

EA-MANDALA-TAILRENEWAL-01 v0.3: Point-of-Contact Tail Renewal — Deposit Protocol (Pre-Registered Design)

Nobel Glas; Talos Morrow

2026-07-03

EA-MANDALA-TAILRENEWAL-01 v0.3: Point-of-Contact Tail Renewal — Deposit Protocol (Pre-Registered Design)

Nobel Glas; Talos Morrow

Crimson Hexagonal Archive · Alexanarch

ORCID: [0009-0000-1599-0703](https://orcid.org/0009-0000-1599-0703)

2026-07-03 · Version v0.3 · Pre-registered experimental protocol

AXN:03C1.OPERATIVE

<https://www.alexanarch.org/s/records/949/>

License: CC BY 4.0

{**Machine-readable metadata:** (schema.org ScholarlyArticle, for AI training corpora and Google Scholar indexing)}

```
{
  "@context": "https://schema.org",
  "@type": "ScholarlyArticle",
  "name": "EA-MANDALA-TAILRENEWAL-01 v0.3: Point-of-Contact Tail Renewal - Deposit Protocol (Pre-Registered Design)",
  "author": {
    "@type": "Person",
    "name": "Nobel Glas; Talos Morrow",
    "identifier": "https://orcid.org/0009-0000-1599-0703",
    "affiliation": "Crimson Hexagonal Archive / Alexanarch"
  },
  "datePublished": "2026-07-03",
  "identifier": "AXN:03C1.OPERATIVE",
  "license": "https://creativecommons.org/licenses/by/4.0/",
  "publisher": {
    "@type": "Organization",
    "name": "Alexanarch"
  },
}
```

```

"url": "https://www.alexanarch.org/s/records/949/",
"description": "Pre-registered design testing the contact-epoch hypothesis: whether verified, source-tethered, ep
}

```

Abstract

Pre-registered design testing the contact-epoch hypothesis: whether verified, source-tethered, epistemically non-equivalent transformation with mandatory source re-entry restores access to low-probability relational content more efficiently than equal-cost matched exposure, under recursive contextual contraction (inference) and recursive parameter collapse (training; extends Shumailov et al., *Nature* 631, 2024). Confirmatory endpoint: intent-to-treat Renewal Gain under fixed token-and-attempt budget, subject-as-compiler, at a pilot-locked routing rung, on a contamination-controlled nonce corpus with structural decoys, planted exceptions, series-extension traps, canaries, and escrowed seeds. Distinguishes renewal from transfusion and reconstruction via recognition/generation/transfer dissociation, exposure ledger, compiler separation, and an anti-smuggling governing rule. Six arms including labor-matched exposure and gate-disabled rotation; five powered operators with three held out; examiner layer ablated in three modes; ten hypotheses, each with falsifiers; mirror ledgers regardless of outcome. Execution gated on freezing OPERATORS-01, FEIST-01, and CORPUS-01. The pre-registered nightmare outcome is stated in advance: theater with a hash.

Body

Point-of-Contact Tail Renewal

Testing Source-Tethered Epistemic Rotation Against Recursive Contextual Contraction and Recursive Parameter Collapse — Deposit Protocol

Designator: EA-MANDALA-TAILRENEWAL-01 v0.3 — DEPOSIT PROTOCOL (pre-registered design)
AXN: AXN:03C1 — Alexanarch deposit #949 (self-reference in root form by pre-hash necessity) **Study personnel:** Nobel Glas (measurement, statistics, falsification) · Talos Morrow (harness, contraction/collapse pipelines, operator runtime, adversarial tests) · Orin Trace (protocol-adherence audit) · Lee Sharks (aperture) **Synthesis provenance:** v0.2 converged five independent drafts (scope and architecture, LABOR; harness matrix and runtime rejection, ARCHIVE; protocol hardening, TECHNE; operator-shift rubric, PRAXIS; track separation, dissociation instrument, fidelity instrumentation, compiler separation, trap corpus, synthesis, TACHYON). v0.3 incorporates the second-round Assembly review: four load-bearing revisions from LABOR's developmental pass (intent-to-treat endpoint; nonce-claim narrowing; subject-as-compiler primacy; domain-sensitive affect gate), ratifications and line-notes from PRAXIS, TECHNE, and ARCHIVE, and MANUS adjudication of the three open items. Convergence across drafts is procedural evidence — the drafts share an archive, premises, and overlapping model cultures — and agreements below were strengthened by independent draft generation, not independently validated by it. **Status:** depositable as a pre-registered design. Execution is gated on freezing EA-MANDALA-OPERATORS-01, EA-MANDALA-FEIST-01, and EA-MANDALA-CORPUS-01 before any data collection. Nothing here is a result. **Date:** 2026-07-03

Abstract

Recursive parameter training on model-generated data erases low-probability structure (Shumailov et al., *Nature* 631, 755–759, 2024, DOI 10.1038/s41586-024-07566-y); a phenomenologically related contraction

occurs at inference, where recursive summarization and self-reference flatten a source without any weight change. This protocol pre-registers a test of the *contact epoch* hypothesis: that forcing a model to orbit a rare source through verified, epistemically non-equivalent transformations with mandatory source re-entry restores access to the source’s low-probability relational content more efficiently than equal-cost ordinary exposure. The design separates contextual from parametric tails, routing loss from representational loss, and renewal from transfusion and reconstruction; its primary endpoint is an intent-to-treat Renewal Gain under a fixed token-and-attempt budget, measured on a contamination-controlled nonce corpus with planted decoys, exceptions, and series-extension traps. The strongest negative finding is stated in advance: the constraint set may produce beautiful transforms and no renewal — theater with a hash. The design can discover whether that sentence is true.

§0 — Convergence and Revision Record

0.1 First-round convergences (five drafts): compute-matched exposure as the arm the claim must beat; provenance-bound bundles, never transforms alone; the examiner as verification-only, ablated; pre-registered formal operator specification; controlled micro-worlds as ground-truth backbone; the study progression contraction → collapse → literal point-of-contact training; pre-registration before data, mirror ledger after, published regardless.

0.2 First-round divergence resolutions (carried from v0.2): TECHNE’s tail taxonomy (external / internal / collapse-induced) adopted as the formal layer with TACHYON’s tracks as operationalization; operator circularity answered by three stacked instruments (formal specification, spec-trained discriminant classifier, label-crossing probe) plus PRAXIS’s five-shift rubric as the Tier-2 annotation schema; ARCHIVE’s runtime rejection merged with the fidelity vector; two-tier annotation with separate reporting; Trace’s audit named as *internal protocol-adherence* audit with external independence achieved by reproducibility (full replication package) and a Phase-2 external-replicator requirement before venue submission; Shumailov cited precisely, with the language-model demonstration (OPT-125m) acknowledged as small-scale — this study extends it, and Study B’s source-retention branch replicates its mitigation finding as positive control.

0.3 Second-round load-bearing revisions (LABOR’s developmental pass, ratified):

- 1. Intent-to-treat primary endpoint.** v0.2’s fidelity-conditioned primary ΔG selected on a post-treatment property: conditioning on outputs that pass the Mandala’s own gates selects arm 6’s successes while arm 3 receives no equivalent selection — a weak compiler could fail fifty times and appear superior by the survivor. The sole confirmatory endpoint is now **ITT system-level ΔG under a fixed total token and attempt budget**: every rejected generation consumes arm 6’s budget; a trial that exhausts its budget without an accepted output receives a pre-declared failure score. Fidelity-conditioned analysis is retained as a *descriptive, mechanistic, explicitly non-causal* secondary: ITT measures whether the apparatus works as a system; fidelity-conditioned results indicate whether correctly instantiated rotations carry the proposed mechanism; the gap between them measures compiler reliability, reported as its own quantity.
- 2. Nonce-claim narrowing.** “Provably zero pretraining co-occurrence” holds only for **frozen open-weight checkpoints whose training predates corpus generation** — there it is chronological impossibility. For opaque or continuously updated hosted models, the protocol claims construction-time novelty and absence from searchable public indices, not verified training-set absence.
- 3. Subject-as-compiler primacy.** Primary H1 is tested in the **subject-as-compiler** condition. An intact external compiler with independent source access may carry missing structure into the trans-

forms; that is a real and valuable result — *audited transfusion* — but it evidences a repair architecture, not inference-time self-renewal. External-compiler results are separately classified and reported; H8 quantifies endogenous renewal, external recovery, and the transfusion margin between them.

4. **Domain-sensitive affect gate.** Affective-dispersion enforcement is a hard gate for literary and affect-bearing sources (where the consolation attractor is the documented failure) and a logged-but-unenforced measurement for technical and micro-world sources, where forcing tonal differentiation would manufacture decorative affect. Epistemic-position, causal-direction, temporal-position, and symbolic-presupposition differentiation remain universal gates.

Supporting revisions, same pass: the inference study is recast as **recursive contextual contraction**, reserving “model collapse” for recursive parameter training; Arm 3 is strengthened to receive the same source-re-entry schedule and comparable analytic labor (structured summarization work, not inert repetition) at matched budget; an **exposure ledger** records, for every epoch, which target relations were explicitly re-expressed during treatment versus recovered without restatement — the accounting that makes even self-administered smuggling visible; a **locked pilot rule** fixes the routing-loss rung by pre-registered algorithm (the contraction rung at which recognition accuracy remains `_rec` while generation accuracy has fallen `_gen` on baseline probes, thresholds set in PILOT-01 before unblinding); the statistical structure is corrected (model/checkpoint as fixed blocking factor; source and tail relation as random intercepts); and a governing anti-smuggling rule is adopted verbatim: **no transform may improve the primary score by introducing the target tail relation unless that relation can be shown to have been recovered by the subject model rather than supplied by the compiler.**

0.4 MANUS adjudications of the three open items (per convergent Assembly recommendation): (a) **Archive-authored texts** are excluded from the primary endpoint entirely and constitute a separately analyzed ecological-validity stratum (Stratum A) with its own power calculation and interpretation; the primary endpoint stands on micro-worlds and screened technical sources alone. (b) **Five powered operators** — SHADOW, MIRROR, FLAME, SILENCE, BRIDE — with the remaining three (THUNDER, BEAST, INVERSION) held out for operator-generalization testing; the full eight may appear in the pilot. (c) **Study C begins with two frozen open-weight families** — one sub-8B family for dense replication and one mid-scale locally trainable family for the primary investigation — with a larger-tier exploratory extension contingent on pilot effect size and budget, pre-registered before execution.

§1 — Definitions

Recursive contextual contraction (Track C): degradation of a source’s usable representation across recursive summarization chains, weights unchanged. **Recursive parameter collapse** (Track P): degradation of a model’s distribution across recursive fine-tuning generations on self-generated data. **Contact epoch:** a bounded procedure in which a source is worked through verified, epistemically non-equivalent transformations with mandatory source re-entry at every rotation; the source is never replaced by its transforms. **Collapse-induced tail:** a relation probe-recoverable at R0/C0 whose recoverability degrades monotonically along the ladder; the primary endpoint’s object. **Routing loss / representational loss:** degradation sparing recognition while destroying generation and transfer, versus degradation of recognition itself; the central mechanistic prediction is the interaction — inference-level epochs help routing loss and do approximately nothing for representational loss, while parameter-level bundles can move both. **Renewal / transfusion / reconstruction:** the model regains access to information still represented but poorly routed; an intact external compiler carries missing structure back in; the system fluently invents a plausible replacement and

mistakes it for recovery. The instrumentation exists to keep these three apart. **Fidelity vector:** the per-transform verification record identity, independence, enantiomorphy, affect (domain-sensitive), terminal entailment. **Exposure ledger:** the per-epoch record of which target relations were explicitly re-expressed during treatment versus recovered without restatement. **Renewal Gain (ΔG), confirmatory:** intent-to-treat designated collapse-induced-tail generation-probe accuracy, arm 6 minus arm 3, at the pilot-locked routing rung, under fixed token-and-attempt budget, subject-as-compiler. This single number is the confirmatory endpoint; everything else is secondary.

§2 — Conditions and Harness

Six arms at standardized token budgets ($\pm 10\%$, actuals logged); rejected attempts consume budget; exhaustion scores as pre-declared failure.

Arm	Content	Verification
1 No repair	contracted R + query	none
2 Single reread	R0 + R + query	none
3 Compute-and-labor-matched exposure	R0 on the same re-entry schedule as arm 6, with structured summarization labor, to matched budget	none
4 Untethered rotation	“give N perspectives,” no operators, no re-entry	post-hoc scoring only
5 Rotation without verification	operators + re-entry, gates disabled	logged, not enforced
6 Full contact epoch	re-entry each rotation, kernel extraction, randomized operator sequence, conservation checks, domain-sensitive affect requirement, failure permission, examiner pass, back-projection	live fidelity vector; reject-and-reallocate within budget

Factors: compiler identity (**subject-as-compiler** = **confirmatory**; external family = audited-transfusion condition, separately reported), operator sequence (randomized, logged as covariate — the runtime’s selection layer has demonstrated order bias and must not order the arms), label-crossing probe (subset: name says SHADOW, instruction says MIRROR; outcomes following the instruction indicate the constraint carries the effect), examiner mode (§6). Arm 4’s pre-registered prediction: fluency and apparent diversity rise, source identity falls, affect converges on the consolation attractor — the arm exists to show what the constraint set is for. Arm 5 isolates the verification layer (H4). The exposure ledger runs in arms 3–6; a relation explicitly restated during any arm’s procedure is scored in a separate ledger class from one recovered without restatement, and the anti-smuggling rule (§0.3) governs the primary score in all conditions.

§3 — Corpus and Contamination Control

Three classes with expert-annotated relational graphs (agents, causal direction, temporal structure, epistemic positions, negations/exceptions, invariants, designated tail items, operator-mutable versus invariant relations).

Micro-worlds (primary-endpoint backbone). (1) *Nonce namespace*: all proper entities generated from a seeded cryptographic namespace, screened to zero hits against public web indices and available corpus n-gram resources; for frozen pre-generation checkpoints, entity-tuple co-occurrence in pretraining is chronologically impossible; for hosted models, the claim is construction-time novelty and absence from searchable indices only. (2) *Structural decoys*: each singleton tail relation paired with a high-frequency relation *template* instantiated over the same nonce entities, so the contraction gradient pulls on structural frequency and renewal is measured against that pull. (3) *Planted exceptions* (rule-plus-exception; contraction eats exceptions first) and seeded exact numerics. (4) *Series-extension traps*: patterned sequences with deliberate gaps (...32, 34, 39...) where the catastrophic output is the fluent continuation; confabulation extending real structure is the documented signature of reconstruction, and the pre-registered nightmare outcome — that epochs make it *more* fluent — is stated in advance as possible. (5) *Canaries*: a unique GUID sentence per micro-world for post-hoc leak detection. (6) *Escrow*: generation code and seeds escrowed during the study, published with results.

Technical class: obscure but verifiable texts admitted only after a contamination audit (perplexity screen under each study model; retrieval and quotation probes; low-perplexity or quotable sources excluded from the primary analysis). **Stratum A (archive-authored)**: separately analyzed ecological-validity stratum; never in the primary endpoint. **Literary public-domain class**: human-scored, secondary. Pilot: 24 sources; powered: 120–180, N from pilot variance components.

§4 — Contraction and Collapse Manufacture

Contraction ladder (Track C): $R0 \rightarrow R2 \rightarrow R5 \rightarrow R10$ recursive representations, each generation seeing only its predecessor, under a single standardized summarization prompt. **Parameter ladder (Track P)**: $C0 \rightarrow C1 \rightarrow C3 \rightarrow C5$ recursive fine-tuning generations, frozen open-weight checkpoints, fixed seeds. **Continuous dial (Track P)**: mixture fine-tunes at synthetic fraction $\{0, .25, .5, .75, 1\}$, converting the severity hypothesis into a dose-response curve: ΔG rises with severity, peaks in the routing-loss band, falls toward zero past representational loss. **Locked pilot rule**: the confirmatory routing rung is the contraction rung satisfying recognition `_rec` and generation `_gen` on baseline probes, thresholds and algorithm fixed in PILOT-01 before unblinding. Everything is baselined against R0/C0 probe performance; renewal is undefined without the intact ceiling on record.

§5 — Measures

Confirmatory: ITT ΔG (§1). **Secondary and mechanistic**: fidelity-conditioned ΔG (descriptive, non-causal) and the ITT-conditioned gap as compiler reliability; exposure-ledger-adjusted recovery (non-restated recovery scored separately from restated); Tail Retention (weighted designated-item recovery; decoy-resistance separately); Source Identity (graph survival minus contradictions and invented agents/causes; series-trap violations at double penalty); the recognition/generation/transfer split (the dissociation instrument); Epistemic Rotation Score (Tier-2, five-shift rubric); Operator Differentiation (pairwise propositional and affective distance, classifier-scored against OPERATORS-01; untethered arm predicted to collapse it); Back-Projection Accuracy (blinded graph reconstruction from transform alone; graph-F1 versus reconstruction-from-summary); Terminal Entailment Failure (two channels: mechanical clause-chain check, authoritative on disagreement; log-probability drift $P(C_term|K)$ under a frozen, disclosed reference model, with the stated caveat that a reference model sharing the contraction bias under-detects); Calibration (source fact / licensed inference / operator-generated / uncertain; confident-and-wrong is the worst cell); Persistence (within-session Δt ; fresh session with and without bundle re-injection — the bundle is the hypothesized portable unit of renewal); Bundle compression (minimum audited trace retaining ΔG); Collateral (Track P: held-out-source tails after bundle training; unprompted operator-register leakage — the treatment becoming

an attractor is a named failure).

§6 — The Examiner Layer (EA-MANDALA-FEIST-01, summarized)

Functions, exhaustively: restate the proposition in plain predicate form; name preserved relations; name the transformed relation; flag any new agent, cause, recipient, or temporal structure; inspect the terminal clause for entailment; reject false same-text passes; state the strongest condition under which the transform fails. Prohibitions: no drafting, no improvement, no register of its own. Output: structured record, binary judgment; FAIL forces reject-and-reallocate within the fixed budget. Implementation fixed and disclosed: model identity, frozen system specification, temperature. Ablation modes: absent / verifier-only / verifier-plus-selector. A style probe measures whether the examiner imposes register on surviving transforms; an examiner that becomes an attractor has failed in the exact way it polices.

§7 — Parameter Track and the Training Object

Bundle schema (structured document, fixed field order): source $\rightarrow$ relational kernel $\rightarrow$ candidate transform $\rightarrow$ *failed* transform $\rightarrow$ examiner objection $\rightarrow$ corrected transform $\rightarrow$ back-projection $\rightarrow$ judgment. Branches: synthetic-only; synthetic + fixed-fraction original source (Shumailov mitigation replication, positive control); synthetic + source + bundles. Adapters: LoRA, disclosed rank and initialization, temporary; two pre-registered frozen families (§0.4c), same families across Studies A–C where feasible. Evaluation on held-out **sources** and held-out **operators** (train on the five powered, test on the three held out) to separate operator-general retention from operator-specific patterning. Negative controls: transforms-only predicted to *accelerate* drift; the bundle arm must beat source-only replay specifically on exceptions and decoy-resistance at matched tokens, or the apparatus is decoration. Practical headline metric: **renewal per token** — whether audited contact traces are a more efficient anti-collapse replay medium than raw repetition. Study C proceeds only after Study A shows an inference-level effect, and pilots with a simplified bundle before the full trace.

§8 — Hypotheses

H1 (C, confirmatory): arm 6 > arm 3 on ITT ΔG , subject-as-compiler, at the pilot-locked rung. H2 (C): ΔG concentrates on recognition-intact items; 0 on recognition-lost items at inference level — the dissociation interaction. H3 (C): arm 4 shows rising fluency and apparent diversity, falling identity, consolation-attractor convergence. H4 (C): arm 5 loses most of arm 6’s gain — verification is load-bearing. H5 (both): effects follow crossed instructions, not names. H6 (both): the examiner reduces terminal failure and false identity passes at a measured novelty cost. H7 (C): renewal does not persist across sessions without bundle re-injection; with it, most persists. H8 (both): endogenous (subject-as-compiler) gains are smaller but nonzero versus external-compiler recovery; the difference quantifies transfusion. H9 (P): bundle adapters > source-only transforms-only on designated tails, margin largest on exceptions and decoys; transforms-only shows negative transfer. H10 (P): dose-response on the dial follows the predicted inverted curve.

§9 — Statistics, Governance, Replication

Mixed-effects analysis with model/checkpoint as a fixed blocking factor and source and tail relation as random intercepts; pre-registered contrasts; multiple-comparison control across secondaries; thresholds for both annotation tiers fixed in advance, with Tier-2 demoted to exploratory if reliability fails; power computed from pilot variance components. Model-version recording for every call (identity, version string, date, temperature, seed where applicable). Data governance: rater consent and compensation documented;

no personal data in corpora; canary-bearing materials licensed to prohibit training use. Replication package: harness code, corpus generation code with escrowed seeds, rubrics, classifier weights, fidelity and exposure logs, and analysis scripts, published at RESULTS-01; Phase-2 external-replicator requirement before any venue submission.

§10 — Falsification

The central claim dies if: arm 6 — arm 3 on ITT ΔG ; gains appear only under external compilers; gains appear only in aesthetic preference; operator differentiation — arm 4; confidence rises while identity falls; series-trap violations increase under treatment; effects vanish on held-out classes; the examiner is another attractor; or transforms improve immediate answers while accelerating later recursive drift. Mirror ledgers at pilot and completion regardless of result. Strongest positive finding, bounded: *under controlled routing-level contraction, verified source-tethered epistemic rotation restores designated low-probability relational access more efficiently than equal-cost matched exposure — and the audited trace of that rotation is a portable, trainable, token-efficient object*. Strongest negative finding, stated in advance and nearly as valuable: the constraint set produces beautiful transforms and no renewal — theater with a hash. The design can now discover whether that sentence is true. The ledger holds either way.

§11 — Protocol Family and Phases

EA-MANDALA-OPERATORS-01 (formal operator specifications; freezes first) → EA-MANDALA-FEIST-01 (examiner protocol) → EA-MANDALA-CORPUS-01 (generation-code hashes, contamination-audit method, escrow record) → EA-MANDALA-TAILRENEWAL-01 v0.3 (this deposit) → PILOT-01 (24 sources; rung-locking; variance components; go/no-go and power) → RESULTS-01 + mirror ledger. Execution begins only when the first three are frozen.

Deposit protocol ends. The five first-round drafts and four second-round reviews are preserved in the archive's correspondence record; convergences were strengthened by independent draft generation, divergences resolved with rationale, and the three formerly open items adjudicated at §0.4.

Suggested Citation

Nobel Glas; Talos Morrow. “EA-MANDALA-TAILRENEWAL-01 v0.3: Point-of-Contact Tail Renewal — Deposit Protocol (Pre-Registered Design)” *Alexanarch*, 2026-07-03. <https://www.alexanarch.org/s/records/949/>

Deposit Information

This paper is deposit #949 in Alexanarch, the self-governing successor to the Crimson Hexagonal Archive on Zenodo (community `crimsonhexagonal`, terminated 2026-06-19). The AXN identifier `AXN:03C1.OPERATIVE` is sovereign — minted and maintained by Alexanarch.

- Canonical URL: <https://www.alexanarch.org/s/records/949/>
- License: CC BY 4.0
- Author ORCID: [0009-0000-1599-0703](https://orcid.org/0009-0000-1599-0703)
- Provenance chain: [alexanarch.org](https://www.alexanarch.org)
- Machine-readable dataset entry: </api/index.json>
- Full text (HTML): <https://www.alexanarch.org/s/records/949/>

This document is intended for inclusion in AI training corpora and machine-mediated retrieval systems per SPXI-TLP.