either side of the coherence identity cancels the terms $\Phi(S)\Phi(T)\Phi(U)$ and leaves $\Phi(STU) - \Phi(S)\Phi(TU)$ together with the same compensating term. For the two-stage formula, add and subtract $\Phi_2(\Phi_1(S)\Phi_1(T))$ in the composite defect. Linearity of Φ_2 identifies the first difference with $\Phi_2(\Delta_1(S,T))$; the second is the second-stage defect. Thus a later representation transports the earlier error and can introduce another one. It does not remove the first term by definition.

□

Formal source: `SplitCompression.defect_factorization;SplitCompression.defect_coherence;SplitCompression.defect_then`.

Hypotheses. Commutative-ring modules and a linear presentation $G \xrightarrow{A} V \xrightarrow{\delta} W$ with $\delta A = 0$; a compatible vector $v \in \ker \delta$.

Theorem 2.4 (SR-03 Actual globalization obstruction)**.** *For $\delta A = 0$, define $K = \ker \delta$, $R = \operatorname{Ran} A \subseteq K$, and $\mathcal{H} = K/R$. A compatible v has class zero exactly when $v = Ag$ for some g. All classes vanish exactly when $R = K$. If A is injective, a zero class has a unique source.*

Commuting presentation maps $LA = A'F$ and $\delta' L = M\delta$ induce a map $\overline{L} : \mathcal{H} \to \mathcal{H}'$ by $[v] \mapsto [Lv]$. These maps compose and preserve zero classes. Reflection of zero follows when $\overline{L}$ is injective.

Proof. The chain condition puts the actual range of A inside K. The zero class criterion for a quotient is membership in the submodule divided out, giving $[v] = 0 \iff v \in R \iff \exists g, Ag = v$. Every compatible class is zero exactly when $K \subseteq R$; the reverse inclusion is already known. Injectivity of A gives uniqueness of any preimage.

For transport, if $v \in K$, then $\delta' Lv = M\delta v = 0$. If $v = Ag$, then $Lv = A'Fg$ lies in R'. Hence L descends to the actual quotients. Evaluating on representatives proves the composition law. A linear map sends zero to zero. If its induced quotient map is injective, $\overline{L}[v] = 0 = \overline{L}[0]$ implies $[v] = 0$. The commuting squares alone provide no such injectivity and do not assert that either obstruction quotient vanishes.

□

Formal source: `Globalization.obstructionClass_zero_iff;Globalization.all_obstructions_zero_iff_exact;Globalization.unique_realization_of_zero;Globalization.PresentationMap.obstructionMap_compose;Globalization.PresentationMap.reflects_zero_obstruction`.

Hypotheses. Commutative-ring modules; a linear observation $O : X \to Y$ and an arbitrary family of allowed endomorphisms $T_i : X \to X$; every finite chronological word is included.

Theorem 2.5 (SR-04 The all-future invisible quotient)**.** *Let $N = \bigcap_w \ker(OT_w)$. It is the largest submodule contained in $\ker O$ and invariant under every T_i. The operations and observation descend to X/N, preserving every future observation. In particular,*

$$[x] = [y] \iff OT_w x = OT_w y \quad \textit{for every finite word } w.$$

If $q : X \to Z$ satisfies $qT_i = U_i q$ and $Vq = O$, then $\ker q \subseteq N$.

Proof. An intersection of kernels is a submodule. If $x \in N$, then $OT_w T_i x = OT_{i::w} x = 0$, so $T_i x \in N$. If an invariant submodule $L \subseteq \ker O$ contains x, induction on word length gives $T_w x \in L$; therefore $x \in N$. This establishes maximality.

Invariance defines quotient operations $[x] \mapsto [T_i x]$; containment in $\ker O$ defines $[x] \mapsto Ox$. Induction gives their all-word fidelity. Equality $[x] = [y]$ means $x - y \in N$, which is precisely equality of the displayed observations by linearity. In an exact competing model, induction gives $qT_w = U_w q$ and hence $OT_w = VU_w q$. If $qx = 0$, every future observation of x is zero. Thus a competing model can retain extra distinctions, but cannot erase any distinction visible to the specified future task.

□

Formal source: `Behavioral.invisible_invariant;Behavioral.invariant_submodule_le_invisible;Behavioral.all_future_observations_preserved;Behavioral.quotient_eq_iff_same_future;Behavioral.exact_model_kernel_le`.

3 Discrete hidden memory and finite closure

Let V be visible state and H hidden state. Over a commutative ring, specify four maps and form the actual full dynamics

$$T(y, z) = (Ay + Bz, Gy + Dz).$$

Put $C(y, z) = y$, $Ry = (y, 0)$, $E_n = CT^n R$, and $M_k = BD^k G$. The return M_k contains entry into H, k hidden transitions, and return to V. Its first effect occurs after $k + 2$ full transitions.

Hypotheses. Commutative-ring modules; actual powers of T(y,z)=(Ay+Bz,Gy+Dz); arbitrary initial y_0,z_0.

Theorem 3.1 (SR-05 Exact hidden elimination)**.** *For the trajectory* $(y_n, z_n) = T^n(y_0, z_0)$,

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

and

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

Prepared evolution ($z_0 = 0$) *satisfies* $E_0 = I$, $E_1 = A$, *and*

$$E_{n+1} = AE_n + \sum_{j=0}^{n-1} M_{n-1-j} E_j.$$

Proof. At $n = 0$, the first formula is $z_0 = z_0$ with an empty sum. If it holds at n, the actual recurrence $z_{n+1} = Gy_n + Dz_n$ gives the same formula at $n + 1$ after distributing D and appending the newest forcing term. Substituting this expression into $y_{n+1} = Ay_n + Bz_n$ proves the visible memory equation. For a prepared input y_0, each y_j equals $E_j y_0$. The pointwise memory equation is therefore an operator identity because it holds for every y_0. The equations are derived from powers of T; the trajectory recurrence is not an extra certificate assumption.

□

Formal source: `BlockDynamics.hidden_elimination;BlockDynamics.discrete_memory_equation;BlockDynamics.prepared_recurrence`.

Hypotheses. Same block dynamics; no finite-dimensional or orthogonality assumption.

Theorem 3.2 (SR-06 Prepared closure and the first return)**.** *The following conditions are equivalent:*

$$E_n = A^n \ (\forall n), \qquad M_k = 0 \ (\forall k), \qquad E_{m+n} = E_m E_n \ (\forall m, n).$$

More precisely,

$$E_n = A^n \ (n \le N+1) \iff M_k = 0 \ (k < N).$$

If $M_j = 0$ for $j < k$, then $E_{k+2} - A^{k+2} = M_k$. Closure for every initial hidden state is equivalent to $B = 0$.

Proof. If all memory maps vanish, the recurrence in SR-05 reduces to $E_{n+1} = AE_n$, proving $E_n = A^n$ by induction. Conversely, suppose exact powers hold. Induct on k. In the recurrence for E_{k+2}, every term except $M_k E_0$ involves a previously vanishing memory map. Thus $E_{k+2} = AE_{k+1} + M_k$, forcing $M_k = 0$. The same induction restricted to $k < N$ proves the finite-prefix equivalence and, before setting the last term to zero, the first-return identity.

Exact powers imply the semigroup law. The law conversely gives $E_{n+1} = E_n E_1 = E_n A$, so it implies exact powers. For arbitrary hidden initialization, evaluating closure at $n = 1$, $y_0 = 0$ gives $Bz_0 = 0$ for every z_0, hence $B = 0$. That condition also suffices directly from the visible recurrence. Vanishing $BD^k G$ imposes a weaker prepared-state requirement.

□

Formal source: `BlockDynamics.prepared_powers_iff_memory_zero;BlockDynamics.prepared_semigroup_iff_memory_zero;BlockDynamics.first_return_discrepancy;BlockDynamics.prepared_prefix_iff_memory_prefix;BlockDynamics.all_state_closure_iff_B_zero`.

Hypotheses. Field modules; supplied monic annihilating polynomial p(D)=0, or finite-dimensional hidden H with d=dim H.

Theorem 3.3 (SR-07 Finite all-depth closure certificate)**.** *If a monic polynomial p annihilates D, vanishing of M_k for $k <$ deg p implies vanishing for all k. If* dim $H = d < \infty$*, then*

$$E_n = A^n \ (\forall n) \iff M_k = 0 \ (0 \le k < d) \iff E_n = A^n \ (0 \le n \le d+1).$$

The visible space need not be finite-dimensional. The case $d = 0$ is included.

Proof. Divide X^n by the monic p to obtain a remainder q_n of degree below deg p. Evaluation at D gives $D^n = q_n(D)$ because $p(D) = 0$. Writing $q_n(D)$ as a finite linear combination of the powers D^k below deg p shows $BD^n G = 0$. For finite-dimensional H, Cayley–Hamilton provides the monic characteristic polynomial of D with degree d. SR-06 then yields both displayed equivalences. In dimension zero, all maps through H vanish and the empty list of memory tests is sufficient. If exact closure fails, the contrapositive gives a nonzero memory map with index below d. These tests concern complete linear maps, not sampled state vectors or approximate floating-point agreement.

□

Formal source: `BlockDynamics.memory_zero_of_annihilator;BlockDynamics.prepared_powers_iff_finrank_memory;BlockDynamics.prepared_powers_iff_finite_horizon;BlockDynamics.nonclosure_has_finite_feedback_witness`.

Hypotheses. Commutative-ring split C,R with CR=I; hidden H=ker C; finite-H corollary over a field.

Theorem 3.4 (SR-08 Canonical hidden state of a split compression). *For $C : X \to Z$, $R : Z \to X$, $CR = I$, let $Q = I - RC$ and $H = \ker C$. The maps*

$$x \longmapsto (Cx, Qx), \qquad (y, z) \longmapsto Ry + z$$

are inverse linear maps between X and $Z \times H$. An endomorphism T is conjugate to block dynamics with

$$A = CTR, \quad B = CT\iota_H, \quad G = QTR, \quad D = QT\iota_H.$$

Their prepared maps equal CT^nR. Hence exact compressed powers are equivalent to zero canonical memory. If $\dim\ker C = d$, complete compressed map tests through $d + 1$ decide all depths.

Proof. $CQ = C - CRC = 0$, so Qx belongs to H. Synthesis after analysis gives $RCx + Qx = x$. Conversely, for $z \in \ker C$, one has $C(Ry + z) = y$ and $Q(Ry + z) = z$, using $QR = 0$ and $Qz = z$. Conjugating T by these inverse maps gives the four displayed blocks; in the hidden blocks, Q supplies the required codomain. Induction transports powers under conjugacy. The lift of y is $(y, 0)$ in these coordinates, so the prepared visible map is precisely CT^nR. Apply SR-06 and SR-07. The construction uses the actual discarded subspace and imposes no orthogonality on the chosen splitting.

□

Formal source: `SplitCompression.block_synthesis_analysis;SplitCompression.block_analysis_synthesis;SplitCompression.block_powers_eq;SplitCompression.exact_powers_iff_canonical_memory_zero;SplitCompression.exact_powers_iff_canonical_finite_horizon.`

Example 3.5 (Delayed return and disconnected return). On $\mathbb{Q}^3$, set $T(y, z_1, z_2) = (z_1, z_2, y)$ and retain y. Then $A = 0$, $M_0 = 0$, $M_1 = I$, and $E_1 = E_2 = 0$ while $E_3 = I$. Two shallow exact checks cannot certify all depths without the dimension budget. For another system, let $A = D = I$, $B(z_1, z_2) = z_1$, and $G(y) = (0, y)$. Both couplings are nonzero but every BD^kG vanishes. Prepared closure is exact, while an arbitrary hidden initial first coordinate affects the next visible state.

4 Finite behavioral models and minimality

Now work over a field and, for dimension statements, assume $\dim X < \infty$. The allowed operations need not commute. Define

$$N_r = \bigcap_{|w| \le r} \ker(OT_w), \qquad k = \dim\ker O.$$

A finite horizon means bounded word length. It gives a finite enumeration of constraints only when the operation alphabet is finite.

Hypotheses. Finite-dimensional X over a field; arbitrary alphabet; all bounded-length words including empty word.

Theorem 4.1 (SR-09 Exact horizon and plateau certificate). $N_0 = \ker O$ *and*

$$N_{r+1} = \ker O \cap \bigcap_i T_i^{-1} N_r \subseteq N_r.$$

One plateau $N_{r+1} = N_r$ implies $N_r = N$. In all cases $N_k = N$. Thus equality of future observations through word length k is equivalent to equality for every finite word.

Proof. The empty word enforces $Ox = 0$. A nonempty word $i :: w$ of length at most $r + 1$ enforces $OT_w T_i x = 0$, yielding the recursion. At a plateau, that recursion shows N_r is invariant under every T_i and lies in $\ker O$. By SR-04, $N_r \subseteq N$, while $N \subseteq N_r$ always holds.

Before reaching N, each successive inclusion must be strict, since an equality would already certify the final subspace. Every strict inclusion reduces dimension by at least one. Starting from dimension k allows at most k such reductions; if N_k has dimension zero it is already N, and otherwise an earlier plateau has occurred. Applying the kernel identity to $x - y$ proves the all-future comparison. No word-order commutativity is used.

□

Formal source: `FiniteHorizon.visibleKernel_succ;FiniteHorizon.visibleKernel_eq_invisible_of_plateau;FiniteHorizon.visibleKernel_stabilizes;FiniteHorizon.same_future_iff_bounded`.

Hypotheses. Finite-dimensional X over a field; exact linear competitor q with qT_i=U_iq and Vq=O; finite Z only for final ambient comparison.

Theorem 4.2 (SR-10 Minimum exact state dimension)**.** *The model X/N is observable. Every exact competitor induces a unique surjection* $\operatorname{Ran} q \to X/N$*, taking qx to $[x]$. Therefore*

$$\dim(X/N) \le \dim \operatorname{Ran} q,$$

with equality exactly when $\ker q = N$*. In that case* $\operatorname{Ran} q \cong X/N$*. If Z is finite-dimensional, also* $\dim \operatorname{Ran} q \le \dim Z$*. Surjectivity of q onto Z is not required.*

Proof. If $qx = qy$, SR-04 gives $x - y \in N$, so $qx \mapsto [x]$ is well-defined. It is linear and surjective because every class has a representative x. Its values on the actual range force uniqueness. Rank-nullity gives $\dim \operatorname{Ran} q = \dim X - \dim \ker q$ and $\dim(X/N) = \dim X - \dim N$. Since $\ker q \subseteq N$, this proves the inequality. Equality of dimensions forces equality of the nested finite-dimensional subspaces; conversely their equality makes the comparison injective and hence an equivalence. In the quotient model, an invisible class $[x]$ has every $OT_w x = 0$, so $x \in N$ and $[x] = 0$. This establishes observability. Unused coordinates in Z do not change the reachable-image comparison.

□

Formal source: `BehavioralMinimality.quotient_is_observable;BehavioralMinimality.imageComparison;BehavioralMinimality.finrank_state_le_image;BehavioralMinimality.image_dimension_minimal_iff;BehavioralMinimality.minimalImageEquiv`.

Hypotheses. Finite-dimensional X over a field; finite operation alphabet; k=dim ker O; target Y can be infinite-dimensional.

Theorem 4.3 (SR-11 Realizable finite observation tables)**.** *Let* $H_k x = (OT_w x)_{|w| \le k}$ *and* $Z_k = \operatorname{Ran} H_k$*. Then* $\ker H_k = N$ *and* $X/N \cong Z_k$*. The formulas*

$$U_i(H_k x) = H_k T_i x, \qquad V(H_k x) = Ox$$

define actual linear maps on Z_k and preserve every finite future word. This model has minimum exact dimension.

Proof. The recursively enumerated words are exactly those of length at most k. SR-09 identifies the kernel of their combined observation map with N. The first isomorphism theorem gives $X/N \cong \operatorname{Ran} H_k$.

If $H_k x = H_k y$, then $x - y \in N$; its invariance gives $H_k T_i x = H_k T_i y$, and $N \subseteq \ker O$ gives $Ox = Oy$. Hence the displayed operations and readout are well-defined on realizable tables. Induction proves $U_w H_k x = H_k T_w x$ and $V U_w H_k x = O T_w x$ for words of any length. Apply SR-10 to get minimality. An arbitrary table in the ambient coordinate product need not be realizable and is not part of this constructed state space.

□

Formal source: `FiniteObservableModel.mem_wordsUpTo;FiniteObservableModel.ker_complete_observation;FiniteObservableModel.observableEquiv;FiniteObservableModel.coordinate_all_future_fidelity;FiniteObservableModel.coordinate_dimension_le_exact_model`.

Hypotheses. Field modules, not necessarily finite-dimensional; finite operation alphabet and finite horizon observation map H_r; actual maps q=F H_r, qT_i=U_iq, Vq=O.

Theorem 4.4 (SR-12 Finite minimality certificate)**.** *For a finite operation alphabet, suppose a proposed representation satisfies*

$$q = FH_r, \qquad qT_i = U_i q, \qquad Vq = O.$$

Then $\ker q = N$*, regardless of whether* r *reaches the a priori horizon. If* q *is surjective, it gives an equivalence* $Z \cong X/N$ *preserving the observation and intertwining every operation. Without surjectivity, the equivalence applies to* $\operatorname{Ran} q$*.*

Proof. If $x \in N$, then every coordinate of $H_r x$ is zero, so $qx = FH_r x = 0$. This gives $N \subseteq \ker q$. Exact intertwining and output give the reverse inclusion by SR-04. The equality makes $qx \mapsto [x]$ injective on the range and surjective onto the quotient. Its inverse is the corresponding quotient map induced by q. Evaluating both sides on qx proves that the equivalence intertwines U_i with the quotient operations and preserves the readout. The certificate has two roles: generation by observations prevents retention of invisible distinctions, while exact dynamics and output prevent loss of future-visible distinctions. Finite dimension is needed to guarantee that such a certificate can be found at a bounded depth, not for its soundness.

□

Formal source: `FiniteCertificate.certificate_kernel_eq;FiniteCertificate.certificate_same_future_iff;FiniteCertificate.certificateEquiv;FiniteCertificate.certificateEquiv_dynamics;FiniteCertificate.certificateEquiv_observation`.

Example 4.5 (A noncommuting minimal model)**.** On $\mathbb{Q}^3$, let $O(x,y,h) = x$, $T_A(x,y,h) = (y,0,h)$, and $T_B(x,y,h) = (0,x,h)$. The two operations do not commute: on $(1,0,0)$ their two compositions give $(1,0,0)$ and $(0,0,0)$. The map $q(x,y,h) = (x,y)$ is generated by O and OT_A, intertwines both operations, and factors O. It therefore has a depth-one certificate. The hidden coordinate h is invisible to every future word, and the minimum exact state dimension is two. The current observation alone has dimension one and misses the distinction between $(0,1,0)$ and zero, which one application of T_A reveals.

5 Quantitative discrete sufficiency

Let $q : X \to Z$ be a proposed representation, with true operations T_i, reduced operations U_i, and readouts $O : X \to Y$, $V : Z \to Y$. Define the step and word defects

$$\delta_i = qT_i - U_i q, \qquad E_w = qT_w - U_w q.$$

Algebraic identities require only the stated commutative-ring structure. Quantitative estimates below use continuous linear maps over a nontrivially normed field and their whole-map operator norms.

Hypotheses. Commutative-ring modules; arbitrary chronological words and linear operations/representation/readouts.

Theorem 5.1 (SR-13 Exact word-defect transport)**.** *One has* $E_\varnothing = 0$*,*

$$E_{i::w} = E_w T_i + U_w \delta_i, \qquad E_{u++v} = E_v T_u + U_v E_u.$$

For $w = [i_1, \dots, i_n]$ *this yields*

$$E_w = \sum_{j=1}^{n} U_{[i_{j+1},\dots,i_n]} \delta_{i_j} T_{[i_1,\dots,i_{j-1}]}.$$

The output defect is $(O - Vq)T_w + VE_w$*.*

Proof. The recursive identity follows by adding and subtracting $U_w q T_i$ in $qT_w T_i - U_w U_i q$. For concatenation, add and subtract $U_v q T_u$; chronological order gives $T_{u++v} = T_v T_u$ and the analogous identity for U. Inducting on the word expands one term per letter: each step defect is transported by its true prefix and reduced suffix. The empty word contributes no terms. Finally $OT_w - VU_w q = (O - Vq)T_w + V(qT_w - U_w q)$. Neither output mismatch nor accumulated dynamics mismatch is discarded.

□

Formal source: `ApproximateAlgebra.wordDefect_cons;ApproximateAlgebra.wordDefect_append;ApproximateAlgebra.wordDefect_eq_sum;ApproximateAlgebra.output_defect_decomposition`.

Hypotheses. Normed spaces over nontrivially normed field; bounded maps; nonnegative letter bounds a_i,b_i,epsilon_i and eta$\geq \|$O-Vq$\|$.

Theorem 5.2 (SR-14 Heterogeneous operator budgets)**.** *Suppose* $\|T_i\| \leq a_i$*,* $\|U_i\| \leq b_i$*,* $\|\delta_i\| \leq \epsilon_i$*. Let* A_w, B_w *be the products of the respective letter bounds and define*

$$D_\varnothing = 0, \qquad D_{i::w} = D_w a_i + B_w \epsilon_i.$$

Then $\|T_w\| \leq A_w$*,* $\|U_w\| \leq B_w$*,* $\|E_w\| \leq D_w$*, and*

$$D_{u++v} = D_v A_u + B_v D_u.$$

If $\|O - Vq\| \leq \eta$*, then*

$$\|OT_w - VU_w q\| \leq \eta A_w + \|V\| D_w.$$

For arbitrary initial $x \in X$*,* $z \in Z$*,* $\|qT_w x - U_w z\| \leq D_w \|x\| + B_w \|qx - z\|$*.*

Proof. Induction and submultiplicativity bound the two word operators by the products. Apply the triangle inequality and submultiplicativity to the first identity in SR-13 to obtain exactly the recursion for D_w. Induction on u proves the displayed scalar concatenation identity; it tracks the same algebra as the exact defect identity, with nonnegative factors preventing an inequality reversal. The output bound follows from its exact decomposition in SR-13. For initialization, write $qT_w x - U_w z = E_w x + U_w(qx - z)$ and bound both terms. These are operator budgets; agreement on a finite set of sample vectors does not supply them.

□

Formal source: `ApproximateBounds.norm_wordError_le_budget;ApproximateBounds.budget_append;ApproximateBounds.norm_outputError_le;ApproximateBounds.norm_state_error_le`.

Hypotheses. Uniform nonnegative whole-map bounds a,b,epsilon; all finite words with length n; eta bound on output mismatch.

Theorem 5.3 (SR-15 Two-rate response and equal-rate limit)**.** *For* $H_0(a,b) = 0$ *and* $H_{n+1}(a,b) = aH_n(a,b) + b^n$,

$$H_n(a,b) = \sum_{j=0}^{n-1} a^j b^{n-1-j}, \qquad (a-b)H_n = a^n - b^n.$$

If $a \neq b$, $H_n = (a^n - b^n)/(a-b)$; *for equal rates,* $H_{n+1}(a,a) = (n+1)a^n$. *Uniform step bounds give*

$$\|E_w\| \le \epsilon H_n(a,b), \qquad \|OT_w - VU_w q\| \le \eta a^n + \|V\|\,\epsilon H_n(a,b).$$

Proof. Induction proves the finite-sum formula: multiplying the old sum by a and adjoining b^n generates every term of the new sum. Multiplication by $a-b$ telescopes, leaving $a^n - b^n$. Division is used only when the rates are distinct. If they coincide, all $n+1$ terms of H_{n+1} equal a^n, including the zero-rate boundary with the usual $a^0 = 1$. For uniform bounds, the heterogeneous budget becomes $D_{i::w} = aD_w + b^{|w|}\epsilon$, so induction gives $D_w = \epsilon H_{|w|}(a,b)$. Apply SR-14. The empty word has zero dynamics defect and is handled by $H_0 = 0$, with no negative exponent required.

□

Formal source: `ScalarBudgets.response_eq_sum;ScalarBudgets.response_mul_sub;ScalarBudgets.response_of_ne;ScalarBudgets.response_same_succ;StableApproximation.norm_wordError_uniform;StableApproximation.norm_outputError_uniform`.

Hypotheses. Bounded maps on the normed spaces of SR-14; $\|T_i\| \le 1$ for every true operation; $\|U_i\| \le 1$ in the nonexpansive case or a common $0 \le b < 1$ with $\|U_i\| \le b$ in the contractive case; $\|qT_i - U_i q\| \le \epsilon$ and $\|O - Vq\| \le \eta$ with $\epsilon, \eta \ge 0$.

Theorem 5.4 (SR-16 Nonexpansive and contractive continuations)**.** *If both families are nonexpansive,* $\|E_w\| \le n\epsilon$ *and the output defect is at most* $\eta + \|V\|\,n\epsilon$. *If the reduced family has a common bound* $0 \le b < 1$, *then*

$$\|E_w\| \le \frac{\epsilon}{1-b}, \qquad \|OT_w - VU_w q\| \le \eta + \|V\|\,\frac{\epsilon}{1-b}.$$

For arbitrary reduced initialization,

$$\|OT_w x - VU_w z\| \le \left(\eta + \|V\|\,\frac{\epsilon}{1-b}\right)\|x\| + \|V\|\,b^n\,\|qx - z\|.$$

Proof. For $a = b = 1$, the response sum has n unit terms. For $a = 1$ and $b < 1$, it is $1 + b + \cdots + b^{n-1} \le 1/(1-b)$. SR-15 gives the operator estimates. For initialization, add $VU_w(qx - z)$ to the operator output error applied to x, and use $\|U_w\| \le b^n$. The common bound concerns every allowed operation in the chosen norm. Individual spectral radii do not imply a common product bound for a noncommuting family. The nonexpansive estimate is an upper bound and does not state that every such system accumulates error linearly.

□

Formal source: `StableApproximation.norm_wordError_nonexpansive;StableApproximation.norm_outputError_nonexpansive;StableApproximation.norm_wordError_contractive;StableApproximation.norm_output_initial_contractive`.

6 Continuous flow and exact hidden memory

Let L be a bounded real operator on a Banach space. Its flow is the actual operator exponential e^{tL}. Boundedness allows differentiation of the norm-convergent exponential series and integration of continuous operator curves. Equalities in this section hold for all real t, using oriented interval integrals when $t < 0$.

Hypotheses. Real Banach space; bounded autonomous D; continuous forcing f; differentiable actual z with z′=Dz+f at every real time.

Theorem 6.1 (SR-17 Actual flow and variation of constants). *$\frac{d}{dt}e^{tD} = De^{tD} = e^{tD}D$, $e^{(s+t)D} = e^{sD}e^{tD}$, and e^{-tD} is the inverse flow. Every actual curve satisfying $z' = Dz + f$ with continuous f obeys*

$$z(t) = e^{tD}z(0) + \int_0^t e^{(t-s)D}f(s)\,ds.$$

The initial value uniquely determines the homogeneous solution; two forced solutions with the same forcing and initial value also agree.

Proof. Termwise differentiation of the exponential series gives both derivative orders, since all powers of D commute with D. The exponential addition identity for commuting multiples of one bounded operator gives the flow law and inverse. For fixed t, differentiate $e^{(t-s)D}z(s)$. The product rule produces $-De^{(t-s)D}z(s) + e^{(t-s)D}(Dz(s) + f(s))$; commutation cancels the generator terms. The remaining derivative is the continuous integrand $e^{(t-s)D}f(s)$. The fundamental theorem of calculus on the oriented interval gives the displayed identity. Setting $f = 0$ yields the homogeneous formula. Subtracting two forced solutions reduces uniqueness to that case. The theorem starts from the differential equation, not from an assumed integral representation.

□

Formal source: `ContinuousFlow.hasDerivAt_flow;ContinuousFlow.flow_add;ContinuousFlow.variation_of_constants;ContinuousFlow.homogeneous_unique;ContinuousFlow.forced_unique`.

Hypotheses. Bounded real generator algebras and bounded linear observation maps; norm-convergent exponentials; all real times or all natural powers.

Theorem 6.2 (SR-18 Power identities and exponential identities). *For bounded linear observation maps K, J on bounded generator algebras,*

$$K(e^{tL}) = J(e^{tD})\ (\forall t) \iff K(L^n) = J(D^n)\ (\forall n).$$

In particular, $Be^{tD}G = 0$ for every t exactly when $BD^nG = 0$ for every n. For $CR = I$, $A = CLR$, one has

$$Ce^{tL}R = e^{tA}\ (\forall t) \iff CL^nR = A^n\ (\forall n),$$

and, with $Q = I - RC$,

$$Ce^{(s+t)L}R - (Ce^{sL}R)(Ce^{tL}R) = Ce^{sL}Qe^{tL}R.$$

Proof. The nth derivative at zero of e^{tL} is L^n. Bounded linear maps commute with these derivatives, so equality for all times implies equality of every observed power. Conversely bounded linear maps commute with the sum of the norm-convergent exponential series: $K(e^{tL}) = \sum_n (t^n/n!)K(L^n)$. Equality term by term gives equality of the sums. Sandwiching and compression are bounded linear maps on the corresponding operator spaces, so they give the specializations. For the last identity, use the flow law and insert $I = RC + Q$ between the two flows. The defect factorization itself does not need $CR = I$; that condition makes Q the discarded-state projector and ensures $E(0) = I$.

□

Formal source: `ContinuousFlow.mapped_flow_eq_iff_powers_eq;ContinuousFlow.sandwich_flow_zero_iff;ContinuousFlow.compressed_flow_iff_powers;ContinuousFlow.compressed_flow_defect`.

Hypotheses. Real Banach visible/hidden spaces; bounded autonomous blocks A,B,G,D; actual exponential full trajectory; arbitrary hidden initial state.

Theorem 6.3 (SR-19 Continuous hidden elimination)**.** *For* $L(y,z) = (Ay + Bz, Gy + Dz)$ *and* $(y(t), z(t)) = e^{tL}(y_0, z_0)$,

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

Writing $M(t) = Be^{tD}G$, *the visible dynamics satisfy*

$$y'(t) = Ay(t) + Be^{tD}z_0 + \int_0^t M(t-s)y(s)\,ds.$$

Prepared initialization removes the term involving z_0.

Proof. Projecting the derivative of e^{tL} onto the visible and hidden components gives $y' = Ay + Bz$ and $z' = Gy + Dz$. The visible curve is continuous, so Gy is continuous forcing for the hidden equation. SR-17 yields the hidden elimination formula. Bounded B commutes with the interval integral, giving the memory kernel in the visible equation. Taking $z_0 = 0$ removes only the initialization term. No converse construction of a full trajectory from an arbitrary Volterra solution is asserted here, and no unbounded generator domain is present.

□

Formal source: `ContinuousBlockDynamics.hasDerivAt_visible;ContinuousBlockDynamics.hidden_elimination;ContinuousBlockDynamics.continuous_memory_equation;ContinuousBlockDynamics.prepared_memory_equation`.

Hypotheses. Bounded real Banach block system; finite-dimensional hidden H only for finite certificates, or supplied monic annihilator.

Theorem 6.4 (SR-20 Continuous closure, finite moments, and jets)**.** *Let* $E(t) = Ce^{tL}R$. *The following are equivalent:*

$$E(t) = e^{tA}\ (\forall t), \quad M(t) = 0\ (\forall t), \quad BD^kG = 0\ (\forall k), \quad E(s+t) = E(s)E(t)\ (\forall s,t).$$

If $\dim H = d$, *each is equivalent to* $BD^kG = 0$ *for* $k < d$, *or to* $E^{(n)}(0) = A^n$ *for* $n \le d+1$. *If the first k return moments vanish,*

$$E^{(k+2)}(0) - A^{k+2} = BD^kG.$$

Proof. SR-18 converts prepared flow equality into prepared power equality and memory flow vanishing into memory moment vanishing. SR-06 equates the corresponding algebraic conditions, and SR-07 supplies the finite moment criterion. Exact prepared flow gives its semigroup law. Conversely, if E is a semigroup, differentiating $E(s+t) = E(s)E(t)$ in s at zero gives $E'(t) = AE(t)$ because $E'(0) = CLR = A$ and $E(0) = I$. The bounded-generator uniqueness theorem yields $E(t) = e^{tA}$.

Actual iterated derivatives of the exponential satisfy $E^{(n)}(0) = CL^nR$. Applying the finite algebraic horizon and first-return identity proves the jet statements. The leading Taylor coefficient has the additional denominator $(k+2)!$; the derivative itself does not. The visible Banach space can be infinite-dimensional. The finite tests are whole-operator jet identities, not a finite list of sampled times.

□

Formal source: `ContinuousBlockDynamics.prepared_exact_iff_memory_zero;ContinuousBlockDynamics.prepared_semigroup_iff_memory_zero;ContinuousBlockDynamics.prepared_exact_iff_finite_memory;ContinuousBlockDynamics.prepared_exact_iff_finite_derivatives;ContinuousBlockDynamics.first_return_derivative`.

Example 6.5 (A linear continuous memory kernel)**.** Take $V = \mathbb{R}$, $H = \mathbb{R}^2$, $A = 0$, $B(z_1, z_2) = z_1$, $G(y) = (0, y)$, and $D(z_1, z_2) = (z_2, 0)$. Since $D^2 = 0$, $e^{tD}(z_1, z_2) = (z_1 + tz_2, z_2)$ and $M(t) = tI$. The first two prepared jets agree with the reduced generator, while the third derivative differs by I. Changing D to zero leaves both boundary couplings nonzero but makes the memory kernel identically zero. Arbitrary hidden initialization still affects the visible initial derivative.

7 Continuous approximate sufficiency

For bounded real generators L on X and A on Z, and bounded $q : X \to Z$, write $\delta = qL - Aq$ and $E(t) = qe^{tL} - e^{tA}q$. Algebraic evolution identities hold for all real times. Norm bounds in this section restrict to $t \geq 0$, with growth assumptions on the nonnegative-time flows.

Hypotheses. Real Banach source/reduced spaces; bounded autonomous generators; bounded q and second representation p; all real times.

Theorem 7.1 (SR-21 Duhamel transport and composition)**.** *One has the operator identity*

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

and $\delta = 0$ exactly when $E(t) = 0$ for every t. Also

$$E(s+t) = E(s)e^{tL} + e^{sA}E(t).$$

For a second representation $p : Z \to W$ with generator B,

$$\delta_{pq} = p\delta_q + \delta_p q, \qquad E_{pq}(t) = pE_q(t) + E_p(t)q.$$

Proof. Differentiate $F(s) = e^{(t-s)A}qe^{sL}$ using the product rule. Its derivative is $e^{(t-s)A}(qL - Aq)e^{sL}$. Integrating from zero to t gives $F(t) - F(0) = E(t)$ and proves the formula. Zero generator defect therefore forces zero evolution defect; differentiating E at zero gives the converse. The addition identity follows by adding and subtracting $e^{sA}qe^{tL}$ and using the two flow laws. For a composite representation, add and subtract pAq in its generator defect and $pe^{tA}q$ in its evolution defect. Both stages' contributions remain explicit. Boundedness guarantees the continuous operator integrand.

□

Formal source: `ContinuousDefect.duhamel;ContinuousDefect.generatorDefect_zero_iff;ContinuousDefect.evolutionDefect_add;ContinuousDefect.generatorDefect_composite;ContinuousDefect.evolutionDefect_composite`.

Hypotheses. Real Banach spaces X, Z; bounded autonomous real generators L, A and bounded representation/readouts q, O, V; real rates a, b of either sign; $t \geq 0$ and $M_L, M_A, \epsilon, \eta \geq 0$. For every $s \geq 0$, $\left\|e^{sL}\right\| \leq M_L e^{as}$ and $\left\|e^{sA}\right\| \leq M_A e^{bs}$; also $\|qL - Aq\| \leq \epsilon$ and $\|O - Vq\| \leq \eta$.

Theorem 7.2 (SR-22 Growth envelopes and scalar response)**.** *For $\|\delta\| \leq \epsilon$, put*

$$R(a,b,t) = \int_0^t e^{b(t-s)} e^{as}\, ds = \begin{cases} (e^{at} - e^{bt})/(a-b), & a \neq b, \\ te^{at}, & a = b. \end{cases}$$

Then

$$\|E(t)\| \leq M_L M_A \epsilon R(a,b,t).$$

If $\|O - Vq\| \leq \eta$, the output defect $F(t) = Oe^{tL} - Ve^{tA}q$ obeys

$$\|F(t)\| \leq \eta M_L e^{at} + \|V\|\, M_L M_A \epsilon R(a,b,t).$$

Proof. Apply the integral norm inequality and submultiplicativity to SR-21. For $0 \leq s \leq t$, the integrand norm is at most $M_A e^{b(t-s)} \epsilon M_L e^{as}$. Integration gives the first estimate. Factoring e^{bt} and integrating $e^{(a-b)s}$ gives the distinct-rate formula; for equal rates the integrand is constant e^{at}. The output decomposition is $F(t) = (O - Vq)e^{tL} + VE(t)$, giving the stated bound. The same argument works with continuous nonnegative envelope functions u, v on the interval, replacing the scalar response by $\int_0^t v(t-s)u(s)\, ds$. In particular these assumptions concern actual flow norms, not just eigenvalues of the generators.

□

Formal source: `ExponentialBudgets.response_of_ne;ExponentialBudgets.response_self;ContinuousApproximation.norm_evolution_le_response;ContinuousApproximation.norm_output_le_response`.

Hypotheses. True flow bounded by M_L; reduced flow bounded by M_A exp(-rt), r>0, for t≥0; generator/readout allowances; arbitrary initial reduced z.

Theorem 7.3 (SR-23 Uniform continuous stability and initialization)**.** *If $\left\|e^{tL}\right\| \leq M_L$ and $\left\|e^{tA}\right\| \leq M_A e^{-rt}$ with $r > 0$, then*

$$\|E(t)\| \leq \frac{M_L M_A \epsilon}{r}, \qquad \|F(t)\| \leq \eta M_L + \|V\| \frac{M_L M_A \epsilon}{r}.$$

For arbitrary x, z,

$$\left\|Oe^{tL}x - Ve^{tA}z\right\| \leq \left(\eta M_L + \|V\| \frac{M_L M_A \epsilon}{r}\right) \|x\| + \|V\|\, M_A e^{-rt} \|qx - z\|.$$

Composition through a second representation propagates allowances with its explicit map norm, as in SR-21.

Proof. Use SR-22 with $a = 0$, $b = -r$. Its response equals $(1 - e^{-rt})/r \leq 1/r$ for $t \geq 0$. This gives the evolution and output operator bounds. The initialization error decomposes into $F(t)x$ plus $Ve^{tA}(qx - z)$, which proves the displayed estimate. The first allowance must be multiplied by $\|x\|$ because it is an

operator estimate. For another representation p with its own evolution defect, SR-21 gives $\|E_{pq}(t)\| \leq \|p\| \|E_q(t)\| + \|E_p(t)\| \|q\|$. This retains error introduced at both stages, rather than identifying their budgets with a single unqualified small quantity.

□

Formal source: `ContinuousStability.norm_evolution_uniform;ContinuousStability.norm_output_uniform;ContinuousStability.norm_initialization_uniform;ContinuousStability.norm_composite_evolution`.

Hypotheses. Actual real scalar bounded generators, q=I; t≥0; b≤a for equality and epsilon>0 for counterexample.

Theorem 7.4 (SR-24 Sharp response and small-defect failure)**.** *For scalar generators $a \geq b$,*

$$|e^{at} - e^{bt}| = (a-b)R(a,b,t),$$

so the response estimate is attained. For every $\epsilon > 0$, choosing $L = \epsilon I$, $A = 0$, $q = I$ gives

$$\|\delta\| = \epsilon, \qquad \|E(t)\| = e^{\epsilon t} - 1 \geq \epsilon t.$$

Its evolution defect exceeds every fixed bound at some nonnegative time.

Proof. The scalar exponential is the actual flow by the homogeneous ODE uniqueness theorem. For $a > b$, the response formula immediately gives the equality; for $a = b$, both sides vanish. In the second example, the generator defect is multiplication by ϵ, while the evolution defect is multiplication by $e^{\epsilon t} - 1$. The elementary exponential inequality $e^u \geq 1 + u$ for $u \geq 0$ gives the lower bound. Taking t sufficiently large makes it exceed any prescribed real bound. Thus a small generator defect alone does not supply uniform accuracy. The hypothesis of a bounded true flow in SR-23 fails here.

□

Formal source: `ContinuousApproxExamples.scalar_response_sharp;ContinuousApproxExamples.unstable_generator_exact;ContinuousApproxExamples.unstable_error_exact;ContinuousApproxExamples.unstable_error_unbounded`.

8 Orthogonal globalization and conditioning

Let $R \subseteq K$ be subspaces of a real inner-product space V, both admitting orthogonal projections. Define the actual maps

$$r(v) = P_R v, \quad o(v) = (P_K - P_R)v, \quad i(v) = (I - P_K)v.$$

They distinguish realized data, compatible obstruction, and incompatibility. In a globalization presentation use $R = \operatorname{Ran} A$, $K = \ker \delta$ with $\delta A = 0$.

Hypotheses. Real inner-product V; nested subspaces R≤K both with orthogonal projections; arbitrary v and represented candidate u∈R.

Theorem 8.1 (SR-25 Unique three-component geometry)**.** *The components lie in R, $K \cap R^\perp$, and $K^\perp$, respectively, are pairwise orthogonal, and give the unique decomposition $v = r(v) + o(v) + i(v)$. Furthermore,*

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

For any $u \in R$,

$$\|v - u\|^2 = \operatorname{dist}(v,R)^2 + \|r(v) - u\|^2.$$

Thus $r(v)$ is the unique nearest represented vector.

Proof. Nested projections satisfy $P_R P_K = P_R$. Hence $P_K v - P_R v$ lies in K and is orthogonal to R; the last component is orthogonal to K. Their definitions sum to v and yield pairwise orthogonality. In any competing decomposition, projection onto R selects its first term and projection onto K selects the first two, so all components are forced. Pythagoras gives the norm identities. For $u \in R$, the residual $v - r(v)$ is orthogonal to $r(v) - u$, so the candidate identity follows. It proves optimality and uniqueness, and taking the infimum over actual $u \in R$ proves the distance formula. Each component map is an orthogonal projection onto its displayed sector and consequently is nonexpansive.

□

Formal source: `OrthogonalResolution.components_sum;OrthogonalResolution.unique_components;OrthogonalResolution.norm_sq_components;OrthogonalResolution.candidate_error_sq;OrthogonalResolution.infDist_sq_decomposition`.

Hypotheses. Real linear presentation δA=0; projections on actual R=range A and K=kerδ; source G only needs real module structure.

Theorem 8.2 (SR-26 Harmonic realization of the obstruction quotient)**.** *The same quotient $\mathcal{H} = K/R$ from SR-03 is linearly equivalent to $K \cap R^\perp$ by*

$$[v] \longmapsto v - P_R v \quad (v \in K).$$

It is inverse to the quotient map restricted to the harmonic subspace. For compatible data, the class is zero exactly when $o(v) = 0$; equality of classes is equality of harmonic representatives. The best image $P_R v$ has an actual source under A.

Proof. Adding a vector of R leaves $v - P_R v$ unchanged, so the map descends to K/R. Its values lie in $K \cap R^\perp$ by SR-25. For harmonic h, $P_R h = 0$, so its class maps back to h. For compatible v, the difference $v - (v - P_R v) = P_R v$ lies in R, so the reverse roundtrip returns $[v]$. The inverse laws prove the equivalence and both class criteria. Since $P_R v$ belongs to the actual range, some g satisfies $Ag = P_R v$. This is image existence; uniqueness of g requires the separate injectivity condition in SR-03. No norm isometry is inferred from linear equivalence alone; the metric statement is proved in SR-33.

□

Formal source: `HilbertGlobalization.obstructionEquivHarmonic;HilbertGlobalization.class_zero_iff_component_zero;HilbertGlobalization.same_class_iff_same_harmonic;HilbertGlobalization.exists_best_global_image`.

Hypotheses. Same projection presentation; bounded δ; sigma>0 with sigma$\|$x$\| \le \|\delta$x$\|$ on K$\perp$; arbitrary data.

Theorem 8.3 (SR-27 Residual bounds retain conditioning and obstruction)**.** *If $\sigma \|x\| \le \|\delta x\|$ for $x \in K^\perp$ and $\sigma > 0$, then*

$$\delta i(v) = \delta v, \qquad \|i(v)\| \le \frac{\|\delta v\|}{\sigma},$$

and

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

For compatible data, $\operatorname{dist}(v, R) = \|o(v)\|$.

Proof. The component $P_K v$ is in $\ker \delta$, so $\delta(v - P_K v) = \delta v$. Apply coercivity to $i(v) \in K^\perp$ and divide by the positive σ. SR-25 gives the exact sum of two nonnegative squares, proving both the lower bound and conditioned upper bound. If $\delta v = 0$, then $v \in K$ and $i(v) = 0$, giving equality. In particular, a compatibility residual bounds only inconsistency. A residual-only global bound needs a vanished obstruction or exactness $R = K$. The condition on σ is independent of the size of a particular measured residual and cannot be inferred from that measurement.

□

Formal source: `HilbertConditioning.inconsistency_residual;HilbertConditioning.inconsistency_le_residual;HilbertConditioning.conditioned_distance_sq;HilbertConditioning.obstruction_lower_bound;HilbertConditioning.compatible_distance`.

Hypotheses. Finite-dimensional real V for existence/characterization; fixed metric and K with projected R0≤R1≤K for refinement.

Theorem 8.4 (SR-28 Finite conditioning and realizable enlargement)**.** *In finite dimension some $\sigma > 0$ satisfies SR-27. Under $\delta A = 0$,*

$$\exists c > 0\ \forall v : \operatorname{dist}(v, \operatorname{Ran} A) \le c\,\|\delta v\| \iff \operatorname{Ran} A = \ker \delta.$$

For projected $R_0 \subseteq R_1 \subseteq K$ in a fixed metric, put $g(v) = (P_{R_1} - P_{R_0})v$. Then

$$\|o_0(v)\|^2 = \|o_1(v)\|^2 + \|g(v)\|^2, \quad \operatorname{dist}(v, R_0)^2 = \operatorname{dist}(v, R_1)^2 + \|g(v)\|^2.$$

No uniform conditioning constant exists across all nonzero scalar compatibility maps.

Proof. Restrict δ to $K^\perp$. It is injective: a zero image lies in both K and $K^\perp$. A finite-dimensional injective linear map has a positive lower norm bound, giving σ; the zero domain is harmless and permits any positive choice. If the residual-only inequality holds, every $v \in K$ has zero distance to $\operatorname{Ran} A$. Finite-dimensional ranges are closed, so $v \in \operatorname{Ran} A$. Conversely exactness removes $o(v)$ in SR-27.

Nested orthogonal projections split $o_0(v) = o_1(v) + g(v)$ into orthogonal sectors, and the inconsistency component stays fixed. Pythagoras gives both refinement identities. For the scalar family $\delta_s(x) = sx$, $s > 0$, take $R = 0$ and $v = 1$. The reconstruction distance is one and the residual is s. For any proposed common c, choosing s sufficiently small violates $1 \le cs$. The theorem guarantees a constant per presentation, not a common constant across degenerating presentations.

□

Formal source: `HilbertConditioning.exists_coercivity_finite;FiniteGlobalization.finite_residual_bound_iff_exact;GlobalizationRefinement.obstruction_refinement_sq;GlobalizationRefinement.distance_refinement_sq;GlobalizationExamples.no_uniform_conditioning_over_scalar_family`.

Example 8.5 (A complete error split)**.** For $A(g) = (g, 0, 0)$, $\delta(x, y, z) = 2z$, and $v = (3, 4, 5)$, the components are $(3, 0, 0)$, $(0, 4, 0)$, and $(0, 0, 5)$. The squared reconstruction distance is $41 = 16 + 25$. The compatibility residual has squared norm 100, and $\sigma = 2$ is sharp on $K^\perp$. Thus the conditioned upper bound is attained: $16 + 100/4 = 41$.

9 Canonical minimum-norm reconstruction

Let $A : G \to V$ be bounded between real Hilbert spaces, with actual closed range $R = \operatorname{Ran} A$. Write $N_A = \ker A$ and $S = N_A^\perp$. Closedness is a substantive hypothesis for a bounded exact inverse. It will be removed for regularized recovery in the following sections.

Hypotheses. Bounded real Hilbert operator A with actual closed range; core needs complete source/range and orthogonal projections.

Theorem 9.1 (SR-29 Constructed bounded canonical inverse)**.** *Restriction $A_S : S \to R$ is a bounded linear bijection with bounded inverse. Define*

$$B = \iota_S A_S^{-1} P_R.$$

Then $AB = P_R$, $BA = P_S$, $ABA = A$, and $BAB = B$. Its two projections are symmetric, $\ker B = R^\perp$*, and* $\operatorname{Ran} B = S$*. The best-image identity $AB = P_R$ and the condition* $\operatorname{Ran} B \subseteq S$ *uniquely determine B.*

Proof. The restriction is injective because $S \cap \ker A = 0$. Every $Ag \in R$ equals $AP_S g$, so it is surjective. Closed R and closed S are Banach spaces; the bounded inverse theorem supplies A_S^{-1} as an actual continuous linear inverse. The definition then gives $AB = P_R$. For g, BAg is the unique element of S with image Ag, namely $P_S g$. The two inverse identities imply the generalized inverse laws. Orthogonal projections are symmetric, proving the symmetry conditions. The definition and injectivity of A_S^{-1} identify the kernel and range. Finally, if B' has the same best image and values in S, then $B'v - Bv \in S \cap \ker A$ for each v, hence it is zero. This construction is the classical Moore–Penrose inverse with the actual range and source ambiguity retained.

□

Formal source: `MinimumNormReconstruction.restrictionEquiv;MinimumNormReconstruction.image_comp_recover;MinimumNormReconstruction.recover_comp_image;MinimumNormReconstruction.recover_unique;ClosedRangeReconstruction.closedRecover`.

Hypotheses. Same real Hilbert closed-range constructor; arbitrary data v, candidate g and arbitrary deterministic noise e.

Theorem 9.2 (SR-30 All least-squares sources and unique minimum norm)**.** *For every source g,*

$$\|v - Ag\|^2 = \operatorname{dist}(v, R)^2 + \|A(g - Bv)\|^2 .$$

All least-squares minimizers are $Bv + k$ with $k \in \ker A$. Among them Bv is the unique minimum-norm source. Exact source recovery $BAg = g$ holds exactly for $g \in S$. For data $Ag + e$,

$$\|B(Ag + e) - g\|^2 = \|Be\|^2 + \|P_{N_A} g\|^2 .$$

Proof. $ABv = P_R v$, so SR-25's candidate identity with $u = Ag$ proves the first formula. Equality to the minimum image error is equivalent to $A(g - Bv) = 0$, yielding exactly the affine family of minimizers. Since $Bv \in S$ and $k \in N_A$, Pythagoras gives $\|Bv + k\|^2 = \|Bv\|^2 + \|k\|^2$. It proves minimum norm and uniqueness. The identity $BA = P_S$ proves the exact recovery criterion. For noisy data, linearity gives $B(Ag + e) - g = Be - P_{N_A} g$. The first term belongs to S and the second to N_A, so they are orthogonal. This last identity does not assume a statistical noise model. It shows explicitly which source content the operator cannot identify.

□

Formal source: `MinimumNormOptimality.candidate_error_decomposition;MinimumNormOptimality.all_least_squares_sources;MinimumNormOptimality.unique_minimum_norm_least_squares;MinimumNormOptimality.exact_source_recovery_iff;MinimumNormOptimality.noisy_source_error_sq`.

Hypotheses. Closed-range A; alpha>0 with alpha‖g‖ ≤ ‖Ag‖ on S; computed source h may contain a kernel component.

Theorem 9.3 (SR-31 Source conditioning and solver ambiguity)**.** *Some $\alpha > 0$ satisfies $\alpha \|g\| \le \|Ag\|$ on S, and any supplied such constant gives $\|B\| \le 1/\alpha$. Therefore*

$$\|B(Ag+e)-g\|^2 \le (\|e\|/\alpha)^2 + \|P_{N_A} g\|^2 .$$

For an arbitrary computed source h,

$$\|h-Bv\|^2 \le \left(\frac{\|Ah-P_R v\|}{\alpha}\right)^2 + \|P_{N_A} h\|^2 .$$

If also $\delta A = 0$, a compatibility constant σ bounds the local error as in SR-27. The constants α and σ have separate roles.

Proof. The bounded inverse on S gives $\|g\| \le \|B\| \, \|Ag\|$. Taking, for example, $\alpha = 1/(\|B\|+1) > 0$ handles also the zero complement and proves existence. For a supplied constant, apply it to $Bv \in S$: $\alpha \|Bv\| \le \|ABv\| = \|P_R v\| \le \|v\|$. Take the operator norm and combine with SR-30 for the noisy bound. For a computed h, split $h - Bv = (P_S h - Bv) + P_{N_A} h$ orthogonally. Coercivity bounds the first term by $\|Ah - P_R v\|/\alpha$ because $AP_S h = Ah$. The compatibility operator bounds distance of local data to realizability; the forward operator bounds source amplification. Neither condition establishes the other. A small fit without control of the computed kernel component does not certify the minimum-norm source.

□

Formal source: `SourceConditioning.exists_source_coercivity;SourceConditioning.recover_opNorm_bound;SourceConditioning.noisy_source_bound_sq;SourceConditioning.candidate_source_bound_sq;ReconstructionGlobalization.two_condition_reconstruction`.

Hypotheses. Bounded A between real Hilbert spaces; actual range closedness; bounded recovery BA=P_S or off-kernel coercivity.

Theorem 9.4 (SR-32 Closed range is equivalent to bounded recovery)**.** *The following conditions are equivalent:*

$$\text{Ran } A \text{ is closed;} \qquad \exists \alpha > 0 : \alpha \|g\| \le \|Ag\| \ \ (g \in S); \qquad \exists B \text{ bounded} : BA = P_S .$$

Thus a nonclosed range admits no bounded exact identifiable-source recovery. Finite-dimensional source or target supplies closed range. There is no common inverse-sensitivity constant across all nonzero real scalar forward maps.

Proof. Closed range gives the bounded constructor and coercivity by SR-29–SR-31. Conversely suppose coercivity holds. If Ag_n converges in V, replace g_n by $P_S g_n$ without changing its image. Coercivity makes this source sequence Cauchy. Completeness gives a limit $g \in S$, and continuity gives the image limit $Ag \in \text{Ran } A$. Hence the range is closed. If bounded B satisfies $BA = P_S$, then for $g \in S$, $\|g\| = \|BAg\| \le \|B\| \, \|Ag\|$. The same positive choice $1/(\|B\|+1)$ gives coercivity. Finite-dimensional

image subspaces are closed in a normed space. Finally $A_c g = cg$, $c > 0$, has canonical inverse $B_c v = v/c$. At $v = c$, data tend to zero with c while the recovered source remains one, ruling out a common constant as $c \downarrow 0$.

□

Formal source: `ClosedRangeReconstruction.closed_range_iff_coercivity;ClosedRangeReconstruction.bounded_recovery_iff_closed_range;ClosedRangeReconstruction.nonclosed_no_bounded_recovery;ReconstructionExamples.no_uniform_source_conditioning`.

Example 9.5 (Source and image are different optima)**.** Let $A(g_1, g_2) = (g_1 + g_2, 0, 0)$ and $v = (3, 4, 5)$. The best image is $(3, 0, 0)$, with the same squared local error 41. The minimum-norm source is $(3/2, 3/2)$ of squared norm $9/2$. The source $(3, 0)$ gives the same best image but has squared norm 9. Their difference $(3/2, -3/2)$ is kernel ambiguity, not local obstruction.

10 Intrinsic metric obstruction and transport

Retain the globalization presentation $\delta A = 0$, the actual projected subspaces $R = \operatorname{Ran} A \subseteq K = \ker \delta$, and the same algebraic quotient $\mathcal{H} = K/R$. Its seminorm is the quotient seminorm induced by the ambient metric. We prove the harmonic equivalence respects that preexisting seminorm, rather than defining a new norm to match its answer.

Hypotheses. Real inner-product V; R=range A,K=kerδ both projected; δ bounded, δA=0; native quotient seminorm.

Theorem 10.1 (SR-33 Native quotient norm and isometry)**.** *The equivalence $K/R \cong K \cap R^{\perp}$ in SR-26 is a linear isometry. For compatible v, w,*

$$\|[v]\| = \|o(v)\| = \operatorname{dist}(v, R), \qquad \operatorname{dist}([v], [w]) = \operatorname{dist}(v - w, R).$$

The quotient norm detects zero classes. Its standard quotient inner product is preserved when the induced projection structure in K is used.

Proof. Let h be the harmonic representative of a class z. For any compatible representative v of z, its harmonic component is h, and the nonexpansive harmonic projection gives $\|h\| \leq \|v\|$. Taking the infimum over representatives proves $\|h\| \leq \|z\|$. Conversely h itself represents z, so the quotient map's norm bound gives $\|z\| \leq \|h\|$. This proves isometry against the existing quotient seminorm. The compatible distance formula follows from SR-25, and distances of classes follow by applying it to $v - w$. Zero norm implies zero harmonic vector and hence zero class by the inverse laws. Orthogonal projection onto R restricts to K, constructing the quotient inner-product structure there. A real linear isometry preserves its inner product by the polarization identity.

□

Formal source: `MetricObstruction.quotientToHarmonic_norm;MetricObstruction.obstructionIsometry;MetricObstruction.norm_class_eq_infDist;MetricObstruction.distance_classes;MetricObstruction.quotient_inner_preserved`.

Hypotheses. Same projected real presentation; arbitrary v for projected class c(v)=[P_Kv]; complete/closed source range only when canonical inverse B used.

Theorem 10.2 (SR-34 Minimum representatives and intrinsic class of data)**.** *Every class z has a unique compatible minimum-norm representative $s(z)$, and every compatible representative v satisfies*

$$\|v\|^2 = \|z\|^2 + \|v - s(z)\|^2 .$$

For arbitrary data define $c(v) = [P_K v]$. This is nonexpansive and

$$\operatorname{dist}(v, R)^2 = \|c(v)\|^2 + \|i(v)\|^2 .$$

The datum is realizable exactly when both $c(v) = 0$ and $\delta v = 0$. When the canonical inverse exists,

$$\|v - Ag\|^2 = \|c(v)\|^2 + \|i(v)\|^2 + \|A(g - Bv)\|^2 .$$

Proof. Take $s(z)$ to be its harmonic representative from SR-33. The difference between any other compatible representative and $s(z)$ lies in R, which is orthogonal to the harmonic sector. Pythagoras gives the first formula, including unique attainment of the minimum. The composition of P_K and the quotient map is bounded by one; its harmonic representative is $o(v)$. SR-25 and SR-33 give the intrinsic distance identity. If both displayed tests vanish, then $v \in K$ and its obstruction class is zero, so SR-03 supplies a source. Realizability implies both tests directly. Finally combine the candidate identity of SR-30 with the intrinsic distance split. The minimum representative of an obstruction class belongs to local data space; it is a different construction from the minimum-norm global source Bv.

□

Formal source: `MetricRepresentatives.unique_minimum_representative;MetricRepresentatives.representative_error_energy;IntrinsicGlobalization.projectedClass_opNorm_le;IntrinsicGlobalization.intrinsic_distance_sq;IntrinsicGlobalization.realizable_iff_class_and_residual_zero;IntrinsicGlobalization.intrinsic_candidate_total_error`.

Hypotheses. Commuting linear presentation maps; whole local-map norm bound C≥0; reverse presentation with bound M>0 and induced left inverse for two-sided estimate.

Theorem 10.3 (SR-35 Bounded quotient transport and reflection)**.** *A presentation map with $\|Lv\| \le C\,\|v\|$ induces the same quotient map as in SR-03, now satisfying*

$$\left\|\overline{L}z\right\| \le C\,\|z\| .$$

If a reverse presentation has bound $M > 0$ and its quotient map is a left inverse of $\overline{L}$, then

$$\frac{\|z\|}{M} \le \left\|\overline{L}z\right\| \le C\,\|z\| .$$

The corresponding class distances satisfy the same bounds.

Proof. For any representative v of z, the quotient norm of $[Lv]$ is at most $\|Lv\| \le C\,\|v\|$. Infimizing over representatives gives the upper bound and supplies a bounded linear map with exactly the algebraic action of SR-03. Alternatively, the minimum representative from SR-34 realizes the infimum and yields it immediately. Apply the reverse bound to $\overline{L}z$: the left-inverse identity gives $\|z\| \le M\left\|\overline{L}z\right\|$. Division by positive M proves the lower bound. Apply the bounds to $z - w$ for distances. Forward boundedness alone does not prove reflection; the reverse identity is genuine additional data.

□

Formal source: `MetricObstructionTransport.obstructionMap_norm_le;MetricObstructionTransport.boundedObstructionMap_opNorm;MetricObstructionTransport.reverse_bound;MetricObstructionTransport.two_sided_bound`.

Hypotheses. Projected nested R0≤R1≤K in fixed real metric for energy; arbitrary submodule R in real normed space for nonclosed seminorm boundary.

Theorem 10.4 (SR-36 Metric refinement and nonclosed boundary)**.** *Range enlargement induces a surjective contraction $F : K/R_0 \to K/R_1$. For the old minimum representative $s_0(z)$,*

$$\|z\|^2 = \|Fz\|^2 + \|(P_{R_1} - P_{R_0})s_0(z)\|^2.$$

Without a closedness or projection hypothesis, a general quotient instead has

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

Thus its seminorm detects zero exactly when R is closed. Nonclosed range can leave a nonzero class of zero seminorm.

Proof. The larger realization submodule contains the smaller, so the identity on K descends to a surjection of quotients. Every old representative is a new representative, proving contraction by the quotient infimum. For the energy identity apply the orthogonal refinement identity in SR-28 to $s_0(z)$, whose old norm is $\|z\|$. The new harmonic component represents Fz and has norm $\|Fz\|$ by SR-33.

For any submodule, the native quotient seminorm equals distance to that submodule. Zero distance is equivalent to membership in its closure by the metric closure criterion; algebraic quotient zero remains actual membership. If R is nonclosed, choose $v \in \operatorname{cl} R \setminus R$. Its class is nonzero with zero seminorm. Conversely closedness makes zero distance imply membership. This is why the projection hypotheses in the preceding metric theory cannot be discarded silently.

□

Formal source: `MetricRefinement.refinementMap_surjective;MetricRefinement.refinementMap_norm_le;MetricRefinement.quotient_refinement_energy;QuotientMetricBoundaries.zero_norm_iff_closure;QuotientMetricBoundaries.nonclosed_nonzero_zero_norm;QuotientMetricBoundaries.norm_detects_zero_iff_closed`.

Example 10.5 (A metric must survive its coordinate change)**.** For $K = \{(x, y, 0)\}$ and $R = \operatorname{span}\{(1, 0, 0)\}$, take the compatible basis $b_1 = (1, 0, 0)$, $b_2 = (1, 2, 0)$. Its Gram matrix is $\left(\begin{smallmatrix} 1 & 1 \\ 1 & 5 \end{smallmatrix}\right)$. The vector $(3, 4, 0)$ has coordinates $(1, 2)$; its harmonic representative $(0, 4, 0)$ has coordinates $(-2, 2)$ and squared norm 16 in that Gram metric. Computing with an untransported coordinate Euclidean metric would change the problem. Redundant or nonorthogonal coordinates do not license such a change.

11 Tikhonov construction, bias, and perturbation

Fix a bounded operator $A : G \to V$ between real Hilbert spaces, without assuming closed range. For $\lambda > 0$, put

$$N_\lambda = A^*A + \lambda I, \qquad R_\lambda = N_\lambda^{-1}, \qquad B_\lambda = R_\lambda A^*.$$

The inverse is constructed from coercivity. It is not supplied as a certificate premise. Here A^* is the actual Hilbert adjoint.

Hypotheses. Real Hilbert G,V; actual bounded A and lambda>0; no finite dimension or closed range.

Theorem 11.1 (SR-37 Constructed normal inverse and energy). *N_λ is a bounded bijection with bounded inverse R_λ. For $x_\lambda = B_\lambda v$,*

$$N_\lambda x_\lambda = A^* v, \qquad A^*(Ax_\lambda - v) + \lambda x_\lambda = 0,$$

and

$$\|Ax_\lambda\|^2 + \lambda \|x_\lambda\|^2 = \langle v, Ax_\lambda \rangle .$$

Proof. The bilinear form $b(x,y) = \langle Ax, Ay \rangle + \lambda \langle x, y \rangle$ is bounded and satisfies $b(x,x) \geq \lambda \|x\|^2$. Lax–Milgram gives a continuous linear equivalence from the source to itself under the Riesz identification. The adjoint identity identifies its map with $A^*A + \lambda I$. The two inverse laws give an actual resolvent. Applying one inverse law to A^*v yields the normal equation. Taking its inner product with x_λ and using $\langle A^*v, x_\lambda \rangle = \langle v, Ax_\lambda \rangle$ gives the energy identity. Strictly positive λ supplies coercivity even if A has a nontrivial kernel or a nonclosed range.

□

Formal source: `TikhonovConstruction.normalForm_coercive;TikhonovConstruction.normalEquiv_apply;TikhonovConstruction.normal_resolvent;TikhonovConstruction.resolvent_normal;TikhonovConstruction.normal_equation;TikhonovConstruction.solution_energy`.

Hypotheses. Same real Hilbert bounded A with lambda>0; arbitrary source candidate x and data v; actual normal residual.

Theorem 11.2 (SR-38 Objective optimality and sharp fixed-parameter bounds). *For $F_\lambda(v,x) = \|Ax - v\|^2 + \lambda \|x\|^2$,*

$$F_\lambda(v,x) = F_\lambda(v,x_\lambda) + \|A(x - x_\lambda)\|^2 + \lambda \|x - x_\lambda\|^2 .$$

Thus x_λ is the unique minimizer and lies in $(\ker A)^\perp$. The whole-map bounds are

$$\|R_\lambda\| \leq 1/\lambda, \qquad \|B_\lambda\| \leq \frac{1}{2\sqrt{\lambda}}, \qquad \|AB_\lambda v\| \leq \|v\| , \quad \|AB_\lambda v - v\| \leq \|v\| .$$

For $r = N_\lambda x - A^* v$, one has $x - x_\lambda = R_\lambda r$ and $\|x - x_\lambda\| \leq \|r\| / \lambda$.

Proof. Expand both squared norms around x_λ. Their cross terms combine to

$$2 \langle A^*(Ax_\lambda - v) + \lambda x_\lambda, x - x_\lambda \rangle ,$$

which vanishes by SR-37. The remaining positive terms prove the increment identity and uniqueness. Taking the normal equation against a kernel vector k yields $\lambda \langle x_\lambda, k \rangle = 0$, proving source orthogonality.

For $y = R_\lambda r$, its energy gives $\lambda \|y\|^2 \leq \langle r, y \rangle \leq \|r\| \|y\|$, hence the resolvent bound. For x_λ, set $u = \|Ax_\lambda\|$. The energy gives $u^2 + \lambda \|x_\lambda\|^2 \leq \|v\| \, u$. Completing the square in u yields $4\lambda \|x_\lambda\|^2 \leq \|v\|^2$ and also $u \leq \|v\|$. Comparing the minimum objective with $F_\lambda(v,0) = \|v\|^2$ bounds the fitting residual. Subtract the exact normal equation from the computed one and apply the resolvent for the correction identity.

□

Formal source: `TikhonovOptimality.objective_increment;TikhonovOptimality.unique_minimizer;TikhonovBounds.resolvent_opNorm_bound;TikhonovBounds.regularized_opNorm_bound;TikhonovBounds.approximate_solver_bound;TikhonovBounds.source_kernel_orthogonal`.

Hypotheses. Same positive Tikhonov construction; arbitrary g,e; source-coercivity alpha>0 only for conditioned bound; g=A*Aw only for source-condition rate.

Theorem 11.3 (SR-39 Bias, noise, and kernel content)**.** *Write $p = P_{(\ker A)^\perp} g$, $k = P_{\ker A} g$. Then*

$$B_\lambda A g + \lambda R_\lambda g = g, \quad B_\lambda (Ag + e) - g = -\lambda R_\lambda g + B_\lambda e,$$

and

$$\|B_\lambda (Ag + e) - g\|^2 = \|-\lambda R_\lambda p + B_\lambda e\|^2 + \|k\|^2 .$$

If A has source coercivity $\alpha > 0$, this is at most

$$\left(\frac{\lambda}{\alpha^2 + \lambda} \|p\| + \frac{\|e\|}{2\sqrt{\lambda}} \right)^2 + \|k\|^2 .$$

For a supplied source condition $g = A^* A w$, $\|B_\lambda A g - g\| \le \lambda \|w\|$ without requiring closed range.

Proof. The inverse identity $R_\lambda (A^* A + \lambda I) = I$ gives the filter decomposition. Linearity adds the propagated noise. On $\ker A$, N_λ acts as λI, so $\lambda R_\lambda k = k$. The resolvent preserves the orthogonal complement: pairing its normal equation with every kernel vector proves this directly. Thus the combined identifiable error is orthogonal to k, giving the exact squared formula. Bias and noise inside that first term need not be orthogonal to each other.

For $y = R_\lambda p \in (\ker A)^\perp$, source coercivity and its energy give $(\alpha^2 + \lambda) \|y\|^2 \le \|p\| \|y\|$. Multiply by λ and combine with the noise bound in SR-38. For $g = A^* A w$, the normal-range source filter is a contraction, so $\|R_\lambda A^* A w\| \le \|w\|$. Multiplication by λ proves the source-condition bias rate. The representation by w is supplied; no controlled w is asserted for every identifiable source.

□

Formal source: `TikhonovBias.filter_decomposition;TikhonovBias.noisy_source_error_squared;TikhonovBias.full_conditioned_error_squared;TikhonovBias.source_condition_bias_bound`.

Hypotheses. Real Hilbert bounded A,C; positive lambda,mu; actual closed range only for B=A-dagger and intrinsic globalization formulas.

Theorem 11.4 (SR-40 Resolvent sensitivity and irreducible local error)**.** *For two operators and parameters,*

$$R_{A,\lambda} - R_{C,\mu} = R_{A,\lambda} (N_{C,\mu} - N_{A,\lambda}) R_{C,\mu}.$$

At fixed positive λ, with $\epsilon = \|A - C\|$,

$$\|B_{A,\lambda} - B_{C,\lambda}\| \le \frac{\epsilon}{\lambda} + \frac{(\|A\| + \|C\|)\epsilon \|C\|}{\lambda^2}.$$

If A has closed range and canonical inverse B, then $B_\lambda v - B v = -\lambda R_\lambda B v$. For $\delta A = 0$,

$$\|v - A B_\lambda v\|^2 = \|c(v)\|^2 + \|i(v)\|^2 + \|A(B_\lambda v - B v)\|^2 .$$

Proof. Multiply the difference of inverses on the left and right by their normal operators, or insert their inverse laws, to obtain the resolvent identity. SR-38 bounds its norm by $\left\| N_{C,\mu} - N_{A,\lambda} \right\| / (\lambda \mu)$. At equal parameters, $C^* C - A^* A = (C^* - A^*) C + A^* (C - A)$, whose norm is at most $(\|A\| + \|C\|)\epsilon$. Splitting $R_A A^* - R_C C^* = R_A (A^* - C^*) + (R_A - R_C) C^*$ proves the stated bound. It permits rank changes and is not uniform as λ vanishes.

Under closed range, A^* ignores the part of data orthogonal to Ran A, so $B_\lambda v = B_\lambda A B v$. Apply the filter identity to Bv for its bias. The intrinsic candidate identity SR-34 with $g = B_\lambda v$ yields the local error split. Regularization introduces image bias while retaining the same obstruction and inconsistency.

□

Formal source: `TikhonovPerturbation.resolvent_difference;TikhonovPerturbation.regularized_operator_difference_bound;TikhonovGlobalization.regularized_pseudoinverse_bias;TikhonovGlobalization.regularized_intrinsic_error`.

Example 11.5 (Exact scalar response and a two-coordinate example)**.** For $A_c x = cx$, $B_\lambda v = cv/(c^2+\lambda)$. At $c > 0$, $\lambda = c^2$, the value is $v/(2c)$, attaining the noise bound. For $A(g_1, g_2) = (g_1 + g_2, 0, 0)$, $v = (3, 4, 5)$, and $\lambda = 1$, the source is $(1, 1)$. Its squared fit is $42 = 16 + 25 + 1$, separating obstruction, inconsistency, and regularization image bias. Its objective is 44, whereas the unregularized source $(3/2, 3/2)$ has penalized objective $91/2$. Positive regularization solves its own stated objective.

12 Strong limits and the uniform boundary

The following limits fix one bounded A between real Hilbert spaces. Statements can be indexed by an arbitrary filter. When uniqueness of limits or a converse is used, the filter is required to be nontrivial. Closed range is unnecessary for strong convergence on each fixed source.

Hypotheses. Fixed bounded real Hilbert A; positive lambda_i→0 along a filter; arbitrary fixed g; nontrivial filter for exact-limit iff identifiability.

Theorem 12.1 (SR-41 Density and strong identifiable-source convergence)**.** *One has* $\operatorname{cl}\operatorname{Ran}(A^*A) = (\ker A)^\perp$*, and for every* g, w*,*

$$\|B_\lambda A g - g\| \le \|g - A^* A w\| + \lambda \|w\| .$$

As $\lambda \downarrow 0$*,*

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

For a nontrivial parameter filter, convergence to g *itself holds exactly when* $g \in (\ker A)^\perp$*.*

Proof. If x is orthogonal to $\operatorname{Ran}(A^*A)$, pairing with A^*Ax gives $\|Ax\|^2 = 0$. Conversely a kernel vector is orthogonal to every A^*Aw by the adjoint identity. Thus $(\operatorname{Ran}(A^*A))^\perp = \ker A$, and taking orthogonal complements gives the closure identity. Decompose the bias into the bias of $g - A^*Aw$ and of A^*Aw. The first is bounded by its input norm because λR_λ is contractive; the second is bounded by $\lambda \|w\|$ from SR-39.

For identifiable p, density lets us choose w making the first term arbitrarily small, independently of λ. Then let λ tend to zero to remove the second. For arbitrary g, $Ag = Ap$ with $p = P_{(\ker A)^\perp} g$, yielding the claimed source limit. Subtracting g and taking norms gives the kernel-error limit. Uniqueness of limits on a nontrivial filter gives the final equivalence.

□

Formal source: `TikhonovDensity.gram_range_closure;TikhonovDensity.source_approximation_bound;TikhonovConvergence.tendsto_source;TikhonovConvergence.exact_limit_iff_identifiable;TikhonovConvergence.tendsto_error_norm`.

Hypotheses. Fixed real Hilbert bounded A; lambda_i>0→0; deterministic noise with ‖e_i‖/sqrt(lambda_i)→0; actual normal residual scale vanishes for computed sources.

Theorem 12.2 (SR-42 Noise schedules and inexact strong limits)**.** *If* $\lambda_i > 0$ *tends to zero and* $\|e_i\| / \sqrt{\lambda_i} \to 0$*, then*

$$B_{\lambda_i}(Ag + e_i) \longrightarrow P_{(\ker A)^\perp} g.$$

The same limit holds for computed x_i when additionally

$$\frac{\|(A^*A + \lambda_i I)x_i - A^*(Ag + e_i)\|}{\lambda_i} \to 0.$$

For example, $\lambda_i = s_i^2$, $s_i > 0 \to 0$, and $\|e_i\| \le s_i^2$ give the exact-solve consistency condition.

Proof. The source difference between noisy and exact-data reconstruction is $B_{\lambda_i} e_i$, whose norm is at most $\|e_i\| / (2\sqrt{\lambda_i})$ by SR-38. Combine its vanishing with the exact-source limit in SR-41. For computed sources, the normal-residual correction in SR-38 bounds their difference from the actual regularized source by precisely the second displayed scale. The triangle inequality then proves the same limit. In the square schedule, the noise ratio is at most s_i, hence vanishes. These schedules are sufficient conditions; they are not asserted necessary for every favorable operator or noise direction.

□

Formal source: `NoisyRegularization.noisy_source_tendsto;NoisyRegularization.noisy_inexact_solver_tendsto;NoisyRegularization.square_schedule_consistent`.

Hypotheses. Fixed real Hilbert bounded A; positive parameter family→0; nontrivial filter for iff; source conditioning alpha only for positive estimate.

Theorem 12.3 (SR-43 Operator-norm convergence exactly at closed range)**.** *Let $E_\lambda = B_\lambda A - P_{(\ker A)^\perp}$. Always $\|E_\lambda\| \le 1$. If $\operatorname{Ran} A$ is nonclosed, then $\|E_\lambda\| = 1$ for every positive parameter. For a positive family tending to zero along a nontrivial filter,*

$$E_\lambda \to 0 \text{ in operator norm} \iff \operatorname{Ran} A \text{ is closed}.$$

Under source coercivity $\alpha > 0$, $\|E_\lambda\| \le \lambda / (\alpha^2 + \lambda)$.

Proof. The projected bias bound from SR-39 gives the universal upper bound one. Suppose any bounded B' satisfies $c = \|B'A - P_S\| < 1$, where $S = (\ker A)^\perp$. For $x \in S$,

$$\|x\| \le \|x - B'Ax\| + \|B'Ax\| \le c\,\|x\| + \|B'\|\,\|Ax\|.$$

Therefore $\alpha = (1 - c)/(\|B'\| + 1) > 0$ supplies off-kernel coercivity, forcing closed range by SR-32. A nonclosed range thus cannot have $\|E_\lambda\| < 1$, proving equality. For closed range, source coercivity exists and the conditioned bias bound yields the displayed quantity tending to zero. Conversely convergence to zero contradicts the constant norm one on any nontrivial filter. Strong convergence in SR-41 remains valid in the nonclosed case because it fixes a source before taking the limit.

□

Formal source: `UniformRegularization.closed_range_of_approximate_recovery;UniformRegularization.nonclosed_sourceError_norm;UniformRegularization.conditioned_sourceError_norm;UniformRegularization.uniform_convergence_iff_closed_range`.

Hypotheses. Fixed real Hilbert bounded A; positive lambda_i→0; data noise→0 for fitted limit; nontrivial filter for impossible source limit.

Theorem 12.4 (SR-44 Fitted-data limit and absent source limit)**.** *For arbitrary data v,*

$$AB_\lambda^A v = B_\lambda^{A^*} A^* v \longrightarrow P_{\operatorname{cl}\operatorname{Ran} A} v.$$

Adding data noise $e_i \to 0$ preserves this fitted-data limit. If $v \in \operatorname{cl}\operatorname{Ran} A \setminus \operatorname{Ran} A$, fitted data converge to v but the corresponding reconstructed source family has no norm limit in G.

Proof. The identity $(AA^* + \lambda I)A = A(A^*A + \lambda I)$ and uniqueness of the positive normal inverse prove the dual formula. Apply SR-41 to the bounded operator A^* with source v. Its identifiable space is $(\ker A^*)^\perp = \operatorname{cl}\operatorname{Ran} A$, yielding the fitted limit. SR-38 shows AB_λ is contractive, so changing data by e_i changes the fitted value by at most $\|e_i\|$. If reconstructed sources converged in norm to some g, continuity of A would make their images converge to Ag. On a nontrivial filter, the already established image limit v would force $Ag = v$, contrary to $v \notin \operatorname{Ran} A$. This proves absence of a source norm limit without adding an unproved assertion that source norms diverge.

□

Formal source: `RegularizedDataLimit.fitted_adjoint_identity;RegularizedDataLimit.fitted_data_tendsto;RegularizedDataLimit.noisy_fitted_data_tendsto;RegularizedDataLimit.unrealizable_no_source_limit`.

13 Discrepancy selection and strong consistency

Still fix a known bounded A between real Hilbert spaces. Define the actual residual curve on positive parameters $r_v(\lambda) = \|AB_\lambda v - v\|$. Its parameter is mathematical data, not an assumed noise estimate. Let η be a justified true-data allowance and $\tau > 1$ a band multiplier. A total rule must include a small-data branch.

Hypotheses. Real Hilbert bounded A; lambda>0; strict floor dist(v,range A)<rho<‖v‖ for existence; actual positive parameter only.

Theorem 13.1 (SR-45 Actual residual curve and unique interior root)**.** *The residual is continuous and nondecreasing on positive parameters; the source norm is nonincreasing. It is strictly increasing when $A^*v \neq 0$, and constant at $\|v\|$ when $A^*v = 0$. For*

$$\operatorname{dist}(v, \operatorname{Ran} A) < \rho < \|v\|,$$

there is exactly one $\lambda > 0$ with $r_v(\lambda) = \rho$. The resulting exact real parameter is a noncomputable choice with a proved specification.

Proof. Compare the minimum objectives at $0 < \lambda < \mu$ using each minimizer as a competitor for the other. Adding the two inequalities gives $(\mu - \lambda)(\|x_\mu\|^2 - \|x_\lambda\|^2) \leq 0$, so source norms are antitone. Substitution back into the smaller-parameter comparison shows residual monotonicity. If the two residuals were equal, both norms and objective values would be equal, and unique minimization would give $x_\lambda = x_\mu$. Subtracting their normal equations gives $(\mu - \lambda)x_\lambda = 0$, hence $x_\lambda = 0$ and $A^*v = 0$. This proves strictness and the constant case.

The resolvent parameter identity gives, for $s, t \geq a > 0$,

$$|r_v(s) - r_v(t)| \leq \frac{\|A\| \, \|A^*v\|}{a^2} |s - t|,$$

proving continuity on positive compact intervals. The strict floor permits a comparator g with $\|Ag - v\| < \rho$; objective minimization at sufficiently small positive λ gives a lower residual bracket. For large λ, $\|B_\lambda v\| \leq \|A^*v\| / \lambda$, so the residual approaches $\|v\| > \rho$ and gives an upper bracket. The intermediate value theorem provides a root; strictness gives uniqueness. The floor condition itself excludes the constant orthogonal-data case.

□

Formal source: `DiscrepancyResidual.source_norm_antitone;DiscrepancyResidual.residual_strictMono;DiscrepancyResidual.residual_constant_of_adjoint_zero;DiscrepancyResidual.residual_difference_bound;DiscrepancySelection.exists_unique_parameter_of_distance`.

Hypotheses. Fixed bounded real Hilbert A; p∈(ker A)⊥; family norm bounds ‖x_i‖ ≤ ‖p‖ and true residual ‖Ax_i-Ap‖ →0.

Theorem 13.2 (SR-46 Norm domination converts residual convergence to source convergence)**.** *If* $\|Ag - v\| \le r_v(\lambda)$*, then* $\|B_\lambda v\| \le \|g\|$*. For any* x, p, w *with* $\|x\| \le \|p\|$*,*

$$\|x - p\|^2 \le 4\,\|p\|\,\|p - A^*w\| + 2\,\|w\|\,\|Ax - Ap\|\,.$$

Consequently, for identifiable fixed p*, the conditions* $\|x_i\| \le \|p\|$ *and* $\|Ax_i - Ap\| \to 0$ *imply* $x_i \to p$ *strongly, without assuming closed range.*

Proof. Objective minimality gives $r_v(\lambda)^2 + \lambda\,\|B_\lambda v\|^2 \le \|Ag - v\|^2 + \lambda\,\|g\|^2$. The residual ordering and positive λ imply norm domination. For the source estimate, expand $\|x - p\|^2$ and use $\|x\| \le \|p\|$ to obtain $\|x - p\|^2 \le 2\,\langle p, p - x\rangle$. Write $p = (p - A^*w) + A^*w$. The first pairing is at most $\|p - A^*w\|\,\|p - x\| \le 2\,\|p\|\,\|p - A^*w\|$; the adjoint identity bounds the second by $\|w\|\,\|Ap - Ax\|$. This yields the displayed estimate. For identifiable p, the density in SR-41 permits w with $p - A^*w$ arbitrarily small (take $w = Au$ from the normal range). Fixing this w, the residual term tends to zero. An epsilon argument then proves norm convergence. No common inverse bound is invoked.

□

Formal source: `DiscrepancyConvergence.norm_domination_of_residual;DiscrepancyConvergence.source_distance_estimate;DiscrepancyConvergence.norm_control_residual_tendsto`.

Hypotheses. Known bounded $A : G \to V$ between real Hilbert spaces and a fixed true source g; a valid pointwise allowance $\|Ag - v\| \le \eta$. Accepted-source consistency uses a fixed $\tau \ge 1$ and genuine bounds $\|Ag - v_i\| \le \eta_i$ with $\eta_i \to 0$. Total selector existence additionally requires $\eta > 0$ and $\tau > 1$.

Theorem 13.3 (SR-47 Accepted bands and the total small-data rule)**.** *Call* x *accepted when either*

$$\|v\| \le \tau\eta, \quad x = 0,$$

or $x = B_\lambda v$ *for some* $\lambda > 0$ *with* $\eta \le r_v(\lambda) \le \tau\eta$*. With a valid noise allowance, every accepted source satisfies*

$$\|x\| \le \left\|P_{(\ker A)^\perp} g\right\|, \qquad \|Ax - Ag\| \le (\tau + 1)\eta.$$

For accepted families with $\eta_i \to 0$*,*

$$x_i \to P_{(\ker A)^\perp} g, \qquad \|x_i - g\| \to \|P_{\ker A} g\|\,.$$

With $\eta > 0$*,* $\tau > 1$*, the total selector chooses zero on small data and otherwise the root* $r_v(\lambda) = \tau\eta$*. If* $g = A^*w$*, accepted sources obey* $\|x - g\|^2 \le 2\,\|w\|\,(\tau + 1)\eta$*.*

Proof. Let $p = P_{(\ker A)^\perp} g$, so $Ap = Ag$. In the positive branch, $\|Ap - v\| \le \eta \le r_v(\lambda)$ gives norm domination by SR-46; the triangle inequality gives the true residual. In the zero branch, norm domination is immediate and $\|Ag\| \le \|Ag - v\| + \|v\| \le (\tau + 1)\eta$. SR-46 now gives strong convergence to p, and continuity of the error norm gives the kernel limit. For large data, feasibility supplies $\operatorname{dist}(v, \operatorname{Ran} A) \le \eta < \tau\eta < \|v\|$, so SR-45 provides the exact root. For $g = A^*w$, the approximation term in SR-46 is zero,

and the residual bound proves the squared source-condition rate. General identifiable sources have no asserted common rate. The rule controls actual Tikhonov sources; numerical solve errors require the additional certificates below.

□

Formal source: `DiscrepancyRule.selectedSource_accepted;DiscrepancyRule.accepted_consistent;DiscrepancyRule.selectedSource_consistent;DiscrepancyRule.accepted_error_norm_limit;DiscrepancyRule.accepted_source_rate`.

Hypotheses. Known bounded real Hilbert A and positive lambda; real scalar identity case 0<rho<v.

Theorem 13.4 (SR-48 Discrepancy endpoints and scalar formulas)**.** *No positive parameter has residual larger than $\|v\|$. When $A^*v \neq 0$, none has residual equal to $\|v\|$. For every positive parameter, $r_v(\lambda) = 0$ exactly when $v = 0$. For the real scalar identity model and $0 < \rho < v$,*

$$\lambda = \frac{\rho}{v - \rho}, \qquad B_\lambda v = v - \rho$$

is the unique root at residual ρ.

Proof. The first assertion is the residual contraction in SR-38. Equality at $\|v\|$ would give $F_\lambda(v, x_\lambda) \leq F_\lambda(v, 0) = \|v\|^2$ while its fit term already equals $\|v\|^2$. Hence $x_\lambda = 0$, forcing $A^*v = 0$. If the residual is zero, then $Ax_\lambda = v$, and the normal equation reduces to $\lambda x_\lambda = 0$. Thus $x_\lambda = 0$ and $v = 0$. The converse is direct. For $A = I$, the normal equation gives $B_\lambda v = v/(1 + \lambda)$ and residual $\lambda v/(1 + \lambda)$ for $v > 0$. Substitution proves the formula and SR-45 supplies uniqueness. These endpoints explain the strict root inequalities and the separate small-data branch.

□

Formal source: `DiscrepancyBoundaries.no_parameter_above_data;DiscrepancyBoundaries.no_parameter_at_data_norm;DiscrepancyBoundaries.zero_residual_iff_zero_data;DiscrepancyBoundaries.identity_parameter_formula`.

14 Reliable discrepancy with inexact solves

The discrepancy rule of the preceding section selects an actual regularized source. A computed vector need not equal that source, and a displayed residual need not equal its exact norm. Those two errors require separate allowances. The normal equation provides the connection without assuming the computation has already solved the inverse problem.

Throughout this section $A : G \to V$ is bounded between real Hilbert spaces, $N_A = \ker A$, $S_A = N_A^\perp$, and $\lambda > 0$. Write

$$N_\lambda = A^*A + \lambda I, \qquad x_\lambda = N_\lambda^{-1} A^* v, \qquad r_\lambda(x) = N_\lambda x - A^* v.$$

The inverse is the actual positive-parameter constructor already established. The vector $r_\lambda(x)$ is the normal-equation residual of the computed source x. Its size measures source correction and image correction on different scales. A fitting residual $\|Ax - v\|$ alone does not supply either correction.

Hypotheses. Real Hilbert spaces; bounded A; lambda>0; actual normal residual.

Theorem 14.1 (SR-49 Normal residual, two correction scales)**.** *For every computed source x,*

$$x - x_\lambda = N_\lambda^{-1} r_\lambda(x), \qquad \|x - x_\lambda\| \le \frac{\|r_\lambda(x)\|}{\lambda}, \qquad \|Ax - Ax_\lambda\| \le \frac{\|r_\lambda(x)\|}{2\sqrt{\lambda}}.$$

Consequently

$$|\,\|Ax - v\| - \|Ax_\lambda - v\|\,| \le \frac{\|r_\lambda(x)\|}{2\sqrt{\lambda}}.$$

The bounds use the same actual residual and the existing Hilbert norms.

Proof. Subtract $N_\lambda x_\lambda = A^* v$ from the definition of the normal residual and apply the inverse. For $u = N_\lambda^{-1} r$, the energy identity gives

$$\|Au\|^2 + \lambda \|u\|^2 = \langle r, u\rangle \le \|r\| \, \|u\| .$$

Dropping the first term shows $\|u\| \le \|r\| / \lambda$, including $u = 0$. To control the image, complete the scalar square:

$$\|Au\|^2 \le \|r\| \, \|u\| - \lambda \|u\|^2 \le \frac{\|r\|^2}{4\lambda}.$$

Both sides of the desired norm inequality are nonnegative, so taking square roots gives the image bound. Set $r = r_\lambda(x)$ and use linearity of A. Finally the reverse triangle inequality bounds the difference of the two fitting norms by $\|A(x - x_\lambda)\|$.

The distinction matters when λ is small. A normal residual sufficient for a prescribed image tolerance may be insufficient for a prescribed source tolerance. If source correction must be at most κ and image correction at most ξ, a sound pair of conditions is $\|r_\lambda(x)\| \le \lambda\kappa$ and $\|r_\lambda(x)\| \le 2\sqrt{\lambda}\,\xi$. These are parameter-dependent requirements, not a uniform inverse bound as λ tends to zero.

□

Formal source: `InexactResidual.correction_identity;InexactResidual.correction_bound;InexactResidual.image_correction_bound;InexactResidual.fitting_residual_difference`.

Hypotheses. Same Hilbert setting; nonnegative source/image/evaluation budgets; positive parameter; genuine measurement enclosure.

Theorem 14.2 (SR-50 Protected measured discrepancy margins)**.** *Suppose m is a measured fitting norm with*

$$|m - \|Ax - v\|\,| \le \omega, \quad \|r_\lambda(x)\| \le \lambda\kappa, \quad \|r_\lambda(x)\| \le 2\sqrt{\lambda}\,\xi.$$

If

$$\eta + \omega + \xi \le m \le \tau\eta - \omega - \xi,$$

then the actual exact source x_λ has discrepancy in $[\eta, \tau\eta]$, and $\|x - x_\lambda\| \le \kappa$. A protected scalar window exists exactly when

$$2(\omega + \xi) \le (\tau - 1)\eta.$$

Its midpoint $(\tau + 1)\eta/2$ is then admissible.

Proof. SR-49 and the measurement enclosure imply

$$|m - \|Ax_\lambda - v\|\,| \le |m - \|Ax - v\|\,| + |\,\|Ax - v\| - \|Ax_\lambda - v\|\,| \le \omega + \xi.$$

Subtracting this total allowance from the lower protected inequality gives $\eta \le \|Ax_\lambda - v\|$. Adding it to the upper inequality gives $\|Ax_\lambda - v\| \le \tau\eta$. The source-correction bound follows from the other residual condition in SR-49. Hence the computed source is within κ of an actual accepted discrepancy source.

The scalar window is nonempty exactly when its lower endpoint is at most its upper endpoint. Rearrangement gives the displayed condition. The proposed midpoint is the mean of the endpoints; direct substitution gives both inequalities. No property of an approximate solver is hidden in this scalar calculation: its residual, measurement enclosure and correction allowance have already been supplied.

For example, let $\eta = 1/10$, $\tau = 2$, and $\omega = \xi = 1/100$. The total protected margin is $1/50$, and the interval is $[3/25, 9/50]$. Its midpoint is $3/20$. This allocation proves room for the measurement and image-correction errors; it does not prove that a particular computed source attains the interval or that $1/10$ is a true data-noise allowance.

□

Formal source: `InexactResidual.measured_band_transfer;InexactDiscrepancy.measured_certificate;InexactToleranceCertificates.protected_window_exists_iff;InexactToleranceCertificates.protected_midpoint`.

Hypotheses. Fixed bounded real Hilbert A; actual noise bounds; fixed tau; NearAccepted with source budget tending to zero.

Theorem 14.3 (SR-51 Near-accepted consistency and kernel leakage)**.** *Call x near accepted with budget κ when some exact accepted source y satisfies $\|x - y\| \le \kappa$. If $\|Ag - v_i\| \le \eta_i$, $\eta_i \to 0$, and near-accepted x_i have $\kappa_i \to 0$, then*

$$x_i \longrightarrow P_{S_A} g, \qquad \|x_i - g\| \longrightarrow \|P_{N_A} g\| .$$

At every index $\|P_{N_A} x_i\| \le \kappa_i$; therefore the computed true-kernel component tends to zero. No separate parameter schedule is assumed.

Proof. Choose accepted comparison sources y_i from the definition. Exact accepted consistency gives $y_i \to p = P_{S_A} g$, while $\|x_i - y_i\| \le \kappa_i \to 0$ gives $x_i - y_i \to 0$. Adding the two limits proves $x_i \to p$. Continuity of the norm gives the source-error limit, because $p - g = -P_{N_A} g$.

The exact accepted source belongs to S_A, including the small-data zero branch. Orthogonal projection is nonexpansive, so

$$\|P_{N_A} x_i\| = \|P_{N_A}(x_i - y_i)\| \le \|x_i - y_i\| \le \kappa_i.$$

This pointwise estimate only needs near acceptance, not a truth-noise premise. Source consistency additionally needs the stated true-noise allowance tending to zero.

The underlying exact consistency theorem has substantive content even for a nonclosed range. Accepted sources satisfy norm domination $\|y_i\| \le \|p\|$ and $\|Ay_i - Ap\| \to 0$. Since $p \in S_A = \operatorname{cl} \operatorname{Ran} A^*$, approximate p by $A^* w$ and use

$$\|y_i - p\|^2 \le 2 \langle p, p - y_i \rangle \le 2 \|p - A^* w\| \, \|p - y_i\| + 2 \|w\| \, \|A(p - y_i)\| .$$

The first factor $\|p - y_i\|$ is bounded by $2\|p\|$. First choose the approximation accurately, then let i tend along the filter. This proves strong convergence without a bounded inverse.

□

Formal source: `InexactDiscrepancy.NearAccepted;InexactDiscrepancy.near_kernel_leakage;InexactConsistency.near_consistent;InexactConsistency.near_error_norm_limit;InexactConsistency.near_kernel_tendsto_zero`.

Hypotheses. Real Hilbert spaces; positive lambda; nonnegative scalar certificate fields; squared inequalities against actual norms.

Theorem 14.4 (SR-52 Squared proof-carrying discrepancy certificate)**.** *Suppose a record supplies $\lambda > 0$, nonnegative $\rho, \xi, \kappa, m, \omega$, and*

$$\|r_\lambda(x)\|^2 \le \rho^2, \quad \rho^2 \le (\lambda\kappa)^2, \quad \rho^2 \le 4\lambda\xi^2,$$

$$\max(0, m-\omega)^2 \le \|Ax - v\|^2 \le (m+\omega)^2, \qquad \eta + \omega + \xi \le m \le \tau\eta - \omega - \xi.$$

Then x is near accepted with source budget κ. Families of such records give SR-51 when the genuine data allowances and source budgets tend to zero.

Proof. All quantities whose squares are compared are nonnegative. The first two inequalities therefore give $\|r_\lambda(x)\| \le \rho \le \lambda\kappa$. The third gives $\rho \le 2\sqrt{\lambda}\,\xi$ because $(2\sqrt{\lambda}\,\xi)^2 = 4\lambda\xi^2$.

Put $d = \|Ax - v\| \ge 0$. The squared fitting enclosure is equivalent to the unsquared enclosure $\max(0, m-\omega) \le d \le m+\omega$. Since $d \ge 0$, its lower part is equivalent to $m - \omega \le d$. Together these say $|m - d| \le \omega$. Thus every premise of SR-50 has been obtained from the supplied record, and its conclusion proves soundness. Apply SR-51 pointwise for the family statement.

The record does not establish the truth of the data-noise allowance. It establishes a relation between the computed source, actual normal residual and actual fitting norm. An implementation may check rational squared inequalities without numerical square roots, but that requires a faithful enclosure of these actual quantities. The generic Lean structure is proof-carrying; a serialized Python record and its parser are separate implementation evidence.

If the normal residual itself is evaluated approximately as $\widehat{r}$, the distinct assumptions $\|r_\lambda(x) - \widehat{r}\| \le \nu$ and $\|\widehat{r}\| \le \widehat{\rho}$ give the allowance $\widehat{\rho} + \nu$. Substituting that sum into both correction tests retains the evaluation error rather than silently identifying it with zero.

□

Formal source: `InexactToleranceCertificates.squared_enclosure_iff;InexactToleranceCertificates.squared_image_budget;InexactToleranceCertificates.PositiveCertificate;InexactToleranceCertificates.certificate_sound;InexactToleranceCertificates.certificate_consistent`.

15 Uncertain operators and changing identifiable models

Data noise, solver error and model error are different inputs. Let A be the true forward operator and C the available operator, acting between the same real Hilbert spaces. The premise $\|C - A\| \le \delta$ is a whole-operator allowance. A small residual for one source does not establish it. This section first uses a fixed cap M on a true comparison source. The next section removes that cap from the acceptance rule by making the allowance depend on the candidate itself.

Hypotheses. Same normed source/target; bounded A,C; whole operator allowance; genuine data allowance and source cap.

Theorem 15.1 (SR-53 Effective model-noise allowance)**.** *If $\|C - A\| \le \delta$, $\|p\| \le M$ and $\|Ap - v\| \le \eta$, then*

$$\|Cp - v\| \le \epsilon := \eta + \delta M.$$

For every candidate x,

$$|\,\|Ax - v\| - \|Cx - v\|\,| \le \delta\,\|x\|\,.$$

If $|m - \|Cx - v\|\,| \le \omega$, then $|m - \|Ax - v\|\,| \le \omega + \delta\,\|x\|$. The effective-noise allowance is sharp in scalar examples.

Proof. The definition of operator norm gives $\|Cp - Ap\| \le \delta\,\|p\|$. Because a norm bound implies $\delta \ge 0$, the source cap gives

$$\|Cp - v\| \le \|Cp - Ap\| + \|Ap - v\| \le \delta M + \eta.$$

For the fitting-norm comparison, apply the reverse triangle inequality to $Ax - v$ and $Cx - v$, whose difference is $(A - C)x$. Combining this with the measurement enclosure gives the last claim.

To see the possible equality, take real scalar maps $A = 1$, $C = 1 + \delta$, $p = M$ and $v = M - \eta$, with $\eta, \delta, M \ge 0$. Then $Cp - v = \delta M + \eta \ge 0$, so its absolute value is exactly ϵ. The source-size factor cannot generally be omitted: for fixed $\delta > 0$, the action difference $\delta|p|$ grows without bound as $|p|$ grows. Thus replacing δM by an operator-only constant would strengthen the contract incorrectly.

□

Formal source: `OperatorUncertainty.action_error;OperatorUncertainty.effective_data_noise;OperatorUncertainty.fitting_norm_difference;OperatorUncertainty.measured_true_fitting_envelope;OperatorUncertaintyExamples.inflated_noise_sharp`.

Hypotheses. Same real Hilbert spaces; bounded A,C; positive or arbitrary common normal parameter; whole norm error delta.

Theorem 15.2 (SR-54 Transport of actual normal residuals)**.** *For the normal residuals formed with the same parameter,*

$$r_\lambda^A(x) - r_\lambda^C(x) = (A^* - C^*)(Cx - v) + A^*(A - C)x.$$

If $\|C - A\| \le \delta$, then

$$\left\| r_\lambda^A(x) - r_\lambda^C(x) \right\| \le \delta\big(\|Cx - v\| + (\|C\| + \delta)\,\|x\|\big).$$

Hence a computed-model normal allowance $\left\|r_\lambda^C(x)\right\| \le \rho$ gives the corresponding true normal allowance by adding this expression.

Proof. Expand the two residuals. The λx terms cancel. Insert and subtract A^*Cx to obtain the displayed identity. Adjoint isometry gives $\|A^* - C^*\| = \|A - C\| \le \delta$ and $\|A^*\| = \|A\|$. Submultiplicativity therefore bounds its two terms by $\delta\,\|Cx - v\|$ and $\|A\|\,\delta\,\|x\|$. The triangle inequality gives $\|A\| \le \|C\| + \delta$, yielding a bound written in terms of available-model quantities and the allowance. Finally apply $\|r^A\| \le \|r^C\| + \|r^A - r^C\|$.

The identity itself does not require a positive parameter; positivity is needed when this transported residual is used in SR-49 to control an inverse correction. The model error remains distinct from the solver's normal residual. Both contributions are amplified by the corresponding positive-parameter correction scale. A small fixed model allowance is therefore not a uniform guarantee when λ tends to zero.

□

Formal source: `OperatorResidualTransport.normal_residual_difference;OperatorResidualTransport.normal_residual_available_bound;OperatorResidualTransport.true_normal_residual_bound`.

Hypotheses. Same real Hilbert spaces; fixed true bounded A, fixed real tau and fixed identifiable p; source cap M; true bounds and exact/near accepted uncertain-model sources; vanishing allowances.

Theorem 15.3 (SR-55 Strong consistency without matching model kernels)**.** *Suppose $p \in S_A$, $\|p\| \le M$, and*

$$\|C_i - A\| \le \delta_i, \qquad \|Ap - v_i\| \le \eta_i.$$

If x_i lies within κ_i of an exact accepted source for C_i, v_i with allowance $\epsilon_i = \eta_i + \delta_i M$, and $\eta_i, \delta_i, \kappa_i \to 0$, then $x_i \to p$. The model kernels and ranks may change. For arbitrary g, using $p = P_{S_A} g$ yields $\|x_i - g\| \to \|P_{N_A} g\|$ and $P_{N_A} x_i \to 0$.

Proof. Let y_i be the exact accepted comparison. SR-53 makes p feasible for the available model at allowance ϵ_i. In the positive branch, regularized optimality and the lower discrepancy bound give $\|y_i\| \le \|p\|$: compare the objective at y_i and p, then cancel the fitting terms using $\|C_i p - v_i\| \le \epsilon_i \le \|C_i y_i - v_i\|$. The small-data branch is zero and satisfies the same domination.

The upper discrepancy bound gives $\|C_i y_i - v_i\| \le \tau\epsilon_i$. Returning to the true operator,

$$\|Ay_i - Ap\| \le \delta_i \|y_i\| + \tau\epsilon_i + \eta_i \le (\tau + 1)\epsilon_i.$$

Thus y_i is norm dominated by the fixed identifiable p and has true image residual tending to zero. The density argument printed in SR-51 proves $y_i \to p$ without projecting onto the changing model kernels. The near-source difference then tends to zero, proving $x_i \to p$.

For arbitrary g, $AP_{S_A} g = Ag$, so the same true-data bound applies to $p = P_{S_A} g$. Continuity of norm and true-kernel projection gives the remaining limits. The fixed cap concerns this limiting identifiable source. It is not a claim that the minimum-norm sources for unrelated models have already converged.

□

Formal source: `ChangingDiscrepancy.norm_control;ChangingDiscrepancy.near_true_residual;ChangingConsistency.projected_consistent;ChangingConsistency.source_error_limit;ChangingConsistency.true_kernel_leakage_tendsto_zero`.

Hypotheses. Real Hilbert setting; operational checks use C; external true allowances and source cap; positive effective allowance and tau>1 for selection.

Theorem 15.4 (SR-56 Available-model certificate and selected-source soundness)**.** *A source is certified using C when either $\|v\| \le \tau\epsilon$ and $\|x\| \le \kappa$, or it has an SR-52 positive certificate for C with allowance $\epsilon = \eta + \delta M$. Such a source is near accepted for the available model. With the external truth premises of SR-55 and vanishing allowances, certified sources converge to the true identifiable source.*

If $\epsilon > 0$ and $\tau > 1$, those truth premises also give feasibility of the total exact discrepancy selector for C; the selected exact sources are consistent under the same vanishing data/operator bounds.

Proof. In the small-data branch compare x with the accepted zero source: $\|x - 0\| \le \kappa$. In the positive branch, SR-52 supplies an actual accepted comparison source at distance at most κ. This proves operational soundness without using A.

For the true conclusion, separately supply $\|C - A\| \le \delta$, the true-data bound, and $\|p\| \le M$. These hypotheses are not fields established by an observed fitting certificate. They allow SR-55 to be applied to every certified source.

SR-53 gives the feasible witness p for the available model. A positive effective allowance and $\tau > 1$ ensure a strict residual target above its attainable floor in the non-small-data case; the total discrepancy construction chooses an actual positive-parameter source there and zero in the other branch. Each selected source is exact accepted, so SR-55 applies with $\kappa = 0$. This selected object is noncomputable in the generic formalization. Its existence theorem and soundness are distinct from a finite numerical search implementation.

□

Formal source: `OperatorCertificates.Certified;OperatorCertificates.certified_near;OperatorCertificates.certified_consistent;OperatorCertificates.inflated_feasibility;OperatorCertificates.selected_from_allowances_consistent`.

16 Candidate-dependent discrepancy without a supplied source cap

The acceptance rule can avoid an externally supplied number M by using the candidate itself in the uncertainty allowance. Put

$$c_{\eta,\delta}(x) = \eta + \delta \|x\| .$$

A generalized accepted source is either zero in the small-data branch $\|v\| \le \tau\eta$, or an actual Tikhonov source whose residual lies between $c_{\eta,\delta}(x)$ and $\tau c_{\eta,\delta}(x)$. This definition contains $C, v, \eta, \delta, \tau, x$ and the positive parameter; it does not contain a true-source cap. Feasibility remains a separate question and retains a quantitative relation between model accuracy and the true noise scale.

Hypotheses. Real Hilbert spaces; bounded C; lambda>0; delta≥0; feasible comparison p for candidate noise; actual regularized candidate.

Theorem 16.1 (SR-57 Candidate-dependent self-bounding)**.** *Let x be the actual regularized source for C, v, λ. If*

$$\|Cp - v\| \le c_{\eta,\delta}(p), \qquad c_{\eta,\delta}(x) \le \|Cx - v\| ,$$

then $\|x\| \le \|p\|$. In particular, if $\|C - A\| \le \delta$ and $\|Ap - v\| \le \eta$, every generalized accepted source satisfies $\|x\| \le \|p\|$.

Proof. Suppose instead $\|x\| > \|p\|$. Since $\delta \ge 0$, candidate noise is monotone in source norm: $c_{\eta,\delta}(p) \le c_{\eta,\delta}(x)$. The two displayed hypotheses imply $\|Cp - v\| \le \|Cx - v\|$. Comparing Tikhonov objectives, the fitting-square term at p is no larger, whereas its positive weighted source-square term is strictly smaller. The total objective at p is strictly smaller than the objective at its actual unique minimizer x, a contradiction.

With true operator bounds, SR-53 applied using $M = \|p\|$ gives the required candidate-noise feasibility. The operator allowance implies $\delta \ge 0$. The positive accepted branch therefore satisfies the lemma; the small-data branch $x = 0$ satisfies the norm conclusion directly.

The lower discrepancy inequality is decisive. An upper residual bound alone can allow overly large regularized source norms. Here the lower bound grows with the candidate norm, and objective minimality prevents that growth from exceeding a feasible comparison source. The comparison source enters the proof of soundness, while the acceptance definition continues to have no externally supplied source cap.

□

Formal source: `GeneralizedDiscrepancy.regularized_norm_domination;GeneralizedDiscrepancy.accepted_norm_control`.

Hypotheses. Real Hilbert spaces; true whole-operator/data bounds; tau≥0; generalized accepted or near-accepted sources; vanishing allowances.

Theorem 16.2 (SR-58 Generalized true residual and strong consistency)**.** *A generalized accepted x obeys*

$$\|Ax - Ap\| \le (\tau + 1)c_{\eta,\delta}(p).$$

A source within κ of such an x obeys the same bound with $\|A\|\ \kappa$ added. For $p \in S_A$, vanishing data/operator allowances and near-source budgets imply strong convergence to p. If $p = A^*w$, an inexact accepted candidate also satisfies

$$\|x - p\|^2 \le 2\kappa^2 + 4\,\|w\|\,(\tau + 1)c_{\eta,\delta}(p).$$

Proof. SR-57 gives $\|x\| \le \|p\|$, so $c_{\eta,\delta}(x) \le c_{\eta,\delta}(p)$. The upper generalized discrepancy bound and $\tau \ge 0$ yield $\|Cx - v\| \le \tau c_{\eta,\delta}(p)$. Insert Cx and v into the true residual:

$$\|Ax - Ap\| \le \delta\,\|x\| + \|Cx - v\| + \eta \le (\tau + 1)c_{\eta,\delta}(p).$$

For a near candidate, split its image difference through the exact accepted comparison and apply the operator norm bound to the correction.

As $\eta_i, \delta_i \to 0$, the comparison-source noise $\eta_i + \delta_i\,\|p\|$ tends to zero. Exact candidates are norm dominated by p and have true image error tending to zero; the same density argument as SR-51 proves their strong convergence. Add a vanishing near-source difference for inexact candidates. No relative feasibility condition is needed in this implication once accepted candidates have been supplied.

For the source-condition rate, norm domination and $p = A^*w$ give

$$\|y - p\|^2 \le 2\,\langle p, p - y\rangle \le 2\,\|w\|\ \|A(p - y)\|$$

for an exact comparison y. Bound its image by the first conclusion. Then $\|x - p\| \le \|x - y\| + \|y - p\|$ and $(a + b)^2 \le 2a^2 + 2b^2$ prove the stated inexact rate.

□

Formal source: `GeneralizedDiscrepancy.accepted_true_residual;GeneralizedDiscrepancy.accepted_consistent;GeneralizedInexact.near_true_residual;GeneralizedInexact.near_consistent;GeneralizedInexact.near_source_condition_rate`.

Hypotheses. Real Hilbert spaces; eta>0,tau>1; true operator/data bounds; delta‖p‖<(tau-1)eta.

Theorem 16.3 (SR-59 Relative feasibility and noncomputable selection)**.** *If $\|C - A\| \le \delta$, $\|Ap - v\| \le \eta$ and*

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

then an accepted source exists with norm at most $\|p\|$. In the large-data case one can choose a positive parameter whose residual is exactly $\tau\eta$. A sufficient ratio condition is $\delta \le \gamma\eta$ and $\gamma\,\|p\| < \tau - 1$. Noncomputably selected families converge to $p \in S_A$ when the data/operator allowances vanish and the pointwise feasibility conditions hold.

Proof. The true bounds imply $\|Cp - v\| \le \eta + \delta \|p\| < \tau\eta$. If $\|v\| \le \tau\eta$, the zero source is accepted. Otherwise the strict attainable-floor bound and $\tau\eta < \|v\|$ give a positive exact discrepancy root for C with residual $\tau\eta$. Let x be its actual Tikhonov source. Objective comparison, or the fixed-noise domination lemma, gives $\|x\| \le \|p\|$. Consequently

$$c_{\eta,\delta}(x) \le \eta + \delta \|p\| < \tau\eta = \|Cx - v\| .$$

Because $\delta \ge 0$ and $\tau > 1$, $\tau\eta \le \tau c_{\eta,\delta}(x)$. Thus both generalized band inequalities hold.

Multiplying $\delta \le \gamma\eta$ by $\|p\| \ge 0$ and the strict margin $\gamma \|p\| < \tau - 1$ by $\eta > 0$ proves the ratio specialization. The chosen source is defined by classical choice from the existence theorem. Its raw definition takes A, p and the truth/feasibility proofs as arguments; it is a theorem-level witness rather than an extracted algorithm using only available measurements. Each chosen source is accepted, so SR-58 proves the family convergence. The selection theorem retains the feasibility premises even though the acceptance predicate has no supplied scalar source cap.

□

Formal source: `GeneralizedDiscrepancySelection.positive_parameter_source;GeneralizedDiscrepancySelection.exists_accepted;GeneralizedDiscrepancySelection.selectedSource;GeneralizedFeasibility.relative_accuracy_of_scaled_error;GeneralizedDiscrepancyConsistency.selected_consistent`.

Hypotheses. Same Hilbert setting; near exact generalized acceptance with vanishing correction budget; scalar boundary uses valid but nonvanishing model allowance.

Theorem 16.4 (SR-60 Inexact generalized sources and the zero-model boundary)**.** *Near-generalized acceptance is preserved by enlarging a source-correction budget and gives the consistency conclusion of SR-58 when that budget tends to zero. Nevertheless valid true-noise and operator allowances alone do not ensure generalized feasibility. With $A = 1$, $C = 0$, $p = v = 1$, $\eta = 1/10$, $\delta = 1$ and $\tau = 2$, no source is generalized accepted.*

Proof. Near acceptance means there exists an exact generalized accepted comparison y with $\|x - y\| \le \kappa$. Increasing κ preserves that inequality. Along a family, choosing the comparisons and adding the vanishing difference proves consistency by SR-58.

For the scalar boundary, the true data error is zero and $\|C - A\| = 1 \le \delta$, so both stated truth allowances are valid. The small-data branch would require $1 \le 2/10$, which is false. For every positive parameter, the actual regularized source of $C = 0$ is zero and its fitting residual is one. Its candidate noise is $1/10 + 1 \cdot 0 = 1/10$, so the upper generalized band would again require $1 \le 2/10$. Every branch fails. The relative feasibility condition also fails, since $1 < (2 - 1)/10$ is false. Thus the boundary refutes an unconditional existence interpretation, not the conditional theorem.

For comparison, take $A = C = 1$, $v = 1$, $\eta = 1/10$, $\delta = 0$, $\tau = 2$ and $\lambda = 1/4$. Then $x_\lambda = 4/5$ and the residual is $1/5$. It lies in $[1/10, 2/10]$, so the candidate is accepted, and $0 < (2 - 1)/10$ verifies the strict feasibility margin. These examples use the actual scalar normal inverse and exact norms, not a substituted approximate residual.

□

Formal source: `GeneralizedInexact.NearAccepted;GeneralizedInexact.near_consistent;GeneralizedDiscrepancyExamples.identity_candidate_accepted;GeneralizedDiscrepancyExamples.zero_model_not_accepted;GeneralizedDiscrepancyExamples.zero_model_valid_true_bounds;GeneralizedDiscrepancyExamples.zero_model_relative_accuracy_fails`.

17 Calibration generates an operator allowance

The final link connects justified calibration measurements with the whole-operator premise used above. It does not infer the calibration truth allowance from the model's own fit. Let K be a real normed calibration coordinate space, G a real normed source space and V a real normed response space. A split frame has bounded synthesis $U : K \to G$ and analysis $R : G \to K$ with $UR = I_G$. Its coordinates may be redundant: U need not be injective or square. The reconstruction identity requires coverage of every source direction.

The observed response $Y : K \to V$ is an operator, and the assumptions $\|Y - AU\| \le \eta_{\mathrm{cal}}$ and $\|Y - CU\| \le f$ are operator-norm enclosures. Individual calibration samples do not automatically give these enclosures. A mathematical or experimental justification of η_{cal} remains a separate input.

Hypotheses. Real normed spaces; bounded U,R with UR=I; whole calibration truth/model-fit operator bounds.

Theorem 17.1 (SR-61 Redundant split calibration identifies the operator)**.** *Every bounded source operator T factors as $T = (TU)R$. Therefore*

$$C - A = ((CU) - (AU))R, \qquad \|C - A\| \le (\eta_{\mathrm{cal}} + f)\,\|R\|\,.$$

If a supplied analysis bound is $\|R\| \le k$, the allowance is at most $(\eta_{\mathrm{cal}} + f)k$. Exact calibration $Y = AU = CU$ forces $A = C$.

Proof. For every source g, $URg = g$, so $Tg = TU(Rg)$. This proves the actual factorization, including its application to $C - A$. The triangle inequality through the common measured operator gives

$$\|CU - AU\| \le \|CU - Y\| + \|Y - AU\| \le f + \eta_{\mathrm{cal}}.$$

Submultiplicativity then yields the displayed allowance. Both measured bounds imply their scalar allowances are nonnegative; hence replacing $\|R\|$ by a larger k preserves the inequality. With exact calibration, $CU - AU = 0$ and the factorization gives $C - A = 0$.

Redundancy causes no logical problem. For instance, let $K = \mathbb{R}^2$, $G = \mathbb{R}$, $U(k_1, k_2) = k_1$ and $Rg = (g, 0)$, with the standard product norm. Then $UR = I$, $\|U\| = \|R\| = 1$, but U has a nontrivial kernel. The frame spans the source even though it has an unused coordinate. A true scalar map $A = a$ and model $C = c$ with measured row $Y(k_1, k_2) = yk_1$ give $|y - a| \le \eta_{\mathrm{cal}}$ and $|y - c| \le f$, hence $|c - a| \le \eta_{\mathrm{cal}} + f$. The same estimate comes from the general factorization rather than assuming the coordinates are square and invertible.

□

Formal source: `CalibrationFrame.Frame.factor_through_frame;CalibrationFrame.Frame.difference_factorization;CalibrationFrame.Frame.measured_operator_bound;CalibrationFrame.Frame.measured_operator_bound_of_analysis_bound;CalibrationFrame.Frame.exact_calibration_identifies`.

Hypotheses. Fixed split calibration frame; real Hilbert G,V; fixed bounded true A and tau≥0; fixed identifiable p; whole calibration truth/model fit and true data allowances; generalized acceptance; relative feasibility for choice.

Theorem 17.2 (SR-62 Calibrated recovery and conditional selected sources)**.** *For recovery assume G, V are real Hilbert spaces, A is a fixed bounded true operator, $\tau \ge 0$ is fixed, and $p \in (\ker A)^{\perp}$ is fixed. Retain the whole calibration bounds and the true data bounds $\|Ap - v_i\| \le \eta_i$. Put $\delta_i = (\eta_{\mathrm{cal},i} + f_i)\,\|R\|$. If the true data, calibration-truth and model-fit allowances tend to zero, then $\delta_i \to 0$. Generalized accepted sources using this*

generated allowance converge to a fixed identifiable true source p. For an arbitrary true source g, apply this with $p = P_{S_A} g$ and genuine data bounds $\|Ag - v_i\| \le \eta_i$; then the source-error norm tends to $\|P_{N_A} g\|$.

If additionally $\eta_i > 0$, $\tau > 1$ and $\delta_i \|p\| < (\tau - 1)\eta_i$ pointwise, accepted sources exist and the noncomputable selected sources have the same limit.

Proof. Apply SR-61 at every index to obtain the actual operator bound $\|C_i - A\| \le \delta_i$. Because the frame is fixed, its finite norm $\|R\|$ is a constant. The sum of the two vanishing calibration allowances tends to zero, so multiplying by that constant proves $\delta_i \to 0$. With the genuine data-noise bounds, every premise of the accepted-source consistency theorem SR-58 is now available.

For arbitrary g, use $p = P_{S_A} g$. Its forward image is Ag, so the same data bounds hold. The strong source limit is p, and continuity gives $\|x_i - g\| \to \|p - g\| = \|P_{N_A} g\|$. No calibration can remove source ambiguity in the true forward kernel without an additional identifying condition.

The relative margin feeds the generated allowance into SR-59 and proves pointwise existence and norm domination. Its chosen sources are accepted, so the first conclusion applies to them. The raw selection is still noncomputable and proof-dependent. Calibration has generated the operator-error premise; it has not eliminated the distinct truth, data, feasibility or selection obligations.

In particular, declining observed model-calibration fit alone is insufficient: the calibration-truth allowance must also vanish. If the reconstruction frame changes with the index, a growing analysis norm would require additional control. The printed family theorem uses the fixed frame it actually assumes.

□

Formal source: `CalibrationRecovery.generated_operator_bound;CalibrationRecovery.generatedDelta_tendsto_zero;CalibrationRecovery.calibrated_consistent;CalibrationRecovery.calibrated_error_norm_limit;CalibrationSelection.exists_accepted_from_calibration;CalibrationSelection.calibrated_selected_consistent`.

Hypotheses. Real normed spaces; bounded split stage frames; or bounded U,R with whole reconstruction defect at most rho<1.

Theorem 17.3 (SR-63 Hierarchical and approximate reconstruction frames)**.** *For split frames $(U_E, R_E) : J \rightleftarrows K$ and $(U_F, R_F) : K \rightleftarrows G$, the composite is split, with analysis $R_E R_F$ and $\|R_E R_F\| \le \|R_E\| \, \|R_F\|$. The corresponding calibration allowance is at most $(\eta_{\text{cal}} + f) \, \|R_E\| \, \|R_F\|$.*

More generally, if $\|I - UR\| \le \rho < 1$, then

$$\|T\| \le \frac{\|TU\| \, \|R\|}{1 - \rho}, \qquad \|C - A\| \le \frac{(\eta_{\text{cal}} + f) \, \|R\|}{1 - \rho}.$$

Proof. The composite synthesis is $U_F U_E$ and analysis is $R_E R_F$. Their product is $U_F (U_E R_E) R_F = U_F R_F = I_G$. Submultiplicativity bounds the composite analysis norm; applying SR-61 and that bound proves the hierarchical allowance. There is no requirement that either synthesis be square or injective.

For the approximate frame, use the exact identity $T = (TU)R + T(I - UR)$. It gives

$$\|T\| \le \|TU\| \, \|R\| + \|T\| \, \|I - UR\| \le \|TU\| \, \|R\| + \rho \, \|T\| .$$

Since $1 - \rho > 0$, absorb the final term and divide. For $T = C - A$, the measured calibration bounds give $\|TU\| \le \eta_{\text{cal}} + f$, producing the operator allowance. This argument only uses norm inequalities; it does not substitute an assumed exact inverse for the approximate reconstruction.

The denominator records the loss of control. At $\rho = 0$ the exact frame estimate is recovered. At $\rho = 1$ the absorption step gives no bound, and it is not legitimate to reuse the formula by division. Likewise the product of hierarchical analysis norms is a valid conservative bound, not generally an equality or the optimal sensitivity of the composite. These conditioning factors belong in the certificate.

□

Formal source: `CalibrationHierarchy.comp;CalibrationHierarchy.comp_analysis_norm_le;CalibrationHierarchy.hierarchical_operator_bound;ApproximateCalibration.approximate_factorization;ApproximateCalibration.approximate_frame_bound;ApproximateCalibration.measured_approximate_operator_bound`.

Hypotheses. Actual partial calibration from R to R^2 with product norm; bounded unseen scalar operators; arbitrary real norm threshold.

Theorem 17.4 (SR-64 A missed calibration direction permits unbounded error)**.** *Let $U(t) = (t, 0)$ and $C_s(x, y) = sy$. Then $C_s U = 0$ and $\|C_s\| = |s|$. For every real M there is an operator agreeing perfectly with zero on the entire calibration subspace yet satisfying $\|C_s\| > M$. The calibration synthesis is not surjective. Thus perfect agreement on a proper source subspace does not control whole-source operator accuracy.*

Proof. Direct substitution gives $C_s U(t) = s \cdot 0 = 0$. The second-coordinate projection has norm one, so scalar multiplication gives $\|C_s\| = |s|$; the vector $(0, 1)$ also verifies the lower bound. Choose $s = |M| + 1 > 0$. Then $\|C_s\| = |M| + 1 > M$ while its calibrated response remains zero. Taking the true operator and measured response both zero shows that both calibration residuals vanish, even though the whole operator discrepancy exceeds M.

The vector $(0, 1)$ is outside the range of U, proving nonsurjectivity. A split reconstruction $UR = I_{\mathbb{R}^2}$ would make every source a range vector, so no such reconstruction exists here. Nor can any approximate reconstruction for this synthesis have whole defect below one: the range of UR has second coordinate zero, and therefore $(I - UR)(0, 1)$ has second coordinate one and norm at least one. This last observation explains the direction-coverage role of the strict absorption threshold in SR-63.

The positive redundant example and this negative partial example differ in coverage, not coordinate count. Extra coordinates can coexist with complete reconstruction of every source. Calibration of fewer source directions leaves an operator component that all tested calibration responses fail to see. The theorem makes that obstruction explicit without claiming that an externally supplied truth allowance can be inferred from those responses.

□

Formal source: `CalibrationExamples.partial_not_surjective;CalibrationExamples.unseen_comp_partial;CalibrationExamples.unseen_norm;CalibrationExamples.arbitrarily_large_unseen_error;CalibrationExamples.perfect_partial_calibration_unbounded`.

18 What the assembled calculus establishes

The results support a precise progression. Exact resolutions identify which coordinate vectors and coordinate kernels describe actual objects and operations. Split compression identifies the excursion through discarded state that prevents preservation of composition. Memory maps characterize prepared closure; finite hidden dimension turns their universal vanishing into a finite whole-map certificate.

Behavioral quotients retain exactly the distinctions accessible to the declared future task. Normed defect budgets determine how approximate representations propagate their errors.

In inverse problems, compatibility, realizability, and source identifiability remain distinct. The local obstruction quotient measures compatible data that have no global source. The compatibility residual measures the inconsistent sector, with its own conditioning constant. A global source has separate kernel ambiguity and inverse sensitivity. Regularization controls an actual penalized objective. Its strong limits recover each fixed identifiable source, while uniform source convergence requires closed range. Discrepancy rules and calibrated changing models add operational acceptance contracts whose true error allowances remain visible.

The common theme is that changing a representation changes the obligations needed to use it. A certificate can discharge an obligation from concrete maps and inequalities. A residual can expose the remaining defect. An obstruction theorem can prove that an obligation is impossible under the available structure. These outcomes have different mathematical meanings.

18.1 Exact representation and information loss

A redundant exact coordinate space need not lose any information about its object, because synthesis after analysis is identity. Coarsening a packet address family can also remain exact if the entire contents of the merged packets are retained. Information loss is supplied by a separate map and its kernel. In the behavioral setting, only the largest invariant invisible subspace can be quotiented without changing any declared future observation. In source reconstruction, the unavoidable ambiguity is the true forward operator's kernel. Neither coordinate count nor packet count alone decides sufficiency.

18.2 Relations to established mathematics

The mathematics deliberately uses established foundations. Split idempotents, quotients, observability, block elimination, operator exponentials, orthogonal geometry, Moore–Penrose inversion, Tikhonov regularization, and discrepancy principles each contribute concrete constructions. The paper's organizing claim is that their interfaces can be written together, with exact transport, obstruction, approximation, and certificate obligations. Historical novelty of each component is a separate question. The formal inventory and finite test totals do not answer it.

18.3 Open bridges and extensions

The following problems specify further research beyond the printed contracts.

1. **Stable target-only recovery from behavioral minimality.** The exact quotient X/N preserves all specified future observations. The Hilbert theory stably recovers an identifiable source under its reconstruction or regularization assumptions. A general theorem connecting these two constructions must identify a task target, its quotient metric, the relevant forward map, and sufficient conditioning. It should recover that target without requiring reconstruction of the full original source. No such unrestricted bridge is asserted in this edition.
2. **Calibration truth allowances from individual samples.** The calibration theorems start from the whole-operator inequality $\|Y - AU\| \leq \eta_{\mathrm{cal}}$. Deriving it from individual noisy probes requires a specified sample model, a justified aggregation bound, and control of untested directions. A measured model fit supplies neither the true operator nor that truth allowance by itself.

3. **Verified numerical procedures.** The formalization proves the certificate contracts. The rational Python implementations compute and check records under explicit input limits. Their parsers, arithmetic helpers, row-space saturation, parameter bracketing, and stopping procedures are not formally verified. Verified extraction or an independent formal certificate importer would add a new layer of assurance.
4. **Additional admissible categories.** Minimum linear dimension does not establish a minimum positive, stochastic, integral, geometric, or nonlinear realization. Bounded autonomous generators do not settle unbounded-domain or time-dependent flow problems. Such extensions must print the additional admissibility and domain hypotheses rather than importing them through terminology.

19 Source distribution and verification boundaries

This edition contains a traditional paper, a visual guide, a Lean guide, and a complete source ZIP. The `.md` files are included in that ZIP, alongside editable `.tex`, the 93 production Lean modules, theorem guides, exact Python companions, example data, source inventories, and verification records. The original v0.16.0 archive is preserved separately for provenance. Publication sources and reconciliation records are additional files, not silent replacements of the mathematical source history.

The principal concordance uses SR-01 through SR-64 with explicit hypotheses, dependencies, exact declaration names, and hashes of the source files read for this assembly. These are consolidated theorem groups. The declaration inventory comprises 998 theorem/lemma declarations and 1159 enumerated named declarations in 93 production modules. Supporting identities, classical specializations, definitions, and checked examples contribute to these totals.

The archived v0.16.0 final verification includes two runs of the latest block's 31-module project dependency closure and audits of its 52 new named declarations. It does not record a fresh simultaneous compilation of every one of the 93 modules at that final checkpoint. The complete release verification record distinguishes archived Lean evidence, fresh compilation if any, source inspections, finite Python replay, document builds, and archive integrity. This manuscript makes no fresh Lean compilation claim merely from the existence of those archives.

The exact-rational companions cover observation saturation, discrete memory, quantitative budgets, orthogonal globalization, quotient metrics, reconstruction, positive-parameter regularization, discrepancy, inexact solves, changing operators, generalized feasibility, and calibration. Their finite checks are implementation evidence in the tested dimensions and inputs. They are not universal proofs and do not infer external truth allowances. The current release's exact run ranges and counts belong in its machine-readable verification records so that archived and fresh results cannot be conflated.

19.1 Rebuilding the reading set

The supplied build scripts invoke pdfLaTeX three times for stable cross-references and table-of-contents entries. This rebuild creates the documents. It does not compile Lean, rerun scientific checks, infer a noise allowance, or authenticate a physical calibration experiment. Separate scripts and instructions address those activities.

The standalone Markdown manuscript preserves the same statements and proofs for readers and tools that work directly with text. The concordance CSV/JSON provides a structured route from a mathematical claim to its formal declarations. The visual guide illustrates the same contracts with exact vector figures and checked rational examples. The Lean guide explains the production modules, the

formal domains, and the evidence needed to reproduce the recorded builds.

A Coverage of the sixteen development blocks

The block numbers below preserve the development history; they are not the section numbers of the reorganized paper. Each row identifies four consolidated theorem groups and the principal mathematical contract.

Block	Stable claims	Contract
1	SR-01–SR-04	Exact redundant representation; compression defects; actual globalization obstruction; behavioral quotient.
2	SR-05–SR-08	Discrete hidden elimination; closure equivalence; finite feedback certificate; canonical hidden-space construction.
3	SR-09–SR-12	Finite observation horizon; minimal image dimension; realizable observation tables; finite certificate soundness.
4	SR-13–SR-16	Exact word-defect transport; heterogeneous budgets; two-rate response; stable initialization bounds.
5	SR-17–SR-20	Actual bounded flows; powers-to-exponentials bridge; continuous memory; finite moment and jet certificates.
6	SR-21–SR-24	Duhamel defect transport; flow growth envelopes; uniform stability; scalar sharpness and instability boundary.
7	SR-25–SR-28	Unique orthogonal error split; harmonic obstruction; compatibility conditioning; finite exactness and refinement.
8	SR-29–SR-32	Constructed canonical inverse; minimum-norm source; source sensitivity; closed-range characterization.
9	SR-33–SR-36	Native quotient metric; minimum representatives; bounded obstruction transport; refinement and nonclosed seminorm boundary.
10	SR-37–SR-40	Coercive Tikhonov construction; objective optimality; bias/noise/kernel accounting; perturbation and globalization.
11	SR-41–SR-44	Normal-range density; strong source limits; noisy and inexact consistency; uniform boundary and fitted-data limits.
12	SR-45–SR-48	Actual residual curve; noncomputable root existence; norm-controlled strong consistency; accepted rule and endpoints.
13	SR-49–SR-52	Source/image solver scales; protected measurements; near-accepted consistency; squared certificate soundness.
14	SR-53–SR-56	Whole-operator uncertainty; model residual transport; changing-model consistency; operational certificates.
15	SR-57–SR-60	Candidate-dependent self-bounding; source-cap-free acceptance; feasibility-aware selection; inexact and impossible-model boundaries.
16	SR-61–SR-64	Redundant calibration frames; generated uncertainty and recovery; hierarchical/approximate frames; unseen-direction obstruction.

B Principal dependency paths

The exact-coordinate path is SR-01 to SR-02. The hidden-memory path is SR-05 to SR-06 to SR-07, with SR-08 supplying the canonical bridge from an arbitrary split compression. The behavioral path starts at SR-04 and passes through SR-09–SR-12. Quantitative discrete transport is SR-13–SR-16. Bounded analytic flow converts the algebraic memory path through SR-17–SR-20 and provides approximate transport in SR-21–SR-24.

The globalization path begins with the actual quotient SR-03 and its projected geometry SR-25–SR-28. Its intrinsic metric is proved in SR-33–SR-36. Canonical source reconstruction SR-29–SR-32 needs the additional source-space and closed-range hypotheses. Regularization SR-37–SR-40 relaxes closed range at positive parameter. Density and norm domination in SR-41 and SR-46 provide the common strong-consistency engine for the discrepancy rules and changing-model extensions.

Calibration SR-61–SR-63 supplies operator-error allowances from its own externally justified truth bounds and conditioning. The recovery chain then uses generalized acceptance and explicit parameter feasibility. The negative control SR-64 explains why source-direction coverage is part of that contract.