What survived: a claims register with falsifiers for a pre-registered AI research program

W. Schulz

Jul 29, 2026

Abstract

Every surviving claim from a pre-registered, run-once program, with its evidence and its falsifier. Alive: pool-relative verifier exploitability, test-time recurrent depth, the AIMC window-policy rule, narrow optical interconnect. Dead: wave/complex substrates, optical compute for inference, frontier-datacenter AIMC, and the ε 'universality class'.

Top-level reconciliation of the two research arcs, plus a claims register: every surviving claim, its evidence, and what would falsify it. Written to be the single legible entry point. The project's real asset is its method — pre-register, run once, report failures, baseline at equal budget — so this document foregrounds status and falsifiers, not narrative.

Last updated after the 10-step roadmap (`docs/ten-step-roadmap.md`).

STATE OF THE PROJECT — 2026-07-29 update (read this first; the body below is the older register)

The body of this register predates programs 3–4, the reckoning, the codesign campaign, and the 2026-07 foundations strands. It remains accurate for arcs 1–2 as far as it goes; this block is the current one-page state, with claims and falsifiers. Per-strand canonical docs are linked from `README.md`.

Alive (build on these):

1. Verifier exploitability (the lead thread). Best-of-N is bounded by verifier *exploitability* ε, not accuracy; N*(ε) shipped as `verifier-budget`; ε is pool-relative (two channels: exploit + coverage; swings 6.2× with candidate pool). Falsifier: a verifier whose selection survives adversarial pool shift without a clean prefilter. (`verifier-bottleneck/FINDINGS.md`)
2. Test-time recurrent depth. Real, parameter-efficient (1/8 params, +0.28 OOD), extrapolating (8×) compute axis; domain governed by update contractivity — externally reproduced at LLM scale (STARS/FPRM). Verifier-gated halting beats best fixed depth (0.991 vs 0.960, ~2.3×±0.8 less compute, N=10). Patent line closed (ACT/PonderNet/DEQ own the wins). (`TTC-recurrent-depth-synthesis.md`)
3. The AIMC window law (ship-ready defensive publication). On hybrid LMs, the attention window policy — not the attention fraction — decides AIMC's decode win (sliding-window ~5.6–6.8× flat in context; full attention collapses to ~1.2× at long context); measured quality boundary on RecurrentGemma-2B (recall 1.0 within window, 0.0 beyond, effective range ~0.5·W at L=2·W). Lane = power-bound 1–8B edge + power-capped $\lesssim$8k disaggregated prefill (K-series all

closed; produced KV portable at ≥5-bit). Falsifier: gain-cell analog attention maturing through the ITO+GQA+OSFET corner (quantified, three immaturities deep). (`PAPER-aimc-window-codesign.md`, `DIRECTION-prefill-flank.md`)

4. Optical interconnect, narrow. Bandwidth-density/capacity fabric for pooled KV at rack reach and extreme scale ONLY; energy case flips under honest laser wall-plug × utilization; the niche is incumbents' shipping thesis (Celestial) — our contribution is boundaries + negatives. (`../codesign-analog-moe/MASTER-SYNTHESIS.md`)
5. The method (the project's durable asset). Novelty gate (5 checks + meta-check), the soundness axis (not-novel ≠ false; landing on a named theorem = corroboration), floor arithmetic for registered thresholds, and gate-first same-day strand closure — calibrated on ~15 deflations and 5 independent reconvergences of "interconnect not compute; strong baseline beats clever structure." (`NOVELTY-GATE.md`, `SOUNDNESS-AXIS.md`)

Dead (do not reopen without new evidence): waves/complex as load-bearing substrate (equal-budget falsified everywhere; only real-OK non-commutativity survives in cognition); optical *compute* for inference (worse-PIM: AIMC beats it 11–20×); quantum/exotic substrates for learning or biology at equal budget; frontier-datacenter AIMC (watts-Amdahl cap ~1.4×); the ε "universality class" (Thurstone identity); all foundation-flavored metaphysics objects (grounding 5/5, ontogeny 14/14, eschatology 5/5, resolution=perennialism — each deflated or reduced to owned mathematics).

The 2026-07 foundations strands (all closed, gate-first, same-day): their durable residue is method plus a few measured facts — self-critique fails via *shared error structure* (same-context critic 12/12 misses; sign FLIPS in the dominated regime); creation/eschatology dissolve into certificate theory + the flux law (expansion manufactures the gap, the gap feeds the flux, the flux makes structure; relics invert the coupling — Layzer + Prigogine + Zel'dovich compose; arrow-capability of a law = recurrence + invariants, residue 0; IC ≡ post-selection = 0 bits).

The two arcs in one paragraph

The project has two independent arcs that share a substrate (complex/wave vectors) and a method. Arc 1 — continual learning & cognition asked whether representing information as waves yields continual learning without forgetting (it does not — that comes from a decorrelating *rule*, not waves) and whether the same substrate models human reasoning (it does, in one specific corner: non-commutativity). Arc 2 — photonic inference asked whether an optics-native SSM beats a GPU on latency (it does, but inside a now five-walled box). Both arcs have reached honest, defensible termini that are far narrower than the original ambitions — and correct.

Arc 1 — continual learning & cognition

What's dead (project a — continual learning)

Three independent roads (capacity bake-off 5-i, online forgetting 5-★, weight-only 5-★b) reached the same place: continual learning is delivered by a decorrelating projection/pseudoinverse rule — a classic 1980s associative-memory result with nothing to do with waves — and a real-valued substrate does it at least as well at equal budget. "No backprop ≠ no forgetting." Terminus; nothing left to chase that isn't goalpost-moving.

What survived, then narrowed (the representational phase edge)

Exps 8/9/13 established a *real* representational edge: a complex/phase encoding separates amplitude-correlated concepts better than an equal-budget real encoding when the data has directional/asymmetric structure — but only with synthetic, injected asymmetry, and only above A≈0.3. Roadmap step 6 closed this thread negatively: on natural corpus directed co-occurrence (asymmetry ≈0.46, *above* the synthetic crossover) the complex encoding loses to equal-budget real at every discriminating noise level, and amplifying the corpus's own asymmetry up to 8× does not flip it. The edge is a property of the synthetic antisymmetric *structure*, not of natural directional data. Surviving claim, fully qualified: *real on synthetic directional data, absent on natural corpus data.*

Where the substrate is genuinely necessary (project b1 — non-commutativity)

This is the one corner where "build it as waves" is justified by evidence, not by the name. The substrate is the Busemeyer/Bruza quantum-cognition formalism, and the roadmap grounded it in our own concept-waves: - Conjunction fallacy (exp 16): `p(F∧B) > p(B)` from incompatible concept-waves, 100% of seeds; vanishes at zero incompatibility (classical control). - Disjunction fallacy (step 1): the *bundled* co-prediction, same geometry unchanged — under the same representativeness principle, `p(F∨B) < max(p(F),p(B))`, 100% of seeds. - Feature-emergent (step 3): the effect survives with no hand-inserted axis — incompatibility emerges from genuinely shared features; null control (no shared features) kills it. - Sure-thing violation (step 5): a *decision* error (second paradigm) — interference drops p(play|unknown) below both conditionals; requires the phase. - QQ-equality on real survey content (step 4): the parameter-free invariant holds to machine precision on the Clinton/Gore pair.

What did not transfer cleanly (reported, not buried): - Quantitative magnitude fit (step 2): the minimal 2-knob model is over-constrained — it gets the fallacy and the human *ordering* but cannot match a 3-value probability triple (one knob couples representativeness to incompatibility). Needs a richer geometry. - Order-effect SIGN (step 4): the naive correlated geometry does not reproduce the Gallup *assimilation* direction (both ratings fall when asked second); matching the sign needs a *fitted* geometry, as Wang & Busemeyer did. The parameter-free QQ invariant is the strong result; the sign is geometry-dependent.

Arc 2 — photonic inference

Verdict: SUPPORTED, but now five-walled (was four). An optics-native static-S4 SSM beats an H100 on decode latency (~55× at d=256, reproduced exactly in step 9) only inside: 1. decode-only (prefill loses ~500×), 2. low-precision (>10^{-4}; the cliff voids it otherwise), 3. small d_model (55×→2.2× over d=256→2048; the size wall), 4. zero dynamic-weight matmuls (the static S4 structure), 5. NEW — task-gated (roadmap steps 7–9): the win requires a workload the static S4 can do at iso-accuracy. On position-selective recall the static SSM cannot reach attention accuracy at any width (step 7: ~16% vs 100%, flat in d), so on recall-heavy decode (the important LLM case) the win is void — no finite-d re-scoring recovers it. The free square-law gate is *not* the problem (step 8: D3 holds); the static *mixing* is. Recovering recall needs a selective (Mamba) SSM — the riskier photonic story the cost model deliberately avoided.

The win's mechanism remains data-residency (weights resident, no per-token HBM traffic), not cheap multiplies (optical energy is ~0–1% everywhere). Cost model is calibrated to literature within an order of magnitude (Prompt 9). The whole 10-part build passes 164 tests.

Claims register

Status: ✓ supported · ⚠ partial/qualified · ✗ falsified · closed terminus.

#	Claim	Evidence	Status	Falsifier
C1	Wave substrate yields continual learning without forgetting	exps 5-i, 5-★, 5-★b	✗	— (falsified: it's the rule, not waves)
C2	Continual learning is achievable online, backprop-free	exp 11 (projection rule, cap 6→60)	✓	a decorrelating rule still shows the cliff
C3	…but the wave substrate is the *source* of C2	exp 12 (real 37 > complex 26)	✗	complex beats real at equal budget
C4	Complex/phase gives a representational separability edge on directional data	exps 8/9/13	⚠	edge gone at decoupled ρ=0 (it held)
C5	…and that edge appears on natural data	step 6 (complex < real on corpus)	✗	complex ≥ real at natural asymmetry
C6	Substrate reproduces quantum-cognition signatures (order, QQ, interference)	exp 14; step 4 (QQ on real pair)	✓	classical control fails to vanish (it vanishes)
C7	Conjunction fallacy emerges from concept-waves	exp 16 (2.7×, 100% seeds)	✓	fallacy persists at zero incompatibility
C8	Disjunction fallacy is a bundled co-prediction (same geometry)	step 1 (100% seeds, ¬B-first)	✓	needs geometry re-tune (it didn't)
C9	The effect needs no hand-inserted axis (feature-emergent)	step 3 (null control kills it)	✓	fallacy needs the artificial salience axis

#	Claim	Evidence	Status	Falsifier
C10	Interference produces a decision error (sure-thing violation)	step 5 (phase region 120–240°)	✓	no phase makes unknown < both
C11	The model fits human probability magnitudes quantitatively	step 2 (max err 0.155 > 0.05)	⚠	(already over-constrained; ordering ok)
C12	The model predicts the order-effect SIGN parameter-free	step 4 (assimilation sign wrong)	✗	sign matches without fitting (it didn't)
C13	Optics-native SSM beats GPU on decode latency	official run; step 9 (55.7×)	✓	win gone in-corner (it holds in-corner)
C14	…at iso-accuracy for the static S4	step 7 (recall fails at any width)	⚠	static S4 matches attention (it doesn't, on recall)
C15	The free square-law gate is an adequate nonlinearity (D3)	step 8 (matches SiLU both tasks)	✓	square-law ≪ SiLU on a task SiLU handles

Roadmap #2 updates (see **`docs/ten-step-roadmap-2.md`**)

Status changes to the rows above, and new claims:

#	Claim	Evidence	Status	Falsifier
C11′	Conjunction model fits the human magnitude triple with a decoupled geometry	r2-step 4 (max err 0.0000)	✓	target Gram not PSD (it is) — C11 over-constraint was the parameterization, not the model
C12′	The order-effect SIGN is reproducible with a fitted geometry (+ QQ still 0)	r2-step 5 (sign ✓, q≈1e-18; magnitudes ~0.05)	⚠	no geometry gives sign+QQ (it does); exact magnitudes need >2D
C14′	A selective SSM recovers recall accuracy the static S4 lacks	r2-step 1 (100% vs static 30%)	✓	selective also fails (it doesn't)

#	Claim	Evidence	Status	Falsifier
C16	Selectivity keeps the photonic latency win (0 dynamic-weight matmuls)	r2-step 2 (45× decode; 0 dyn-wt)	✓	selective adds a data×data matmul (it doesn't)
C17	A selective photonic SSM is Pareto-optimal at iso-accuracy	r2-step 3 (60× vs accurate transformer)	✓	no SSM on the iso-accuracy front (selective is)
C18	Conjunction & disjunction fallacies are positively linked (one incompatibility)	r2-step 6 (r=0.69; both rise w/ incompat)	✓	r≈0 in-model (it's 0.69) — novel human-testable prediction
C19	The natural-data null (C5) is explained by antisymmetric effective rank	r2-step 7	✗	eff-rank predicts it (it doesn't) — cause relocated to the M_sym correlation regime

Two headline shifts from roadmap #2: - Photonic arc UPGRADED. The roadmap-1 "fifth wall" (task-gating: static S4 can't recall) is navigable — a *selective* SSM recovers recall accuracy (C14′) while keeping zero dynamic-weight matmuls and a 45–60× decode win (C16, C17). The arc's verdict moves from "SUPPORTED, task-gated (recall = dead end)" to "SUPPORTED; recall reachable via a selective SSM that keeps the win." - Cognition (b1) strengthened but contextualized. C11/C12 resolve once the geometry is decoupled / fitted (C11′, C12′), and a genuinely novel falsifiable prediction emerges (C18). BUT a literature cross-check (r2-step 10) surfaced Boyer-Kassem et al. 2016, "Quantum-like models cannot account for the conjunction fallacy" — the QP account is *contested*, and our own fitting strains (C11 original, C12) are *consistent with that critique*. The honest position: the substrate reproduces each signature cleanly, but jointly satisfying all constraints is genuinely hard and debated — cite the critique, don't overclaim.

One-line bottom line: waves earn their keep in exactly one place — non-commutativity in human judgment (b1) — where the conjunction/disjunction/sure-thing errors fall out of incompatible concept-waves; everywhere else (continual learning, natural-data representation, general photonic inference) the honest verdict is narrower than the name suggests, and now precisely bounded.

What's genuinely open (worth doing; not goalpost-moving)

Both items previously listed here were CLOSED by Program 4 (`docs/program4-closing-the-gates.md`): - ~~Cognition magnitude fit (C11/C12)~~ → CLOSED (P4 s6). A real-valued 2-D projective model fits the Moore-2002 Clinton/Gore order-effect magnitude (RMSE 0.052) and satisfies the

parameter-free QQ-equality to 1e-15. Magnitude and sign are real-representable; complex adds nothing the QQ constraint doesn't already bind. The last clean cognition number is closed — *complex decorative*. - ~~Photonic selective/Mamba SSM relevance~~ → CLOSED (P4 s1, decisive negative). The latency win needs zero dynamic-weight matmuls; only pure static/diagonal-SSM keeps it (≈0% of deployed decode), and attention breaks it outright. The photonic story does not extend to real LLM decode — structurally, calibration-independent.

The actually-open items now live on the surviving (substrate-agnostic) tracks — not photonics or cognition, both of which are closed: - Does the TTC hybrid scheduler generalize off Kuramoto slow-dynamics? (Program 3: bounded so far.) - A deployable verifier-ε estimator for real reward models (ε-budget CLI shipped; real-RM estimation open). - Does "worse PIM" hold for large/high-precision models? (LightML, ISCA 2025, contests it at scale.)

Program 4 re-charter recommendation (the fork)

The founding premise (wave-based continual learning / Thesis B) is dead; the two applied bets resolved (photonic compute DEAD, interconnect ALIVE-but-bounded). Recommendation: PIVOT the project from "*wave-based continual learning*" to "inference-time compute + verifier reliability tooling" — the TTC hybrid scheduler and the verifier ε-budget tool are the only upward, substrate-agnostic survivors and depend on none of the dead theses. Carry optical interconnect as an *optional hardware sidecar*, and retire the wave/continual-learning thesis with a clear, well-evidenced obituary. See `docs/program4-closing-the-gates.md` step 10.

Published by NeuroGenesys (AW3 Technology, Inc.). Canonical version: https://neurogenesys.si/research/claims-register