

One Theory, Two Hemispheres

Right-Brain AI (November 2025) Read Through the Vacuum.net Theory, and What Has Since Been Built and Measured

**J. Konstapel — Constable Research, Leiden —
September 2026**

This article establishes a correlation that has not been written down before: between "Left and Right Brain AI" (Konstapel, constable.blog, 23 November 2025 — the primary reference of this article) and the vacuum.net theory as it took its vocabulary in 2026. It then measures the 2025 claims against what stands running and proved in September 2026: MAZE.AI. The article is self-contained; every number used appears in the text.

1. Two documents, ten months apart

In November 2025, "Left and Right Brain AI" laid out an architecture in the vocabulary of fields: contemporary AI as a hypertrophied left hemisphere, and a complementary Resonant Stack — oscillatory substrate, nilpotent coherence kernel, multi-scale coupling — as the missing right hemisphere. The essay was written from intuition and from physics, before the vacuum.net theory had its present vocabulary of strand, winding, closure, mesh and tension.

In September 2026, MAZE.AI stands running: an answer loop over 8.19 million addressed pieces, with machine-checked invariants, a signed ledger, and a pre-registered placement measurement of 15 out of 20. It was designed in the vocabulary of discrete addresses.

The claim of this article is that these are not two projects. They are one theory read at two winding depths. Right-Brain AI reads the knowing layer as a continuous field — the deep, fast windings. The maze reads it as discrete addresses — the shallow, slow windings. Identity across scale, not analogy. The correlation below is term-by-term, and each term ends in something checkable: a theorem, a measurement, or a named open instrument.

2. What the 2025 article claimed

Six claims carry the original essay. They are restated here faithfully, in its own terms.

C1 — The left brain is a discrete stack with three constraints. The dominant AI of 2025 is discrete at every joint: bits, tokens, steps. It suffers temporal fragmentation (fixed context windows; long-range coherence simulated by bookkeeping), loss-function myopia (behaviour determined by a human-chosen scalar loss), and an energy ceiling that the essay calls a physical constraint, not a solvable engineering problem.

C2 — The alternative computes by coherence. The Resonant Stack replaces the bit by phase and frequency. Computation is the self-organisation of a coupled oscillator field, Kuramoto-style, ideally photonic.

C3 — The kernel is nilpotent. Above the field sits a coherence kernel after Rowlands: only states with $N^2 = 0$ — states respecting conservation and zero-totality, the whole summing to nothing — are admissible. Incoherent states are not rare failures to be corrected; they are impossible.

C4 — Learning is not gradient descent. Admission of a new configuration is algebraic: propose, compute the nilpotent vector, accept only on $N^2 = 0$.

C5 — The hemispheres meet in a designed corpus callosum. Interface pattern 5.1 of the essay: the language model receives the request, decomposes and phrases; the coherent core is consulted before anything executes; the model proposes, the core disposes.

C6 — The safety argument changes class. A stack with nilpotent constraints can show that whole classes of destructive states are impossible — not rare, not filtered at 99.9%, but impossible. The essay names this as a different class of guarantee than "we tested the model and it performed well."

3. The net vocabulary, in one paragraph

The vacuum.net theory describes everything as one strand that winds. A winding that closes carries no net tension; closure is the stable form. Windings of windings form a mesh; position in the mesh is depth. What persists is selected by fit: only what returns in phase remains. Where two quantities are one state read twice, the theory writes a constitutive relation — simultaneous, no causal arrow. The register of the net is the trit: +1, 0,

−1, the balanced form in which the zero is lawful and the two signs carry each other.

4. The correlation, term by term

Left brain = the long chain in the language layer. The 2025 essay says: discrete at every joint, brilliant emissary, blind to the whole. The net says what that is: computation held in the talkative, shallow layer, with assigned addresses (a vocabulary table someone maintains) and a deep continuous register (16-bit floats). Everything that requires an institution. One statement, two depths.

Energy ceiling = register depth, as a constitutive relation. The essay calls the ceiling physical. The net gives the number and refuses the arrow. Register depth $16/1.58 \approx 10.1$. Measured energy ratio of a full-precision against a ternary model of the same class: $0.258/0.028$ joules per token ≈ 9.2 . The same magnitude read twice — the ceiling *is* the register depth. No mechanism story is needed, and none is offered.

Coherence = selection by fit. The Resonant Stack's principle — only what synchronises persists — and the net's core operation — only what returns in phase remains — are one operation. The transformer already performs it in crutch form: the top-k gate keeps what fits and drops the rest, but still pays a learned score to decide fit. Deterministic routing work shows the score can go; the operation stays.

Nilpotent kernel = balanced closure = the trit. $N^2 = 0$ says: admissible states sum to nothing. Balanced ternary is the discrete object with exactly that property: +1 and −1 carry each other, the zero stands lawful in the middle, and a closed winding carries no net tension. The number 43 is written +—+ : one step out, three steps back, one step out — and the address is the record of the walk, not a label pasted on it. The nilpotent kernel of 2025 and the ternary kernel of 2026 are the same closure form in two registers: continuous field, discrete trit.

Optics = the carrier assumption lifted. The 2025 essay chose photonics because light computes in phase. The 2026 theory derives why that choice was right: there is no separate carrier. "Hardware" is deep, slow winding treated as frozen background; computing in light is computing in the windings themselves, without the background/field split. What was a hardware preference in 2025 is a derivation in 2026.

The master's mode = the mesh. The essay asks for broad attention: holding a whole in view, context and margin with the centre. The address kernel supplies the discrete form of exactly that: a prefix is a coarser class, a neighbour is a relative, and truncation is proved stable under deepening — so the coarse view is never invalidated by the fine one. A hash has no such metric; the balanced address does. Knowing-by-place is not an add-on to the maze; it is the maze.

5. The 2025 claims, measured in 2026

Each claim of Section 2 now has a status.

C6 first, because it is the sharpest. The 2025 essay promised a new class of safety argument: impossible, not filtered. That class now exists, machine-checked — delivered first by the discrete hemisphere. In the running MAZE.AI loop, fabrication is not rare; it is impossible by the output type: no constructor exists without an address, the gate theorem holds as an equivalence on the finite-map store (the loop emits exactly what the store holds, both directions), and the proofs are bound to production by a differential test: 30,572 addresses, zero differences, under one control command. The promise of C6 was written for the field kernel; it was first cashed by the ternary kernel. Same closure form, so the same class of guarantee — identity across scale, demonstrated.

C5 is running. Interface pattern 5.1 — the model proposes, the core disposes — is the deployed design of maze.swarp.nl/ai: the language model's answer enters the loop as a typed proposal, may phrase and may not assert, and leaves only as store-backed content or as a vacancy. The corpus callosum of the 2025 essay is in production, with the coherent core cast discretely as store, gate and ledger. And the boundary between the hemispheres is not a metaphor in this construction: it is one named predicate — the total meaning relation over open language, unformalisable after Tarski — and everything on the machine side of it is proved.

C1's three constraints, answered discretely. Temporal fragmentation: in the address kernel, depth is timescale — the shallow prefixes are the slow modes, proved never invalidated by refinement, so long-term memory is intrinsic to the mesh rather than simulated by bookkeeping; the ledger adds the slow mode of settled meaning, append-only and signed. Loss-function myopia: the answer loop has no loss function at all — admission is by certificate, selection by prefix, settlement by a named human. The energy ceiling: confirmed from outside, by the field's own measurements (the 9.2 ratio above) and by its behaviour — the industry is dismantling its

own tables, scores and register depth under exactly the pressure the essay named.

C4, measured in the negative direction. One experiment tried to keep gradient descent while adopting the resonant register: a byte-level model with computed addresses, prefix routing and native ternary weights, trained by backprop under two hashed plans. Result: 2.96 bits per byte against 1.39 for a conventional small model of the same size. Booked honestly and archived. Read through the correlation, the miss confirms C4 rather than refuting the register: the loss function is the last left-brain remnant, and a right-brain register trained by a left-brain method does not cross. The field's own ternary successes still use backprop; nobody yet trains by coherence. That is the open instrument this claim still waits for.

C2 and C3 as substrate: open, and arriving from outside. The oscillatory field itself — photonic Kuramoto modes, the Virtual Resonant Being, multi-scale coupling — is unbuilt here, and the 2025 essay's own hedge holds: photonics is coming either way, as interconnect if not as computer. Every photonic-computing result published is free measurement for this claim. Meanwhile the discrete hemisphere gives the right-brain functions provisional homes: neighbourhood ranking within the image of an address (apprehension by nearness), a weakness flag at a fixed coherence threshold (0.32), and human settlement recorded as an attributed event. Shadows of the field, standing in until the field exists.

One archive-level measurement joins the two. Merging is proved as union in the discrete kernel, and measured on the running net: of 4.08 million distinct addresses built by independent sources, 0.06 percent are shared. Federation without translation — the essay's planetary ambition, demonstrated at the scale of one archive.

6. What the correlation changes

Three things follow.

First, Right-Brain AI stops being an earlier, separate proposal. It is the vacuum.net theory applied to knowing, written before the theory had its words. Every 2025 term has exactly one net term as its translation, and the translation is where the 2026 measurements land in place.

Second, the build order explains itself. The left hemisphere was built first not by accident but because the split-brain literature the essay cites predicts it: a right hemisphere without an articulate left partner perceives

but cannot state. MAZE.AI is the articulate partner — exact, provenance-bound, incapable of confabulation by construction. The essay's own pathology of the isolated left hemisphere — it confabulates and denies — is precisely what the claim discipline removes.

Third, the roadmap of the 2025 essay acquires a measured first milestone. Its 2026–2027 phase asked for a kernel released, a first lattice, and hybrid operation of the two stacks. What stands in September 2026: a proved discrete kernel (740 lines of Lean, three standard axioms, no additions), the proposer/gate hybrid in production, one placement measurement (15/20 against a pre-registered threshold of 14), and one failed, honestly booked training experiment that locates the next open instrument exactly: training by coherence instead of loss.

7. Status, stated once

Theorems: closure and uniqueness of the balanced address, truncation stable under deepening, merge as union, the gate equivalence, claim completeness by type — machine-checked. Measurements: the 9.2 energy ratio (external), 15/20 placement, 2.96 against 1.39 bits per byte, 0.06 percent archive overlap, zero kernel differences. Readings: the term-by-term correlation of Section 4 is a reading — offered as the diagnosis this article exists to state, not as a theorem. Open instruments: training by coherence; the oscillatory substrate; the token-axis question on ternary registers raised by Ouyang et al. The boundary of Section 5 is classical logic applied and is not open.

8. Conclusion

The 2025 essay asked for a whole-brain infrastructure and warned against an emissary ruling alone. The 2026 construction answers with the half that could be proved: a left hemisphere that no longer confabulates, holding the mesh as its memory and one named predicate as its boundary. The field hemisphere remains the open half, and the correlation says where it will attach when it exists: at the proposal interface that is already running, under the same closure form that is already proved. One strand, two windings, one theory.

Annotated References

1. Konstapel, J., "Left and Right Brain AI", constable.blog, 23 November 2025. *Why read?* The primary reference of this article: the full field-side architecture — left-brain stack with its three constraints, Resonant Stack in five layers, the corpus-callosum interface patterns, the roadmap. Every claim measured in Section 5 is stated there. <https://constable.blog/2025/11/23/left-and-right-brain-ai/>

2. **Konstapel, J., "The Resonant Stack: A Paradigm Shift from Discrete Logic to Oscillatory Computing", constable.blog, 19 November 2025.** *Why read?* The core exposition of the field hemisphere: Kuramoto dynamics, KAYS functions, oscillatory computation. The stack whose kernel Section 4 translates into the trit.
3. **Konstapel, J., "Right Brain AI — Technical Architecture", constable.blog, November 2025 (PDF).** *Why read?* The technical companion to reference 1; the specification level of the field hemisphere.
4. **Rowlands, P., "A Universal Algebra and Rewrite System Approach to Physics" (2002); Rowlands, P. & Diaz, B., "Aspects of a Computational Path to the Nilpotent Dirac Equation", Foundations of Physics 37 (2007).** *Why read?* The nilpotent principle: only zero-totality states are stable. Section 4 identifies its discrete counterpart in the balanced register. *Advice:* the 2007 paper is the more readable entry.
5. **McGilchrist, I., "The Master and His Emissary" (2009).** *Why read?* The hemispheric asymmetry synthesis behind the 2025 essay's design heuristic, including the split-brain pathology that Section 6 uses: the isolated left hemisphere confabulates.
6. **Kuramoto, Y., "Self-Entrainment of a Population of Coupled Non-Linear Oscillators" (1975).** *Why read?* The canonical synchronisation model of the field hemisphere — selection by fit in continuous form.
7. **Knuth, D. E., "The Art of Computer Programming", vol. 2, §4.1.** *Why read?* Balanced ternary: range, uniqueness, and the arithmetic of the closure form that Section 4 identifies with the nilpotent kernel.
8. **Tarski, A., "The Concept of Truth in Formalized Languages" (1936).** *Why read?* The theorem behind the one-predicate boundary where the hemispheres meet. *Advice:* a modern exposition suffices.
9. **Konstapel, J., "Where Computation Comes From" (2026).** *Why read?* The transformer read as a net: the four operations already approximated, the two remnants, the constitutive relation register-depth $\approx$ energy-ratio ($10.1 \approx 9.2$) used in Section 4. <https://claude.ai/artifact/2KUPC9w6iqpWM4LxtLZurv>
10. **Konstapel, J., "Provable AI — A Claim-Bounded Criterion, a Semantic Boundary, and a Machine-Checked Construction" (2026).** *Why read?* The criterion and boundary that make the left hemisphere provable, with the Lean core as appendices. <https://claude.ai/artifact/PSG2144zBPND2jQvnfMszE>
11. **Konstapel, J., "MAZE.AI: A Claim-Bounded Answer Loop with Machine-Checked Provenance", whitepaper, 20 September 2026.** *Why read?* The running system and every 2026 number used in Section 5: the gate equivalence, 30,572/0, 15/20, 2.96/1.39, 0.06 percent. Live at maze.swarp.nl/ai.
12. **Ma, S. et al., "BitNet b1.58 2B4T" (Microsoft Research, 2025).** *Why read?* The measured ternary register: 0.4 GB, 0.028 J/token against 0.258 for a comparable full-precision model — the external half of the constitutive relation.
13. **Ouyang et al., "Low-Bit Quantization Scaling" (ACL 2025).** *Why read?* The honest counter-measurement: the low-bit gap grows with training tokens. The token-axis question it raises is listed open in Section 7.
14. **DeepSeek-AI, "DeepSeek-V3 Technical Report" (2024) and its open inference code.** *Why read?* The left hemisphere examined at source: rotary phases, top-k gate, low-rank truncation, one residual stream — the four net operations in crutch form.
15. **Klein, G. et al., "seL4: Formal Verification of an OS Kernel" (2009).** *Why read?* The methodological model for proving a running system and stating exactly what the proof does not cover — the discipline Section 5 applies.