

title: "Progress Report: The MAZE Project" subtitle: "What was investigated, proved and built, 8 August – 23 September 2026" author: "J. Konstapel, Leiden, 23 September 2026"

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

0. Summary

MAZE applies the Vacuum.Net theory to the human world and to knowing. It started on 8 August 2026. In seven weeks it went through seven phases. It produced an address system, three instruments, a model of the human, a knowledge net of 8.19 million addressed pieces, a machine-checked answer machine that runs live, a composition rule measured in production, and today an exact theory of closure dynamics.

This report records every step in detail. Each statement carries one status:

- **Proved** — a mathematical result with its proof in the text or machine-checked.
- **Built and measured** — running software with numbers. The runs were performed by Claude Code; the figures are self-reported until independently reproduced.
- **Corresponded** — a declared identification between two vocabularies. It can fail.
- **Predicted / tested** — a claim put against data.
- **Open** — named, not yet done.

The fixed image is the fishnet. The vacuum is a net of closed meshes. A mesh is the smallest closed unit. Strands carry winding. Tension lives on the net and relaxes; topology does not. What persists is what closes.

Phase	Dates	Main result
1. Name and ground	8–10 Aug	Address system in balanced ternary; kernel in code
2. Instruments	10 Aug – 5 Sep	Tension Gauge live; space-weather and business layers
3. The human	12–13 Aug	The Living Winding: model, mapping, protocols, self-calibration
4. The knowledge net	24–28 Aug	Wikipedia, music and art in one address space; public MAZE
5. maze.ai, left hemisphere	18–20 Sep	Machine-checked answer loop, live, 15/20 test
6. Maxwell and the right hemisphere	20–22 Sep	R_Q derived and measured (AUC 0.704), in the gate
7. Closure without alignment	23 Sep	Exact closure dynamics; Maxwell bridge; first data

1. What MAZE is

The name comes from the labyrinth of Chartres. It is not a maze of dead ends. It is one path, ring after ring. There is no wrong turn, only a next ring at a computable distance. In Dutch the word sounds like *maas*: the mesh, the smallest closed unit of the net.

The programme was stated on 8 August in *MAZE: Applying the Vacuum Net Model*. Vacuum.Net is the model of the whole net. MAZE works out its human consequences: collectively through markets, individually through the blueprint, and for knowing through addressed knowledge.

Two architectural facts were fixed at the start. MAZE is built outside SWARP and tested separately. And every result is reproducible from public data or from the store itself. There is no office and no gatekeeper; verification runs through the ledger and through practice.

2. The ground: balanced ternary and the address

2.1 The register

A self-crossing of the strand has three states: over (+1), under (−1), none (0). The digits of balanced ternary are exactly these three values. MAZE writes every place in the net as a string of such trits.

Worked example: 43. Divide by 3 repeatedly and choose remainders in {−1, 0, +1}:

- $43 = 3 \cdot 14 + 1 \rightarrow \text{digit } +1$
- $14 = 3 \cdot 5 - 1 \rightarrow \text{digit } -1$
- $5 = 3 \cdot 2 - 1 \rightarrow \text{digit } -1$
- $2 = 3 \cdot 1 - 1 \rightarrow \text{digit } -1$
- $1 = 3 \cdot 0 + 1 \rightarrow \text{digit } +1$

Read from most significant: + − − − +. Check: $81 - 27 - 9 - 3 + 1 = 43$.

2.2 Theorem T7 (Unique Address) — proved

Every integer has exactly one finite balanced-ternary representation with a nonzero leading trit.

Existence. Write the number in ordinary ternary. Replace every digit 2 by a carry of 1 and a digit −1, since $2 \cdot 3^i = 3^{i+1} - 3^i$. Each step removes a 2 at or below the current position, so the procedure ends.

Uniqueness. Suppose two strings give the same number. Their difference gives $\sum d_i \cdot 3^i = 0$ with $d_i \in \{-2, \dots, 2\}$, not all zero. Let k be the highest index with $d_k \neq 0$. Then $|d_k \cdot 3^k| \geq 3^k$. The rest is at most $2 \cdot (3^0 + \dots + 3^{(k-1)}) = 3^k - 1$. The sum cannot vanish. ■

2.3 Consequences — proved

- **Mirror.** Negation is trit-wise: swap +1 and −1. The address of −43 is − + + + −. No minus sign is imported.
- **Capacity per depth.** The largest value at depth n is $(3^n - 1)/2$: 1, 4, 13, 40, 121, ... Exceeding it forces the next depth.
- **Truncation is rounding.** A tail of k trits is at most $(3^k - 1)/2$ in absolute value, less than half the next unit. Cutting a string is rounding to nearest. A deeper address refines a shallower one and never invalidates it. The address layer never has to be re-indexed.
- **Identity over scale.** Leading zeros change nothing. A shift of the string is one scale step.

- **Radix economy.** The cost of radix r is proportional to $r/\ln r$. For $r = 2$: 2.885. For $r = 3$: 2.731. For $r = 4$: 2.885. The continuous optimum is $e = 2.718$. Three is the nearest integer.

2.4 The kernel in code — built

On 10 August the MAZE kernel entered the SWARP repository as `server/maze-kernel/maze.ts`: 114 lines, no database dependency, BigInt arithmetic so the bijection stays exact at every depth. It has 96 lines of tests that recompute the exact figures of the paper. It holds turn, route, the bijection route $\leftrightarrow$ integer, mirror and truncation. The rewrite dynamics was deliberately left out: that is open problem O2.

3. The instruments

3.1 The Tension Gauge (Weather Instrument) — built, live

The first instrument reads the tension of the human sphere from markets. It went live on weer.swarp.nl.

- **Input.** 17 indices: the VIX and 16 regional exchanges in three equally weighted regions.
- **Score.** $S = \text{clamp}(50 + \Delta \cdot 15)$, from 0 to 100, refreshed every five minutes.
- **Labels.** Calm, unsettled, turbulent, storm.
- **Discharge law.** A falling price is loading; a rising price is discharge. Thresholds are calibrations, treated as state functions.
- **Test engine.** A paper-trading simulation with €10,000 virtual capital: buy when $C < 42$, sell with profit when $C > 58$, damped by the VIX.

A structural insight followed on 12 August. The score is a deterministic formula on public data. Its whole history can therefore be reconstructed without an app log. The app is only the live end of the formula.

3.2 Space weather — designed

The working paper *MAZE Space Weather* (12 August) extends the gauge into a layered early-warning system:

- tension (the existing gauge) as the base layer;
- the VIX;
- market loading;
- space loading: the geomagnetic A_p index with threshold $A_p > 29$, and a six-day memory with decay weights from 1.00 to 0.25;
- human sentiment from GDELT (tone, pessimism, attention);
- a discharge detector.

The test design is four nested models, A to D (base, plus sentiment, plus space, plus both), with a 60/20/20 out-of-sample split. One lesson was recorded. An external AI first fabricated correlations ($r = -0.58$) and then retracted them. The own-data run is designed and not yet performed.

3.3 The institutional layer — designed

On 4 September a slow layer was added: seven public Dutch series (bankruptcies and start-ups, vacancy rate, building permits, house price to income, the NL–DE spread, energy spot price, AEX-

relative), a frozen reference period 1990–2018, a monthly cadence, the same four weather states, and a nested test variant E with a 1–12 month horizon.

3.4 Scope

On 13 August the scope was fixed. The instruments are scaffolding. MAZE is the application of Vacuum.Net to the human and to knowing. The instruments stay as the live end of the model on public data.

4. The human: The Living Winding

4.1 The model — proved within its definitions

The model was rebuilt from the base model in mathematical language: eight definitions and one postulate (diffusion on the graph). A human is a throughflow closure. Four quantities describe it:

- **N** — the identity pattern: an equivalence class of chord diagrams on the strand. The realisation (the actual set of contacts) is replaced; the class persists.
- $\bar{\chi}$ — the mean state (tension).
- τ — the retention time.
- **J** — the throughflow.

Life is closure within throughflow: $dC/dt = \mu J - \lambda C$. Drift and repair give a duty cycle $f = vJ/(vJ + r)$. The steady-state capacity is:

$$C^*(J) = \mu r J / [\lambda(vJ + r)]$$

It rises with J and saturates at the ceiling $\mu r/(\lambda v)$. Its slope is $dC^*/dJ = \mu r^2/[\lambda(vJ + r)^2] > 0$. There is a sustainable window of throughflow, with a lower edge $J_{\min} = \lambda C_{\min} \cdot r/(\mu r - \lambda v C_{\min})$ and an upper edge set by the supply of the environment.

A tower lemma places the retention time of a group above the capacity-weighted harmonic mean of its members.

The human scale does not follow from Vacuum.Net alone. It needs the closure law, a spectral condition, the window, and one chosen temporal selection criterion (choice 3b). That dependency is stated in the definition.

4.2 How it was tested

Eight editions went through seven adversarial stress tests. Each round ended in a correction written on the record: a normalisation error, a low-load approximation, a missing reservoir equation, a nesting problem in the model comparison. After round seven no algebraic inconsistency remained. The theory phase was closed.

4.3 The mapping onto the sciences of the human — corresponded

Each discipline received a table of model term against its own term, with status exact, partial or empty:

- autopoiesis = throughflow closure;
- Schoenheimer's dynamic state = the split between class and realisation;
- allostatic load = accumulated drift;
- emotional inertia in affect dynamics = τ , measured for twenty years;
- Borbély's Process S = the duty cycle f ;
- Gibson = perceiving without a representation step;
- Simon's near-decomposability (1962) = choices 1 and 3b;
- Durkheim's collective effervescence = the group repair phase;
- critical slowing down = the retention ratio approaching 1.

Four terms have no preimage: representation, cause between things, the isolated subject, disease as an entity.

4.4 Protocols and self-calibration — designed

Five predictions T1–T5 were written as preregistrable tests in two tracks. Track A re-analyses existing data: T1 on experience-sampling data, T3 on sleep and workload data, T5 on dyad data. Track B designs new tests: T2 load titration, T4 multi-scale registration over 90 days.

The *Self-Calibration Guide* turns the model into an instrument anyone can use: one tension score twice a day, a personal baseline from two ordinary weeks, and five procedures to measure one's own τ , rest curve, window and pair coupling. It has six honesty rules; the first is that thresholds are set before the data.

Status: the theory is closed. None of T1–T5 has been run.

5. Waste and alchemy

5.1 The exponent law — proved

A production chain is a stretched mesh. Each transition keeps a fraction q . Over k transitions, q^k remains. With $q = 0.85$: ten links keep 20%, twenty links keep 4%. The loss sits in the number of steps. Optimising a step does not touch the exponent. Unpriced transition costs land at the nodes held below their window.

5.2 The criterion and three doors — open

On 13 August a hard criterion was set: either the vacuum can be manipulated, or the line stops. Manipulation means a commanded, measured change of the local state function $g(\chi)$. Three places were named where the medium may yield, each with an instrument and a kill criterion:

1. **Threshold softness.** Near $\chi \approx 1$, as in the radial-acceleration relation; optical clock pairs below 10^{-18} as meter.
2. **Phase selection.** Distance to bifurcation as the design variable.
3. **Resonance.** Sonoluminescence as a natural soft spot where no state-function reading has ever been done.

Status: speculative, with kill criteria. Not measured.

6. The knowledge net

6.1 Reengineering the Universe V3 — built

MAZE became a trit-addressed knowledge net. A route stops where its period closes. A periodic route has the closed form $v_p/(3^p - 1)$. The English Wikipedia, all music and all art were loaded into one address space. The public MAZE runs on maze.swarp.nl, with a memory-palace interface in which every address is a URL. The census of 28 August counted 6,691 inhabited windings and 3.6 million pieces from Wikipedia (65 fields as the first round), the Metropolitan Museum, three music corpora and constable.blog. A daily reasoning loop records vacancies.

6.2 The carry as rewrite rule — proved as candidate for O2

Balanced-ternary addition has a carry. That carry is the candidate rewrite rule R. The residue stays; the overflow moves one weight deeper. Properties:

- $R \cdot R = 0$ follows from the digit range;
- nilpotent up to four simultaneously loaded addresses;
- termination in at most $n + 1$ rounds;
- depth grows by at most one ring;
- closing happens before coarsening.

Two numbers describe a rewrite: σ , the initial overload, and ϱ , the number of rounds. Worked sums in the essay *How to Calculate Science*: $43 + 25 = 68$, $40 + 40 = 80$, $1 + 1$, $1 + 1 + 1$.

6.3 Art as route — corresponded

A work is a route, not a knot. A tradition is a periodic route. The coding rule per transition is over, straight or under: where the new ordering principle stands relative to the old. Western music in six transitions gives $81 - 27 - 9 + 3 - 1 = 47$ at depth 5. Its mirror -47 is the cheapest test. The carry $47 + 47 = 94$ gives $\sigma = 5$ and $\varrho = 1$.

A blind-coding test by two independent readers is designed. It has not been performed. Disagreement on the first transition costs $2 \cdot 3^{(d-1)}$; on the last, 2. Blind coding must fix the first transitions above all.

6.4 The memory palace — built

The interface became a modern memory palace, taking up the proposal in the author's 2010 blog *About Magic and the Memory Palace*. Camillo's chronological order is completed by the number system. Plato's objection in the *Phaedrus* is answered by walking instead of retrieving. The dialogue has five verbs: where, neighbour, mirror, together, up.

6.5 A finding about the coder

The production field coder is a layer vector, not the ring rule of the blog. The self-test in *The Address of a Proof* gave that essay the signature -144719454 at depth 19 and showed the coder to be too coarse for proofs. This finding shaped the next phase.

7. maze.ai: the left hemisphere

7.1 The design principle

The language model is the counter clerk, never the source. The system claims one relation only: that its output was derived from its store. It claims nothing about world truth or meaning. An assertion without an address becomes a vacancy, not an answer. The vacancy is the enforced "I do not know".

7.2 The six theorems — proved

The Six Theorems of the MAZE (19 September) proves in full:

1. unique address (T7);
2. truncation is rounding, so no re-indexing is ever needed;
3. merging stores is set union;
4. the gate invariant: every emission is content at an address, or a vacancy;
5. prefix partition into 3^m classes is balanced by construction;
6. radix economy: $2.731 < 2.885$.

It also proves a limit. The step from language to address cannot be proved correct, because meaning is not a formal object. It can be measured. In a weight-based AI the unprovable part is the whole machine. In MAZE it is one measurable function.

7.3 The Lean kernel — proved by machine

`MazeAI.lean` and `check.lean` were accepted by Lean 4 on its three standard axioms. The obligations:

- O1: $\text{decode}(\text{encode } n) = n$;
- O2: the truncation lemma, both directions;
- $\text{merging} = \text{union}$;
- O3: the gate in both directions (delivered content is exactly in the store; a vacancy only at a truly empty address);
- O5': the ledger chain step;
- O6: completeness per output type.

The executable self-test gives $\text{encode } 43 = + - - - +$.

7.4 The article and its review

Provable AI — A Claim-Bounded Criterion, a Semantic Boundary, and a Machine-Checked Construction went through four AI reviews. All points of the strongest review were adopted. The fixed vocabulary is non-fabricating, store-faithful and claim-bounded. The phrase "cannot lie" was removed: membership of the store is not world truth. The nearest existing work (attributed and verifiable generation) is placed in a comparison table. A Dutch blog, *De AI die niets verzint*, and a specification followed.

7.5 Built and measured

Claude Code built the system. It runs live on maze.swarp.nl/ai, with one output path, including the home page. Figures from the whitepaper of 20 September:

- 8.19 million addressed pieces, including Mathlib theorems and EPO patents;
- Lean 4, about 740 lines in six files, three standard axioms;
- a finite-map store with the gate invariant as an equivalence and write-once as a theorem;
- a sha256 + HMAC chained ledger;
- differential testing of the production kernel against the Lean kernel: 30,572 addresses, zero differences;
- a preregistered twenty-question placement test: 15 of 20 against a threshold of 14 fixed beforehand; the weakness flag caught 4 of the 5 misses;
- merging 4.08 million addresses from independent sources: 0.06% overlap, so federation is plain union;
- one failed attempt: training an own three-in-one model gave 2.96 against 1.39 bits per byte; archived as candidate host.

Open: the formal proof of O4 (verifier soundness), personal keys, and the gap between image address and exact address.

8. Maxwell and the right hemisphere

8.1 The record — historical

Maxwell wrote a constitutive theory of a medium. In Arts. 591–619 of the *Treatise* (1873) the fundamental equations stand as one group in quaternions. That form carried a scalar register and medium moduli. In 1884–1885 Heaviside and Gibbs cut both away. The six editions of *Clerk Maxwell and the Vacuum.Net Theory* document this and rewrite the equations into the net. The speed of light entered physics as a computed modulus: 193,088 miles per second against the light-speed figures of the day.

8.2 The architecture — derived

From the two-layer law (address and tension) follow two hemispheres and one boundary:

- **LAI**, the address hemisphere: frozen, exact, provable. This is maze.ai.
- **RAI**, the tension hemisphere: continuous, phase-carrying, selection by fit.
- **The gate**: the only relation between the layers is a reading. The model proposes; the store supplies; one output path.

8.3 Why the cosine is not enough — proved

The running system ranked by cosine: AUC 0.87 on 200 blind questions. The cosine is the scalar part of the quaternion product alone. Two identities close the side doors:

- **The 4c identity**. For unit vectors, $|q + a|^2 - |q - a|^2 = 4c$. "Energy minus momentum" is four times the cosine.
- **Pure-quaternion collapse**. Pure quaternions on one shared frame reduce every product to pairwise cosines or a near-zero determinant. The scalar path score $R = (1/k) \cdot \sum \cos(q, n_i) \cdot \cos(n_i, a)$ is the amputation per step.

Two shortcuts failed as the theory said they would: arithmetic on addresses, AUC 0.41 (an address is place, not tension); pairwise closure, AUC 0.53 (closure is a property of a ring).

8.4 The composition rule R_Q — derived and measured

Per neighbour a frame; each item as a full rotation quaternion; the relative winding $D_i = (Q^{(i)})^* \cdot A^{(i)}$. The ring measure is sign-invariant (Markley): $B = (1/k) \cdot \sum D_i \cdot D_i^T$, $R_Q = \lambda_1 \in [1/4, 1]$. Self-test: flipping $D_i \rightarrow -D_i$ must leave R_Q unchanged. Worked example: $R_Q = 0.826$.

Measured by Claude Code:

- the neighbour-centre frame gave $R_Q \equiv 1$: the flat-mesh lemma, observed;
- replaced by the region frame per inhabited sister cell (depth-5 prefix, 24 gate axes);
- result: AUC 0.704 [0.61–0.79] on 196 cases, beside the cosine at 0.886;
- threshold 0.70 reached; sign self-test $\Delta R_Q = 0$;
- since plan 17, R_Q stands in the gate as referee (blind sets: 73 and 78 of 100);
- field coherence scored AUC 0.52 and stays with the field, not the gate;
- $K_0 = 1$ fixed.

8.5 Phase of a piece and learning rule — written

The phase of a piece is $\phi(x) = \text{atan2}(x \cdot e_2, x \cdot e_1)$ in its region frame. The orientation of the frame is fixed by the earliest-admitted piece: memory fixing chirality. The learning rule is Hebbian–Kuramoto plasticity driven by the ledger: on each settlement $K_{uv} \leftarrow K_{uv} + 0.1 \cdot (s - K_{uv})$; daily decay toward $K_0 \cdot c_{uv}$ with $\lambda = 0.01$. A 90-day Spearman test is fixed in advance. It is open.

8.6 Further steps

- **Register depth.** 16 bits against 1.58 bits per trit gives 10.1. A published ternary-versus-full-precision energy comparison gives 0.258 against 0.028 joules per token, ratio 9.2. The two models differ; the pair is an external comparison, not an isolation.
- **World specification.** Measurement M1 as a per-domain switch: field against external model through the same gate, blind, $n = 200$. Rollout in three rings: finance, then mathematics and engineering, then language after M1.
- **The Interpreter Behind the Gate.** Gazzaniga's left-hemisphere interpreter, which never says "I do not know", corresponds to a language model without gate. The vacancy is that sentence. The left/right placement is a correspondence, not a derivation.

9. Closure without alignment (23 September)

9.1 Setting

A ring of N positions carries trits T_i . The window is (T_{i-1}, T_i, T_{i+1}) . Its charge is $D_i = T_{i-1} + T_i + T_{i+1} \bmod 3$. The window is closed when $D_i = 0$. Of 27 triples, 9 are closed: 6 heterogeneous (all three values) and 3 homogeneous (a, a, a). Each position updates with probability p . The closure rule sets $T_i \leftarrow -(T_{i-1} + T_{i+1}) \bmod 3$. The copy rule sets T_i to one of its neighbours.

9.2 Exact results — proved

- **First-step law.** From a random ring: $C_1(p) = 1/3 + (2/3) \cdot p(1 - p)^2$. Only when the centre moves and both neighbours stand still is the window certainly closed; every other case leaves it uniform. The maximum lies exactly at $p = 1/3$, where $C_1 = 35/81 \approx 0.432$.
- **Synchronous updating is chance.** At $p = 1$ the new window sum contains $T_{\{i-2\}}$ with coefficient -1 , which makes it uniform. $C = 1/3$.
- **The self-closing register.** An update at i with charge q sets D_i to 0 and subtracts q from both neighbouring charges. A charge splits into two charges $-q$; in balanced ternary $(-1) + (-1) \equiv +1$. Charge is conserved mod 3. Noise adds δ to three windows: $\delta + \delta + \delta \equiv 0$. Checked on 20,000 random updates: 20,000 exact for mod 3; 3,023 failures for an integer-sum rule.
- **The limit of one dimension.** The charge of a block depends only on four boundary trits: $T_{\{a-1\}} + 2T_a + 2T_b + T_{\{b+1\}}$. There is no bulk invariant on a line.

9.3 Measured in simulation

$N = 1000$, $p = 1/3$, noise 0.05, five seeds:

	R (alignment)	C	C_{het}
random	333	333	222
closure rule	323	484	333
copy rule	681	512	38

Closure and alignment are independent order parameters. C_{het} separates them; total C does not. The exact stationary state on a ring of 7 (2,187 states) places the long-run optimum on a plateau at $p \approx 0.2-0.35$. A control test: a plain integer zero-sum rule also raises C_{het} (0.366). What is unique to mod 3 is the self-closing register, not the closure level.

9.4 Derived from the axioms — proved

- Among linear symmetric window conditions $a + b + c \equiv k$, only $k = 0$ is invariant under shifting all trits (no preferred state) and under mirroring (no preferred orientation).
- Fox's crossing condition $a + c \equiv 2b \pmod{n}$ is symmetric in all three arms only when $3b \equiv 0$ for every b , so $n = 3$. For one strand crossing itself, three is not only the cheapest register; it is the only one without a privileged arm.
- With fixed neighbours exactly one value closes a window (N2). A crossing is a separate event (N1), which gives asynchronous updating.

9.5 The Maxwell bridge — proved in the trit register

Since $3T \equiv 0$, the window sum equals $T_{\{i-1\}} - 2T_i + T_{\{i+1\}}$: the discrete Laplacian. With T as scalar potential, D is charge, as in Poisson. Since $\frac{1}{2} \equiv -1 \pmod{3}$, the closure rule is averaging over neighbours: relaxation toward a harmonic state. The shift of all trits leaves D unchanged: gauge freedom. The block-charge boundary term is Gauss's law. The linearised nearest-neighbour Kuramoto coupling carries the same Laplacian, so the trit register and the phase register share one operator.

The closure condition is also the colouring rule in the proof of theorem T1 of the Foundational Paper: at each crossing one colour or all three.

9.6 First test on real swarm data — tested, negative

Data: the open fish-school trajectories of Jhawar et al. (15, 30 and 60 fish, 11 trials). Trit per fish per frame: left, straight or right, with tertile thresholds. Window: a fish and its two nearest neighbours. Null model: trits shuffled within each frame. Result: alignment above chance in 11 of 11 trials; heterogeneous closure below chance in 11 of 11 (8 intervals exclude zero). These fish follow the copy rule. The swarm prediction fails for this data and this trit definition.

9.7 Consequences for the build

- The daily relaxation of all couplings K_{uv} at once is a synchronous step. The first-step law suggests relaxing a random fraction of about a quarter to a third per round.
- The cosine is alignment. A closure column on triples (question, answer, neighbour) measures something else. It is a candidate for the existing protocol.

10. The human layer

Three results reach people directly. The answer machine gives an honest "I do not know" where others invent; a reader can trace every answer to its address. The self-calibration guide lets anyone measure their own retention time and rest curve with a notebook, without an institution in between. And the memory palace gives one shared ground plan of knowledge, music and art, with a wing each person can furnish with an own archive.

11. Status ledger

Status	Items
Proved	T7 and its consequences; six theorems; Lean kernel O1–O3, O5', O6; carry properties; Living Winding algebra; exponent law; 4c identity and collapse lemmas; closure first-step law; synchronous chance; charge register; boundary term; axiom derivation; $n = 3$
Built and	kernel; Tension Gauge; public MAZE; maze.ai live; 8.19 M pieces; 0 differences on 30,572 addresses; 15/20; 0.06% overlap; R_Q 0.704 in the gate; closure simulations
Corresponded	the human mapping; art as route; Maxwell dictionary; hemispheres; Vacuum.Net as a swarm of relations
Tested	fish schools: copy, not closure (negative)
Open	O2 (full rewrite dynamics); O4 formal; image-vs-exact gap; 90-day plasticity test; M1; blind coding; T1–T5; models A–D; the three doors; O1; a 2D knot test of closure; closure

12. What follows

In the fixed working order:

1. **Publish.** This report; an updated *Closure Without Alignment* with sections 9.4–9.6; the Dutch blogs owed for the Maxwell article and *The Interpreter Behind the Gate*.
2. **Build.** Claude Code evaluates asynchronous relaxation in the learning rule and the closure column as a candidate under the frozen protocol.

3. **Measure.** The 90-day plasticity test; M1 per domain; the blind-coding test; the formal O4 proof.

Each MAZE is one step further. This report fixes where the step now stands.

Annotated references

Konstapel, J. (2026). *The Vacuum.Net Theory: Foundational Paper*, sixth edition. constable.blog. Why read? The axioms N1–N5, theorem T7 with its proof, the five epistemic statuses and the open calculations O1–O4. Everything in this report builds on it. *Reading advice:* Parts I–III first. Part II answers "why three".

Konstapel, J. (2026). *Clerk Maxwell and the Vacuum.Net Theory*, sixth edition. Why read? The derivation of LAI, RAI and the gate, and the composition rule R_Q with its measurement. *Reading advice:* Part III for the architecture, Part IV for R_Q, Part V for the learning rule.

Konstapel, J. (2026). *MAZE.AI: A Claim-Bounded Answer Loop with Machine-Checked Provenance*. Whitepaper, 20 September. Why read? The measurement report of the running system: every number in section 7.5. *Reading advice:* Read beside the Provable AI article. The article is the criterion; the whitepaper is the measurement.

Konstapel, J. (2026). *Provable AI — A Claim-Bounded Criterion, a Semantic Boundary, and a Machine-Checked Construction*. Why read? The exact line between what can be proved in an AI and what can only be measured. *Reading advice:* The status table and the claim ladder carry the argument.

Konstapel, J. (2026). *The Living Winding: One Model, Seven Variables, and a Mapping onto the Sciences of the Human*. Why read? The model of the human in full, with the mapping onto eight disciplines. *Reading advice:* Start with the mapping tables; they show where existing measurements already fit.

Konstapel, J. (2026). *Closure Without Alignment*. Why read? The exact closure theory of section 9. *Reading advice:* Theorems 1–4 are short and complete.

Maxwell, J. C. (1873). *A Treatise on Electricity and Magnetism*, Vol. II, Arts. 591–619. Why read? The quaternion grouping of the field equations: the scalar register the net restores. *Reading advice:* Read Maxwell's own names for the quantities: elasticity, displacement, momentum.

Fox, R. H. (1970). "Metacyclic invariants of knots and links." *Canadian Journal of Mathematics* 22. Why read? The source of the colouring condition that turns out to be the closure rule. *Reading advice:* The definitions suffice.

Knuth, D. E. (1997). *The Art of Computer Programming*, Vol. 2, §4.1. Why read? Balanced ternary: uniqueness, truncation as rounding, radix economy. *Reading advice:* A few pages; the examples are worth doing by hand.

Markley, F. L. et al. (2007). "Averaging Quaternions." *Journal of Guidance, Control, and Dynamics* 30(4). Why read? The sign-invariant eigenvector form used in R_Q. *Reading advice:* The short derivation of the accumulator matrix is the core.

Gazzaniga, M. S. (2000). "Cerebral specialization and interhemispheric communication." *Brain* 123. Why read? The left-hemisphere interpreter that never says "I do not know": the clinical

counterpart of a model without gate. *Reading advice*: The split-brain observations are the relevant part.

Simon, H. A. (1962). "The Architecture of Complexity." *Proceedings of the American Philosophical Society* **106**. *Why read?* Near-decomposability: the strongest existing motivation for the time-scale choices of the Living Winding. *Reading advice*: Read the parable of the two watchmakers.

Jhavar, J. et al. (2020). "Noise-induced schooling of fish." *Nature Physics* **16**. *Why read?* The open fish data of section 9.6, and the authors' own finding that the fish copy pairwise. *Reading advice*: Compare their pairwise-copying result with the negative closure test.

Konstapel, J. (2010). *About Magic and the Memory Palace*. **constable.blog**. *Why read?* The programme text for the memory-palace interface, sixteen years earlier. *Reading advice*: The closing proposal is the one MAZE carries out.