Complex amplitudes are decorative: a pre-registered falsification of the wave/quantum thesis for learning and cognition

W. Schulz

Jun 27, 2026

Abstract

Across continual learning, cognition, and compositional binding, the complex/wave degree of freedom is load-bearing nowhere at equal real-parameter budget. The geometric intuition survives in real-valued form; the transferable contribution is the equal-budget falsification methodology itself.

Preprint draft, 2026-06-27. A negative-results + methodology paper. Every claim below is from a run-once, pre-registered experiment in this repository; the internal audit trail is `RECKONING.md`, `synthesis-and-claims-register.md`, and the per-track `FINDINGS.md`.

Abstract

A recurring idea holds that representing information as complex/wave amplitudes (Hilbert-space states, Born-rule readout, phase) buys something a real-valued vector cannot — non-forgetting continual learning, or a uniquely quantum account of human judgment. We built a numpy substrate for this thesis and tested it as a falsification machine: pre-register, run once, baseline at equal real-parameter budget, report failures. Across three domains — continual learning, cognition, and compositional binding — the complex/wave degree of freedom is load-bearing nowhere at equal budget. Continual learning comes from a decorrelating *rule* (a 1980s associative-memory result), not from phase; the famous "quantum cognition" order/conjunction effects need only real non-commutative projectors, not complex amplitudes; and holographic binding (FHRR) merely ties its real-valued counterpart (HRR) at matched budget. The idea's *geometric* intuition — isolate new knowledge in a stable representation — survives, but in real-valued form; phase was the wrong implementation of a right idea. We report the negative result in full and argue the methodology (equal-budget falsification) is the transferable contribution.

1. The thesis and the test

Thesis (Project B): information represented as normalized complex vectors, evolved by interference and attractor settling, yields (a) continual learning without catastrophic forgetting and (b) a substrate for human-like reasoning. Null: an equal-real-budget real-valued vector does as well. A complex vector of dimension D/2 carries the same real parameter count as a real vector of dimension D; all comparisons hold this fixed. Each experiment is pre-registered with a kill criterion and run once.

2. Continual learning — the phase degree of freedom does not help

Probe	Result
Capacity bake-off (holographic vs replay)	Capacity is not the edge; a 32-slot replay buffer dominates a single holographic trace.
Phase stability–plasticity (static, structured data)	Phase gives a *representational* separability gain only when data is directionally structured — and only statically.
Online Hebbian, real vs complex (equal budget)	Phase headroom does not survive the shared-weight online rule; forgetting is catastrophic for both. "No backprop ≠ no forgetting."
Iterative self-training / model collapse	Phase does not delay collapse — it collapses faster.
Natural directed co-occurrence	Even with strong measured asymmetry (A≈0.94), complex loses; real directionality is worse than random.

Mechanism (the real-valued lesson): Class-IL forgetting is, mechanically, recency bias in the shared classifier head; the large lever is a bias-free head (NCM: 0.195→0.634), not a fancier rule, more capacity, or smarter buffers. Stability comes from putting new knowledge in a *fresh, isolated, sparse, or frozen* region of a *stable* representation — representation isolation, which real orthogonal/sparse/ frozen subspaces deliver. The intuition behind the wave program was right; complex phase was the wrong implementation.

3. Cognition — real non-commutative projectors suffice

The substrate reproduces the canonical "quantum cognition" signatures — the Linda conjunction fallacy ($p(F \wedge B) > p(B)$, 2.7×, robust over 200 seeds, switched on by projector *incompatibility* with an a=0 control that kills it), order effects, and QQ-equality. These are genuine and the substrate models them well. But the load-bearing ingredient is non-commutativity of measurement projectors, which is available to real-valued projectors; the *complex* amplitude adds nothing the effects require. The earlier belief that waves were load-bearing here (via Busemeyer's formalism) is superseded: it validates a *reasoning* model built on non-commuting real projectors, not the complex substrate, and not the *learning* thesis.

4. Binding — FHRR ties HRR at equal budget

Holographic compositional binding (role⊗filler, unbind, analogical inference) works — but the complex Fourier HRR (FHRR) ties the real-valued HRR at matched real budget. The binding algebra is real; the phase is bookkeeping.

5. What survives (and it isn't complex)

Three results survived adversarial, equal-budget scrutiny — and none rests on complex amplitudes: 1. Test-time-compute dynamics — depth-vs-width-vs-verifier structure (Tunnel Vision, width-escapes) is a property of *slow dynamics*, not of phase. (Caveat: degrades on learned models; see the TTC thread.) 2. The verifier exploitability bound — $N^*(\varepsilon)$, substrate-free, the project's most externally relevant result (separate paper). 3. Photonic interconnect/energy — a hardware story about data residency, orthogonal to the learning thesis, and itself narrow and conditional.

6. The methodology is the contribution

The result was produced by a discipline, not a model: pre-register the prediction and a kill criterion; run once; baseline at equal budget; report failures and self-falsifications as first-class. Of ~26 probes across the project, the high-value bits were overwhelmingly negatives and self-falsifications — cheap to produce, and the actual product. A hyped idea (quantum/wave cognition and learning) was given its best equal-budget shot and found decorative. We argue this equal-budget falsification protocol is the transferable contribution, and that publishing the clean negative is worth more than another regime-hunt for a phase win.

7. Honest scope

This falsifies *complex amplitudes as a load-bearing computational primitive* for these tasks at equal budget on a CPU/numpy substrate. It does not speak to physical quantum computers, to optical hardware's *energy* story (a separate, narrow positive), or to whether phase helps at non-equal budgets (it can — but that is a different, weaker claim than the thesis made).

Published by NeuroGenesys (AW3 Technology, Inc.). Canonical version: https://neurogenesys.si/research/negative-results-wave-cognition