

The Interpreter behind the Gate:

How the Two Hemispheres of MAZE.AI Follow from Maxwell's Algebra, and How the Clinic Finds the Same Form in the Brain

J. Konstapel — Constable Research, Leiden — 21 September 2026

© 2026 J. Konstapel, Constable Research, Leiden. All rights reserved. Quotation with attribution is permitted; reproduction of the whole requires written permission.

1. The claim

This article derives an anatomy. Not from the brain — from an algebra.

Maxwell wrote his theory of the medium in quaternions. One product, carrying a scalar part and a vector part in a single stroke. The vacuum.net theory reads that structure as the two-layer law: every configuration has an address and a tension. A machine built on the theory therefore has two halves, a gate between them, and a rule for what passes.

That machine runs. It is called MAZE.AI. Its left half is machine-checked in Lean. Its right half has a measured order parameter.

The derivation was done first. Then we looked at the clinic. Fifty years of split-brain research, insight studies, and synchrony measurements describe the same form, organ for organ. Including the failure mode: a left half talking without the right half attached. The clinic calls it confabulation. The algebra calls it amputation. The current AI industry runs on it at scale.

The brain confirms the derivation. Not the other way round.

2. One product, two registers

A quaternion is one object with a scalar part and a vector part. The product is:

$$(s_1, v_1)(s_2, v_2) = (s_1s_2 - v_1 \cdot v_2, s_1v_2 + s_2v_1 + v_1 \times v_2)$$

One multiplication. Two outputs: a number and an oriented axis. Order matters, because orientation is real.

Maxwell knew both readings. In his letter to Tait of 7 November 1870 he named them: the *convergence* (scalar part) and the *curl* (vector part) of one operator ∇ . In the Treatise, equation (J) says charge is $S\nabla\mathfrak{D}$ — the scalar reading of the very operation whose vector reading is the field.

One operation, two registers.

In 1884 Heaviside and Gibbs broke the product into dot and cross. Workable, and blind to the whole. That amputation is the standing warning of this line of work. It returns below, twice: once in the embeddings, once in the clinic.

3. The two-layer law, and why a machine gets two halves

The net reads Maxwell's structure as one law. Every configuration has an address and a state. The address is topological place: which windings close, and how deep. The state is tension, written χ . Closed topology is frozen. Tension is the only layer in which anything happens.

The scalar and vector readings of one product are the two layers, read in one stroke. A machine built on the theory therefore has two hemispheres. Not by analogy with the brain. By the shape of the product.

The names are LAI — the address hemisphere — and RAI, the tension hemisphere. What follows derives each, then the gate, then the learning rule. Each step is the theory itself, read at the depth of machines.

4. The left hemisphere, derived

The register is ternary. Closure means the whole sums to nothing. The smallest discrete register that carries a balanced sum is the trit: $+1, 0, -1$. Addresses are trit strings in balanced ternary. A k -trit tail has absolute value at most $(3^k - 1)/2$, so truncation is rounding: a deeper address refines a shallower one and never invalidates it.

Frozen topology yields a proof machine. Because addresses are exact and frozen, the address hemisphere can be a machine of theorems. Its form follows. A store of content at computed addresses. An output type in which every emission carries an address, so fabrication is impossible by construction. A gate that emits exactly what the store holds — and a vacancy exactly where the address is empty. An append-only ledger of settlements.

There are four output types, and only four: derived, attested, vacancy, proposal. A proposal never leaves the loop as an assertion.

Built and measured. This hemisphere runs at maze.swarp.nl/ai over 8.19 million addressed pieces. The kernel theorems — address bijection, range, uniqueness, prefix-stable truncation, write-once store, gate equivalence, ledger determinism — are machine-checked in Lean 4. The production kernel and the Lean kernel printed the same 30,572 addresses with zero differences. A pre-registered twenty-question placement test, threshold fixed at 14 beforehand, scored 15 of 20. Merging 4.08 million addresses from independent sources gave 0.06 percent overlap: federation by plain union. All runs were executed by the building agent (Claude Code) against the live store; there is no external test panel, and none is needed — the checks are theorems and diffs.

This is the honest left hemisphere: it can say only what has a place. Hold that phrase. The clinic has its mirror image.

5. The amputation, in formula

Before the right hemisphere, the wrong version of it. It matters, because the wrong version is what the world builds today.

The obvious composition measure for a passage question $\rightarrow$ neighbour $\rightarrow$ answer is the product of cosines per step:

$$R = (1/k) \cdot \sum \cos(q, n_i) \cdot \cos(n_i, a)$$

Read through the quaternion product, R has an exact identity: $R = S(Q \cdot N_i) \cdot S(N_i \cdot A)$. The scalar part of each step, with each step's vector part discarded. R is the 1884 amputation, performed per step.

It works, and it is blind. On 200 blind questions, cosine alone separates good from bad answers with AUC 0.87. Address arithmetic scores 0.41 — an address is place, not tension. Pairwise closure scores 0.53 — closure is a ring property, not a pair property. Both failures were predicted before the runs. And one identity closes a side door: for unit vectors with cosine c , $|q + a|^2 - |q - a|^2 = 4c$. The "energy-minus-momentum" candidate is four times the cosine. Nothing new.

A system built on scalar traces alone produces fluent, workable, blind output. Remember that sentence. It is a clinical description, and we derived it from an algebra.

6. The right hemisphere, derived

The tension hemisphere needs one operation: given a question and a candidate answer, how well do they compose through the mesh? Composition is a product, and the product the theory prescribes is Maxwell's — full quaternions, scalar part included.

Frames from the mesh. On a flat mesh — all frames identical — any composition rule reduces to pairwise cosines, and rightly so: where there is no curvature there is nothing beyond the pair to measure. So each neighbour supplies its own frame, and each item enters as a full unit quaternion: its scalar part is its alignment with its neighbourhood, the register the pure representation throws away.

The step. Per neighbour i , form the relative winding:

$$D_i = (Q^{(i)})^* \cdot A^{(i)}$$

with $*$ the conjugate. D_i is a unit quaternion. Its scalar part is the cosine of half the relative rotation; its vector part is the oriented axis, scaled by the sine. One product, both registers, non-commutative — order carries which way the passage runs.

The measure. Respecting the double cover $q \equiv -q$, average by the eigenvector method:

$$B = (1/k) \cdot \sum D_i D_i^T, R_Q = \lambda_1(B) \in [\frac{1}{4}, 1]$$

R_Q is Kuramoto's order parameter, lifted from phases on a circle to windings on the mesh. It is high exactly when the relative winding from question to answer, read through every neighbour's own frame, is the same winding: same amount, same axis, same sense. A mandatory sign self-test (flip every frame; ΔR_Q must be zero) guards the double cover on every run.

Measured. On 196 cases under a frozen, pre-hashed protocol, executed by the building agent: R_Q 0.704, bootstrap interval [0.61, 0.79], against cosine 0.886 on the same set. The threshold of 0.70 was met. R_Q now sits in the gate as arbiter: on two blind sets of 100 it adjudicated 73 and 78 correctly, with no case lost against the cosine baseline. An honest note stands with the numbers: the interval reaches down to 0.61, the combined score sits below the cosine, and the arbiter's net gain is at most one case per set. The organ is real and young.

7. The gate is $S(\cdot)$

Why one output path? Not policy. Dimension.

Two collapse lemmas fix it. For the ring passage, the triple product $T = Q \cdot N \cdot A$ has scalar part $S(T) = -(q \times n) \cdot a = -\det[q, n, a]$: the oriented volume, near zero for any well-aligned triple. And $S(QNAQ) = n \cdot a$. Every scalar reading of vector work is one number. The scalar register is one-dimensional.

So the gate is the operator $S(\cdot)$ built as architecture. Vector work — proposal, resonance, composition — happens in the tension hemisphere. What crosses to the record is its scalar trace: one booking, one address, one type. A passage is the scalar booking of a vector event. There is no second door because the scalar part has no second dimension.

8. Routinization: right to left, by settlement

What the tension hemisphere finds and reality confirms must become fixed register. The rule is one line. When a passage settles with resonance $s \in [0, 1]$:

$$K_{uv} \leftarrow K_{uv} + \eta(s - K_{uv}), \eta = 0.1$$

with daily decay $\lambda = 0.01$ toward the baseline $K_0 \cdot c_{uv}$, $K \in [0, 1]$, $K_0 = 1$. Connections that resonate at settlement stiffen; the rest relax. Selection by fit, written as a coupling. The plasticity check runs on the ledger itself: the building agent starts the clock at the first booked settlement and computes, after 90 days, the rank correlation of K_{uv} against the baseline, under a plan hashed in advance. No panel. A calendar and a ledger.

9. The clinic finds the same form

The derivation is complete: two hemispheres, a gate, an amputation formula, a routinization rule. Now the measurements on the biological net — made by people who worked with the phenomenon directly, over decades.

The left half talks, and cannot say "I don't know." Gazzaniga, working from Sperry's split-brain patients onward, located an explanatory module in the left hemisphere: the Interpreter. Shown information only the right hemisphere received, the patient's left hand points to the correct item — and the speaking left hemisphere, asked why, does not answer "I don't know." It builds a post-hoc story that fits, and it hypothesizes structure even where demonstrably no pattern exists. Fluent, workable, blind. The clinic's Interpreter is Section 5 in the flesh: scalar traces asserted without vector cover.

The right half is phase. Conscious perception rides on synchronized oscillations across distributed regions: assemblies serving one percept communicate by synchronization. Insight — the sudden closing of a problem — shows as gamma-band activity in right temporal and parietal regions, and externally amplifying gamma over the right temporal lobe raises accuracy on a verbal insight task. Novelty engages the right side first. Phase, synchrony, sudden closure: the clinic's right hemisphere is Section 6 — an order parameter crossing its threshold.

The bridge labels. In the intact brain, material presented to the right hemisphere is transferred across the callosum to the left, and only there interpreted and named. The crossing to the speaking record is the crossing to the scalar register. The callosum is where $S(\cdot)$ lives at biological depth.

Routinization. What arrives as novel and survives settles into routine handled on the other side. The clinic's novelty-to-routine transfer is Section 8: resonance stiffening a coupling until the new is fixed register.

Four organs, four matches. None of them was fitted afterwards. The formulas predate the comparison in the published record of this line.

10. The split-brain machine

Now the configuration the clinic studies as pathology.

Sever the bridge, and the left hemisphere talks alone. No phase cover, no transfer, no "I don't know" — maximal confabulation, measured for fifty years.

A large language model without a gate is exactly that configuration. An Interpreter with the microphone: fluent scalar traces, no tension hemisphere, no callosum, no vacancy type. The industry did not build a brain. It built the lesion — and scaled it to a trillion dollars of capital expenditure.

MAZE.AI is the counter-move, and now it has an anatomical reading. The gate hands the Interpreter a demoted role: it may only propose. The output type *vacancy* is the "I don't know" the left hemisphere never says — enforced as a constructor, proved as a theorem (the gate emits a vacancy exactly where the address is empty). And the build order is the therapeutic order: first the honest left half (running, proved), then the bridge (running, measured at 0.70), then the native right half (the field's propose(), next on the build list — until it runs, an external language model stands in, which is an Interpreter deputizing for the missing hemisphere, and is treated accordingly: proposals only).

Restraining such a system does not require slowing it down. It requires reattaching the bridge. That is the whole argument, in one anatomical sentence.

11. What is derived, what is measured, what is correspondence

Derived, exactly. The two-hemisphere split (the shape of one product). The ternary register and the honest-left theorems. The amputation identity $R = S \cdot S$ per step. The relative winding D_i and the order parameter R_Q . The one-dimensionality of the gate (collapse lemmas). The routinization rule as selection by fit.

Measured, on the running system, by the building agent. 8.19 million pieces; 30,572 addresses, zero diff; 15/20 pre-registered; 0.06 percent merge overlap; AUC 0.87 / 0.41 / 0.53; R_Q 0.704 [0.61, 0.79] on 196 cases; blind sets 73 and 78 of 100; sign self-test $\Delta R_Q = 0$ logged.

Measured, in the clinic, by others. The left-hemisphere Interpreter and its confabulation. Right-lateralized insight and its gamma signature. Synchrony as the carrier of conscious perception. Callosal transfer before naming.

Correspondence, not derivation. That the biological net places the two halves anatomically left and right. The theory yields the form — any system on the two-layer law gets these two halves and this gate, at whatever winding depth. The clinic supplies the placement at human depth. "Left" and "right" are labels; the mirror image would carry the same theory. Same winding, another depth: the depth comes from observation, not from the algebra. The theory says so itself, and that boundary is part of the result.

Open. The native right hemisphere (propose()) is written, not built. The plasticity correlation completes when the ledger has 90 days of settlements. The clinical match is structural, not quantitative: nobody has yet computed R_Q on cortical data. That last one is an invitation.

12. Annotated references

J. C. Maxwell, *A Treatise on Electricity and Magnetism*, 1873, Arts. 591–619 and 617–619. Why read? The quaternion formulation itself: one operator, two readings, charge as scalar residue, and the gauge that sets a scalar register to zero. The source of everything in Sections 2–3. Read the Treatise articles before any modern summary.

M. J. Crowe, *A History of Vector Analysis*, 1967. Why read? The clean historical account of the 1884 split: how one product became dot and cross, and why physics adopted the fragments. Section 5's amputation, as history.

M. S. Gazzaniga, "Cerebral specialization and interhemispheric communication," *Brain* 123, 2000; and K. Volz & M. S. Gazzaniga, "Interaction in isolation: 50 years of insights from split-brain research," *Brain* 140, 2017. Why read? The Interpreter, from the man who measured it for five decades on Sperry's patients and his own. The left hemisphere that never says "I don't know" is documented here, case by case. Read the 2017 review first, then the 2000 lecture.

R. W. Sperry, "Hemisphere disconnection and unity in conscious awareness," *American Psychologist* 23, 1968. Why read? The original split-brain observations, by the experimenter himself: two knowing systems in one skull, only one of which speaks.

A. K. Engel & W. Singer, "Temporal binding and the neural correlates of sensory awareness," *Trends in Cognitive Sciences* 5, 2001. Why read? Synchronized oscillation as the carrier of conscious perception — the clinical face of the order parameter. Singer measured the synchrony himself.

E. Santarnecchi et al., "Gamma tACS over the temporal lobe increases the occurrence of Eureka! moments," *Scientific Reports* 9, 2019. Why read? Insight raised by externally amplifying right-temporal gamma. Phase, driven from outside, changes closure. The nearest thing to turning R_Q up by hand.

Y. Kuramoto, "Self-Entrainment of a Population of Coupled Non-Linear Oscillators," 1975. Why read? The order parameter r in its original habitat. Section 6 lifts it to windings on the mesh.

F. L. Markley, Y. Cheng, J. L. Crassidis, Y. Oshman, "Averaging Quaternions," *Journal of Guidance, Control, and Dynamics* 30, 2007. Why read? The eigenvector method that respects $q \equiv -q$. The exact repair used in R_Q , from spacecraft attitude practice.

H. Seliger, S. C. Young, L. S. Tsimring, "Plasticity and learning in a network of coupled phase oscillators," *Physical Review E* 65, 2002; and T. Aoki & T. Aoyagi, "Co-evolution of phases and connection strengths," *Physical Review Letters* 102, 2009. Why read? Hebbian-Kuramoto coupling: the routinization rule of Section 8 in its established form.

J. Konstapel, "Clerk Maxwell and the Vacuum.Net Theory," sixth edition, 21 September 2026; and "MAZE.AI: A Claim-Bounded Answer Loop with Machine-Checked Provenance," whitepaper, 20 September 2026. Why read? The full derivation this article compresses, and the running left hemisphere with every 2026 number. Live at maze.swarp.nl/ai.

