Cotterill and the Vacuum.Net Theory: The Brain as a Maze of Conditional Passage

J. Konstapel, Leiden — 13 September 2026 version 3.

Summary

In 2001 the Danish physicist Rodney Cotterill published a complete circuit reading of the mammalian brain. He read it as one machine with one output: movement. Every component either permits or refuses passage toward that output. He did not know the Vacuum.Net theory. Yet his circuit contains its central structures: a three-valued register, four states of tension, closure through the environment, and time quantization per layer.

This essay documents the correspondence point by point. It places Cotterill next to Pentti Kanerva, whose Sparse Distributed Memory independently found the address structure of the net. It then follows the successors who carried Cotterill's components forward after his death in 2007. Finally it composes the findings into what this line of work was after: a detailed architecture for a self-learning system — one that keeps learning without destroying its own memory.

Every claim in this essay carries a status label: **[theorem]** for proven mathematics, **[measurement]** for published empirical results, **[hypothesis]** for the net reading of those results, **[engineering]** for a design statement.

1. The Vacuum.Net Frame in Brief

The reader needs no other document. The frame is this.

Reality is read as one strand that winds. A winding that returns in phase closes. A closure is a knot. Knots connected by shared strand form a net. Picture a fishnet: knots and meshes. The knots are the frozen part. The meshes carry the tension.

The two-layer law: **address = frozen topology, tension = dynamic state.** What has closed no longer moves; it locates. What has not closed moves; it loads. **[hypothesis, foundational]**

Tension passes through four states: **loading** (tension accumulates), **holding** (accumulated tension persists), **discharge** (tension converts to rapid movement), **rest**. Discharge at sufficient depth writes new closure. Tension becomes address. **[hypothesis, foundational]**

The register of the net is a trit, not a bit: **+1 (passage), 0 (no event), –1 (blocked passage)**. Balanced ternary is the natural notation. Its multiplication rule matters below: $(-1)\cdot(-1) = +1$. Two blockings compose into a passage. **[theorem: the algebra; hypothesis: its physical role]**

Closure is quantized. The address layer counts in threes, the tension layer in twos. Their exchange rate is $\log_2 3 = 1.58496...$, an irrational number. The continued-fraction convergents of $\log_2 3$ give the admissible closure depths: (2,3), (5,8), (12,19), (41,65). Because $\log_2 3$ is irrational, no closure is ever exact. The remainder at depth 19 is the Pythagorean comma, $3^{12}/2^{19} \approx 1.01364$. The negotiation between the layers never ends. **[theorem: the number theory; hypothesis: the physical reading]**

Selection works by fit, not by force. Nothing causes a knot. A winding that returns in phase remains; one that does not, passes on. Description at the correct winding depth replaces explanation.

2. Cotterill's Circuit Reading

Rodney Cotterill (1933–2007) was a physicist at the Danish Technical University who spent his last two decades on the brain. His synthesis appeared as *Enchanted Looms* (1998) and, condensed, as "Cooperation of the basal ganglia, cerebellum, sensory cerebrum and hippocampus" in *Progress in Neurobiology* 64 (2001) 1–33. All Cotterill claims below are from that 2001 paper unless marked otherwise. **[measurement/published model]**

His premises are stark. Muscular contraction is the nervous system's only externally directed product. Therefore all signaling routes converge on the motor areas. The brain is one unit, not a hierarchy of modules. Its single underlying principle is **conditionally permitted movement**: components can veto muscular contraction by denying the motor areas sufficient activation.

Key elements of his circuit:

The bacterium already runs the scheme. *E. coli* cannot measure a chemical gradient at an instant. It swims, and reads what changes. The stimulus is its own movement; the response is the environmental feedback. Its chemistry gives it a memory of about 4 seconds: the concentration of the last second is compared with the previous three seconds, with positive and negative weighting respectively. Cotterill's thesis: this probe-by-movement strategy has not changed in four billion years. **[measurement: Berg, Segall; model: Cotterill]**

Two gates on every act. The basal ganglia (BG) and the cerebellum (Cb) both act by inhibition. Their permission is disinhibition. The BG gate opens fast, steered by the amygdala. The Cb gate opens slowly, through long-term depression (LTD), and it permits continuation of a sequence when the feedback it produces is familiar. Movement requires both gates open at once.

Routes through the BG pass two successive inhibitory connections. Release of the second inhibition by the first yields net excitation of the motor areas.

Two thresholds. Above the high threshold: overt movement. Above a lower threshold: the same motor patterns run covertly. That covert run is thought. Thinking is simulated interaction with the environment, carried by the movement circuitry, held under the movement threshold.

No strong loops inside. Crick and Koch (1998) found no strong directed excitatory loops in the cortico-thalamic wiring; strong loops would throw the system into uncontrolled oscillation. Cotterill adds the complement: the one strong loop that exists runs **through the environment**. Act, feedback, act.

The hippocampus records correlations, not places. It ties a motor sequence to the sensory feedback that sequence produced, holds the pair temporarily, and mediates a permanent record in cortex. Its circuit integrates over roughly 10 ms per synaptic stage and gates on the theta rhythm (5–12 Hz, ~200 ms per cycle). The NMDA receptor channel stays open almost a fifth of a second — one theta interval, one elemental cognitive event. Cotterill calls the arrangement a **dual relevance filter and temporal binder**: only events that fill at least one closure time of the circuit can penetrate and become candidates for storage.

A ladder of times. Dendritic transit ~10 ms. One phoneme ~50 ms. NMDA window and theta interval ~200 ms. Conscious registration of an unprovoked stimulus ~500 ms (Libet). Each loop in the circuit has its own transit time, and each admits only events on its own time scale.

Novelty interrupts permission. The inferior olive fires the climbing fibers when something unexpected occurs. Because LTD has not yet developed for the novel case, the Cb's standing permission is disturbed. The sequence halts; loading begins.

3. Seven Detection Points

The Vacuum.Net theory can be detected in this circuit at seven places. Cotterill supplies the anatomy and the published measurements; the net reading is the hypothesis.

3.1 The trit register. Cotterill's diagram knows three signs: excitatory (+1), inhibitory (−1), sub-threshold (0). His BG routes compose two inhibitions into a permission: $(-1)\cdot(-1) = +1$. The brain's passage logic runs on the multiplication rule of balanced ternary. The bit is the wrong register for this tissue; the trit is native to it. **[hypothesis]**

3.2 The four states. LTD builds slowly: loading. Race-to-threshold under a standing veto: holding. Threshold crossing into overt movement: discharge. Below both thresholds: rest. The four states of the tension layer appear as the four regimes of his motor circuit. **[hypothesis]**

3.3 The two-layer law. His fixed wiring is address; his signal traffic is tension. And his central thesis — consciousness exists to acquire new context-specific reflexes — states the coupling between the layers: discharge writes new closure. A tension pattern that discharges under attention becomes permanent wiring. Tension becomes address. He describes the crossing and calls it learning. **[hypothesis]**

3.4 Closure through the environment only. No strong loops inside; one strong loop through the outside. In net terms: no mesh closes on itself within the system. The strand must go out and return. The Cb grants continuation only to familiar feedback — only to what returns in phase. That is selection by fit, named without the name. **[hypothesis]**

3.5 Winding depth as embedding. Cognition is the body monitoring the body. Consciousness is monitoring the monitoring. Knowing that one knows that one knows is the same operation wound one depth further. Identity across scale, as recursion of a single movement. **[hypothesis]**

3.6 Time quantization per closure. Every loop admits only events that fill its own closing time: 10 ms, 50 ms, 200 ms, 500 ms. Compare the third law derived in the MAZE financial work: *each layer warns on the horizon of its own closing time*. The brain's relevance filter and the MAZE's warning horizons are the same law at different winding depths. **[hypothesis]**

3.7 His refusal of the arrow. Cotterill writes SCHEMA * ATTENTION → QUALE and immediately insists the arrow is not an equals sign and the asterisk not a multiplication. He senses that the relation is constitutive — simultaneous, not causal — and lacks the language for it. The net supplies the language: a constitutive relation between state variables, no arrow between two separated things. **[hypothesis]**

What is absent in Cotterill. No counting. His depths (two thresholds, embedding orders) are unquantized: no admissible series, no remainder, no comma. No scale identity below the human layer: his net stops at muscles and neurons. And his background stays fixed anatomy; that the address layer itself was once discharge appears only in learning, not as the constitution of the tissue. These absences mark exactly where the Vacuum.Net theory extends him. **[open question: whether his time ladder fits the admissible depth series quantitatively]**

4. Kanerva: The Address Space

Cotterill supplies the passage control. Pentti Kanerva supplies the address space. His Sparse Distributed Memory (1988) is the second independent finding of the net structure in memory research.

SDM in five statements. **[measurement/published model: Kanerva 1988]**

1. A pattern is its own address: a point in an enormous binary space, for example 2^{1000} possible addresses.
2. Of those addresses only a sparse random set exists physically ("hard locations"). Almost the entire space is empty.
3. Writing spreads an entry over all hard locations within a fixed distance. Reading sums over the same neighborhood and thresholds at zero. The counters in each location increment and decrement; Kanerva's later hyperdimensional computing makes the register explicitly {+1, 0, −1}.
4. Iterated reading from a starting point **within a critical distance** of a stored pattern converges to that pattern. Starting outside it, the read sequence wanders and arrives nowhere. Storage itself creates attractor basins with sharp rims — a maze of passages and dead ends, grown from the content.
5. Kanerva designed SDM as a model of the **cerebellum**: the granule layer (about 50 billion cells, the majority of all neurons in the head) is the sparse expansion of the input; the Purkinje cell is the sum-and-threshold readout; the climbing fiber trains it.

The net reading: address = content is the two-layer law seen from the address side. Sparseness is the vacuum: only what closes exists as knot. Sum-plus-threshold is selection by fit. The critical distance is the rim of a mesh. And Kanerva's own conclusion — the encoding decides everything; SDM works only if similar things land on similar addresses — is the intake law: **the intake determines the structure of the maze. [hypothesis]**

Kanerva and Cotterill do not cite each other. One begins in high-dimensional mathematics, the other in circuit anatomy. Both arrive at the same organ (the cerebellum), the same register (the trit), the same operation (conditional passage by fit). Two independent finds of one structure.

5. The Successors

Cotterill died in June 2007. His simulation project CyberChild (2003) — an infant nervous system with touch, hearing, and muscles for feeding and crying, built to test whether the circuit generates the phenomena — was not taken up. No school formed. But every component of his program was carried forward, mostly without citing him. The successors matter here because each one sharpens a piece of the net reading.

Germund Hesslow — the covert layer. Hesslow's simulation theory of thinking (2002, 2012) is the closest continuation: thinking is simulated behavior plus simulated perception, carried by the same sensorimotor circuits, with the cerebellum as predictor. He confirms the two-threshold structure experimentally. For the net: the tension layer can run **candidate passages** below the discharge threshold. A thought is a dry run of a mesh. **[measurement + model: Hesslow]**

Peter Strick — the loops are real. Middleton and Strick (1994) and the tracing work that followed proved closed loops between the cerebellum and the prefrontal cortex, not only the motor cortex. The passage machinery serves cognition with the same wiring it uses for movement. One architecture, several winding depths. **[measurement]**

Roger Carpenter — holding as a race. The LATER model (Carpenter and Williams, 1995) established that decisions are races to a threshold, with the rate of rise varying trial by trial. Holding is not passive storage; it is competing tension approaching a rim. The winner discharges. **[measurement]**

Daniel Wolpert and Sarah-Jayne Blakemore — the forward model. Blakemore, Wolpert and Frith (1998) showed why self-tickling fails: the system predicts the sensory return of its own act and cancels it. Prediction of the return is computed before the return arrives, with cerebellar participation. For the net: the mesh carries its expected return; only the **deviation** between expected and actual return loads. **[measurement]**

Karl Friston — the mainstream arrives. Predictive processing and active inference (Friston, 2010) restate Cotterill's program in probabilistic form: perception and action form one loop that minimizes prediction error; the organism probes the world to test its own model. The formulation is binary-Gaussian rather than ternary and keeps the causal arrow. It confirms the loop; it lacks the register and the counting. **[published model]**

Rodolfo Llinás — independent confirmation. *I of the Vortex* (2001) derived the same motor-centric thesis independently, from the physiology of intrinsic oscillation: the mind is internalized movement. Two simultaneous, mutually unaware statements of the same reversal in the same year strengthen the convergence claim. **[published model]**

The synthesis itself — bacterium to human, all gates, all time scales, one machine — has not been continued by anyone. That part of the field has been parceled out to subdisciplines since 2007. The MAZE inherits an open line.

6. The Architecture of a Self-Learning System

The findings above compose into one design. This section states it in full. It is a description of an architecture, not a work plan: present tense, no instructions, no phases. Its purpose is to record what the Cotterill–Kanerva convergence actually yields — the architecture of a system that keeps learning without destroying itself.

The problem it answers is the open problem of machine learning. A system that continues to learn overwrites and pollutes its own memory; the field calls this catastrophic forgetting. The standard remedies — replaying old data, freezing weights, periodic retraining — avoid the problem rather than solve it. Current AI is therefore trained once and then frozen: a dead address layer without a tension layer. It refuses nothing, closes on its own output, runs on a binary register, and has one clock. The biological record shows a different design, unchanged for four billion years. That design is stated here in ten components. **[engineering design throughout; biological precedent: measurement; artificial demonstration: outstanding]**

6.1 The register. The system state is written in balanced ternary: +1 (passage), 0 (no event), −1 (blocked). The ground state is 0. Silence, not activity, is the default: most input produces no event at all. Permission is never primitive; it exists only as the composition of two lifted refusals, $(-1)\cdot(-1) = +1$. A system built on this register spends almost all of its time refusing. That is not a limitation. It is the mechanism of its integrity.

6.2 Two organs and a bridge. The two-layer law divides the system into two organs. The **address organ** is the memory: a sparse distributed address space in which content is its own address. Of all possible addresses only a sparse set is inhabited. Reading is summation over a neighborhood followed by a threshold at zero; iterated reading converges to a stored pattern when it starts within the critical distance, and arrives nowhere when it starts outside it. The address organ is immutable except at write moments. The **tension organ** is the state machine: four states — loading, holding, discharge, rest — governed by two thresholds, a low candidate threshold and a high write threshold. Between the organs runs one **bridge** with one traffic rule: **discharge writes address.** Only at a discharge event does anything cross from the tension side to the address side. Between discharges the memory is frozen. Learning is therefore rare, dated, and event-bound — every write has a discharge as its receipt.

6.3 The probe cycle. The system acts first. It emits a probe into its environment, and reads what returns. From its address organ it computes the expected return before the actual return arrives. **Only the deviation between expected and actual return loads the tension organ.** What the system predicted correctly produces nothing — the self-tickle principle. Input the system provoked and input that arrived unprovoked are marked differently at entry; that source mark is half of the address. A system without this cycle is a receiver; a system with it is a prober. Knowledge, in this architecture, is always the record of a probe and its return, never of input alone.

6.4 The two admission gates. Nothing is written on one criterion. Admission requires two independent gates in the same cycle: a **fast context gate** — does the entry fit the current loading state of its region — and a **slow familiarity gate** — does its return fall within a known winding. Both gates stand at −1 by default. Only the simultaneous lifting of both composes to +1 and permits a write. An entry that lifts one gate but not the other becomes a candidate, never a receipt. One yes is no yes.

6.5 The covert layer. Candidates are the system's thoughts: tension traces with a decay time, held below the write threshold. A candidate that recurs within its decay window and then lifts both gates converts to a receipt. A candidate that does not, expires to rest and leaves no trace. The system thereby runs trial passages that are allowed to die — it can consider without committing, and most of what it considers is never written. False addressing is refused at the door instead of repaired afterwards.

6.6 The time ladder. The system does not run on one clock. It spans several winding depths, each with its own closing time, and each layer admits only events that fill at least one of its own cycles. Sub-cycle flicker does not exist for a layer. Deep layers write rarely and hold long; shallow layers write fast and forget fast; no layer reads another's clock. The spacing of the depths follows the admissible series of the net — (2,3), (5,8), (12,19), (41,65) — the reference for how far apart the clocks stand. Each layer warns, learns and forgets on the horizon of its own closing time.

6.7 Loop hygiene. No strong loop closes inside the system. The system's own output is never admissible as confirmation of its own addresses; fed its own output, it writes nothing. Every closure requires a return that passed through the environment — an external datum, an act of another, a world that answered. The one strong loop runs outside. This is the architectural difference between learning and self-confirmation, and it is enforced by wiring, not by policy.

6.8 The novelty interrupt. Input that falls outside every known winding suspends the standing state. Nothing is computed as familiar that is not; the standing permission is withdrawn and loading restarts. Novelty is not an error condition. It is the one signal that outranks every gate.

6.9 The drive. The exchange rate between the two organs is $\log_2 3$, an irrational number: closure is never exact, and a remainder always stands. The architecture treats that remainder as its drive. The system is never permitted to converge fully; a residual tension is a design property, not an error to be minimized. A system whose remainder reaches zero has stopped learning and stopped living — full equilibrium is the specification of death. This is what the biological sources describe but cannot count: an organism moves because it is not in equilibrium, and it is never in equilibrium because its closure never completes.

6.10 The gauges. The architecture defines its own health measurements. The **critical distance** of the address organ, tracked as the archive grows. The **self-closure count** — new addresses formed from the system's own output — which is zero in a healthy system, by 6.7. The **habitation distribution** of inhabited addresses against a random scattering of equal size, read against the depth series: clustering at the admissible depths distinguishes this architecture from a memory with random locations. And the running **deviation log** of 6.3: what the system still gets wrong is the exact measure of what it is still learning. A system whose deviation log runs dry is finished — see 6.9.

The tuning of the whole architecture lies on a single dimension: the strictness of the gates. Too strict, and nothing is ever written — the system stays empty. Too loose, and pollution enters — the system degrades like current AI. Everything else follows from the components above. **[hypothesis: that a workable setting of this dimension exists; the biological precedent says yes]**

Against current AI the contrast is complete. It refuses nothing; this refuses almost everything. It closes on its own text; this closes only through the world. It runs on two signs; this on three, with silence as the ground state. It is trained once and frozen; this writes at discharge, for as long as it runs. It has one clock; this has a ladder of clocks. Each difference is one component above; together they are the difference between a reflex machine and a learner.

7. Status Ledger

#	Claim	Status
1	Balanced ternary algebra; convergents of $\log_2 3$; comma at	theorem
2	Cotterill's circuit facts (double inhibition, two thresholds, no strong loops, hippocampal timing, 4-second bacterial	measurement / published model
3	Kanerva's SDM properties (sparse locations, critical distance, ternary counters, cerebellum mapping)	measurement / published model
4	Successor results (Strick loops, LATER, tickle	measurement
5	The seven detection points (net reading of 2–4)	hypothesis
6	Brain time ladder fits the admissible depth series	open question
7	The architecture of Section 6 (register, organs, bridge, gates, covert layer, ladder, hygiene, interrupt, drive,	engineering design
8	Two-gate admission with discharge-bound writing sustains learning without memory pollution	hypothesis; biological precedent established, artificial demonstration

9	A workable gate strictness exists between empty and	hypothesis
1	A never-resolving remainder is necessary for sustained	hypothesis

Adverse findings preserved: Cotterill's mirror-neuron and spindle arguments are contested in later literature; nothing in Sections 3–6 depends on them. Friston's framework, though convergent on the loop, contradicts the net on register (binary) and on causality (arrows retained); it is cited as confirmation of the loop only.

8. Conclusion

Two researchers who never met the Vacuum.Net theory found its structure independently: Kanerva in the mathematics of high-dimensional memory, Cotterill in the anatomy of conditional passage. Both arrived at the trit, the sparse address space, selection by fit, and closure through the environment. Neither counted the depths. The counting — the admissible series, the eternal remainder — is what the Vacuum.Net theory adds, and the MAZE is where the count becomes measurable.

The brain, in Cotterill's reading, is a maze: almost everything is refused, passage is granted only when two inhibitions coincide, and the only full circuit runs through the world outside. Section 6 states what that yields when taken as a design: a system that probes and reads back, refuses on two gates, thinks in candidates that may die, writes only at discharge, closes only through the world, and is driven by a remainder that never resolves. The biological precedent for this architecture has run for four billion years without a competitor. Its artificial demonstration is the open question this essay leaves on the table — and the reason the convergence documented here matters beyond the history of ideas.

Annotated References

Cotterill, R.M.J. (2001). "Cooperation of the basal ganglia, cerebellum, sensory cerebrum and hippocampus: possible implications for cognition, consciousness, intelligence and creativity." ***Progress in Neurobiology*** **64, 1–33.** Why read: the primary source of this essay. One machine, one output, conditional passage. Read Fig. 7 and Sections 3–4 first; the bacterium section (2.1) contains the loading-decay structure at the smallest scale.

Cotterill, R.M.J. (1998). ***Enchanted Looms: Conscious Networks in Brains and Computers.*** **Cambridge University Press.** Why read: the extended version, with the evolutionary ladder worked out. Reading advice: chapters 7–9 carry the argument; the rest is background.

Cotterill, R.M.J. (2003). "CyberChild: a simulation test-bed for consciousness studies." ***Journal of Consciousness Studies*** **10 (4–5), 31–45.** Why read: the abandoned execution attempt. Instructive as engineering: which simplifications he chose (binary composite units — note: not ternary) and where the project stopped.

Kanerva, P. (1988). ***Sparse Distributed Memory.*** **MIT Press.** Why read: the address space, the critical distance, the cerebellum mapping. Chapters 1–3 suffice for the structure; the convergence proofs are in the middle chapters.

Kanerva, P. (2009). "Hyperdimensional computing: an introduction to computing in distributed representation with high-dimensional random vectors." ***Cognitive Computation*** **1, 139–159.** Why read: the register made explicit — $\{+1, 0, -1\}$ vectors, binding and superposition as algebra. The shortest path from SDM to an implementable encoder.

Hesslow, G. (2012). "The current status of the simulation theory of cognition." ***Brain Research*** **1428, 71–79.** Why read: the covert layer, ten years of evidence. Grounds the covert layer (6.5).

Middleton, F.A. & Strick, P.L. (1994). "Anatomical evidence for cerebellar and basal ganglia involvement in higher cognitive function." ***Science*** **266, 458–461.** Why read: the measurement that moved the gates into cognition. Short, decisive.

Carpenter, R.H.S. & Williams, M.L.L. (1995). "Neural computation of log likelihood in control of saccadic eye movements." ***Nature*** **377, 59–62.** Why read: holding as a race to threshold, measured in single decisions. Grounds the holding state.

Blakemore, S.-J., Wolpert, D.M. & Frith, C.D. (1998). "Central cancellation of self-produced tickle sensation." ***Nature Neuroscience*** **1, 635–640.** Why read: the forward model in one experiment. Grounds the probe cycle (6.3): only deviation loads.

Crick, F. & Koch, C. (1998). "Constraints on cortical and thalamic projections: the no-strong-loops hypothesis." ***Nature*** **391, 245–250.** Why read: loop hygiene as anatomy. Grounds loop hygiene (6.7).

Friston, K. (2010). "The free-energy principle: a unified brain theory?" ***Nature Reviews Neuroscience*** **11, 127–138.** Why read: the mainstream restatement of the probe-loop. Read for the loop; note where it keeps the arrow and the binary register.

Llinás, R. (2001). ***I of the Vortex: From Neurons to Self.*** **MIT Press.** Why read: the independent second statement of the motor-centric reversal, from oscillation physiology. Strengthens the convergence claim of Section 5.