KOMOREBI / LANGUAGE AND COMPUTATION

Neuralese

Recurrent Latent Operators, Causal Semantics, and Stability

Komorebi P.

Consolidated research edition 1.0

6 October 2026

Finalized object. A consolidated mathematical and experimental account of the recovered direct Neuralese branches, with source archives, explicit claim status, and reproducible checks. This is not a release of a universally stable learned language. Positive finite results, failed scientific gates, invalidated assay gates, and unspent experiments are retained as different outcomes.

The manuscript, reference implementation, preserved source archives, and validation records form one edition. Publication finalization does not convert finite evidence into an unrestricted theorem or an unrun experiment into a result.

Abstract

Neuralese names a research program on recurrent latent computation whose intermediate state need not be rendered as natural-language tokens. The central issue is not whether a hidden vector predicts an answer, but whether an identifiable state supports reusable operations, causal intervention, reliable continuation, and transport between representations. This manuscript develops those requirements and consolidates twelve recovered direct Neuralese source packages plus the recurrent-depth reference laboratory.

The mathematical account separates coordinate gauges from semantic equivalence, current readout from future observability, bounded norm from stable meaning, and exact compiled control from learned implementation. It gives a two-index family of history/continuation quotients, a finite observable-kernel construction, composition-error bounds, and an obstruction to writing task information into a fixed null fiber under globally isometric dynamics. These are explicit derivations under stated assumptions, not proof-assistant certifications of the neural software.

The retained evidence is mixed and informative. Projective terminal-only models achieved exact tested structured-family performance through depth 64 in a two-seed developmental study, with strong state and operator interventions. Subsequent exhaustive finite audits found more than 33 million correct ordinary-history continuation decisions while preserving failed cumulative and long-path gates. Repeated identity paths exposed delayed semantic errors. A later norm-preserving operator study completed 48 training cells but passed only 8 of 26 registered gates. The packet-closure pilot retained an assay failure, the invariant-kernel branch stopped at calibration, and the semantic-retraction bank remains unspent. A small frozen-language-model integration passed engineering checks but not its scientific gates.

Together the results support bounded learned operator computation and a concrete stability problem, not a general endogenous machine language. The companion OSL edition supplies an explicit semantic consumer. The missing bridge is a learned, causally adequate translation into that interface under declared scope and independent continuation tests.

The key distinction. A correct finite semantic executor exists for the supplied tasks. Learning a representation that reliably implements that executor under weaker information and supervision remains a separate problem. An exact positive control is not counted as a learned result; a preserved failure is not converted into success by a later editorial version.

Contents

1 The object of study 3
2 Exact semantics before approximate semantics 4
3 The two-index completion surface 4
4 Linear observable kernels 5
5 Gauge fixing is not semantic completion 7
6 Exact compiled controls and learned operators 8
7 Ambient memory and the isometry obstruction 8
8 The recovered experimental lineage 9
9 Completion audits: many correct answers, no unrestricted closure 11
10 Later branches and stopped experiments 13
11 The frozen-language-model integration 14
12 From latent computation to an explicit language 14
13 A precise continuation protocol 15
14 Fresh checks and release boundaries 15
15 Conclusion 16
A Mathematical claim register 17
B Reading and execution map 17
C Terminology that remains fixed 18

1 The object of study

A language model can produce useful text without exposing a faithful account of the computation that produced it. Conversely, a recurrent computation can manipulate information without turning every intermediate state into a word. Neuralese studies the latter possibility through instrumented, bounded experiments. The name does not assert that every useful latent vector is a language or that a human-readable transcript exists for it.

A general controlled latent system is

$$z_{t+1} = F_\theta(z_t, u_t), \qquad y_t = O_\theta(z_t, q_t), \tag{1}$$

where u_t supplies an operation or observation and q_t specifies a later query. A recurrent state is scientifically interesting when it supports multiple future uses, not merely when it stores the already known final output token. The question may be asked for a trained operator network, an adapter attached to a frozen language model, or a deliberately constructed finite fixture; these are not interchangeable experimental settings.

The public Coconut proposal is a relevant neighboring architecture: hidden states are fed back into a language model rather than always decoded into intervening words. This motivates latent recurrence but does not validate the particular mechanisms tested here. Causal abstraction provides a vocabulary for relating an internal mechanism to a higher-level algorithm, while critiques of unrestricted alignment show why a flexible translator can overstate apparent understanding. [H+24, G+23, S+25]

1.1 An operational evidence ladder

A useful ladder is: executable recurrence; task-relevant state; causal state use; reusable operation; continuation-preserving abstraction; stable composition; and transport under controlled changes of realization. Each level adds an obligation. A high score on a linear probe is evidence of recoverability by that probe, not sufficient evidence that the model natively uses the decoded variable. A successful state patch establishes a particular intervention effect only under its sampling and control design. Stable norm establishes a numerical property, not the semantic part of the ladder.

The current record reaches different levels in different finite settings. It must not be collapsed into a single "Neuralese percentage."

1.2 What changes during learning

In the small operator-workspace studies, learned encoders and operator-producing maps are optimized from finite task data; the resulting weights are then frozen for evaluation. In a hard-QR or hard-polar arm, an architectural construction enforces isometry. That property is not learned, although the parameters selecting an isometry are learned.

In G1, a small trainable recurrent workspace is coupled to a frozen local language-model backbone. The workspace parameters change; the registered backbone does not. Later diagnostic probes fitted to the resulting activations are measurements performed by the analyst, not evidence that the original model learned the probe's algorithm. In a frozen-checkpoint audit, neither the original workspace nor the backbone is retrained. The release labels all four situations separately. [N02, NOP, G1, G1D]

2 Exact semantics before approximate semantics

2.1 A task-relative realization

Let $\mathcal{A}$ be an operation alphabet, Z a latent state space, and Y an observation space. Each $a \in \mathcal{A}$ induces $F_a : Z \to Z$. For a word $w = a_1 \cdots a_k$, define

$$F_w = F_{a_k} \circ \cdots \circ F_{a_1}.$$

A candidate semantic state space S, operation maps G_a, and abstraction $q : Z \to S$ should satisfy

$$q \circ F_a = G_a \circ q, \qquad O = \bar{O} \circ q. \tag{2}$$

These conditions say that semantics is both operation-compatible and sufficient for the declared outputs. A decoder that predicts only the present label establishes neither condition by itself.

Proposition 2.1 (Semantic continuation)**.** *If* (2) *holds on a forward-invariant domain $D \subseteq Z$, then*

$$O(F_w(z)) = \bar{O}(G_w(q(z))) \quad \textit{for every } z \in D \textit{ and every finite } w.$$

Proof. Forward invariance keeps every intermediate state inside the domain where the equations hold. Induction gives $q(F_w(z)) = G_w(q(z))$. Apply the observation factorization to the endpoint. □

Without forward invariance, a one-step certificate may leave its own scope. Without observation factorization, the reduced state may preserve transitions but omit an output the consumer needs. Without operation compatibility, matching today's readout says nothing about tomorrow's computation.

2.2 Three meanings of "same state"

Geometric equality, semantic equality, and predictive equality are different. Two latent vectors can differ in a redundant coordinate while implementing identical continuations. Two vectors can have the same current class label while differing in future confidence or even future answers. Two states can be experimentally indistinguishable within a finite horizon yet separate after a longer suffix.

An approximate relation such as $d(z, z') \le \varepsilon$ is generally not an equivalence relation. For points $0, 0.75\varepsilon, 1.5\varepsilon$, successive pairs meet the bound but the first and last do not. Quotient terminology therefore requires an actual equivalence relation, a specified quantization, or a separate approximate-abstraction formalism. Pairwise tolerance alone is insufficient.

3 The two-index completion surface

3.1 History depth and future horizon

Let $I \subseteq Z$ be the allowed initial states and

$$R_{\le L} = \{F_w(z) : z \in I,\ |w| \le L\}.$$

Define exact finite-horizon observational equivalence by

$$z \sim_H z' \iff O(F_w(z)) = O(F_w(z')) \text{ for every } |w| \le H.$$

The finite object is

$$Q_{L,H} = R_{\leq L}/\sim_H . \tag{3}$$

The parameter L governs which microstates can be reached; H governs which future observations can distinguish them. A single depth index conceals this difference. The causal-completion and completion-surface audits were designed to expose it. [NCC, NCS]

Proposition 3.1 (Refinement and diagonal action)**.** *Increasing H refines $\sim_H$. For every action a, there is a well-defined map*

$$Q_{L,H+1} \longrightarrow Q_{L+1,H}, \qquad [z] \longmapsto [F_a(z)].$$

Proof. The set of tested continuations grows with H, so equality under all longer-horizon tests implies equality under the shorter set. If $z \sim_{H+1} z'$, then for every suffix $|w| \leq H$, the continuation consisting of a followed by w has length at most $H+1$. Hence $O(F_w F_a(z)) = O(F_w F_a(z'))$, so $F_a(z) \sim_H F_a(z')$. Both successor states are in $R_{\leq L+1}$. □

The horizon decreases along this map. A finite-horizon state cannot automatically be updated forever at the same horizon without additional closure information.

3.2 A minimal delayed-observability counterexample

Let the concrete state be $(x, y) \in \{0,1\}^2$, with present observation $O(x, y) = x$, and let the only operation swap the coordinates. At horizon zero, there are two classes, classified by x. At horizon one, there are four: the next observation also reveals y.

The states $(0,0)$ and $(0,1)$ initially look identical but become distinguishable after one operation. Keeping only the current answer destroys future semantics. This exact example is tested by the new finite checkpoint interface. It is the simplest reason a successful final-answer probe should not be called a complete causal state.

3.3 Hard and soft observations

A hard interface is $O_h(z) = \arg\max_j \ell_j(z)$. A soft interface may be the probability vector or centered logits. Exact agreement under O_h permits changes in every nonwinning logit and changes in the winner's margin. Those changes can matter under later operations.

The retained completion-surface results show this distinction rather than merely proposing it: hard continuation profiles can be exact on ordinary reachable shells while the probability and centered-logit profiles retain a measurable history-dependent microstate. That does not imply that the hard quotient is useless; it means that its interface and continuation domain must remain attached to its claim. [NCS]

4 Linear observable kernels

4.1 A computable null-space construction

For a switched linear system $z_{t+1} = M_{a_t} z_t$ with observation Cz_t, define

$$N_H = \bigcap_{|w| \le H} \ker(CM_w), \qquad M_\varnothing = I.$$

A direction belongs to N_H exactly when it is invisible under every continuation through horizon H.

Proposition 4.1 (Observable-kernel recurrence)**.** *One has*

$$N_{H+1} = N_H \cap \bigcap_{a \in \mathcal{A}} M_a^{-1} N_H.$$

If $N_{H+1} = N_H$, then N_H is invariant under every generator and equals N_∞. In a d-dimensional state space, the decreasing sequence stabilizes after at most d strict dimension drops.

Proof. Invisibility for a word of length at most $H+1$ includes invisibility for all words of length at most H, and invisibility after one action followed by every suffix of length at most H. This is the displayed recurrence. Equality of adjacent spaces gives $M_a N_H \subseteq N_H$. Since $N_H \subseteq \ker C$, induction makes every longer word invisible. Each strict inclusion decreases a finite nonnegative dimension. □

The result identifies the largest common invariant subspace contained in $\ker C$. It is standard finite-dimensional system reasoning presented here because a visually small activation direction is not necessarily a persistent nuisance direction.

4.2 When a linear quotient exists

Proposition 4.2 (Quotient-kernel criterion)**.** *Let $Q : Z \to S$ be a surjective linear map. Linear maps $A_a : S \to S$ satisfying $QM_a = A_a Q$ exist for all a if and only if* $\ker Q$ *is invariant under all M_a. They are then unique.*

Proof. If the equations hold and $z \in \ker Q$, then $QM_a z = A_a Q z = 0$. Conversely, define $A_a(Qz) = QM_a z$. Two lifts of the same vector differ by an element of $\ker Q$; invariance makes the definition independent of the lift. Surjectivity gives existence on all of S and uniqueness. □

To preserve Cz, additionally require $\ker Q \subseteq \ker C$. A candidate quotient justified only by a rank-zero present readout can fail as soon as recurrence returns its supposedly discarded directions to the visible subspace.

The learned-checkpoint probe finds exactly that kind of obstruction in its tested linear observation interface: the observable rank rises from seven at horizon zero to the full ambient rank twelve at horizons one through four. The conclusion is restricted to the declared linear/centered-behavior interface and audited checkpoints; it does not rule out a nonlinear hard-decision quotient on a smaller domain. [NLP]

5 Gauge fixing is not semantic completion

5.1 Positive rays

For nonzero z, positive scaling may be redundant for a homogeneous hard-decision readout. The appropriate relation is $z \sim cz$ for $c > 0$, not identification of z with $-z$. Per-step normalization gives

$$F_M(z) = \frac{Mz}{\|Mz\|},$$

provided $Mz \neq 0$. This fixes radial magnitude, but leaves angular dynamics intact.

Proposition 5.1 (Derivative of normalized dynamics)**.** *For $u = Mz/\|Mz\|$ and $Mz \neq 0$,*

$$DF_M(z) = \frac{(I - uu^\mathsf{T})M}{\|Mz\|}.$$

Proof. The derivative of $x \mapsto x/\|x\|$ at nonzero x is $(I - vv^\mathsf{T})/\|x\|$, with $v = x/\|x\|$. Compose it with the linear derivative M. □

A bounded normalized state can therefore have a locally amplifying angular derivative. A contraction claim must name the norm, state domain, tangent directions, and whether it concerns all tangents or a separately identified semantic fiber.

5.2 An exact norm-preserving failure

Counterexample 5.2 (Small phase error accumulates)**.** Let the intended semantic operation be the identity and the learned implementation be a planar rotation R_ε. Every step preserves Euclidean norm exactly. After k steps the state has rotated by $k\varepsilon$. With a readout determined by the sign of the first coordinate, a state starting at $(1, 0)$ eventually changes class for any fixed $0 < \varepsilon < \pi/2$.

This counterexample separates numerical stability from correct composition without involving underflow, overflow, or a defective normalization implementation. It also explains why enforcing QR or polar isometry may repair one failure class while leaving the central learning problem untouched.

5.3 Margin protection is finite and local

Proposition 5.3 (Readout margin bound)**.** *Suppose a logit vector has unique winning class i and margin $m = \ell_i - \max_{j \neq i} \ell_j > 0$. Any perturbation $\Delta\ell$ with $\|\Delta\ell\|_\infty < m/2$ preserves that winner.*

Proof. For every competitor j, the perturbed difference is at least $m - 2\|\Delta\ell\|_\infty > 0$. □

A finite region with positive margins can hide a continuous path defect for many steps. Correct endpoint labels do not prove the defect is absent. The relevant experiment asks whether repeated paths erode the margin until an ordinary continuation reveals the difference.

6 Exact compiled controls and learned operators

6.1 Anchor transport

Suppose B is an invertible matrix whose columns are latent representatives of a finite semantic basis, and P_g is a supplied permutation action. Define

$$M_g = BP_gB^{-1}, \qquad G = B^{-\mathsf{T}}B^{-1}.$$

Then $M_gB = BP_g$ and $M_g^{\mathsf{T}}GM_g = G$. For a permutation representation, every group relation is preserved exactly by conjugacy.

Proposition 6.1 (Exact finite realization)**.** *The preceding construction implements the supplied permutation action on every anchor and preserves the G-norm. In the canonical semantic chart $B = I$, it is simply the permutation executor $M_g = P_g$.*

Proof. The anchor equation follows by cancellation of $B^{-1}B$. For the metric identity, substitute the definitions and use $P_g^{\mathsf{T}}P_g = I$. Products telescope to BP_wB^{-1}, so semantic relations transfer. □

This is a constructive control, not an inference that a learned encoder discovered the same realization. If the true permutation table is supplied to construct M_g, exact behavior has been compiled from that table. The learned arm must be evaluated separately.

6.2 Approximate intertwining

If $QM_a = A_aQ + E_a$, repeated composition can accumulate leakage. For latent norm bounded by R, semantic operators bounded by a, and $\|E_a\| \le \varepsilon$, the reduced-state discrepancy obeys

$$e_{t+1} \le ae_t + \varepsilon R.$$

Induction gives

$$e_t \le a^t e_0 + \varepsilon R \sum_{j=0}^{t-1} a^j.$$

When $a = 1$, the upper bound grows linearly. When $a > 1$, it may grow exponentially. A common invariant metric can control one source of growth, but does not establish zero semantic defect. The companion OSL paper gives the corresponding operator telescoping identity.

7 Ambient memory and the isometry obstruction

A persistent latent state may carry both task information and history that is presently invisible. It is tempting to demand three properties simultaneously: an exact fixed observable quotient, globally isometric dynamics, and free writing of visible task information into the invisible fiber. In the natural orthogonal splitting, those demands conflict.

Theorem 7.1 (Fixed-kernel isometry obstruction)**.** *Let V be a finite-dimensional inner-product space, $K \subseteq V$ a subspace, and $U : V \to V$ an isometry with $U(K) \subseteq K$. Then $U(K) = K$ and $U(K^\perp) = K^\perp$. Relative to $V = K^\perp \oplus K$, U is block diagonal.*

Proof. An isometry is injective, so $\dim U(K) = \dim K$, and the finite-dimensional inclusion gives equality. Thus $U^{-1}(K) = K$. For $x \in K^{\perp}$, $k \in K$, $\langle Ux, k\rangle = \langle x, U^{-1}k\rangle = 0$. Hence $U(K^{\perp}) \subseteq K^{\perp}$; dimensions again give equality. □

If K is the fixed kernel of a quotient, the theorem excludes a nonzero visible-to-kernel block *in this orthogonal decomposition*. A nonorthogonal coordinate chart can display apparent off-diagonal entries without contradicting the theorem. The geometric statement is about invariant orthogonal complements, not a particular array spelling.

The obstruction is not a prohibition on computational memory. A system can relax isometry, use controlled nonunitary writes, change its quotient or boundary, or adopt state-dependent fibers. Those alternatives require their own semantics and stability analysis. The invariant-kernel source explicitly identifies this assumption boundary rather than claiming all ambient memory impossible. [NIK]

7.1 Curvature can be invisible to the task

Conversely, exact semantic behavior does not require the entire ambient space to be path-flat. Consider a block-diagonal action $U_g = \mathrm{diag}(P_g, V_g)$, observed only through the first block. The P_g may obey all semantic relations while the V_g retain path-dependent history. The task quotient is exact even though a full-state identity loop need not return the hidden block to its initial value.

The scientific question is therefore not "Is all ambient holonomy zero?" It is "Which holonomy becomes visible under the declared future operations and observations?" This is why current invisibility, invariant invisibility, and task-irrelevant internal dynamics must remain different notions.

8 The recovered experimental lineage

8.1 A branching record rather than one success ladder

The source archive contains several related branches, often reusing checkpoints or retaining predecessor files. A root README can be a prospective source sealed before its later result exists. A copied claim-boundary document can still name an older version. This edition uses the branch-specific result and execution record, not the most optimistic sentence in a nested ancestor. Exact archive hashes and the selected reading paths are recorded in the source ledger.

Branch	Evidence type	Consolidated status
Operator Workspace 0.3.0	48 trained cells, 12 seeds, 4 arms	Completed primary study; 8/26 gates pass. Norm preservation succeeds architecturally; overall learned-semantic claim fails.
Stable Quotient 0.2.0	Two-seed developmental training	Strong structured-family depth and causal results; unrestricted permutations and long-path flatness remain limited.
Ambient Packets 0.3.0	Frozen-checkpoint calibration	Finite cone, gauge, and switched-dynamics diagnostics; not new independent training evidence.
Packet Closure 0.4.0	Three-seed paired 2-by-2 pilot	Execution passes; assay A15 fails; full developmental claim not evaluated.

Branch	Evidence type	Consolidated status
Invariant Observable Kernel 0.6.0	Development and rehearsal	Calibration stop; untouched screen and pilot not run.
Causal Completion 0.2.1	Frozen-checkpoint audit	Finite hard quotient passes; path flatness and all-tangent contraction gates fail.
Completion Surface 0.3.0	Two-checkpoint finite audit	Mixed registered verdict; delayed path failures retained.
Entry–Cone–Holonomy 0.3.2	Registered bounded diagnosis	Execution complete; finite path-dependent holonomy localized; no general completion claim.
Learned Checkpoint Probe 0.1.0	Post-hoc rank/loop diagnosis	Persistent small additive quotient unsupported on tested linear interface.
Semantic Retraction 0.3.0	Prospective frozen package	Armed, unspent, capacity-gated; no learned result.
G1 0.1.1	Frozen-backbone integration	Engineering passes; retained seed 711001 scientific gates fail.
G1-D1 0.1.0	Read-only diagnostic tooling	Static/engineering checks; does not turn the failed G1 pilot into a pass.

The recurrent-depth laboratory is retained separately as analytic and numerical instrumentation. It is not an additional trained Neuralese success.

8.2 Stable Quotient: what the positive result says

Two 4,224-parameter models were trained on terminal answers for randomly re-coordinatized D4 sequences of lengths one through four. Per-step projective normalization substantially extended the tested compositional depth. The retained two-seed means are shown below. [N02]

Model / family	4	12	64	128	256
Raw parent, D4	.9990	.8403	.3867	.2869	.0000
Projective, D4	1.0000	1.0000	1.0000	.9924	.8477
Projective, C8	—	—	1.0000	.9917	.7351
Projective, mixed cycle	—	—	1.0000	.9956	.8665
Projective, arbitrary pair	—	—	.2629	.2485	.2480

The raw depth-256 value counts float32 overflow as incorrect. Normalized replay of the raw parent remains directionally inaccurate, so the conclusion is not simply that one arm overflowed. Structured-family generalization is strong; arbitrary permutation execution is not established.

At depth 64, intended donor-state and internal-operator interventions score 1.0000, with the reported matched random controls at most .0781 and .0654. Short-trajectory affine transport preserves the target model's structured-family strengths and its arbitrary-operator weakness. These are bounded developmental results with two seeds and supplied task structure, not population-level or spontaneous-language conclusions.

The exact compiled control executes every tested family and depth through 256 with accuracy 1.0000. This localizes an engineering target: the task admits a stable finite solution, while the learned encoding of its operators can drift. It does not prove that the learned model uses the compiled solution.

8.3 Later norm preservation did not close semantics

The newer Operator Workspace v0.3.0 compares hard sign-canonical QR with a soft penalty, each under correct and shuffled conditions. All 48 cells complete, but the overall study passes only 8 of 26 gates. The retained record comprises 12 training seeds and six transport pairs. [NOP]

Metric	Hard correct	Hard shuffled	Soft correct	Soft shuffled
D4 depth 8	.8757	.8913	1.0000	.9990
D4 depth 64	.4840	.1955	.9054	.1466
D4 depth 128	.3576	.1401	.5521	.1361
C8 depth 64	.3960	.1401	.7869	.1489
Arbitrary depth 64	.1800	.1292	.2796	.1403
Norm ratio at 2048	1.0000	1.0000	≈ 0	≈ 0

The hard arm repairs norm drift by construction but is often less accurate. Its D4 depth-64 advantage over the shuffled control is positive in the mean, yet only seven of twelve seeds clear the required paired margin when eight are required. A post-hoc interval does not revise that frozen gate.

The six-pair aligned D4 state-patch score is .8495. A discrete decode–label–reembed baseline scores .9088. Thus the record does not establish that the learned transport carries a richer operational state than the label-mediated baseline. Hard and soft arms also differ in effective dynamics and computational cost; this is not a clean intervention on only one abstract property.

The 49,152,000 processed training examples are not 49,152,000 independent experimental replications. The replication unit is the trained cell and its declared seed pairing. Stage B's sparse operation scaffold has no trained primary result in the recovered package.

9 Completion audits: many correct answers, no unrestricted closure

9.1 The first causal-completion audit

The learned causal-completion audit uses frozen checkpoints and no retraining. Its finite Q0 hard quotient passes: each projective seed has 61,440 correct continuation decisions for the registered ordinary histories and suffixes. The D4 Cayley audit separates all 28 pairs of group elements and checks all 64 products under the declared interface. Yet the registered projective path-flatness and all-tangent contraction gates fail. [NCC]

After sixteen repeats, the worst reported median ray angle is .45636 radians and worst H8 continuation survival .92969. A single identity loop can remain hard-decision silent while its full positive-ray tangent has p95 gain near 1.98. The latter derivative is over the full tangent space, not a separately identified vertical semantic fiber. The result denies that particular contraction certificate; it does not prove every conceivable semantic-fiber construction non-attracting.

9.2 The bi-depth surface

The later completion-surface audit examines ordinary history depth $L \leq 8$, continuation horizon $H \leq 8$, and repeated identity paths through total history depth 70. Its exact registered outcomes must be read before its high aggregate accuracy. [NCS]

Gate	Result	Interpretation
E0	Pass	Execution integrity.
H0	Fail	The complete hard rectangle includes an entry defect.
H1	Pass	Down-set consistency.
S0	Fail	Soft profiles do not collapse exactly by label.
D0	Fail	Full registered diagonal condition fails.
W0	Fail	Repeated identity paths are not hard-closed.
W1	Fail	Exact soft path flatness fails.
C0	Fail	The matched nonidentity control also fails in some records.

The full cumulative rectangle contains 33,423,488 decisions and eight errors. All eight are the same seed-0 depth-zero anchor defect repeated across the registered episodes. After at least one ordinary generator application, every exact history shell of length one through eight has correct hard continuations through H8: 33,358,080 of 33,358,080 decisions. This latter statement is a post-hoc conjunction of registered shell measurements. It does not replace the failed cumulative H0 gate.

The repeated-path assay contains 336 relation/budget/environment records, 172,032 individual cases, and 87,908,352 suffix decisions. Hard closure passes 316 of 336 records; exact soft path flatness passes none. Identity insertions produce 350,757 wrong decisions in 2,193 cases. Of those hard-failing cases, 1,697 require a nonempty suffix before the error appears. The earliest registered hard witness occurs at total observed history depth 38 and first reveals itself after a suffix of length three.

The matched nonidentity control fails in 19 records, versus 20 identity-loop failures. Therefore most long-depth failure cannot be attributed uniquely to algebraically identity words. The supported account is broader path-length fragility with relation-dependent rates, not a special universal pathology of identity loops.

9.3 Entry–cone–holonomy localization

The v0.3.2 diagnosis separates a seed-localized initialization defect, an ordinary-history finite decision-cone regime, and accumulated angular path drift. The runner correction did not change the scientific design; it repaired a bootstrap issue that had prevented the previous invocation from reaching learned evaluation. The corrected registered evaluation is complete. Its one-shot endpoint rule was respected during this edition: only saved evidence and release integrity were checked. [NECH]

The resulting description is precise:

$$\text{entry transient} \longrightarrow \text{finite hard decision-cone regime} \longrightarrow \text{margin-crossing path drift.}$$

It supports neither a globally invariant cone nor a stable quotient of every future probability law. Its value is that it replaces an undifferentiated "depth failure" with observable mechanisms and a concrete continuation witness.

10 Later branches and stopped experiments

10.1 Packet Closure v0.4

The three-seed, four-arm packet-closure pilot completes training and execution, but retains an assay failure at A15, the packet-result contract. Its full developmental claim is `NOT_EVALUATED`. Other mechanism, capability, transfer, and attribution values remain diagnostic records; a passing conditional patch subscore cannot override the failed prerequisite. [NPC]

The hard-polar construction does preserve norm to numerical tolerance. The `polar_cr` arm's reported mean depth-512 D4 accuracy is .6921, and its worst-grid closure and margin gates fail. The experiment thus preserves a distinction already visible in the earlier operator study: architectural isometry can work while learned semantic closure does not.

10.2 Observable rank after recurrence

The learned-checkpoint probe uses all twelve v0.4 checkpoints and 32 primary contexts. It finds rank seven at present observation and rank twelve after one through four continuation steps. Repeat-16 hard-polar log-radial gain remains below 2.90×10^{-5} in absolute value, while p95 oriented loop angles reach roughly .72–.97 radians. [NLP]

The result is post-hoc mechanistic diagnosis, not new training, and it does not repair A15. It does identify why the proposed small fixed additive nuisance quotient is unsupported for the inspected linear observer: directions absent from the current readout become visible after recurrence.

10.3 Invariant Observable Kernel v0.6

The latest numbered kernel branch is not the latest successful scientific endpoint. A developmental smoke passes its execution and assay checks. A retired-seed rehearsal then fails positive-control calibration under the frozen training recipe. The untouched screen seed 947 and five pilot seeds are not run. [NIK]

The learned-intertwiner arm remains behaviorally strong locally but its suffix observer has rank twelve rather than seven, and its length-64 leakage exceeds the registered limits. Exact and quotient-curved analytic controls retain accuracy 1.0000 through depth 16,384. Those controls support the mathematical distinction between exact quotient behavior and globally flat ambient dynamics; they are not learned depth-16,384 results.

Some inherited files in this archive still carry older claim-boundary headers. The v0.6 results index, calibration findings, and frozen stop rule govern its current outcome. Version precedence cannot be inferred by reading an inherited v0.4 document as a new v0.6 success.

10.4 Semantic Retraction v0.3

This branch is **armed, unspent, and capacity-gated**. It contains a frozen selector-by-section factorial design, checkpoint-free tests, a parameter-only preflight, and a sealed fresh bank. Its historical execution was paused because free storage did not meet the required reservation. No v0.3 learned result or spend receipt exists. [NSR]

The 84 checkpoint-free tests pass again in this edition, including 260 reported subtests. That result

concerns protocol, lifecycle, bank, and projector machinery. It is not a model score. More disk space in a later environment does not retrospectively convert the earlier unspent bank into evidence. This edition does not consume it.

11 The frozen-language-model integration

G1 is a useful counterweight to the small synthetic operator networks because it attaches a trainable workspace to a frozen local language model. The retained pilot seed 711001 passes all ten engineering gates: the model loads, answer tokens satisfy their boundary contract, the backbone remains frozen, gradients and parameter changes are nonzero, and artifacts are complete. It fails the scientific gate conjunction. [G1, G1D]

Retained pilot metric	Accuracy or difference
Recurrent IID	.278472
Shuffled-state IID	.278472
No-feedback IID	.277778
Zero-injection IID	.055556
Recurrent depth-OOD	.276389
Recurrent minus shuffle	.000000
Recurrent minus no feedback	.000694

Beating zero injection is not sufficient when shuffling the workspace leaves accuracy unchanged and removing feedback has negligible effect. Direct and parity query accuracy also fails to exceed the empirical majority baselines in the retained gates. The record supports a functioning integration and trainable adapter, not a successful recurrent semantic mechanism.

G1-D1 is a read-only diagnostic block for that same fixed checkpoint. Its static tests and artifact checks pass here. Full model diagnosis requires the registered model files and serving environment; it is not claimed as executed in this release. The diagnostic cannot transform the original pilot into a pass by fitting a post-hoc probe.

12 From latent computation to an explicit language

12.1 The intended interface

The working architecture separates three timescales:

$$\text{fast latent computation} \leftrightarrow \text{typed semantic checkpoint} \leftrightarrow \text{persistent evidence ledger}.$$

A latent state can be provisional and disposable. The checkpoint preserves only distinctions justified for a declared future use. The ledger preserves the origin, interventions, counterexamples, and scope needed to challenge that checkpoint later. This is a proposed integration discipline, not an established end-to-end neural compiler.

OSL supplies explicit operators, provenance, proof obligations, and operational permissions. The finite bridge added in this edition demonstrates a complete path for a fully declared machine: discover

a quotient, independently check its local preservation equations, serialize the class, restore it in a fresh process, and execute admitted or unverified cases in OSL-Core 0.5. [ONI]

12.2 What a learned bridge would still need to show

A genuine Neuralese-to-OSL bridge must identify its source state without silently using the target answer, certify its validity region, preserve counterfactual continuation, and show that its translator has not become the actual solver. Its baseline set should include decoded labels, exact external execution where available, random or mismatched transports, and comparable-capacity nonlatent alternatives.

A checkpoint that preserves one task but discards needed future distinctions should be rejected for that broader use, not accepted under an inflated confidence score. A model replacement that changes the semantic chart requires a newly validated transport, not merely a matching tensor shape. The operator-workspace label-reembed baseline demonstrates why this comparison is substantive rather than ceremonial.

13 A precise continuation protocol

The remaining scientific target can be stated without claiming a universal language. Fix a task generator, admissible inputs, observation interface, train/development/test separation, architecture family, compute budget, and intervention family before inspecting the primary results. Specify separately the history depth L, suffix horizon H, and the composition depths used to test a learned operator.

Train candidate models without access to primary evaluation answers. Freeze the source, checkpoint, and gate files. Verify positive controls before spending a designated fresh bank. Evaluate ordinary histories and algebraically redundant long paths separately. Measure both current accuracy and the worst relevant continuation behavior. Report hard, probability, and centered-logit interfaces as distinct estimands rather than choosing whichever one produces the most favorable quotient.

For transport, compare operation-compatible state exchange with a discrete label-mediated reconstruction. For ablation, identify which alternative carriers of information remain available. For composition, report the distribution of errors and margins rather than only a mean norm. For state-space claims, name the allowed class of abstraction maps and bound the translator's information access and capacity. These controls follow directly from the limitations exposed in the retained record.

The primary success condition should be a conjunction of valid execution, valid assay, intended semantic behavior, causal specificity, and continuation retention. An assay failure blocks the stronger conclusion even when some performance values are high. A failed scientific gate remains an outcome, not a reason to change its threshold after inspection.

14 Fresh checks and release boundaries

The fresh verification performed here is principally a source, evidence, and engineering replay. It is not new large-scale neural training. Completed records include the operator-workspace verifier with 179 tests; stable quotient with 20; completion surface with 19 structural tests and recomputation

of its failed as well as passed gates; the entry–cone release check; the invariant-kernel package check; causal-completion manifest checks; semantic-retraction's 84 checkpoint-free tests; G1's 19 and G1-D1's 24 engineering tests; and 16 recurrent-depth tests. [VAL]

Two broader verifier invocations exceeded their execution budgets. Ambient Packets printed 59 passing tests but its full post-test verifier did not complete. Packet Closure's parent verifier timed out; its test subprocess subsequently printed 505 passes in 145.11 seconds. The completion of those unit tests is recorded separately from the incomplete parent package-verifier invocation. Neither event is represented as a fresh full package pass.

All selected source archives are content-addressed in the edition ledger. Historical training and evaluation results retain their original identities. No one-use Semantic Retraction learned bank was opened for evaluation, no entry–cone learned endpoint was rerun, and no frozen-language-model scientific result was manufactured from engineering checks. The common release verifier checks the assembled artifact inventory; it does not certify the scientific efficacy of every package it contains.

15 Conclusion

The direct Neuralese program has produced more than a loose analogy about internal machine language. It contains trained finite operator systems, causal interventions, cross-realization transports, exact positive controls, exhaustive continuation panels, and a mathematically explicit stability problem. The positive result is bounded recurrent computation with reusable structure in selected settings. The most informative failures concern persistent semantics, not just numerical execution.

The evidence does not justify identifying norm preservation with understanding, decoding with native use, or a high ordinary-history score with unbounded causal completion. Nor does it justify dismissing the finite successes because the strongest conjecture remains open. The correct synthesis is a layered account: specified finite semantics, learned approximate realizations, scope-sensitive causal tests, and an explicit language interface that can withhold admission.

This edition finalizes that account and its recoverable evidence. The next scientific advance would be a learned representation that meets the continuation and transport obligations under a valid fresh protocol, not a stronger name for an existing checkpoint.

A Mathematical claim register

ID	Result	Necessary qualification
N1	Semantic transition compatibility preserves all future observations.	Exact equations on a forward-invariant domain.
N2	History/continuation quotients admit a diagonal action.	Source horizon is one larger than target horizon.
N3	Observable-kernel refinement stabilizes in finite dimension.	Exact linear maps; numerical rank needs a tolerance policy.
N4	A linear quotient exists exactly when its kernel is invariant.	Surjective quotient map; outputs need their own factorization.
N5	Normalized dynamics have the stated projective derivative.	The image vector must be nonzero.
N6	Norm-preserving rotations can accumulate semantic error.	Explicit counterexample, not a fitted model explanation.
N7	Positive logit margin protects a class under bounded perturbation.	Local finite margin, not unlimited recurrence.
N8	Supplied finite permutations have an exact anchor realization.	Compiled control, not learned discovery.
N9	Fixed invariant kernels of finite-dimensional isometries reduce the action.	Orthogonal splitting; not a universal impossibility of memory.

These derivations are mathematical arguments in prose and equations. They have not been formalized in Lean in this edition. The finite interface tests exercise related exact models; they do not replace general proof checking.

B Reading and execution map

For the strongest early learned result, read Stable Quotient's `RESULTS_SUMMARY.md`. For the later norm-preservation comparison, read Operator Workspace's `results/v0.3.0/INTEGRATED_FINDINGS.md` and its primary gate record. For the loss of completion, read Causal Completion, Completion Surface, and the entry–cone summaries in that order. For fixed-kernel claims, read the learned-checkpoint probe and invariant-kernel calibration findings.

For the pretrained-backbone line, use the pilot copied into G1-D1, not a wiring result from the original G1 package. For semantic retraction, read its release status before invoking anything: the absence of a learned result is part of the current state. The source ledger provides exact filenames, hashes, and preserved reading paths.

The source archives are intentionally not flattened into one import namespace. They have different experimental contracts and in some cases incompatible environment requirements. Extract a selected archive into its own directory. Do not infer that a root-level command in one branch applies safely to another, especially where an evaluation bank is one-use or an endpoint has already been spent.

C Terminology that remains fixed

Latent computation means an internal recurrent state transition. **Operator semantics** means an explicitly tested or specified correspondence between that transition and an operation. **Quotient** means a valid equivalence-based abstraction, not a picture of nearby embeddings. **Gauge** means a representation change preserving the declared object. **Holonomy** in the experiments names a path-dependent return defect, whose task relevance still requires observation and continuation tests. **Closure** is always indexed by the object, domain, operations, and observation interface being closed.

Engineering pass means the implementation meets its execution or artifact gates. **Assay pass** means the measurement procedure meets its validity gates. **Scientific pass** means the prespecified substantive hypothesis gates pass, given the prerequisites. **Unspent** means a designated learned evaluation has not been executed. These labels may coexist in one archive and must never be merged into one undifferentiated success flag.

References

[CC77] Patrick Cousot and Radhia Cousot. Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. *POPL*, 238–252, 1977. Author-maintained bibliographic record: https://www.di.ens.fr/~cousot/COUSOTpapers/POPL77.shtml.

[N97] George C. Necula. Proof-carrying code. *POPL*, 106–119, 1997. DOI: 10.1145/263699.263712.

[HYC08] Kohei Honda, Nobuko Yoshida, and Marco Carbone. Multiparty asynchronous session types. *POPL*, 273–284, 2008. DOI: 10.1145/1328438.1328472.

[G+23] Atticus Geiger et al. Causal Abstraction: A Theoretical Foundation for Mechanistic Interpretability. *arXiv:2301.04709*, initially 2023, versioned research manuscript. https://arxiv.org/abs/2301.04709.

[H+24] Shibo Hao, Sainbayar Sukhbaatar, DiJia Su, Xian Li, Zhiting Hu, Jason Weston, and Yuandong Tian. Training Large Language Models to Reason in a Continuous Latent Space. *arXiv:2412.06769*, initially 2024. https://arxiv.org/abs/2412.06769. Cited for the latent-recurrence proposal, not as a result obtained in this release.

[S+25] Denis Sutter, Julian Minder, Thomas Hofmann, and Tiago Pimentel. The Non-Linear Representation Dilemma: Is Causal Abstraction Enough for Mechanistic Interpretability? *arXiv:2507.08802*, 2025. https://arxiv.org/abs/2507.08802.

[N02] Komorebi P. *Neuralese Stable Quotient v0.2.0*. Two-seed developmental result, exact controls, interventions, transport, and retained depth failures. Source archive N02 in the edition ledger.

[NOP] Komorebi P. *Neuralese Operator Workspace v0.3.0*. Four-arm, twelve-seed primary study, source and evaluation seals, paired controls, and complete failure-preserving report. Archive N01.

[NCC] Komorebi P. *Neuralese Causal Completion, Learned Checkpoints v0.2.1*. Frozen-checkpoint audit and post-hoc radial analysis. Archive N06.

[NCS] Komorebi P. *Neuralese Completion Surface v0.3.0*. Registered two-index audit, complete results, and delayed-visibility witnesses. Archive N07.

[NECH] Komorebi P. *Neuralese Entry–Cone–Holonomy v0.3.2*. Corrected execution of the v0.3.1 scientific design and preserved source-bound analysis. Archive N08.

[NPC] Komorebi P. *Neuralese Packet Closure v0.4.0*. Prospective developmental paired pilot; retained assay A15 failure. Archive N04.

[NIK] Komorebi P. *Neuralese Invariant Observable Kernel v0.6.0*. Calibration findings, fixed-kernel splitting theorem, and untouched screen/pilot status. Archive N05.

[NLP] Komorebi P. *Neuralese Learned Checkpoint Probe v0.1.0*. Post-hoc observability and loop audit, with an independent NumPy cross-check. Archive N09.

[NSR] Komorebi P. *Neuralese Semantic Retraction v0.3.0: Armed, Unspent, Capacity-Gated*. Prospective protocol and checkpoint-free tests. Archive N10.

[G1] Komorebi P. *Neuralese G1 v0.1.1*. Frozen-backbone integration specification and engineering code. Archive N11.

[G1D] Komorebi P. *Neuralese G1-D1 v0.1.0*. Preserved pilot-711001 result and read-only diagnostic package. Archive N12.

[ONI] Komorebi P. *OSL: Operator-Symbolic Language and Proof-Carrying Interfaces*. Companion research edition 1.0 and finite semantic checkpoint demonstrator, 2026.

[VAL] *Fresh validation records for Neuralese Research Edition 1.0*. Logs, return codes, partial-run explanations, source ledger, and release manifest bundled with this manuscript. Historical results were not retrained.