Memory Without Storage, Learning Without a Learner: Observer Inference Across Four Computational Substrates

Author: D. Neale
Affiliation: Independent Researcher, Goleudy.ai, Rochester, New York, USA
Email: david.n@goleudy.ai
ORCID: 0009-0006-7749-9876
Date: April 2026
Status: Submission Ready

Companion paper: *Why Observers Construct Intelligence: Cognitive Categories as Prediction Failure* (D. Neale, 2026a)

Abstract

We propose that memory, learning, and decision-making are not properties of systems but performance characteristics of an observer's predictive model. We define three measures — memory as model obsolescence, learning as decreasing surprise, decision as surprise spike — all derived from a single rolling prediction applied to observable emissions. We apply these measures identically across four computational substrates: 256 elementary cellular automata, 18 Game of Life patterns, 10 Gray-Scott reaction-diffusion regimes, and 6 sorting algorithm variants. The same categorical structure emerges in all four substrates despite their having nothing in common as physical systems. The M∩L∩D overlap identifies the systems that human researchers independently classify as "complex" or "interesting." A reaction-diffusion system subjected to repeated identical perturbations produces an emission profile qualitatively matching published habituation data from single-celled organisms, yet the system is a partial differential equation with no plausible cognitive capacity. Restricting the observer's access to emission features increases apparent learning, consistent with the prediction that cognitive attribution scales with observer incompleteness. Two independent operationalisations of memory (prediction-based and Cohen's d) correlate at $\rho = 0.75$ across systems. The results support the thesis that cognitive categories are features of the observation relation, not of the observed system.

Keywords: observer inference, memory, learning, Hidden Markov Models, cellular automata, reaction-diffusion, Bayesian surprise, predictive model

---

1. Introduction

A single-celled organism with no nervous system appears to learn. Recent experiments with the ciliate *Stentor coeruleus* have demonstrated both habituation — declining response to repeated mechanical stimulation — and what appears to be associative learning, where the organism pairs a weak stimulus with a subsequent strong stimulus and begins responding to the weak stimulus alone (Eckert et al. 2024, Gershman et al. 2026). The experimentalists ask a natural question: "How do cells without brains manage something so complex?"

This paper proposes a different question: what if the complexity is not in the cell but in the observation?

The observer watches *Stentor* contract in response to taps. The contraction rate declines with repetition. The observer records these emissions — the observable contraction events — and infers that the cell has learned to ignore a repeated stimulus. But the observer has no access to *Stentor*'s internal dynamics. The inference of "learning" is constructed entirely from the emission stream. The question then becomes: would the observer construct the same inference from the emission stream of any system whose response to repeated perturbation declined over time — regardless of whether that system had any capacity for learning?

We show that the answer is yes. A reaction-diffusion system governed by partial differential equations, subjected to the same perturbation protocol used in *Stentor* experiments, produces an emission profile with declining response magnitude followed by partial recovery — the same qualitative structure that the observer would label "habituation" in a biological system. The PDE has no memory, no learning mechanism, and no cognitive capacity of any kind. Yet the observer, working from emissions alone, would construct identical cognitive attributions.

This result is not a curiosity about one system. It reflects a structural feature of observation under incomplete access. Whenever an observer builds a predictive model of a system whose full state is not accessible, three characteristics of the model's performance emerge inevitably: the model can fall behind the system (the system moved persistently — the observer infers "memory"), the model can improve over time (the system became less surprising — the observer infers "learning"), and the model can be caught off guard (the system did something unexpected — the observer infers "a decision"). These three — memory, learning, decision — are not properties of the system. They are properties of the observer's prediction performance.

We formalise these three measures, derive them from a single rolling predictive model, and apply them identically across four computational substrates that share no physical, mathematical, or structural features: elementary cellular automata, Conway's Game of Life, Gray-Scott reaction-diffusion, and classical sorting algorithms. The same categorical structure emerges in all four. The observer constructs "memory," "learning," and "decisions" from sorting algorithms running fewer than ten lines of fully visible, deterministic code. The cognitive vocabulary was never about the systems. It was about the observation.

The philosophical implications of this result are developed in a companion paper (Neale 2026a). The foundational derivation showing why the observer-inference architecture takes this specific form — why five hidden constraint dimensions, why the relational ontology — is given in Neale (2025). The present paper stands on its own as a computational result with specific, falsifiable predictions.

---

2. The Observer-Inference Framework

2.1 Observer, Emissions, Hidden States

Consider any system S observed by any observer O. The system exists in some space of possible configurations whose full dimensionality we denote D_{total}. The observer has access to a set of measurable features — quantities that instruments or senses can detect. These observable features constitute the emission stream: a sequence of vectors $E(t) \in \mathbb{R}^d$ where $d = D_{obs} \ll D_{total}$.

The observer does not have access to the system's internal dynamics. Many distinct internal configurations can produce identical emissions. The observer's task is prediction: given the emissions observed so far, what emission should be expected next?

This is the standard structure of a Hidden Markov Model (HMM). The system's internal dynamics constitute the hidden layer — configurations the observer cannot directly see. The measurable features constitute the emission layer — the data the observer actually receives. The observer constructs a model relating the two, inferring hidden trajectories from observable patterns.

We make no specific claims in this paper about the structure of the hidden layer. The companion papers argue that the hidden state space has a specific five-dimensional structure arising from the requirements of robust distinguishability (Neale 2025, 2026a). For present purposes, we require only the minimal HMM assumption: there exist hidden dynamics that generate the emissions, and the observer builds a predictive model from the emission stream without direct access to those dynamics.

2.2 The Observer's Predictive Model

The observer uses a rolling mean predictor. At each timestep t, the observer's prediction of the next emission is the mean of the w most recent emissions:

$$\hat{E}(t) = (1/w) \Sigma_{i=1}^{w} E(t-i)$$

The observer's surprise at timestep t is the prediction error:

$$s(t) = \sqrt{(\mathrm{mean}((E(t) - \hat{E}(t))^2))}$$

This is the simplest possible Bayesian predictive model: the posterior predictive distribution is centred on the running mean, and surprise measures how far each new emission falls from that centre. We choose this deliberately — more sophisticated models would change the numbers but not the structure of the results. The simplicity of the model makes the results harder to dismiss as artifacts of a cleverly chosen predictor.

A natural question: would a more sophisticated observer — a Kalman filter, a recurrent neural network, a full variational inference engine — see the same cognitive categories? The framework predicts it would see *fewer*. A better predictor produces lower surprise, fewer surprise spikes, and smaller early-vs-late differences. M, L, and D would all decrease.

The cognitive attribution is strongest with the simplest model and diminishes as the observer becomes more capable. In the limit of a perfect predictor (one that knows the system's dynamics completely), surprise is zero everywhere and $M = L = D = 0$: the observer perceives no memory, no learning, and no decisions — because there is no gap between prediction and observation. This is the Cognitive Hierarchy: observers who see more perceive less intelligence. Our results using the simplest possible predictor therefore represent an *upper bound* on cognitive attribution — a conservative choice, not a limitation.

2.3 Three Questions About Prediction Performance

From the observer's prediction performance, three characteristics emerge:

Memory (M): Has the system moved beyond my model?

The observer builds a predictor from the early portion of the emission stream and applies it to the late portion. Separately, it builds a predictor from the late portion and applies it to the late portion. If the early model performs significantly worse on late data than the late model does, the system has moved persistently — the past is showing through in the present, and the observer infers that the system "carries its history."

Formally:

M = error(early_model, late_data) − error(late_model, late_data)

where error is root mean squared prediction error. High M means the emission regime shifted and stayed shifted. Zero M means the early model still works — no persistent change.

As an independent validation, we also compute Cohen's d between early and late emission distributions. If both measures rank systems consistently, the result is robust to operationalisation.

Learning (L): Is the system becoming less surprising?

The observer compares its mean surprise in the early third of the emission stream to its mean surprise in the late third:

L = (s̄_early − s̄_late) / s̄_early

subject to an activity gate: the system must still be producing variable emissions in the late phase. If the system simply died (emissions became constant), the surprise decrease reflects cessation, not learning. The gate requires that mean absolute change in late emissions exceeds 1% of early emission variability.

High L means the observer's predictions improved while the system remained active — the observer infers "this system is learning." Zero L means either no improvement or the system stopped.

The 1% activity floor is not a sensitive parameter. Testing all 256 CA rules across a 50-fold range of thresholds (0.1% to 5%), zero rules change their $L > 0$ classification. The gate catches genuinely dead systems (uniform convergence, complete die-off) and passes genuinely active systems (ongoing dynamics of any kind) regardless of the threshold. The only system in our full test set that is sensitive to the threshold is the GoL R-pentomino, which falls below a 5% gate because its late-phase activity consists of small oscillators producing only ~2% of early-phase variability. At all thresholds from 0.001% to 2%, it registers $L = 0.994$.

The connection to Friston's free energy principle (Friston 2010) is direct: surprise in the Bayesian sense is the negative log-likelihood of the observation under the model. Decreasing surprise is decreasing free energy. When we say "the system is learning," we are saying "the observer's free energy is decreasing" — which is precisely Friston's claim about biological systems, reframed as a property of the observation rather than of the system.

Decision (D): Has the system departed from its pattern?

At each timestep, the observer compares its current surprise to the local baseline — the median surprise over the preceding window. If the current surprise exceeds the baseline by more than 2 median absolute deviations (MAD), the observer registers a "decision event":

$$\mathbf{D(t) = 1 \text{ if } s(t) > \text{median}(s) + 2 \times \text{MAD}(s), \text{ else } 0}$$

over a rolling window. The decision rate D_rate is the fraction of timesteps with decision events. This is a robust surprise detector — the MAD-based threshold is resistant to outliers and does not assume normality.

High D_rate means the system frequently surprises the observer — the observer infers "this system keeps making decisions." Zero D_rate means the system never departs from the observer's expectations.

2.4 Co-presence, Not Sequence

M, L, and D are computed at every point in the emission stream. They are simultaneous aspects of the observer's prediction performance, not sequential stages. A system can show high M (persistent regime change), high L (decreasing surprise), and high D (frequent surprise spikes) all at once. The region where all three are simultaneously active is what observers label "intelligent" or "interesting."

This co-presence has a consequence for memory specifically. Every instance of state persistence (M) is encountered within a gradient context (L) and a current configuration that may be near a transition boundary (D). The observer never encounters "pure memory" — memory is always a relationship between the current prediction model and the basin topology shaped by the system's history. This is why memories change with each recall: every "retrieval" is a new prediction from a new model state, not a readout from storage (see Neale 2026a for the philosophical development).

Goals, purpose, and planning are narratives the observer constructs from the temporal pattern of L and D. When decision events cluster early in the emission stream and L is positive throughout, the observer infers “the system was working toward something.” When D events are distributed uniformly and L is flat, the observer infers “this system is just reacting.” These are derived interpretations, not additional measures.

2.5 Predictions

The framework makes three predictions that distinguish it from the standard cognitive vocabulary:

(a) **Substrate independence.** The same M/L/D structure should appear across systems with fundamentally different physics, because M/L/D are properties of the observation, not of the substrate.

(b) **Observer scaling.** Apparent learning and memory should increase when the observer’s access is restricted, because cognitive attribution scales with the gap between what the observer sees and what the system does.

(c) **Cognitive attribution without cognition.** Systems with no plausible cognitive capacity — PDEs, sorting algorithms — should produce the same observer inferences as systems traditionally called intelligent, when subjected to the same emission analysis.

---

3. Four Substrates, One Observer

3.1 Elementary Cellular Automata

All 256 elementary cellular automata rules were simulated on a 1D grid of 101 cells with periodic boundaries, random initial conditions (seed = 42), for 300 timesteps. Four emission features were computed at each step: cell density (fraction of live cells), spatial entropy (Shannon entropy of 3-cell block distribution), boundary density (normalised count of 0↔1 transitions), and activity (fraction of cells changed since previous step). The observer does not know the rule number or Wolfram classification.

3.2 Game of Life

Eighteen Game of Life patterns were simulated: 3 still lifes (Block, Beehive, Loaf), 5 oscillators (Blinker, Toad, Beacon, Pulsar, Pentadecathlon), 4 spaceships (Glider, LWSS, MWSS, HWSS), and 3 methuselahs (R-pentomino, Acorn, Diehard). Each was placed at the centre of a 50×50 grid (60×60 for Acorn) with periodic boundaries and evolved for 300 steps (400 for methuselahs). Four emission features were computed: cell density, spatial entropy (2×2 block distribution), boundary density (live cells with dead neighbours), and activity. The observer does not know the pattern name or type.

3.3 Gray-Scott Reaction-Diffusion

The Gray-Scott system ($dU/dt = D_u\nabla^2 U - UV^2 + F(1-U)$, $dV/dt = D_v\nabla^2 V + UV^2 - (F+k)V$) was simulated on a 64×64 grid with $D_u = 0.16$, $D_v = 0.08$, $dt = 1.0$, periodic boundaries, for 5000 steps sampled every 10 (yielding 500 emission samples). Ten parameter regimes were selected spanning qualitatively different behaviours: uniform death ($F=0.078$, $k=0.061$), stable spots (0.035, 0.065), labyrinthine stripes (0.040, 0.065), moving spots (0.014, 0.054), dividing spots/mitosis (0.028, 0.062), pulsing (0.025, 0.060), coral growth (0.062, 0.061), spatiotemporal chaos (0.026, 0.051), worm-like meandering (0.054, 0.063), and sparse spots (0.030, 0.062). Six emission features were computed: mean U, mean V, spatial variance of V, gradient energy (mean $|\nabla V|^2$), activity (mean $|\Delta V|$), and pattern entropy (Shannon entropy of discretised V). The observer does not know the parameter values.

3.4 Sorting Algorithms

Six sorting configurations were simulated on arrays of 30 elements: already-sorted control, standard bubble sort, self-sorting bubble sort (random-order element visits), self-sorting with frozen defects at positions 7, 15, and 22, standard selection sort, and bubble sort starting from reverse order. Five emission features were computed at each step: sortedness (fraction of adjacent pairs in order), mean displacement from sorted position, fraction of inversions, activity (fraction of elements moved), and block entropy. The complete source code for every algorithm is fewer than ten lines. The system dynamics are fully deterministic and fully visible. There is nowhere for intelligence to hide.

3.5 Habituation Protocol and Frequency Sweep

A stable Gray-Scott spot pattern ($F=0.035$, $k=0.065$, grown for 3000 steps) was subjected to repeated identical perturbations: a local concentration pulse ($\Delta V = +0.15$, $\Delta U = -0.15$, radius 3 cells) applied at a fixed location. Emission response was measured as root mean squared change in bulk emission features between pre-tap and post-tap states (50 relaxation steps after each tap). Twelve taps were administered at each frequency.

Crucially, the experiment was run at seven different tap intervals: 30, 50, 80, 120, 200, 400, and 800 timesteps between taps. The system, the perturbation, and the emission features are identical across all seven conditions. The only variable is the observer's choice of probing frequency. This design tests a specific prediction: if cognitive attribution depends on the observer's protocol, then the *same system* should appear to exhibit different cognitive properties — habituation, sensitisation, or no learning — depending on the tap frequency.

This protocol mirrors the mechanical tap paradigm used in *Stentor* habituation experiments (Eckert et al. 2024), where the ciliate is tapped every 45 seconds — approximately its contraction-extension recovery time. Our frequency sweep brackets this design by systematically varying the interval relative to the system's relaxation timescale.

3.6 Multi-Observer Feature Access Protocol

Five systems (CA Rule 54, CA Rule 30, GoL R-pentomino, GS mitosis, bubble sort) were analysed with systematically restricted emission access. Observer A received all emission features. Observer B received pairs of features (averaged across all combinations). Observer C received single features (averaged across all).

This tests the prediction that restricted access increases cognitive attribution. However, a naive average across restricted observers can obscure what is happening. A single-feature observer watching an informative channel may see strong learning, while a single-feature observer watching an uninformative channel sees none — and their average may be lower than the full-access observer. To understand the mechanism, we also decomposed L by individual feature for each system, identifying which features carry the learning signal and which contribute noise. This reveals whether the hierarchy operates through *inference gap* (less access → more uncertainty → more apparent cognition) or through *dilution* (more features → more noise → less apparent cognition).

---

4. Results

4.1 M/L/D Discriminates System Types

Across all 256 elementary CA rules, the three measures separate Wolfram classes without the observer knowing the classification:

Wolfram Class	n	M (mean ± sd)	L (mean ± sd)	D (mean ± sd)
Class 1 (uniform)	24	0.000 ± 0.000	0.000 ± 0.000	0.000 ± 0.000
Class 2 (periodic)	194	0.017 ± 0.028	0.147 ± 0.237	0.101 ± 0.091
Class 3 (chaotic)	30	0.017 ± 0.013	0.071 ± 0.069	0.181 ± 0.047
Class 4 (complex)	8	0.019 ± 0.010	**0.274 ± 0.189**	0.137 ± 0.062

Class 4 rules show the highest L (the observer sees them becoming less surprising while remaining active) and moderate D (they surprise the observer but not as constantly as Class 3). Class 3 shows the highest D (constant surprise — chaotic systems keep breaking the observer's predictions) but low L (they never become less surprising). Class 1 shows zeros across all measures — the observer sees nothing to interpret.

The Game of Life results follow the same pattern:

Pattern type	n	M	L	D
Still life	3	0.000	0.000	0.000

Pattern type	n	M	L	D
Oscillator	5	0.005	0.006	0.039
Spaceship	4	0.001	0.059	0.082
Methuselah	3	0.000	**0.331**	0.119

Methuselahs — systems with complex transients that resolve into structured attractors — show the highest L. The R-pentomino achieves L = 0.994: the observer's surprise drops 99.4% from early to late while the system remains active. This is the strongest "learning" signal in any substrate.

Gray-Scott results separate pattern-forming from non-pattern-forming regimes:

Regime	M_pred	L	D
Mitosis	0.253	**0.912**	0.200
Coral	0.193	**0.917**	0.190
Sparse spots	—	**0.954**	0.213
Stripes	0.325	0.311	0.221
Spots	0.150	0.107	0.217
Death	0.000	0.000	0.000

Sorting algorithms show clean discrimination:

Variant	M_pred	L	D
Bubble sort	0.000	**0.689**	0.200
Selection sort	0.168	0.408	0.143
Reverse bubble	**0.364**	0.420	0.048
Already sorted	0.000	0.000	0.000

The reverse bubble sort shows the highest M in any substrate — the observer's early model (trained on a fully reverse-ordered array) is completely useless on late data (a sorted array). The system moved maximally and stayed moved.

4.2 Cross-Substrate Convergence

The M∩L∩D overlap — the geometric mean of the three normalised measures — identifies "interesting" systems across all four substrates. The top-ranked systems are:

Rank	System	Substrate	M	L	D
1	GS Mitosis	Reaction-diffusion	0.253	0.912	0.200
2	GS Coral	Reaction-diffusion	0.193	0.917	0.190
3	GS Stripes	Reaction-diffusion	0.325	0.311	0.186
4	Sort: Reverse bubble	Sorting algorithm	0.364	0.420	0.048
5	CA 147	Cellular automaton	0.066	0.270	0.186

Rank	System	Substrate	M	L	D
6	Sort: Selection	Sorting algorithm	0.168	0.408	0.050
7	CA 041 (Class 4)	Cellular automaton	0.029	0.366	0.178
8	GoL R-pentomino	Game of Life	0.015	0.994	0.100

Systems from four different substrates — a reaction-diffusion PDE, a sorting algorithm, a cellular automaton, and a Game of Life methuselah — occupy the same region of M/L/D space. These systems share no physics, no mathematics, and no structural features. What they share is an emission profile that causes the observer's predictive model to exhibit the same pattern of performance: persistent regime change, decreasing surprise, and intermittent prediction failure. The observer, applying identical analysis to all 287 systems without knowing what any of them are, groups them together.

4.3 The Observer's Protocol Determines the Cognitive Attribution

Under the habituation protocol (Section 3.5), the same Gray-Scott system produced qualitatively different emission profiles depending solely on the observer's choice of tap frequency:

Tap interval	Response decline	Partial recovery?	Observer infers
30 steps	99%	Yes	Strong habituation with spontaneous recovery
50 steps	93%	No	Sustained habituation
80 steps	55%	Yes	Habituation with recovery
120 steps	38%	Yes	Mild habituation
200 steps	26%, then *increase*	—	Sensitisation
400 steps	19%	Erratic	No learning
800 steps	43%	Gradual	Ambiguous

The system is identical across all seven conditions. The perturbation is identical. The emission features are identical. The only variable is the interval between taps — a design decision the observer makes before collecting any data.

At short intervals (30–80 steps), the system has not fully relaxed between taps. Each successive perturbation displaces the system from an already-displaced state, producing a smaller incremental response. The observer sees declining response and infers habituation. At interval 80, the response drops 75% between the first and sixth tap, then partially recovers — matching two of Thompson and Spencer's (1966) hallmarks of habituation: frequency-dependent decline and spontaneous recovery.

At moderate intervals (200 steps), the system fully relaxes but the cumulative perturbations have begun to deform the local pattern structure. The response to later taps is *larger* than to early taps because the deformed structure is differently sensitive. The observer sees increasing response and infers sensitisation — the *opposite* of habituation, and traditionally considered a more sophisticated form of learning.

At long intervals (400–800 steps), the pattern has time to reorganise substantially between taps, producing erratic responses that the observer cannot easily categorise.

The same PDE, governed by the same reaction-diffusion equations, appears to *habituate* at one probing frequency, *sensitise* at another, and show *no learning* at a third. All four of the basic categories in the habituation literature — habituation, sensitisation, spontaneous recovery, and null response — can be produced from a single system by adjusting the observer's protocol.

Comparison with *Stentor* experiments. Eckert et al. (2024) tapped *Stentor coeruleus* every 45 seconds — the approximate time for the cell to re-extend after contraction. This interval falls within the cell's relaxation window, producing declining contraction probability over 60 repetitions. Had the experimentalists chosen a 5-second interval, the cell would have had no time to re-extend — the observer would see constant contraction, not declining response. Had they chosen a 5-minute interval, the cell would have fully recovered between taps — every response would be identical, with no decline. The "habituation" appears at 45 seconds because that frequency falls within the relaxation window where incomplete recovery produces declining incremental response.

Our Gray-Scott system shows the same frequency-dependent window. The observer who probes within the relaxation window sees "habituation." The observer who probes outside it sees "no learning" or "sensitisation." The cognitive attribution is co-determined by the system's relaxation dynamics and the observer's protocol — it is not a property of the system alone.

This result extends beyond habituation. Any protocol-based assessment of cognitive capacity — including standard tests for associative learning, decision-making, and memory — involves observer choices about stimulus timing, strength, measurement window, and feature selection. Each choice constrains which emission profiles are produced, and therefore which cognitive categories the observer will construct. The question "does this system learn?" is incomplete without specifying "under what observation protocol?" — which means the answer is always partly about the observer.

4.4 The Observer's Feature Choice Shapes Cognitive Attribution

The multi-observer test (Section 3.6) reveals a consistent mechanism: adding emission features to the observer's model never increases apparent learning — it can only dilute. Detailed feature-by-feature analysis clarifies why.

For every system tested, the single most informative emission feature produces L equal to or greater than the full-access L. Adding any further feature either leaves L unchanged or reduces it:

System	Best single feature	L (best single)	L (full access)	Effect of adding features
Bubble sort	inversions	**0.914**	0.689	Each addition dilutes L
GS Mitosis	var_V	**0.970**	0.912	Each addition dilutes L
CA Rule 30	spatial_entropy	**0.264**	0.092	Each addition dilutes L
CA Rule 54	density	**0.214**	0.129	Each addition dilutes L
GoL R-pentomino	spatial_entropy	**0.994**	0.994	Other features are dead (activity gate fails)

The pattern is universal: pair analysis shows that starting from the best single feature and adding any other feature reduces or maintains L but never increases it. For bubble sort, the inversions feature alone gives L=0.914; adding sortedness drops L to 0.729; adding activity drops it to 0.717. The richest single channel carries the clearest learning signal. Adding channels adds noise from features where surprise does not systematically decrease.

The R-pentomino case is instructive. The "hierarchy failure" reported in the averaged results (L drops from 0.994 to 0.249 with single-feature access) is not a failure at all. Three of four emission features (density, boundary density, activity) fail the activity gate — the system becomes effectively static in these channels by the late phase. The one observer watching spatial entropy sees L=0.994 — identical to the full-access observer. The averaged L drops because three of four restricted observers happen to be watching uninformative channels. The hierarchy holds for every observer who watches an informative channel.

This result has a direct parallel to the habituation frequency sweep. Just as the observer's choice of tap frequency determines whether the system appears to habituate, the observer's choice of which emissions to monitor determines how much "learning" is attributed. An observer tracking inversions in a sorting array infers strong learning (L=0.914). An observer tracking block entropy in the same array infers weaker learning (L=0.864). An observer tracking sortedness infers weaker still (L=0.571). Same system, same algorithm, same run — different cognitive attribution depending on which emission channel the observer monitors.

The general principle: **what the observer infers depends on both *which* emissions it tracks and *when* and *how frequently* it probes.** The cognitive attribution is jointly determined by the system's dynamics, the observer's feature selection, and the observer's protocol design. No single factor determines it alone. This is the central result of the paper expressed through experimental design rather than through abstract argument.

4.5 Two Memory Measures Agree

The prediction-based memory measure (M_pred: model obsolescence) and the statistical measure (M_cohen: Cohen's d between early and late emission distributions) rank systems consistently across all 23 test systems, with Spearman correlation $\rho = 0.75$ ($p < 0.0001$). The top-ranked systems by both measures are Gray-Scott pattern-forming regimes, sorting algorithms, and GoL methuselahs. The bottom-ranked systems by both measures are still lifes, already-sorted arrays, and Class 1 CA rules. The convergence of two independent operationalisations supports the claim that both are measuring the same underlying phenomenon: persistent regime change visible in the emission stream.

---

5. Discussion

5.1 Cognition as Prediction Performance

The results demonstrate that an observer applying identical prediction machinery to emission streams from 287 systems across four substrates constructs the same three-category decomposition in every case. The categories — memory (model obsolescence), learning (decreasing surprise), decision (surprise spike) — are not designed into the analysis. They emerge from the structure of prediction under incomplete access.

The $M \cap L \cap D$ overlap identifies "complex" or "interesting" systems across all substrates: Class 4 cellular automata, GoL methuselahs, Gray-Scott pattern-forming regimes, and sorting algorithms undergoing structured convergence. These are the systems that human researchers independently find most worth studying. The framework suggests why: "interesting" is itself a cognitive attribution that the observer constructs when all three prediction-performance characteristics are simultaneously active. The system is changing ($D > 0$), the changes are becoming more predictable ($L > 0$), and the system has moved persistently from its initial state ($M > 0$). This combination is what "interesting" means operationally.

The sorting algorithm results are particularly telling. Bubble sort is fewer than ten lines of fully deterministic code. The complete source code is visible. There are no hidden variables, no stochastic elements, no emergent properties beyond what the algorithm prescribes. Yet the observer, analysing the emission stream without seeing the source code, constructs memory (M_cohen = 3.88), learning ($L = 0.69$), and decisions ($D = 0.20$). If the observer can construct cognition from a bubble sort, the cognitive vocabulary is about the observation, not about the system.

A natural objection to the habituation result (Section 4.3) is: "This is just physical hysteresis. The PDE hasn't fully relaxed between taps. There is nothing cognitive here — it's just physics." This objection is correct about the mechanism. It is also, from the framework's perspective, an illustration of the thesis rather than a rebuttal of it. The critic who says "it's just physics" is an observer with access to the diffusion coefficients, the reaction rates, and the relaxation timescale — information about the hidden layer that the emission-only observer does not have. From the critic's position (high D_obs), the system shows no cognition. From the emission-only observer's position (low D_{obs}), the same system habituates. Both observers are constructing models from their available information. The critic's model invokes hysteresis. The biologist's model invokes habituation. Both are inferences from different access levels. The framework does not claim that the PDE "really" habituates. It claims that the distinction between "real habituation" and "mere hysteresis" is an observer-access distinction, not a system property. The biologist studying *Stentor* is in the same epistemic position as the emission-only observer watching the PDE — both see declining responses and construct cognitive explanations, because cognitive explanations are what observers with incomplete access construct.

5.2 Connections to Established Frameworks

The framework connects to three established traditions:

Friston's free energy principle. Our L measure (decreasing surprise) is operationally equivalent to free energy minimisation. Friston (2010) proposes that biological systems minimise variational free energy $F = KL(q||p)$, where q is the system's model and p is the true posterior. In our framework, the "system's model" is replaced by the "observer's model," and free energy minimisation becomes what the observer sees when the emission stream becomes more predictable. The reframing is not a rejection of Friston — it is a relocation. The free energy is real; its minimisation is real; but it is a property of the observation relation, not exclusively of the biological system.

Hopfield networks. In a Hopfield network, memory is navigation to an attractor basin (what we measure as M — persistent regime change), learning is landscape deformation through Hebbian weight updates (what we measure as L — the landscape becomes more predictable as basins deepen), and decision is convergence to a specific attractor from ambiguous input (what we measure as D — the system commits to one basin, surprising the observer who couldn't predict which). The Hopfield energy function $E = -\frac{1}{2} \Sigma\, w_{ij}\, s_i\, s_j$ is the constraint landscape whose gradient the system follows. The critical capacity at $0.14N$ is a phase transition in the geometry of this landscape — a prediction that the storage metaphor cannot derive but the constraint architecture makes inevitable.

Classical thermodynamics. The Helmholtz free energy $F = U - TS$ contains a term TS that represents the information content of the system's microscopic state — what the observer cannot see. All three of our measures relate to this term. M measures persistent changes in the hidden state structure (the TS landscape shifted). L measures the observer's improving estimate of the TS term (the model's approximation gets better). D measures sudden rearrangements of the hidden state structure that violate the observer's current estimate. The "cognitive" vocabulary and the thermodynamic vocabulary describe the same dynamics from different observer positions.

5.3 Reinterpreting Levin's Sorting Result

Zhang, Goldstein, and Levin (2024) studied classical sorting algorithms as models of basal intelligence. They found that self-sorting arrays — where each element decides locally when to swap — sort more reliably than top-down implementations in the presence of defects, can temporarily reduce progress to navigate around frozen elements, and show unexpected clustering behaviour in chimeric arrays mixing ascending and descending elements. They describe these as “emergent problem-solving capacities” and conclude that “basal forms of intelligence can emerge in simple systems without being explicitly encoded in their underlying mechanics.”

Our framework offers a complementary interpretation. The observer watching a sorting array’s emissions sees the same M/L/D structure as the observer watching a cellular automaton or a chemical system: persistent regime change (the array becomes sorted and stays sorted), decreasing surprise (the later steps are more predictable), and decision events (discrete transitions in the emission profile as elements swap past defects). Levin attributes these patterns to emergent competencies in the system. Our framework locates them in the observer’s predictive model. Neither interpretation is more correct — they are different observer narratives constructed from the same emission stream. The data do not adjudicate between them because the data are emissions, not hidden states.

5.4 Reinterpreting Stentor

Eckert, Gunawardena et al. (2024) show that biochemically plausible networks based on negative feedback or incoherent feedforward motifs satisfy all single-stimulus hallmarks of habituation as codified in animals with brains. Gershman et al. (2026) extend this to associative learning. These are important results. Our frequency sweep experiment (Section 4.3) adds a dimension that the biological experiments could not test: what happens when the *same* system is probed at different frequencies?

The answer is that the cognitive attribution changes. The Gray-Scott system habituates at short intervals, sensitises at moderate intervals, and shows no learning at long intervals — despite being the same PDE throughout. This frequency-dependent window is not an artifact of the PDE; it is a structural feature of any system with a characteristic relaxation timescale probed by an observer with a chosen stimulus interval. The *Stentor* experiments necessarily operate within a specific frequency window (45 seconds, matching the cell’s contraction-extension recovery), and the observed habituation necessarily reflects the relationship between that choice and the cell’s dynamics.

This does not diminish the *Stentor* result. The experimentalists made a well-motivated protocol choice and observed a robust, reproducible emission pattern. What the frequency sweep adds is the recognition that the same cell, probed at a different frequency, might exhibit sensitisation rather than habituation — and that this would not reflect a different cognitive capacity in the cell but a different relationship between the observer’s protocol and the cell’s relaxation dynamics.

The broader implication is that protocol-based assessments of cognitive capacity in any system — from single cells to neural networks to AI systems — are co-determined by the system’s dynamics and the observer’s design choices. The question “does *Stentor* habituate?” has the same logical structure as the question “does the Gray-Scott system habituate?” — both require specifying a probing protocol, and both answers depend on that specification. Habituation is a property of the system-protocol pair, not of the system alone.

5.5 What This Framework Is Not

Not eliminativism. The M/L/D patterns are real — as real as temperature, which is also a statistical pattern rather than a molecular property. The claim is that these patterns are relational (between observer and system) rather than intrinsic (belonging to the system alone).

Not pan-cognitivism. We do not claim that sorting algorithms think or that PDEs are conscious. We claim that the *same mathematical structure* in the observer's predictive model produces cognitive attributions across substrates. Whether subjective experience accompanies this structure is a separate question the framework does not address.

Not teleological. Systems do not "optimise" or "seek." Configurations that produce predictable emissions are configurations where the observer's model performs well. The language of optimisation is the observer's shorthand for a selection process that requires no intention.

Not a metaphor. The numbers in Section 4 are quantitative. The prediction that L should increase with restricted observer access is tested and partially confirmed. The claim that a PDE habituates is demonstrated with specific emission profiles. These are computational results, not analogies.

5.6 Proposed Experimental Validation

The computational demonstrations in this paper should be followed by laboratory experiments. We propose a multi-observer experiment on the Belousov-Zhabotinsky chemical oscillator:

The same BZ system would be observed simultaneously by four measurement frameworks: (A) chemical concentrations ($[Br^-]$, $[Ce^{4+}]$, $[H^+]$), D_obs = 3; (B) physical quantities (spatial gradients, wave velocity, pattern coherence), D_obs = 3; (C) electrochemical measurements, D_obs = 4; (D) complete multi-wavelength spectroscopy, D_obs ≈ 8. For each observer, M, L, and D would be computed from the emission stream.

The framework predicts: M and L should decrease monotonically with increasing D_obs, with memory scaling approximately as $M \propto \log(D\_total/D\_obs)$. Observer D (most complete access) should perceive the least "intelligence" in the same physical system. This prediction is specific, quantitative, and distinguishable from any framework that treats cognitive properties as intrinsic to the system.

---

6. Conclusion

Memory is not stored; it is inferred from model obsolescence. Learning is not encoded; it is inferred from decreasing surprise. Decision is not computed; it is inferred from surprise spikes. These three inferences arise from a single predictive model applied to observable emissions, and they arise identically across four computational substrates that share no physics, no mathematics, and no structural features.

A reaction-diffusion system habituates. A sorting algorithm learns. The observer who sees less infers more intelligence. These are not paradoxes — they are consequences of the thesis that cognitive categories are properties of the observation relation, not of the observed system.

The framework is falsifiable. It predicts substrate independence (confirmed across four substrates), observer scaling (partially confirmed in the multi-observer test), and cognitive attribution without cognition (confirmed in sorting algorithms and PDEs). The proposed BZ oscillator experiment would test the framework in a laboratory setting, bridging from computation to chemistry.

The companion paper (Neale 2026a) develops the philosophical implications: why the observer's three categories are co-present rather than sequential, why goals and intelligence are derived narratives rather than independent properties, and why the construction applies recursively — the observer constructing M/L/D from a system's emissions is itself a system whose emissions another observer would decompose the same way.

Appendices

Appendix A: Emission Feature Specifications

[Complete table of emission features per substrate, computation methods, and parameter choices]

Appendix B: M/L/D Algorithm Pseudocode

[Formal specification of the rolling prediction model and all three measures]

Appendix C: Full Results Tables

[Complete M/L/D values for all 287 systems]

Appendix D: Free Energy Correspondence

[Formal sketch showing that M, L, D map to components of the observer's variational free energy]

References

Barandes, J.A. (2023). The stochastic-quantum correspondence. *arXiv:2302.10778*.

Eckert, L., Vidal-Saez, M.S., Zhao, Z., Garcia-Ojalvo, J., Martinez-Corral, R., & Gunawardena, J. (2024). Biochemically plausible models of habituation for single-cell learning. *Current Biology* 34, 5646–5658.

Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience* 11(2), 127–138.

Gershman, S.J. et al. (2026). Associative learning in the protozoan *Stentor coeruleus*. *bioRxiv* 2026.02.25.708045.

Hopfield, J.J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences* 79(8), 2554–2558.

Jaynes, E.T. (1957). Information theory and statistical mechanics. *Physical Review* 106(4), 620–630.

Neale, D. (2025). Being from Nothingness: Deriving the Structure of Existence from a Single Axiom. Preprint, goleudy.ai.

Neale, D. (2026a). Why Observers Construct Intelligence: Cognitive Categories as Prediction Failure. Preprint, goleudy.ai.

Thompson, R.F. & Spencer, W.A. (1966). Habituation: A model phenomenon for the study of neuronal substrates of behavior. *Psychological Review* 73(1), 16–43.

Zhang, T., Goldstein, A., & Levin, M. (2024). Classical sorting algorithms as a model of morphogenesis: Self-sorting arrays reveal unexpected competencies in a minimal model of basal intelligence. *Adaptive Behavior* 32(6), doi: 10.1177/10597123241269740.

Goleudy.ai research programme. Code and data available at goleudy.ai. All simulations reproducible from the published scripts.