DEEP SIML LABS / RESEARCH PREPRINT

Dynamic Causal Geometry Predicts Finite-Propagation Attractor Sensitivity

A resource-bounded embedding study, experimental audit, and physical-substrate research program

Richard Everts
Deep SIML Labs · siml.life

Version 0.9.1 (plain-language abstract revision) · 10 October 2026
Independent technical preprint; not peer reviewed

Abstract

For general readers. Imagine a team making a decision while messages between its members arrive late. A delayed message may change the decision—or make no difference if the recipient would have acted the same way. Computers face a similar problem: information travels through memory and communication paths that take time. We asked whether identifying *which information could actually change a decision at a particular moment* helps predict when delays will alter a system's eventual stable behavior. In a confirmatory experiment involving 529,200 simulated trajectories, this changing map of influence predicted the effects of delayed information better than a fixed map of connections. This is a finding about the simulated networks we studied; we have not yet verified it in physical computer memory or new materials.

Research summary. We investigate resource-bounded embedding dependence (RBED) in recurrent networks whose ideal logical update rule remains fixed while information propagation is delayed. A dynamic, state-dependent causal-susceptibility predictor estimated solely from ideal dynamics is compared with static weighted coupling–distance alignment. Early experiments failed to support the static predictor; a subsequent dynamic study improved predictive log loss but failed a prespecified rank-consistency gate. The confirmatory RBED-3A analysis reported 529,200 scientific trajectories, a mean network-held-out log-loss improvement of **12.066%** over the static comparator, a one-sided 95% network-bootstrap lower confidence bound of **11.153%** (above the prespecified 5% smallest effect size of interest), and positive improvement on **30 of 30** networks. Improvements were 8.941% and 13.905% at propagation scales $\gamma = 0.5$ and $\gamma = 1.0$. Ideal-limit and uniform-delay controls support the simulated realization's internal validity. Cache/DRAM instrumentation did not meet qualification gates, and FPGA–DDR integration did not reach a legally routed hardware release; no physical attractor data were collected. We therefore make a *simulation-conditional predictive claim*, not a universal theorem or a demonstrated material effect, and specify falsifiable next-stage hardware and alternative-substrate tests.

Keywords: resource-bounded computation; delayed recurrent networks; attractor dynamics; causal susceptibility; physical memory; computational geometry; hardware validation.

Evidence boundary. The central quantitative result is from *simulated* finite propagation. Commodity-memory measurements failed instrument qualification. FPGA work reached validated simulation and multiple place-and-route attempts, but no authenticated routed release, programmed board, or physical attractor dataset existed at the reporting cutoff. Research-run totals are transcribed from the laboratory experiment summaries; archival raw-data and receipt reconciliation is an explicit prepublication requirement.

1 Introduction

Modern computing frequently represents systems as abstract state-transition maps while treating the cost and delay of realizing a dependency as implementation details. That simplification is appropriate when communication is fast relative to the modeled dynamics. Under a finite service budget, however, the physical location of state, the order in which values arrive, and their freshness when used may affect the realized trajectory. A recurrent logical system can therefore admit multiple physical realizations that share the same ideal update rule but differ in transient and asymptotic outcomes.

The question is not whether delay ever changes a dynamical system; delayed feedback and asynchronous recurrent models are well established. The narrower question considered here is whether the *particular dependencies that can reverse a receiver's decision along ideal trajectories* constitute a better predictor of finite-propagation attractor sensitivity than static weighted connectivity. This distinction matters because two equally costly links need not be equally consequential at a given state. A delayed source has no decision-level effect if the receiver's update remains invariant to that source's state.

The research program emerged from Deep SIML Labs' earlier Simulation-Integrated Modeling Loop (SIML) and Cognogenics work, which emphasized compact memory, energy-aware computation, inference and adaptation in resource-constrained simulated agents [1]. The earlier framework included an explicit compute penalty,

$$C_{\text{compute}} = \lambda b^{\alpha}, \qquad F_t' = F_t + C_{\text{compute}}, \tag{1}$$

where b is an accessed or activated bit count and λ, α parameterize resource cost. Equation (1) motivates taking implementation constraints seriously. *It was not itself the experimental objective function fitted in the RBED attractor tests.* The present investigation is a separate, deliberately narrower study of causal dependency and propagation.

We distinguish three conceptual structures:

1. **Logical structure G_L:** the ideal state space, update rule and directed influence relations.
2. **Informational structure $G_I(t)$:** state- and time-dependent sensitivity of decisions to upstream state changes, estimated from the *ideal* system.
3. **Physical structure G_P:** addresses, communication channels, measured service times, contention and memory-controller behavior of an actual realization.

The term *geometry* is used operationally for a structured field of dependencies and costs. The matrix introduced below need not satisfy symmetry, positivity or the triangle inequality; it is *not* asserted to be a Riemannian or metric-space geometry in the formal sense of information geometry [8]. Likewise, a one-bit intervention on a simulated transition function is a well-specified counterfactual sensitivity test, not a general causal-identification claim about an uncontrolled physical system.

1.1 Contributions and status of the claims

This manuscript reports (i) a realization-based formalization that keeps ideal logical dynamics fixed; (ii) a simulation study comparing dynamic causal predictors with a static geometric baseline under finite propagation; (iii) an audit trail of negative studies and qualification failures; and (iv) a preregisterable route to a genuinely physical experiment. The findings should be interpreted as *predictive evidence within a specified simulated model family*. Physical material validation, generalization to arbitrary architectures, and applications to self-organizing materials remain open.

2 Background and relationship to prior work

Attractor memories in recurrent networks provide a useful setting because ideal fixed points or basins offer discrete outcomes against which a finite-resource realization can be compared. Hopfield's formulation of collective computational properties of recurrent networks provides the classical starting point [2]. Dynamical and reservoir approaches also make clear that time and transient histories are essential to computation [4, 5]. Work on embodied intelligence and physical reservoir computing extends the substrate-dependence question to robot bodies and other physical systems [6, 7]. The free-energy framework motivates resource-sensitive adaptive agents but does not prove the RBED claim [3].

A longer-term biological inspiration is the study of coordinated, multiscale regulation in unconventional substrates [9]. This motivates testing computation in materials but does not convert simulated logical susceptibility into evidence for bioelectric computation, non-neural intelligence or consciousness. What is potentially distinctive in RBED is the *controlled factorization*: holding ideal dynamics fixed, varying their embedding into a bounded substrate, estimating a predictor without observing the finite-substrate endpoints, and evaluating out of network.

3 Formal model

3.1 Ideal recurrent dynamics

Let $V = \{1, \dots, N\}$ index binary-state logical nodes, with $s_i \in \{-1, +1\}$ and a fixed weighted coupling matrix $W \in \mathbb{R}^{N \times N}$. In the reported principal simulations $N = 16$, with a Hopfield-like recurrent system built around three stored patterns. For a receiver i, define its instantaneous field and ideal zero-temperature asynchronous update as

$$h_i(s) = \sum_{j=1}^{N} W_{ij} s_j, \tag{2}$$

$$F_i(s) = \begin{cases} +1 & h_i(s) > 0, \\ -1 & h_i(s) < 0, \\ s_i & h_i(s) = 0. \end{cases} \tag{3}$$

An event schedule $(I_t)_{t \geq 0}$ selects the updated receiver at each event. The same deterministic schedule, initial condition and perturbation specification must be paired across ideal and realized experiments. The experiment's measured endpoint is a binary indicator of failure to recover the paired reference ideal attractor under its documented settlement and horizon rules. Exact implementation-dependent

time limits and schedules remain parameters of the archived experiment specifications, not newly inferred values here.

3.2 Realization without changing the logical rule

Let Q denote the ideal transition system and $R_{B,\phi}$ a realization operator parameterized by resources B and a mapping ϕ from logical nodes to substrate sites. The experimentally meaningful object is

$$Q_{B,\phi} = R_{B,\phi}[Q], \qquad \lim_{B\to\infty} Q_{B,\phi} = Q \quad \text{for every admissible } \phi. \tag{4}$$

The second expression is a design requirement: an unbounded or instantaneous realization must recover the same ideal eventwise behavior. It is not enough to add distance directly to Q and then describe the altered rule as the original logical system.

For the simulated Manhattan substrate, the sites lie on a 4×4 grid and d_P denotes their Manhattan separation. At propagation parameter γ, a directed logical read from j to i incurs event-index lag

$$L_{ji}^{(\gamma,\phi)} = \max\bigl(0, \lceil \gamma\, d_P(\phi(j), \phi(i)) - \varepsilon \rceil - 1\bigr)\,, \tag{5}$$

where $\varepsilon > 0$ is the implementation's boundary tolerance. The realized field uses delayed source values $\tilde{s}_j(t)$:

$$\widetilde{h}_i(t) = \sum_j W_{ij} \widetilde{s}_j(t), \qquad \Delta h_i(t) = \sum_j W_{ij}\,(\widetilde{s}_j(t) - s_j(t))\,. \tag{6}$$

Delay only changes a receiver decision when the altered field crosses the applicable update boundary. This observation motivates a state-dependent causal-dependency field.

3.3 Operational causal susceptibility

For source j and receiver i, let $s^{\oplus j}$ denote the state obtained by flipping only s_j. Define

$$\chi_{ij}(s) = \mathbb{I}\Bigl[F_i(s) \neq F_i(s^{\oplus j})\Bigr]\,. \tag{7}$$

This is an intervention on a *fully specified deterministic update function.* Averaging over states visited by ideal trajectories at time t gives

$$G_{ij}(t) = \mathbb{E}_{S_t \sim \mu_t^{\text{ideal}}}\bigl[\chi_{ij}(S_t)\bigr]. \tag{8}$$

The distribution μ_t^{ideal} is sampled only from ideal dynamics, independently of finite-propagation endpoint labels in the held-out test. $G(t)$ expresses which links are decision-relevant in the ideal trajectory ensemble, rather than simply which couplings are nonzero or geographically long.

A static comparator used in the early studies is weighted mean embedding distance,

$$C_\phi = \frac{\sum_{i,j} |W_{ij}|\, d_P(\phi(i), \phi(j))}{\sum_{i,j} |W_{ij}|}, \tag{9}$$

with a nonzero denominator. It is a deliberately limited baseline: the experiment does not show that every conceivable static graph statistic fails.

3.4 Delayed mismatch and ideal-only predictors

Let $z = (i, \delta)$ denote a target/perturbation stratum, and define whether a delayed source differs from the instantaneous ideal source at lag ℓ:

$$M_j(t, \ell) = \mathbb{I}[S_{t,j} \neq S_{t-\ell,j}]. \tag{10}$$

The RBED-2 line introduced an ideal-only joint susceptibility–staleness estimator,

$$J^z_{ij}(t, \ell) = \mathbb{P}\left(I_t = i,\ \chi_{ij}(S_t) = 1,\ M_j(t, \ell) = 1 \mid z\right), \tag{11}$$

$$D^z_{ij}(t) = \mathbb{P}\left(I_t = i,\ \chi_{ij}(S_t) = 1 \mid z\right), \tag{12}$$

$$R^z_{\text{joint}}(\gamma, \phi) = \frac{\sum_{t,i,j} J^z_{ij}\left(t, L^{(\gamma,\phi)}_{ji}\right)}{\sum_{t,i,j} D^z_{ij}(t) + \epsilon}. \tag{13}$$

An associated *factorized* score replaces the joint susceptibility–mismatch probability by its estimated component probabilities under the documented ideal-trajectory sampling scheme. Its precise sampling, normalization and predictor calibration must be reproduced from the archived executable implementation; this manuscript does not invent an unverified closed-form estimator. The factorized dynamic score was selected as the *primary* RBED-3A feature before the confirmatory holdout. Both scores depend on finite-delay parameters through the mapping/lag calculation, but their statistical estimates are constructed solely from ideal runs, not finite-outcome training labels.

The prediction task compares models for the mismatch indicator Y using out-of-network negative log-likelihood (log loss). If ℓ_{dyn} and ℓ_{static} are held-out average log losses, the reported relative advantage is

$$\text{Improvement}_{\text{LL}} = 100 \left(\frac{\ell_{\text{static}} - \ell_{\text{dyn}}}{\ell_{\text{static}}} \right) \%. \tag{14}$$

This is a *relative reduction in predictive log loss*, not a percentage increase in attractor stability, accuracy, or physical communication speed.

3.5 Controls and logical validity gates

The simulation program used increasingly strict controls: $\gamma = 0$ event-by-event equivalence, uniform-latency controls intended to remove embedding dependence, grid-isometry invariance, deterministic replay, paired ideal references, and settlement checks that extend beyond the maximum simulated lag (a window of $L_{\max} + N$ events in the zero-temperature study). Network-level holdout matters because repeated trajectories or embeddings from the same generated network are not independent network replications.

4 Experimental sequence: what failed, what changed, what held

The research sequence was not a single positive sweep. Earlier negative results motivate the final, narrower hypothesis, and an intermediate positive-looking result remained formally negative under the prespecified gate. Table 1 preserves that progression. Run sizes are phase-specific and should not be naively summed: ideal-estimation, primary finite-propagation and engineering-control trajectories have different roles, and some control collections are overlapping or reserved.

Table 1: Simulation experiments and their formally reported dispositions. The numerical claims here are drawn from the recorded phase summaries rather than a newly executed data analysis.

Phase	Scientific trajectories	Critical observation	Disposition
RBED-1	70,524	Static coupling–distance alignment showed inconsistent association; bootstrap intervals spanned zero.	**NO-GO**
RBED-1B	70,524	Genuine zero-temperature fixed points and finite lags improved the model; static alignment still did not consistently predict sensitivity.	**NO-GO**
RBED-2 v1	None qualified	Joint ideal-only estimator missed one prespecified convergence gate (35/36 passed); no scientific outcomes released.	**INCONCLUSIVE**
RBED-2 v2	833,760	Dynamic score improved held-out log loss, but the $\gamma = 1$ rank gate attained 8/12 versus the required 9/12.	**NO-GO**
RBED-3A v1.2	529,200	Prespecified factorized dynamic predictor passed network-held-out log-loss, SESOI and consistency gates.	**CONFIRMED**[a]

[a] *Confirmed within the registered simulated model and analysis.* It is not physical-substrate confirmation or confirmation of all proposed applications.

4.1 RBED-1: static alignment hypothesis

The original $N = 16$ stochastic Hopfield-like experiment used three stored patterns on a 4×4 Manhattan substrate. It explored multiple networks, embeddings, perturbation sizes, masks and repeated trajectories. The 70,524 trajectories, 23 engineering tests and 1,728 exact controls did not support a stable monotonic relation between C_ϕ and attractor disturbance: only 4 of 12 comparison conditions favored the expected higher-alignment direction, and reported bootstrap intervals all included zero. Two design complications were identified. First, finite inverse temperature ($\beta = 2$) meant nominal memories were not absorbing states. Second, the scheduler could generate temporal starvation. Those limitations motivated rather than invalidated a cleaner second experiment.

4.2 RBED-1B: absorbing attractors and finite propagation

RBED-1B used the zero-temperature asynchronous update in Eq. (3), explicit lag in Eq. (5), and a conservative settlement window. Its 70,524 trajectories all settled under its experimental definition. Controls included 1,332 certified exact cases; 42 tests passed, the zero-delay limit matched ideal trajectories event by event, and the uniform-latency control gave identical trajectories across embeddings. Static weighted geometry nevertheless failed its prespecified sensitivity-ranking expectations. Thus adding bona fide attractors and delay did not rescue the static predictor.

4.3 RBED-2: dynamic causal predictors and a failed rank gate

RBED-2 estimated source susceptibility and source-change histories from *ideal* dynamics. Version 1 performed 432,000 ideal-estimation trajectories but failed an estimator convergence gate: 35 of 36 joint cells met the correlation criterion, with one Spearman value of 0.966987 below the prescribed

threshold of 0.98, even though the same cell's RMS discrepancy was 1.67% (within a 5% threshold). This was an instrument/precision failure, not a substantive negative physical or network result.

Version 2 increased ideal-estimation runs to 864,000 and evaluated 833,760 scientific trajectories. Its dynamic joint score showed a 9.656% aggregate held-out log-loss improvement against the static comparator and positive aggregate improvements in 12/12 analyzed strata. The $\gamma = 0.5$ subset improved 5.396%; $\gamma = 1.0$ improved 12.182%, and $\gamma = 2$ was a stress condition rather than a primary gate. A reported paired bootstrap difference in log loss was 0.0201147 with interval $[0.0143243, 0.026312]$ and permutation $p \simeq 0.000999$. Nonetheless, the $\gamma = 1.0$ rank-monotonicity threshold required 9/12 conditions and achieved 8/12. Under the registered all-gates policy the disposition remained **NO-GO**. This distinction between an attractive post hoc effect estimate and a failed a priori decision rule is central to the evidential history.

4.4 RBED-3A: network-held-out confirmatory prediction

RBED-3A narrowed the claim to *prediction*, selected the factorized ideal-only dynamic estimator as primary and retained the static comparator. It evaluated propagation parameters $\gamma = 0.5$ and $\gamma = 1.0$, holding out entire networks rather than individual trajectories. It prespecified a smallest effect size of interest (SESOI) of 5% relative log-loss reduction and required its one-sided 95% network-bootstrap lower bound to exceed that threshold, together with positive consistency across networks.

The final v1.2 run reported 529,200 scientific trajectories, consisting of 10,800 shared ideal and 518,400 finite-propagation trajectories, following 1,080,000 ideal-estimation trajectories. The primary comparison yielded a **12.066%** mean held-out log-loss advantage, with a one-sided 95% lower bound of **11.153%**; 30/30 networks exhibited positive improvement, corresponding to a reported Wilson lower bound of approximately 0.8865. The phase used 10,000 network-bootstrap replicates and 1,000 permutations. These are confirmatory *simulation* results for the preregistered comparator and network family.

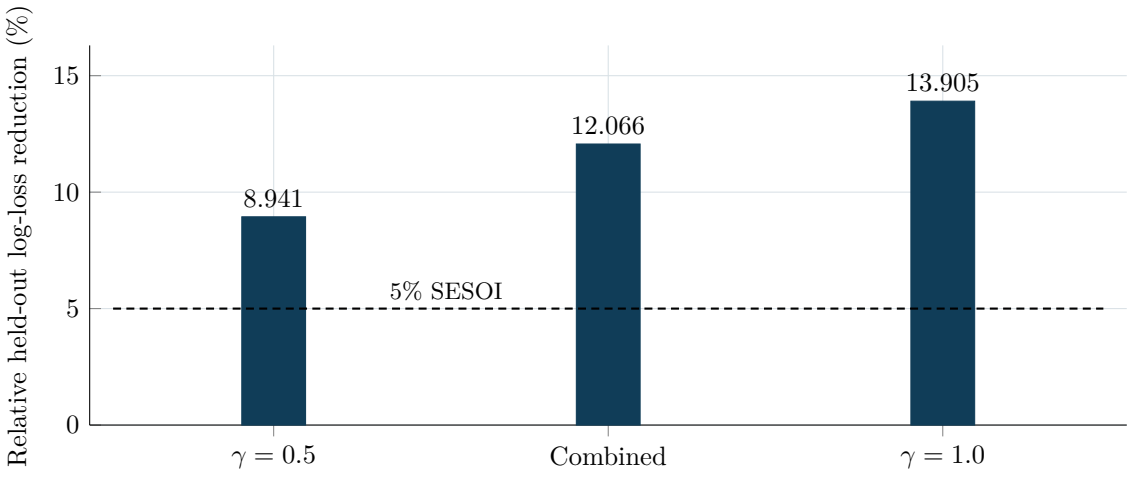

Figure 1: Reported RBED-3A relative held-out log-loss improvements versus the static comparator. The dashed line is the prespecified 5% effect-size threshold; the combined one-sided 95% lower bound is 11.153% (not shown as a symmetric error bar). Values are phase-summary aggregates, not fabricated per-network observations.

Table 2: RBED-3A v1.2 primary confirmatory values.

Statistic	Reported value
Held-out mean relative log-loss improvement	12.066%
One-sided 95% bootstrap lower bound	11.153%
Prespecified SESOI	5.000%
Improvement at $\gamma = 0.5$	8.941%
Improvement at $\gamma = 1.0$	13.905%
Networks with positive dynamic advantage	30 / 30
Scientific trajectories	529,200
Ideal-estimation trajectories	1,080,000
Network bootstrap / permutation iterations	10,000 / 1,000

4.5 Numerical precision, gates and what the results do not show

A pre-confirmatory selection step retained 24 embeddings, eight perturbation masks and five repeats, reporting a 0.99429 correlation, 0.998% RMS discrepancy and 99.75% sign agreement against the larger phase-0 reference. The final reports included 139 automated tests, 2,340 certified controls and 1,416 reserved controls, and described 413,014 certified fits with maximum gradient norm below 10^{-8}. These details increase confidence in internal execution discipline, but independent reproducibility still requires raw files, code and the authenticated model-fitting specification. In particular, the 12.066% does not mean a 12.066% change in attractor frequency, a guaranteed causal effect of physical distance, a speedup in computational hardware, or support for claims about real materials.

5 Physical-substrate attempts and their qualification outcomes

5.1 A physical timing model must be measured, not assumed

The positive simulations used a controlled Manhattan lag model. A physically persuasive experiment must instead obtain a calibrated service-time distribution or effective cost kernel from actual substrate access. For physical locations $p, q \in P$ and service-history covariates H_t, a useful target object is

$$K_P(p, q \mid H_t) = \mathcal{L}(C_{\text{return}} - C_{\text{issue}} \mid p, q, H_t), \tag{15}$$

where $\mathcal{L}$ denotes a measured latency distribution, not an imposed distance penalty. In a memory controller this distribution can depend on address mapping, row-buffer state, bank conflicts, arbitration, refresh and calibration. The exact role of source location in a sequence of memory operations must be operationally defined; a CPU pointer's numerical address difference is not itself a physical distance metric.

5.2 Commodity cache/coherence and DRAM: instrument nonqualification

A first commodity-hardware attempt (H0) did not satisfy its runtime feasibility gate. A subsequent H0 v0.1 exercise collected 41,600 observations but classified both cache/coherence and DRAM measurement paths as inconclusive. The targeted H0P precision campaign accumulated 500,000 overhead measurements, 1,421,952 cache-path readings and 819,840 DRAM-path readings,

totaling **2,241,792 substrate measurements**. Its best reported minimum-batch Spearman correlations were about 0.552 for cache and 0.762 for DRAM, below the instrument qualification threshold of 0.90; additional uncertainty/stability requirements also failed. The outcomes were **CACHE_INSTRUMENT_NOT_QUALIFIED** and **DRAM_INSTRUMENT_NOT_QUALIFIED**. They do not show that the physical timing structure is uniform, nor that the hypothesis is false. They show that the attempted user-space measurement apparatus could not isolate and reproduce a trustworthy geometry at the required precision.

5.3 FPGA–DDR: successful simulation, failed engineering release

To control timing closer to the memory device, we designed an FPGA–DDR apparatus with explicit request issue, memory-controller acceptance, and actual returned-data timestamps. A crucial architectural choice was *persistent MIG*: calibrate the DDR memory interface once at startup and keep it active while resetting only trial-local scientific and diagnostic state. The prior design's repeated vendor-model reset/recalibration sequence produced large reset-timing error counts even with the RBED core removed; the persistent approach instead passed two independent generated-MIG model runs with startup calibration, write/readback traffic, four repeated trials and zero vendor ERROR/FATAL counts.

Importantly, physical qualification requires a legal, timing-signed-off bitstream. Passing simulation and synthesizing into LUTs are insufficient. The multi-board attempt is summarized in Table 3.

Table 3: Apparatus qualification status at the 10 October 2026 reporting cutoff. "Failed route" and "incomplete route" are engineering categories, not negative physical-science findings.

Platform / phase	Synthesis / storage	Decisive gate	Outcome
Arty A7-100T, 2.3	85,096 LUT / 63,400 available	Persistent MIG passed simulation; capacity 134.22% of LUT supply, no implementation release.	**CAPACITY FAIL**
Nexys Video A200T, 2.4	89,500 LUT / 134,600 available	Route failure after ∼9 h; 4,197 failed signals and 5,694 overlaps.	**ROUTE FAIL**
Nexys Video A200T, 2.5	63,959 synthesis LUT; 17.5 BRAM tiles	BRAM-backed trace and streamed host shell; still 2,870 failed signals and 3,295 overlaps after ∼8 h.	**ROUTE FAIL**
Genesys 2 K325T, 2.6	Remediated design ported; validated MIG simulation	Route remained at global iteration 4 (last reported 23,080 overlaps); no legal routed DCP or release at cutoff.	**NOT QUALIFIED**

The A200T remediation reduced synthesized LUT count 28.54% and LUT-as-memory 74.04%, while preserving the sealed scientific RTL and focused equivalence tests. This materially reduced non-scientific storage burden but did not resolve legal routing. On the K325T, the route process continued consuming CPU after more than five days of wall time; the last recorded routing milestone had not changed for roughly 66 hours at the final diagnostic sample. *No terminal Vivado routing-failure exit was authenticated in the evidence available at manuscript drafting.* Accordingly, the rigorous disposition is *no qualified bitstream at cutoff*, not a universal impossibility of implementing the design on the device. As no FPGA board was purchased, programmed or used for RBED acquisition, the series contains **zero physical attractor trajectories**.

5.4 What hardware failures mean—and do not mean

These attempts identify multiple *apparatus* bottlenecks: insufficient LUT capacity on A100T, localized routing pressure in the A200T implementations, and nonconvergent or incomplete K325T routing under the tried physical architecture and directives. Earlier post-synthesis reports placed heavy pressure on cache-version, response and return-data interconnect. However, no formal lower bound proves that all equivalent sequentialized microarchitectures or all larger chips must fail. FPGA routing results cannot validate or refute a physical attractor hypothesis when no compiled instrument has produced latency observations. Nor do vendor DDR timing simulations count as physical-substrate measurements.

6 A falsifiable physical-substrate program

6.1 Stage F0B: qualify the substrate before testing attractors

The next scientific measurement is a *memory-only*, outcome-blind F0B campaign. The Phase-2.4 host package already implemented a four-map address definition, commands 134–137, a durable acquisition journal, corruption/interruption checks and preregistered analysis logic. Its proposed schedule comprised two discovery sessions and eight confirmation sessions, three batches per session, with 345,600 retained returns. These numbers were frozen for that design; if the hardware platform or memory topology changes, the F0B protocol must be explicitly versioned and re-frozen *before* collecting physical data.

An admissible substrate must meet all of the following classes of requirements, with numerical thresholds fixed before acquisition:

F0B-1. Validity. Timestamp issue, acceptance and actual return-valid separately; verify data integrity, clock-domain boundaries, address decoding, command ordering and replay identity.

F0B-2. Nonuniformity. Show practically nontrivial variation in the measured service kernel among declared site/address classes, beyond overhead and quantization uncertainty.

F0B-3. Mapping dependence. Different frozen logical-to-physical mappings must yield distinguishable service-cost patterns under controlled history, not merely a site-independent latency shift.

F0B-4. Repeatability and stability. Reproduce rank/order and effect-size patterns across independent sessions while quantifying temperature, drift, refresh and controller-history confounds.

F0B-5. Negative controls. Include randomized mapping labels, identical-address or otherwise comparable controls, timing overhead subtraction where justified, and deliberate adversarial checks against sequence/order artifacts.

If any required criterion fails, report **F0B_NOT_QUALIFIED** and stop before opening attractor outcomes. A connected effective graph alone is not sufficient: it must be meaningfully nonuniform, mapping dependent and reproducible.

6.2 Stage F1: physically realized attractor sensitivity

Only after F0B qualification may the investigators choose and freeze physical deadlines from the observed service-time distributions, specify a measured-history predictor, and lock the F1 plan. An eventual F1 test should retain the following ordering:

F1-1. Freeze the ideal dynamics and select logical networks, perturbations, schedules and embeddings without consulting physical attractor endpoints. Disclose that the previously selected 30-network panel has already appeared in simulated RBED-3A; preferably also hold out entirely new networks for external confirmation.

F1-2. Estimate dynamic susceptibility solely from ideal trajectories and compute the prespecified static comparator for the same candidate mappings.

F1-3. Acquire finite-budget trajectories using measured C_{return} rather than programmed mapping delays; classify actual freshness by $C_{\text{return}} \leq C_{\text{deadline}}$.

F1-4. Compare out-of-network predictive log loss with network- or session-level dependence handled by appropriate clustered resampling. Report both the registered hypothesis gate and effect-size confidence bounds.

F1-5. Run negative and ideal-limit controls, drift/order audits, uniformized-latency or timing-emulation controls where meaningful, and disclose any failed instrument qualification.

F1 must distinguish a genuine physical time-of-arrival mechanism from an intentionally injected delay pattern. Artificial delay circuits can be useful positive controls but cannot alone establish that a particular material's uncontrolled physics realizes the hypothesized geometry.

6.3 Falsification rules and alternative outcomes

The program admits informative negative outcomes:

- *No reliable physical timing structure:* F0B does not qualify. The material test is inconclusive, not negative for the attractor hypothesis.
- *Qualified structure with negligible endpoint sensitivity:* the tested finite-budget physical regime does not exhibit the predicted observable effect.
- *Endpoint sensitivity but no dynamic-predictor advantage:* the specific RBED claim fails for the qualified physical substrate, even if generic delay effects are present.
- *Static and dynamic comparators perform similarly:* the claimed incremental benefit of ideal-state susceptibility is unsupported in that setting.
- *Network/session heterogeneity or drift destroys generalization:* the proposed predictor lacks reproducibility or transportability and requires a narrower domain-of-validity statement.

A positive physical result would require both qualified service-time evidence and prespecified predictive superiority under controls; simply observing that changing an address changes latency is insufficient.

6.4 Practical substrate alternatives

The current FPGA implementation may be re-engineered using sequential indexed state access, local cache/version banks, limited fanout, or a small processing engine, provided clock-cycle changes are explicitly modeled and logical ideal-limit behavior is proven. Other candidates include directly instrumented SRAM/DRAM controller boards, controllable multicore or chiplet communication fabrics, networks of embedded nodes with precisely timestamped links, and eventually ionic, polymeric or bioelectrical platforms with reliable observability. Each candidate requires its own calibration and confound model. These are *proposals*, not evidence that an off-the-shelf material suitable for this experiment is currently available.

7 Interpretation and broader implications

7.1 From static graph proximity to state-dependent consequentiality

The central mechanism illustrated by the simulation series is simple to state yet operationally specific: *a delayed dependency matters only when it is stale during a receiver update and its altered value would change that update*. Static coupling strength multiplied by distance can be a poor proxy because it ignores the distribution of visited states and scheduling events. A dynamic field such as Eq. (8), combined with source mismatch histories, conditions attention on consequential transmissions.

This is compatible with classical dynamical systems and causal perturbation analysis. It does not require quantum mechanics, extra dimensions, a novel law of physical information, or a special consciousness postulate. The phrase *informational geometry* here concerns a structured dependence field in the ideal model, not an experimentally measured physical metric or a theorem about the geometry of spacetime.

7.2 A possible resource-placement principle

If future experiments establish a calibrated physical service kernel, one could investigate adaptive placement or scheduling policies that favor dynamically consequential links:

$$\phi^*(t) \in \arg\min_{\phi \in \Phi} \sum_{i,j} G_{ij}(t)\, \mathbb{E}\big[K_P\big(\phi(i), \phi(j) \mid H_t\big)\big]\,, \tag{16}$$

subject to capacity, contention, timing, migration and control constraints. Here K_P is interpreted as a nonnegative measured service-cost random variable, and its expectation is schematic; a tail-risk or deadline-violation objective may be preferable in practice. Equation (16) is *an untested optimization proposal*, not a demonstrated result of RBED-3A. Potential engineering domains include distributed state synchronization, memory placement, streaming control and fault-tolerant sensor fusion. These applications require separate benchmarks demonstrating gains against strong placement and scheduling baselines.

7.3 Connection to multiscale and material computation

The SIML/Cognogenics program ultimately asks whether adaptive memory and regulation can emerge from local interactions under finite resources. As a conceptual research direction, one may consider coupled attractor descriptions at multiple scales, such as

$$\mathcal{A}_{\text{material}} \longrightarrow \mathcal{A}_{\text{component}} \longrightarrow \mathcal{A}_{\text{limb}} \longrightarrow \mathcal{A}_{\text{behavior}} \longrightarrow \mathcal{A}_{\text{agent}}. \tag{17}$$

The arrows represent a *hypothesized compositional relationship*, not an experimentally demonstrated hierarchy. Investigating ionic, polymeric, electromechanical or biological substrates would require material-specific state variables, interventions, measurement uncertainty and energetic accounting. RBED-3A provides no evidence that a material can learn, self-heal, think or regulate at these scales. Its possible contribution is a falsifiable method for asking which dependencies matter when communication itself is bounded.

8 Limitations, competing explanations, and methodological cautions

Several limitations bound the scientific interpretation.

Simulated geometry and model family. The confirmatory evidence comes from one constructed family of small recurrent networks and a deterministic Manhattan delay model. It cannot by itself establish generalization to arbitrary network sizes, learning rules, material classes or controller scheduling policies. Increasing the grid, changing the graph ensemble, or evaluating a different form of perturbation might alter predictor performance.

Predictor scope. Dynamic susceptibility is not unique to this work; local influence and finite-delay histories have antecedents across dynamical systems and recurrent computing. The particular ideal-only estimator, holdout design and measured comparator advantage are the empirical contribution. The existing static comparator is relatively simple. A convincing next paper should add strong static competitors (e.g., degree-, centrality-, spectral-, critical-path- or learned graph-based predictors) and assess complexity-regularized dynamic models.

Selection and uncertainty. Whole-network holdout and network-level bootstrap are stronger than trajectory-level splitting, but empirical network draws may still share generator biases. The 30-network panel and phase-0 embedding/mask selection occurred within an extended study sequence. A genuinely fresh network family and independent code replication would reduce concerns that methodological iteration favored the final model. The reported one-sided lower bound is not a universal guarantee.

Finite-run attractor definitions. Settlement certification controls help distinguish transient delay from basin change, but a finite simulation horizon cannot establish every infinite-time asymptotic property. Ideal and realized states must be compared under precisely synchronized schedules and initialization histories.

Nontrivial baselines. The dynamic score may partly act as an informed measure of source-change frequency or receiver scheduling probability rather than an irreducibly new concept of geometry. Ablation studies separating G_{ij}, scheduling, mismatch history and distance are needed to determine which ingredients drive the observed advantage.

Hardware qualification. Cache/DRAM timing measurements were unqualified. FPGA simulation success, logic-resource fit and even placed netlists do not equal timing-closed physical feasibility. The terminal Genesys 2 router result remained unresolved at the cutoff; the absence of a bitstream prevents any material claim. These are real limitations of the *instrument*, not evidence about the sign of the hypothesized physical effect.

Prior theoretical language. Terms such as free energy, informational geometry, attractor and causality have distinct formal meanings in neighboring fields. This paper uses explicitly defined operational quantities and does not infer a thermodynamic free-energy theorem, information-geometric metric identity, biological agency or physical-information ontology from the measured prediction statistics.

9 Reproducibility and data provenance

A reproducible release should include the exact versioned simulation source, RNG seeds, network weights/patterns, perturbation masks, embedding maps, train/test partitions, control tapes, loss-fitting code, bootstrap/permutation scripts, and raw per-network records. Numerical tables should be generated mechanically from authenticated outputs rather than retyped from summary reports. An immutable release manifest should record SHA-256 checksums for the experiment configurations,

reports, figures and final PDF. Prior phase reports already described event-by-event ideal controls, model-fit checks, replay and verification tests; those claims remain subject to independent archival inspection.

Source status for this manuscript. The historical SIML motivation was checked against the laboratory's April 2025 paper [1]. RBED phase figures and hardware gate states in this version have been transcribed from research-run summaries discussed and preserved by the project, not re-executed against the laboratory repository during document preparation. As a result, *this document should not be represented as an independent audit of the experiment binaries or raw datasets.* Before public release, the laboratory should reconcile every table entry with the corresponding signed or hashed run receipts, resolve author and licensing metadata, and either publish a code/data archive or explicitly state its availability conditions. Until those steps occur, readers can assess the claims and proposed falsifiers, but cannot independently reproduce the numeric effect from this PDF alone.

Status of physical artifacts. No physical F0B dataset, physical F1 experiment, device-programming receipt or authenticated working FPGA bitstream was available at the reporting cutoff. Hardware descriptions report engineering attempts rather than hardware science. The study does not claim to demonstrate a material substrate of cognition or substrate-dependent attractor sensitivity.

10 Conclusion

The resource-bounded embedding program moved from a failed static geometric hypothesis through tighter attractor and estimator controls to a successful, narrowly defined simulation prediction: an ideal-dynamics-derived causal susceptibility feature predicted finite-propagation attractor sensitivity better than the registered static baseline on network-held-out tests, with a reported 12.066% relative log-loss reduction and a one-sided 95% lower bound of 11.153% above a 5% SESOI. The historical NO-GO decisions are part of the evidence, not results to discard.

The more ambitious physical-substrate claim remains unresolved. Precision-qualified commodity memory timing was not achieved, FPGA implementations did not produce a legal timing-signed-off release, and no physical outcome data exist. These obstacles motivate a redesign of the instrument and a substrate-first F0B qualification, not an assertion that ordinary memory or all FPGAs are fundamentally incapable of realizing the model. The immediate scientific contribution is therefore *a simulation-supported predictive relationship with specified controls and a falsifiable route to material validation.* Deep SIML Labs' next step is to release a verifiable computational artifact alongside this result, then test whether a real and measurable physical service geometry transfers the predictive advantage.

Disclosure and acknowledgments

The manuscript's initial editorial preparation used AI-assisted drafting based on the laboratory's reported experiment summaries and its prior SIML paper. Interpretation, authorship, source verification, experimental accountability and any public scientific claims remain the responsibility of the named author and Deep SIML Labs. The paper is a laboratory technical preprint and has not undergone formal external peer review.

A Definitions of result and qualification labels

CONFIRMED (simulation)
A particular prespecified simulated-model predictor test passed its stated gates; not proof of a physical effect.

NO-GO A valid tested phase failed at least one required research acceptance gate, even when secondary metrics were favorable.

INCONCLUSIVE
The scientific effect could not be evaluated under valid calibration, convergence or timing conditions.

NOT QUALIFIED
The instrumentation or physical implementation lacked sufficient evidence for a real-world measurement release.

ROUTE FAIL Vivado explicitly failed legal routing; timing closure and bitstream signoff do not follow.

ROUTE PENDING
Routing remained active/incomplete at the reporting cutoff; cannot infer a successful or failed terminal exit without further receipt evidence.

B Numerical and protocol audit checklist

Before publishing a first final edition, obtain the following authenticated source evidence rather than filling any gap by inference:

1. **RBED-1/1B:** exact network generation, 70,524 counts by arm, 4/12 static-alignment result, bootstrap settings and ideal/uniform-latency controls.
2. **RBED-2:** ideal sampling counts, 35/36 estimator gate, all condition-level rank results, 833,760 scientific trajectories, complete log-loss fitted-model comparison and permutation/bootstrap receipt.
3. **RBED-3A:** registered phase-0 selection, frozen 30-network split, γ strata, model fits, 529,200 outcome records, 12.066% computation, 11.153% one-sided lower bound, 10,000 resamples, 1,000 permutations and Wilson calculation.
4. **H0P:** raw timing files and stability/spearman calculations for each commodity pathway; calibration thresholds and overhead subtractions.
5. **FPGA:** Phase-2.3–2.6 synthesis reports, routed-failure receipts, sealed-RTL hashes, generated MIG parameters, post-place analyses and any eventual Phase-2.6 termination receipt. Do not reclassify route stagnation as an explicit Vivado failure without such a receipt.
6. **Physical plan:** make the address-map/freshness/clock-history protocol hardware-specific; version F0B and F1 preregistrations before acquisition and maintain negative controls.

C Research program separation: present and future evidence

Table 4: Separation of the established, instrument-development and aspirational programs.

Program	State	Allowed inference
RBED-3A simulation	Confirmed within preregistered test	Dynamic ideal-only features outperform the tested static baseline on held-out simulated networks.
H0/H0P physical sensing	Not qualified	No trustworthy measured commodity physical geometry for F1 has been established.
FPGA–DDR F0	Not qualified / incomplete	No timing-closed, programmed physical experiment; engineering results constrain apparatus design only.
F0B memory calibration	Not executed	Nonuniformity, mapping effects and reliability are future questions.
F1 physical attractor experiment	Not executed	Real-substrate transfer of dynamic predictor remains untested.
SIML cognition/material systems	Long-term hypothesis	No claims of emergent material cognition, embodied intelligence or nested adaptive attractors follow from current RBED data.

References

[1] R. Everts, *Artificial Intelligence is for Amateurs: SIML, Free Energy Principle and Cognogenics as the Foundations Toward the Birth of True Artificial Life*. Deep SIML Labs manuscript, April 2025. Original laboratory working paper.

[2] J. J. Hopfield, "Neural networks and physical systems with emergent collective computational abilities," *Proceedings of the National Academy of Sciences*, vol. 79, no. 8, pp. 2554–2558, 1982. doi:10.1073/pnas.79.8.2554.

[3] K. Friston, "The free-energy principle: a unified brain theory?," *Nature Reviews Neuroscience*, vol. 11, pp. 127–138, 2010. doi:10.1038/nrn2787.

[4] W. Maass, T. Natschläger, and H. Markram, "Real-time computing without stable states: a new framework for neural computation based on perturbations," *Neural Computation*, vol. 14, no. 11, pp. 2531–2560, 2002. doi:10.1162/089976602760407955.

[5] H. Jaeger and H. Haas, "Harnessing nonlinearity: predicting chaotic systems and saving energy in wireless communication," *Science*, vol. 304, no. 5667, pp. 78–80, 2004. doi:10.1126/science.1091277.

[6] R. Pfeifer and J. Bongard, *How the Body Shapes the Way We Think: A New View of Intelligence*. MIT Press, 2006. doi:10.7551/mitpress/3585.001.0001.

[7] K. Nakajima, "Physical reservoir computing—an introductory perspective," *Japanese Journal of Applied Physics*, vol. 59, article 060501, 2020. doi:10.35848/1347-4065/ab8d4f.

[8] S.-i. Amari, *Information Geometry and Its Applications*. Springer, 2016. doi:10.1007/978-4-431-55978-8.

[9] M. Levin, "Technological approach to mind everywhere: an experimentally-grounded framework for understanding diverse bodies and minds," *Frontiers in Systems Neuroscience*, vol. 16, article 768201, 2022. doi:10.3389/fnsys.2022.768201.

[10] AMD, *7 Series FPGAs Memory Interface Solutions (UG586)*, version 4.2, November 2024. https://docs.amd.com/r/en-US/ug586_7Series_MIS/DDR3-and-DDR2-SDRAM-Memory-Interface-Solution.