


Rules, Memory, and Repair

Local rules, hidden context, and testable recovery



30 September 2026

Team eR33t Science | Research essay

137 / UNUS MUNDUS

RESEARCH ESSAY / 30 SEPTEMBER 2026

Research essay / September 2026

Local rules, hidden context, and testable recovery

On 200 matched eligible damaged bitmaps, a one-entry Life mutant produced ten moving gliders after 64 synchronous updates following damage, compared with eight under exact Conway Life. Yet the same lookup entry conflicts with an undamaged reference path: an identical binary 3×3 neighborhood requires birth for the observed repair and death for healthy preservation. This local-state aliasing establishes a precise limitation of a memoryless local mapping. [E2] We connect that result to Michael Levin's work on morphogenesis and collective control, while separating anatomical regulation, behavioral savings, and theoretical cognition. [S6][S7][S10] Wolfram's cellular-automaton framework and Cook's Rule 110 proof supply a complementary distinction between simple transition rules and complex computational capabilities. [S2][S5] The digital evidence comprises a frozen Jev 1.13.0 rule-table audit, deterministic trajectories, a finite damage census, and bounded synthetic and Rule 30 explorations. [E1][E2][E3] The engineering consequence is an explicit division of labor: code acquires and qualifies evidence, executes exact transitions, and verifies outcomes; typed selection operates over permitted candidates; generative models support research and explanation. A proposed digital experiment tests whether budgeted persistent local state improves unseen recovery without degrading healthy behavior. Additional state can separate histories; a repair claim still requires held-out evidence.

PURPOSE	STATUS	METHOD
Research essay	Editorial edition 1.0	Saved evidence and source review

A note from Justin

For my beautiful mom, my hard-headed, hardworking dad, and my little brother Cody.

Mom gave me heart. Dad gave me grit-and a stubborn streak I'm choosing to call persistence. Cody went into radiology, so apparently he's the one with a proper excuse for seeing through all of us.

This started with a weird idea and a lot of curiosity. Somewhere along the way, it became a reminder: the rules matter, but so does what we carry forward.

Love you guys.
- Justin Hedge

At a glance

The decisive conflict	The ID9 witness asks one identical local input to produce incompatible next outputs. No memoryless 3×3 mapping can meet both one-step targets; a richer observation or a changed target is required. [E2]
Representation changes the audit	The frozen raw table scored 231/512, below an always-dead baseline of 372/512. The counted representation scored 18/18 and expands exactly to 512/512. These are artifact-specific results, not general Jev accuracy estimates. [E1]
Architecture follows authority	Keep acquisition, eligibility, arithmetic, transitions, and verification in code. Let Jev select among permitted candidates or abstain. Preserve receipts and independent review: valid types and confident outputs do not establish truth.
Memory remains a hypothesis	The next experiment should compare persistent local state with matched memoryless, sham, yoked, and explicitly circular oracle controls. Healthy-path preservation and unseen damage recovery must both be preregistered.

Reading map

01 The recovery that contradicts its own rule The most revealing result was a successful recovery with an incompatible healthy-path requirement. A single lookup-table change preserved eight moving outcomes and added two removal recoveries on the same 200 eligible damaged bitmaps. [E2]	02 Auditing a frozen rule Before interpreting a moving pattern, audit the transition function that makes it move. This lab's central artifact is one frozen Jev 1.13.0 recording. [E1]
03 A finite world fills up Small transition errors can produce large trajectory differences. Population, however, is only one projection of a board's state.	04 What the damage census actually counts The recovery comparison becomes meaningful only after its eligibility rule and denominators are visible. The wave-3 census contains 4,000 actual one-cell edits. [E2]
05 Morphogenesis supplies a harder question Michael Levin's relevance is the study of form, physiological state, and collective control. Morphogenesis concerns the organization of anatomy; regeneration concerns its restoration after damage.	06 Simple rules do not settle long futures The computational connection is strongest when it distinguishes a transition rule, a demonstrated capability, and a finite observation.
07 Put typed decisions behind an evidence boundary For assistants such as AlexAI, the practical lesson is to make authority and state explicit before asking a probabilistic system to choose.	08 Proposed experiment: distributed state with a budget Everything in this section is proposed and unrun. The aim is to test whether locally maintained history improves unseen digital recovery while preserving declared healthy behavior.
09 What a system must know to repair The philosophical connection concerns the relation among rules, goals, and available information. Those three can be separated precisely enough to test.	

01 / FIELD NOTE

1

The recovery that contradicts its own rule

The most revealing result was a successful recovery with an incompatible healthy-path requirement. A single lookup-table change preserved eight moving outcomes and added two removal recoveries on the same 200 eligible damaged bitmaps. [E2]

These were wrapped-board outcomes after 64 synchronous updates following damage: ten moving cases for the ID9 mutant, eight for exact Life. A glider is a small pattern whose changing shape carries it across the grid. Here, repair means returning a damaged pattern to a specified undamaged trajectory within a finite test. The word describes a digital target, without assigning the pattern an organism's physiology. [E2]

The mechanism is more interesting than replacement of a missing pixel. In the recorded witness, the deleted cell at (30,17) remains dead in the next phase. Its neighbor at (29,18) instead acquires a different local input: the intact neighborhood is ID13, while deletion changes it to ID9. The mutant births that neighbor. The complete next bitmap then equals the undamaged control's next bitmap. Recovery occurs through a compensating transition elsewhere in the pattern. [E2]

ID9 and ID13 are identifiers for encoded neighborhoods, not cell counts. The difficulty appears when ID9 is encountered on the healthy reference orientation labelled R0. Preserving that healthy step requires death. Thus the same nine visible bits demand output 1 in the damage witness and output 0 in the healthy witness. The mutant's eligibility on some intact paths does not make it faithful on every orientation and phase. [E2]

Write the neighborhood as n_9 and a memoryless local rule as f . The two requirements are $f(n_9) = 1$ and $f(n_9) = 0$. No function satisfies both. This is local-state aliasing: distinct situations become indistinguishable through the chosen observation boundary. Randomizing the answer cannot guarantee both requirements either, because an unchanged input supplies no basis for choosing the appropriate context-dependent distribution.

The conflict belongs to the targets, not to a defective law of Life. Exact Life correctly applies its own rule to the damaged input. Asking it to recover the counterfactual undamaged trajectory adds another objective. Likewise, correcting a recorded table toward Life and modifying Life to recover damage are different operations. A modification may improve one objective while violating the other.

A larger spatial view, phase information, a trustworthy contextual signal, or persistent state could distinguish the cases. A supplied healthy-versus-damaged label would settle this particular contradiction immediately, but would also conceal the hard part: obtaining that distinction without handing over the answer. The witness motivates more context. It does not prove that explicit memory or signaling is necessary for every possible repair model.




Figure 1 The one-step witness: deletion at (30,17) changes the neighborhood at (29,18) from ID13 to ID9; a neighboring birth produces whole-bitmap next-state equality while the deleted site stays dead. The healthy R0 witness requires death for the identical ID9 input. [E2]

02 / FIELD NOTE

2

Auditing a frozen rule

Before interpreting a moving pattern, audit the transition function that makes it move. This lab's central artifact is one frozen Jev 1.13.0 recording. [E1]

Conway Life uses a binary center cell and its eight neighbors. A dead center becomes live with three live neighbors; a live center survives with two or three. All other next states are dead. The complete binary 3×3 input space therefore contains $2^9 = 512$ neighborhoods. The recorded raw table agreed with this exact B3/S23 rule on 231 of them. An always-dead predictor would agree on 372. [E1]

That baseline matters. A raw score without the output distribution could sound like partial mastery of a difficult local task. Here it is weaker than a constant answer under uniform enumeration. Uniform enumeration is useful because it exhausts this finite function; it is not a natural-frequency estimate for neighborhoods encountered during a particular rollout. Those frequencies depend on the seed, boundaries, and rule itself.

The 281 disagreements fall into five families: 59 false births, 47 missed births, four underpopulation survivals, 157 overcrowding survivals, and 14 erroneous deaths. These categories distinguish ways a rule can alter dynamics. A false birth introduces activity where Life would not; an overcrowding-survival error retains a live cell that Life would remove. Their downstream effects depend on which neighborhoods the evolving board visits. [E1]

A counted representation produced a sharply different artifact. Supplying center status and neighbor count gives only 18 possibilities: two center states multiplied by nine possible counts. That table scored 18/18 and expands deterministically to the exact 512-entry Life table. The expansion is an equivalence check, not 512 additional model successes. Counted, expanded, and directly implemented exact tables are not independent replications of the same finding. [E1]

This contrast supports an architectural lesson about representation, but does not isolate a causal explanation for every raw error. The recording alone cannot distinguish failures of counting, prompting, indexing, batching, or another representation detail. Nor does one recording establish the behavior of other prompts, versions, or repeated calls. Exhaustive coverage of one artifact's inputs is different from broad coverage of a model's behavior. [E1]

Once saved, the lookup table becomes an ordinary deterministic program. Each tick indexes recorded answers; it does not ask Jev to reason again. That distinction makes the subsequent experiments inexpensive to replay and precise to inspect. It also limits their interpretation: thousands of trajectories explore consequences of the frozen artifact, rather than thousands of fresh opportunities for the model to succeed or fail.




Figure 2 Correct outputs over all 512 binary neighborhoods: raw recording 231, always-dead baseline 372, counted-table expansion 512. The counted result originates in 18/18 center/count cases; expansion adds no independent model evidence. The raw table has 281 errors across five families. [E1]

03 / FIELD NOTE

3

A finite world fills up

Small transition errors can produce large trajectory differences. Population, however, is only one projection of a board's state.

The experiments used a 64×40 board, with 2,560 sites, under wrapped and dead-edge semantics. Wrapping joins opposite boundaries; dead-edge semantics treats positions beyond the board as dead. These choices belong to the transition contract. A pattern approaching an edge experiences different neighborhoods under the two conventions, so a boundary change is not merely a display preference. [E1]

In the original wrapped rollout, generation 263 contained 442 live cells under the raw table and five under exact Life, with 441 differing cells. Those quantities measure different things: two population counts and a whole-grid disagreement count. The raw trajectory later reached a fixed point at generation 1,273 with 1,918 of 2,560 sites alive, or 74.9%. The board size, timing, and density describe this seed and recording. [E1]

The original raw glider also had three early deaths before its later nondecreasing population behavior. That observation should not be promoted into a monotonicity law. The lookup table includes several mechanisms capable of changing occupancy in either direction. A trace may become trapped in a region of state space where some possibilities stop occurring, even though the full rule still permits them elsewhere. [E1]

The broader deterministic sweep validated 7,312 trajectories with no fresh Jev calls. Its design included 32 family conditions—the on/off combinations of five error families—24 random-edit controls, a declared seed mixture, and a generation cap of 2,048. These experiments examine structured corrections against other table edits. They do not turn the seed mixture into a random sample of every possible initial board. [E1]

In that declared design, correcting overcrowding-survival errors produced the largest average population reduction among the family corrections. The family also supplies a direct persistence mechanism: cells survive neighborhoods in which exact Life requires death. Yet reducing population is not equivalent to restoring correct motion. A corrected table can empty a board, freeze a pattern, or increase activity elsewhere through altered interactions. [E1]

A fixed-point claim requires equality of complete consecutive states. A constant population can conceal continuing rearrangement; a repeated whole state establishes a cycle, with a one-step repeat establishing a fixed point. Runs reaching the cap without a complete repeat remain censored. A finite deterministic board must eventually revisit a state, but that mathematical fact does not identify the recurrence time within an affordable experiment. The sweep appropriately separates observed repeats from unresolved finite traces. [E1]

What the damage census actually counts

The recovery comparison becomes meaningful only after its eligibility rule and denominators are visible. The wave-3 census contains 4,000 actual one-cell edits. [E2]

The design combined eight glider orientations, damage phases 8-11, two boundary conventions, and 25 fixed candidate sites per phase: five removals and twenty additions. Across the tested tables, 160 groups qualified for editing, yielding $160 \times 25 = 4,000$ edited cases. Qualification required finite whole-grid agreement with the healthy reference path from generation 0 through absolute generation 64. It did not establish indefinite healthy fidelity. [E2] The recovery endpoint is separate: 64 updates after damage at absolute generations 8-11 ends at absolute generations 72-75.

A group's eligibility matters because different tables preserve different subsets of intact paths. The healthy-path gate therefore conditions the damage experiment before any recovery score is calculated. An eligible-group count is not a count of independent biological subjects or randomly sampled environments. Phase aliases and duplicate trajectories also remain in the labelled census. The following wrapped-board counts retain the original denominators. [E2]

Read the matched comparison first

For ID9, the appropriate exact comparator is eight moving outcomes on the same 200 eligible damaged bitmaps, not the exact table's aggregate 32/800. The mutant preserves all eight and adds two removal recoveries. That overlap establishes a narrow improvement on a shared test set. Comparing 10/200 directly with 32/800 would mix a rule change with a change in eligible cases. [E2]

The other modifications tell different stories. The raw table patched with the canonical ten support entries produced twenty moving cases in its 200-case wrapped group, but lost the eight exact successes and recovered a different set. The raw table patched with all D4 support produced four moving cases versus exact Life's 32 on the matched 800-case comparison. D4 denotes the eight square-grid rotations and reflections. [E2]

These results separate intact-path support from damage recovery. Adding entries that reproduce healthy motion constrains a particular corridor through state space. An edit pushes the pattern outside that corridor, where other table entries become decisive. Expanding healthy support can therefore remove transitions that happened to help a damaged pattern. More faithful intact behavior need not produce more recovery.

The useful conclusion is conditional and mechanistic. ID9 has a repeatable, whole-bitmap one-step recovery witness, alongside a healthy-path contradiction. The census identifies where that modification helps under the stated gate; it supplies neither independent probability samples nor evidence that a repaired glider heals biologically.

Wrapped table	Moving after 64 updates	Removal recovery	Addition recovery
Exact Life	32/800	0/160	32/640
Exact + ID9	10/200	2/40	8/160

Wrapped table	Moving after 64 updates	Removal recovery	Addition recovery
Raw + canonical ten	20/200	0/40	20/160
Raw + all D4 support	4/800	0/160	4/640

EVIDENCE NOTE / JEV / FINITE EXPERIMENT

Damaged-glider recovery

Wrapped boundary; 64 updates after damage. Every finite denominator is part of the claim.




Figure 3 Matched wrapped-board outcomes after 64 synchronous updates following damage: ID9 preserves eight exact successes and adds two removal recoveries among 200 identical eligible damaged bitmaps. The all-D4 comparison uses 800 matched cases, with four moving outcomes against 32 under exact Life. Counts are labelled cases, not independent probabilities. [E2]

05 / FIELD NOTE

5

Morphogenesis supplies a harder question

Michael Levin's relevance is the study of form, physiological state, and collective control. Morphogenesis concerns the organization of anatomy; regeneration concerns its restoration after damage.

A historical bridge begins with Turing's reaction-diffusion model. His 1952 analysis showed how small disturbances can destabilize an initially homogeneous chemical system and generate spatial structure. [S1] The transferable idea is that local interactions can generate organized patterns without a cell-by-cell spatial script. Similar-looking output, however, does not establish a shared mechanism. A cellular automaton and a chemical model must each be judged by their own variables and equations.

Levin and Martyniuk's review treats membrane potentials, ion-channel dynamics, and bioelectric networks as information-bearing components of growth and pattern regulation. [S6] That framing enlarges the modeling question: what state must be represented to explain a system's response to perturbation? For an engineer, its value is less a metaphor about intelligence than a demand to distinguish visible material configuration from the variables participating in control.

Two primary studies make the physiological claim concrete. Pai and colleagues reported that changes in transmembrane potential could induce eyes outside the anterior neural field in *Xenopus* embryos. [S8] Beane and colleagues reported bioelectric effects on head and organ scaling during planarian regeneration. [S9] These are bounded developmental results. They motivate attention to state-dependent pattern control without identifying a universal voltage code or validating any digital recovery mechanism.

Anatomical pattern memory introduces history. Levin's 2014 synthesis discusses persistent changes in regenerative target morphology after transient physiological interventions, including double-headed outcomes and reset phenomena. It proposes endogenous bioelectric networks as a medium for non-genetic patterning information. [S11] The stable anatomical outcomes and the proposed storage explanation are separate claims. For modeling, the relevant question is whether present anatomy alone predicts the next regenerative response.

Behavioral memory is a different evidential category. Shomrat and Levin reported environmental familiarization and a behavioral savings effect after head regeneration: previous experience was associated with later performance in their training paradigm. The study did not identify the physical storage substrate. [S10] Neither persistence of a behavioral effect nor the regeneration of a head licenses the stronger assertion that a particular memory was located outside the brain.

Theoretical cognition adds a third category. Levin's scale-free cognition article proposes computational boundaries for systems that measure, model, and act across space and time. It is explicitly hypothesis-and-theory work. [S7] Its vocabulary can help specify an engineered controller's observation and action boundary. Using that vocabulary does not experimentally establish cognition in a lookup table, a tissue, or a document assistant.

These distinctions improve the digital question. Anatomical regulation suggests evaluating geometry and proportion rather than activity alone. Behavioral savings suggests testing whether history predicts performance after disruption. Theoretical boundary language suggests declaring which information a controller can access. None supplies a ready-made biological explanation for ID9. Each instead suggests a different control or measurement that could make the computational experiment more discriminating.

Lobo, Beane, and Levin argue for spatial, algorithmic models that expose how cellular activities could produce self-monitoring and growth control. [S12] Our corresponding obligation is to expose target definitions, state variables, update rules, and failure conditions. A glider has no independently demonstrated biological target; its healthy continuation is selected by the experimenter. That explicit target makes the digital problem tractable, while also marking the limit of the analogy.

"Any Self is demarcated by a computational surface"

MICHAEL LEVIN / S7

Simple rules do not settle long futures

The computational connection is strongest when it distinguishes a transition rule, a demonstrated capability, and a finite observation.

Wolfram's 1984 paper organized one-dimensional cellular automata into four broad behavioral classes and explored the possibility of universal computation in complex localized dynamics. [S2] His 1985 work examined undecidable and computationally intractable questions through cellular-automaton examples. [S3] These sources undermine the assumption that a short deterministic specification guarantees a convenient shortcut to every long-run property. They do not classify this particular Life experiment as computationally irreducible.

The Rule 110 result is more specific. Wolfram's book provides the historical simple-rule framing; Cook's 2004 paper proves universality through a constructed simulation of a cyclic tag system on a bi-infinite row with repeating backgrounds and a finite encoded middle segment. [S4][S5] Universality establishes a capability under an appropriate encoding. It does not mean that every starting row performs useful computation, and it supplies no proof that a biological tissue or a Jev selector has the same capability.

There is also an important vocabulary trap. A memoryless local lookup rule can participate in a globally stateful system: information can persist in spatial patterns and their interactions. Our ID9 impossibility concerns the information available at one update boundary. Adding persistent local registers would still permit a deterministic transition function over an enlarged state. The proposed change concerns what the state contains, rather than whether the system obeys rules.

The separate Rule 30 notes provide a bounded example of conjecture testing. The search examined all 8,191 nonzero initial rows supported on positions [-6,6]. For candidate periods 1-32, every tested center-column periodicity from time zero was broken by step 43. The notes also include 4,096 finite-trace reconstruction cases. These are finite computations, not an all-time theorem. [E3]

In particular, breaking a period from time zero does not rule out an eventually periodic tail after a transient. Testing those initial supports and periods does not cover larger supports or other periods. The all-time inverse-tail argument remained unresolved, already at period two. The Rule 30 center-column question must therefore stay separate from Cook's proved Rule 110 universality result. [E3][S5]

The appropriate evidential threshold depends on the statement. Enumerating 512 inputs completely checks a finite Boolean table. Two incompatible constraints on one input suffice to refute a proposed function class. A collection of finite traces cannot, by itself, resolve an unrestricted infinite-time assertion. This distinction also disciplines model-assisted research: a useful conjecture or reconstruction can guide the next test without becoming a proof through repetition or confidence.

A research trail: 110, 137, 42, and 30

For an elementary binary rule, write the three-cell input as (L,C,R), with index $i = 4L + 2C + R$. The output is bit i of the rule number. Reading the eight outputs from 111 down to 000 gives Rule 110 as 01101110 and Rule 137 as 10001001. Their relationship is exact: $F_{137}(1-L, 1-C, 1-R) = 1 - F_{110}(L, C, R)$. Checking all eight inputs exhausts the

local claim. Induction transfers it to a complete trajectory when the starting states, background, and boundary values are complemented together. [S16]

Rule 137 therefore transfers any correctly encoded Rule 110 computation through a change of colors. Cook's periodic-background construction transfers too. This is inherited computational equivalence, not a second independent universality proof. Starting Rule 137 from the same isolated live cell on the same zero background is a different experiment: its 000 input produces 1. The companion atlas pairs the transformed states explicitly. [S5][S16]

Rule 42 gives a useful contrasting control. Its output bits are 00101010, and exhaustive enumeration gives $F_{42}(L,C,R) = R \text{ AND NOT}(L \text{ AND } C)$. On an otherwise zero line, one isolated live cell moves one site left at each synchronous update. Its population stays one while its position changes. A wrapped finite board repeats after traversing the circumference; a dead edge eventually removes the pulse. This simple case catches scoring systems that confuse a steady headcount with a fixed state. The number is also a small wink; the research value comes from the stated rule and test.

Further researchers can test representation invariance: complement a complete experiment, map its boundaries, decode its output, and compare whole states. Canonicalize color and reflection aliases before counting independent findings. Keep the initial background, seed, encoding, boundary, and time horizon in every receipt. Rule 30 remains a useful search problem with the finite limits described above. A space-time image checks an implementation; a universality or all-time claim requires its own proof.

Rule	Bits: 111 to 000	What is established	Next useful test
110	01101110	Cook's encoded universality result	Declare the periodic background
137	10001001	Color conjugacy to 110	Decode complemented full states
42	00101010	$R \text{ AND NOT}(L \text{ AND } C)$	Track a moving pulse, not headcount
30	00011110	The cited search is finite	Target the unresolved tail argument

ELEMENTARY RULE ATLAS / SPACE-TIME PLATE

Three bits. Three views.

Standard Wolfram encoding · index = 4L + 2C + R · synchronous update



256 cells × 96 ticks · Rule 137: inverted seed / background 1 · wrap replay

Figure 4 Computed 256-cell, 96-update wrapped space-time examples. Rules 110 and 42 begin with one central live cell on background 0; Rule 137 uses the complemented seed and background 1. The 110/137 trajectories agree after decoding the colors. The Rule 42 pulse moves left. These plots check implementations, not universality. [S16]

"very much simpler rules should also show universality."

STEPHEN WOLFRAM / S4

07 / FIELD NOTE

7

Put typed decisions behind an evidence boundary

For assistants such as AlexAI, the practical lesson is to make authority and state explicit before asking a probabilistic system to choose.

TypeSafe describes Jev as making typed probabilistic decisions over supplied state. Its documented limitations include arithmetic, counting, indirection, context handling, and structural invariants. [S13][S14] The architecture proposed here uses that bounded role. It leaves exact operations with code and treats the selector's answer as a candidate decision whose semantic validity still requires checking.

Start with source acquisition. Code should preserve the document version, content location, extraction method, and binding between a candidate value and its evidence. OCR and layout acquisition are separate from selection. Metadata can triage material, but cannot substitute for content evidence. Before a candidate reaches Jev, eligibility checks should enforce the requested identity, period, unit, authority, version, and provenance, routing ambiguous extraction or source conflicts for review.

Consider a hypothetical request for an annual quantity in tonnes. One document candidate has the right entity, year, and unit. Another contains a similar number for a different year; a third is only a filename suggesting relevance. Code excludes the latter two before selection. Jev receives the qualified candidate and its supporting content, with permission to choose it, decline the offered candidates, or report insufficient evidence. It does not decide whether an ineligible year is close enough.

A small state contract makes that boundary inspectable: `DecisionState = {task, eligible_candidates, evidence_bindings, checks}`; `Decision = choose(candidate_id) | none | unknown`. Here, `none` declines the offered candidates; `unknown` indicates insufficient evidence to decide. A receipt binds the input-state hash, candidate set, selected identifier, model version, validation outcome, and fallback reason. These fields are useful only if they refer to preserved artifacts rather than decorative labels.

Execution, explanation, and review

After selection, code resolves the identifier, checks it against the permitted set, and executes the allowed transition. For Life, that means exact neighborhood encoding and synchronous updates. For document work, it means returning the bound value and source location without silently changing its unit or scope. A valid type proves membership in an allowed output structure. It does not prove that the source is accurate or that the selected interpretation answers the question.

A generative model belongs in the research and explanation layer: organizing sources, proposing hypotheses, describing traces, and drafting an account that reviewers can check. It should not acquire authority merely by restating a selector's answer fluently. Independent review must return to content evidence and, where applicable, source-owner qualification. Agreement between two models can be informative, but shared mistakes remain possible when both depend on the same defective extraction.

Fallbacks should preserve that distinction. On abstention, timeout, failed validation, or provider unavailability, a deterministic path can return an unresolved status with a reason, or use a previously approved result only if its eligibility still holds. It should never silently replace a failed typed decision with an unconstrained generated answer. Receipts must record whether the outcome came from selection, validation, fallback, or review.

A small supporting experiment

The synthetic Type1 exercise comprised eight live requests and 23/25 primary checks. A three-swap location puzzle scored 2/4 and repeated the same two errors. When ordinary code supplied the completed transition state, both final locations were answered correctly. [E3] This supports the division of labor in a narrow example: supply resolved state for selection rather than asking the selector to reconstruct a multi-step update chain.

The exercise establishes neither generalized accuracy nor a latency advantage. Its checks are small and related, and repeated errors are not independent observations. Future performance reporting should separate acquisition, orchestration, model evaluation, validation, and end-to-end time. Likewise, a saved rule table replaying thousands of ticks is a deterministic engine workload, not evidence of a model making thousands of fast live decisions.

A live qualification: sixteen resolved relations

A one-off server-side qualification on the live eR33t physics host tested this boundary on 30 September 2026. Deterministic code supplied two bit outputs for each of sixteen cases: eight Rule 110/137 complementary-input pairs and eight Rule 42/Boolean-gate pairs. Jev 1.13.0 chose between same and opposite. The expected relation

was withheld from the request, and candidate order was balanced eight ways each. All sixteen responses passed the stated response-shape checks and matched the independently recomputed relation. [E4]

This is a finite qualification of selection over resolved state. It does not show Jev deriving the rule identities, proving universality, or achieving general accuracy. The proof is separate: the eight-row deterministic truth table exhausts each local identity, and induction carries the 110/137 complement relation across a trajectory with complemented states and boundary values. [E5] Provider-fetch latency ranged from 107.75 to 301.67 milliseconds, with a 119.96-millisecond median. Those measurements include the network fetch; they are not isolated model-evaluation times. The sequential sixteen-case run took 2.640 seconds. [E4]

Resolved-state family	Expected relation	Validated Jev selections
110 / 137, complementary inputs	opposite	8 / 8 correct
42 / R AND NOT(L AND C)	same	8 / 8 correct

08 / FIELD NOTE

8

Proposed experiment: distributed state with a budget

Everything in this section is proposed and unrun. The aim is to test whether locally maintained history improves unseen digital recovery while preserving declared healthy behavior.

The ID9 witness supplies a sharp starting point: additional state must actually distinguish the aliased contexts, not merely accompany a more elaborate controller. The biological inspiration is distributed, state-dependent pattern control. The experimental object remains a finite automaton with an externally specified target. No tissue model, biological protocol, or synthetic performance result is implied.

Preregister the targets and observations

Use the same 64×40 geometry under separately reported wrapped and dead-edge semantics. Before fitting a controller, publish healthy reference trajectories, damage windows, allowed transformations, and scoring rules. A proposed 256-generation horizon would require exact whole-grid healthy preservation throughout. For damaged trials, preregister recovery within 32 ticks followed by 64 ticks of continued reference agreement, scheduling damage early enough to observe both intervals.

Score two endpoints separately. First, test the incompatible ID9 one-step requirements, including healthy Ro. Second, test longer-window recovery on held-out damage. A memoryless rule cannot satisfy the first pair with the same local input, but could potentially recover by a different, slower route. Keeping these endpoints distinct prevents the impossibility proof from being overstated as a theorem against every memoryless recovery strategy.

Define alignment before looking at results. The primary visible target should be the complete healthy bitmap at the specified phase and position. Report phase-delayed or translated matches separately, with displacement and motion direction. Population agreement alone is insufficient. Explicitly label local-input aliases, orientation

aliases, and phase aliases so that two records with different names do not masquerade as different underlying challenges.

Constrain memory and compare meaningful controls

Give each site two persistent bits, initialized uniformly and updated during an undamaged warm-up. A site may read only the current 3×3 occupancy neighborhood and the corresponding two-bit memory states: at most 27 input bits, with two stored bits per site. Updates are synchronous. Prohibit coordinates, global clocks, orientation labels, damage flags, hidden buffers, and post-damage access to the reference trajectory. Boundary handling for memory must be specified as exactly as occupancy handling.

Fork damaged and healthy runs from a shared pre-damage history. Never reload the damaged controller's state from its healthy counterpart. Distinguish occupancy-only damage, which leaves site memory intact, from coupled damage, which also resets memory at affected sites. This distinction matters: a persistent register beneath a deleted visible bit could retain exactly the information the test seems to remove. Recovery that survives local memory loss is a stronger distributed-state result.

Include exact Life, the ID9 mutant, and a fitted memoryless 3×3 baseline with matched training cases, search effort, and model-call allowance. Also include a current-only 5×5 observer, reading 25 occupancy bits, to test whether spatial context explains any advantage. Its computational capacity is not identical to the memory model's; it is a deliberate alternative mechanism with a comparable input-width budget.

Add a stored-target oracle that receives the healthy next bitmap. Label it an explicitly circular upper bound: the answer is supplied to the controller. Its purpose is to check target attainability and scoring, not demonstrate learned repair. Sham memory should use the same interface without case-relevant history. Yoked memory should come from a different matched history, breaking the binding between a trial and its own past while retaining comparable state structure.

Corrupt memory deliberately: reset local patches, flip stored bits, and disrupt spatial associations under preregistered conditions. Check undamaged patterns and remote background sites for off-target changes. A controller that repairs one deletion by producing unrelated births elsewhere has failed whole-grid preservation, even if a small crop looks convincing. These interventions test whether useful information is locally stored, redundantly distributed, or merely correlated with an easier cue.

Test transfer, accounting, and failure

Separate genuine damage transfer from symmetry transfer. Split primary damage-history classes before symmetry augmentation, keeping equivalent histories together. Reserve unseen damage topologies, such as selected multi-cell removals or additions, for testing. Run a separate challenge with held-out rotations, reflections, translations, and phase offsets; identify cases that are transformed versions of familiar ones rather than counting them as independent novelty. Include boundary-near placements where geometry changes the task.

Save full occupancy and memory fields. Report visible bitmap distance, phase alignment, displacement, direction, and sustained motion, together with memory-field agreement against the healthy control. A visible recovery does not establish recovery of the augmented state. Detect cycles using the full state relevant to the controller, and label nonrepeating runs at the horizon as censored. Report results by damage family and alias class rather than treating all rows as independent probability samples.

Freeze each controller before final testing. Account separately for generative proposal calls, Jev selections, fitting or search operations, and any live intervention calls. Deterministic replay is not another model call. A receipt should bind controller version, split definition, state budget, seeds, settings, and traces, allowing an independent reviewer to rebuild both successful and failed trajectories.

The hypothesis loses support if improvement disappears on held-out damage, if healthy fidelity degrades, or if sham and yoked state perform comparably to correctly bound history. A wider memoryless observer matching the improvement would favor spatial context over persistence as the sufficient addition. Success requires an advantage that survives these alternatives within the declared finite window; the oracle's performance cannot earn that claim.

09 / FIELD NOTE

9

What a system must know to repair

The philosophical connection concerns the relation among rules, goals, and available information. Those three can be separated precisely enough to test.

A rule determines what happens next. A target determines which outcomes count as acceptable. An observation boundary determines what a controller can distinguish when choosing an action. In the ID9 witness, the target distinctions are finer than the controller's input distinctions. The impossibility follows from that mismatch. Enlarging the state can refine the distinctions, but only a successful design establishes that the added information is maintained and used appropriately.

This also clarifies the role of memory in software. An audit log preserves history for a reviewer; it does not automatically make that history available to the decision process. A source receipt establishes where a candidate came from; it does not by itself establish the candidate's truth. Persistent controller state can affect future behavior, but its relevance must be demonstrated through intervention. Storage, provenance, and useful control memory are related engineering resources with different tests.

Levin's conceptual lens encourages asking where coordinated control begins and ends. Wolfram's framework encourages taking simple rules seriously as computational objects. The digital lab places both questions inside an inspectable boundary: what information is present, which transition is allowed, what target is being pursued, and what observation would falsify the claim? That bridge is ambitious because it connects form and decision-making; it remains bounded because no shared biological mechanism has been demonstrated.

For Team eR33t, the practical commitment is consequently specific. Preserve exact engines and qualified sources. Expose the selector's permitted choices. Keep abstention and fallback visible. Judge proposed memory by unseen perturbations and healthy controls. The missing cell in the ID9 witness did not simply reappear; a neighboring transition restored the next bitmap. Explaining that exact event, including the healthy path it contradicts, is more informative than calling the result intelligent. More context is a justified next variable. Repair remains an earned result.

"We can only see a short distance ahead, but we can see plenty there that needs to be done."

ALAN M. TURING / S15

Continue in the live labs

Open the public article, live experiments, replay, and evidence bundle from this paper.

The illustrated research article	The publication hub: PDF, demonstrations, replay, and evidence.
Live Life Lab	Explore the recorded Jev table and exact Conway comparison.
Method and evidence	Read the original lab method and recording boundaries.
Neural Cellular Automata Lab	A separate growing-cells sandbox for further exploration.
Memory and recovery companion	Run the embedded frozen table, local rule variant, and stored-target control offline.
Elementary rule atlas: 110 / 137 / 42 / 30	Inspect truth tables and synchronized color-conjugate trajectories.
The production Jev qualification	Read the sixteen-case method and open its public result rows.
Watch the 30-second mechanism replay	Computed frames show the neighboring birth and whole-bitmap equality.
Research companion bundle	Download the public data, replay code, figures, and compact evidence supplements.

Appendix

Scope, provenance, and publication checks

This is a sourced research essay with a finite digital lab, not a peer-reviewed empirical biology paper. Section 8 is a proposed digital-only experiment; it reports no completed trials.

Dataset and replay limits

E1 covers the frozen-table audit and deterministic sweep. E2 covers the wave-3 damage census and mechanism clarification. E3 covers the bounded synthetic Type1 and Rule 30 notes. Equivalent tables, repeated trajectories, phase aliases, and related checks must not be counted as independent replications. Finite replay results do not establish general Jev accuracy or infinite-grid behavior.

Source verification

Public primary, publisher, and author-hosted sources were checked on 30 September 2026. The numerical claims, measurement clocks, quotations, and interpretations were reviewed against the saved experiment artifacts and cited sources. These editorial checks are not peer review.

AI-assisted preparation

Research organization and drafting used GPT-6 Astra in pro mode at maximum reasoning effort. Codex coordinated independent source and numerical checks, editorial revisions, and page-by-page visual review. Justin Hedge supplied the project direction and personal dedication. Model confidence and agreement are not verification.

Art and scientific figures

The cover was created with GPT-Image-2.5 Sunburst as editorial illustration. The scientific figures and 30-second replay are computed from saved rules, source-bound counts, and independently checked state transitions. The offline lab embeds the 512 recorded Jev decisions. Its stored-target mode receives the answer as a control; its B23/S23 variant is illustrative. Neither implements the proposed persistent-memory experiment.

Credits and boundaries

Life is credited to Conway and its public exposition to Gardner; Wolfram supplies computational context, and Michael Levin supplies conceptual inspiration. The subject is morphogenesis and regeneration, not origin-of-life biogenesis. No collaboration, endorsement, biological healing, or medical application is claimed.

Explore and reproduce

The companion publication links the offline lab, computed replay, evidence summaries, and research bundle. The bundle contains the original public rule-repair data and replay code, the compact wave-3 census supplement, and the source-bound mechanism witness. The full wave-3 case ledger is hash-bound but omitted from that supplement. Reference E2 supplies the recovery clock and matched denominators; E3 labels the limited Type1 and Rule 30 evidence. E4 records the sixteen live typed selections; E5 supplies the independent deterministic identities and induction argument.

References

- [S1] The chemical basis of morphogenesis. Alan Mathison Turing (1952).
- [S2] Universality and complexity in cellular automata. Stephen Wolfram (1984).
- [S3] Undecidability and intractability in theoretical physics. Stephen Wolfram (1985).
- [S4] A New Kind of Science: The Rule 110 Cellular Automaton. Stephen Wolfram (2002).
- [S5] Universality in elementary cellular automata. Matthew Cook (2004).
- [S6] The bioelectric code: An ancient computational medium for dynamic control of growth and form. Michael Levin; Christopher J. Martyniuk (2018).
- [S7] The Computational Boundary of a "Self": Developmental Bioelectricity Drives Multicellularity and Scale-Free Cognition. Michael Levin (2019).
- [S8] Transmembrane voltage potential controls embryonic eye patterning in *Xenopus laevis*. Vaibhav P. Pai; Sherry Aw; Tal Shomrat; Joan M. Lemire; Colin D. McCaig; Michael Levin (2012).
- [S9] Bioelectric signaling regulates head and organ size during planarian regeneration. Wendy Scott Beane; Junji Morokuma; Joan M. Lemire; Michael Levin (2013).
- [S10] An automated training paradigm reveals long-term memory in planarians and its persistence through head regeneration. Tal Shomrat; Michael Levin (2013).
- [S11] Endogenous bioelectrical networks store non-genetic patterning information during development and regeneration. Michael Levin (2014).
- [S12] Modeling Planarian Regeneration: A Primer for Reverse-Engineering the Worm. Daniel Lobo; Wendy S. Beane; Michael Levin (2012).
- [S13] Jev 1.13 jaggedness. TypeSafe AI (2026).
- [S14] Introducing System One Models & Jev. Diogo Almeida / TypeSafe AI (2026).
- [E1] Rule repair: what a recorded Jev table does to Conway Life. Team eR33t / AlexAI lab (2026).
- [E2] Wave-3 recovery census and one-step mechanism: companion evidence. Team eR33t (2026).
- [E3] Bounded Type1 and Rule 30 research notes: source-labelled evidence summary. Team eR33t (2026).
- [S15] Computing Machinery and Intelligence. Alan M. Turing (1950).
- [S16] A New Kind of Science: Atlas of Elementary Cellular Automata, Equivalent Rules. Stephen Wolfram (accessed 2026).
- [E4] Production Jev qualification: sixteen resolved rule relations. Team eR33t (2026).
- [E5] Elementary-rule identities: exhaustive local checks and induction argument. Team eR33t (2026).

Editorial edition 1.0 · 30 September 2026 · Team eR33t Science. A non-peer-reviewed research essay; finite digital results and proposed experiments are identified separately.