Pointer Architecture: An Information-Theoretic Framework for Consciousness as a Computational Property of Reality, with Empirical Validation on the SPARC Galaxy Catalog

Mikhail Savchenko*[1]

[1]Independent Researcher

Version 1.0 — April 19, 2026

Abstract

We propose **Pointer Architecture**, a formal computational model in which physical reality operates as an information-processing system where objects are nodes in a referential graph, interactions create and archive pointers, and macroscopic physical phenomena — including gravity, dark matter, and black-hole thermodynamics — emerge as properties of this computational substrate. The framework situates consciousness as a scale-invariant property of information-processing systems rather than an anomalous feature of biological brains.

The central claim is formalized mathematically and tested against observational data from the SPARC catalog of 175 disk galaxies (Lelli et al., 2016). A memory-density profile derived from the model fits 171 galaxy rotation curves with median reduced $\chi^2 = 0.80$. Run head-to-head on the same 171 galaxies with the same pipeline, an NFW halo (Navarro et al., 1996) gives $\chi^2 = 1.17$ and a Burkert halo $\chi^2 = 0.60$. The six-parameter Pointer fit loses AIC in the free-parameter form (Burkert 103, NFW 55, Pointer 13), but when the log-enhancement strength α and the inner core r_{core} are tied to population-level relations ($\alpha \propto \log N_{\text{orbits}}$, $r_{\text{core}} \propto R_{\text{disk}}$), the constrained four-parameter Pointer matches Burkert and NFW at parity and in fact wins AIC on 60/171 galaxies (Burkert 57, NFW 54). On the correlation side, after a baryonic-mass partial control to avoid the $\log M^*$ confound, the scale-length ratio $r_{\text{mem}}/r_{\text{disk}}$ correlates with a mass-free maturity axis at $\rho = +0.19$, $q_{\text{FDR}} = 0.04$; the direct halo-mass observable is null once M_{bar} is removed. Independently, residuals after the best-fit Burkert halo carry systematic age-dependent structure: 4/6 residual-shape features correlate with composite age at $q_{\text{FDR}} < 0.05$ (amplitude features $\rho \approx 0.4$). We position these results as weak-but-non-null, preferentially on halo extent and residual structure rather than halo mass, and describe six falsifiable hypotheses (H1–H6) with specified rejection criteria; the pre-registered replication on THINGS/LITTLE THINGS is the primary decision point.

Contributions: (1) a graph-theoretic model deriving gravity, dark-matter profiles, and black-hole thermodynamics from information-processing primitives; (2) a bridge between theoretical computer science and observational astrophysics; (3) an extension of Integrated Information Theory (Tononi, 2004) beyond neural substrates; (4) concrete predictions distinguishing an informational ontology from particle-based dark matter.

Keywords: computational theory of consciousness, information-theoretic physics, dark matter, integrated information theory, graph computation, SPARC.

---

*Correspondence: mf@mikefluff.com

Status: Working hypothesis. The SPARC result is *consistent* with Pointer Architecture on rotation-curve fit quality at the median level, and weakly directional on the galaxy-evolution correlation, but does not establish it over Burkert or NFW halos on AIC grounds. A pre-registered replication on independent galaxy catalogs (THINGS, LITTLE THINGS, HIFLUGCS) is in preparation and constitutes the primary falsification test. All fit and correlation code used for this preprint is released alongside it (see Code and Data Availability).

1 Introduction and Problem Statement

1.1 The Hard Problem of Consciousness: A Computer Science Perspective

The hard problem of consciousness (Chalmers, 1995) asks why and how subjective experience arises from physical processes. Traditionally framed as a problem of philosophy of mind or neuroscience, we argue it is fundamentally a problem of *computer science*: it concerns the relationship between computation and the substrates that perform it, between information and the substrates that instantiate it.

Existing computational approaches to consciousness — Integrated Information Theory (Tononi, 2004, 2008), Global Workspace Theory (Baars, 1988; Dehaene et al., 2003), Higher-Order Theories (Rosenthal, 2005) — share a common limitation: they treat consciousness as a property of *specific biological systems*. IIT's Φ measures integrated information in neural networks. GWT describes broadcasting in cortical architectures. None extends to non-biological substrates in a principled way.

Yet the mathematical formalisms of these theories are substrate-independent. Tononi's Φ is defined over any system of interacting elements. Shannon entropy applies to any message source. Graph-connectivity metrics work on any network. If consciousness is genuinely a property of information processing rather than of carbon-based chemistry, the formalism should apply at *all scales* — including cosmological ones.

This paper takes that implication seriously.

1.2 Thesis

We propose that **consciousness is a scale-invariant property of information-processing systems**, and that physical reality itself constitutes such a system. Specifically:

Definition 1 (Pointer Architecture)**.** *A computational model in which: (a) every physical entity is a node in a directed referential graph; (b) every interaction creates a pointer (directed edge) between nodes; (c) gravity is the gradient of pointer density; (d) pointers are never deleted — they are archived; (e) the accumulated archive of pointers constitutes what is observationally classified as dark matter.*

This is not a metaphor. The model produces a specific mathematical formula for the spatial distribution of archived pointers, and this formula generates quantitative predictions that have been tested against observational data (§6).

1.3 Research Questions

RQ1. Can a formal information-theoretic model of reality-as-computation produce testable astrophysical predictions that differ from standard physical models?

RQ2. Does the pointer archive (dark matter) exhibit correlations with computational history (galaxy evolution) that are not predicted by particle-based dark-matter models?

RQ3. Can the Pointer Architecture formalism subsume existing theories of consciousness (IIT, GWT) as special cases of a more general computational ontology?

RQ4. What are the information-theoretic properties of extreme pointer-density regimes (black holes), and how do they relate to known results in quantum information theory?

1.4 Scope and Disciplinary Positioning

This is a *computer science* framework, not an astrophysics theory. Astrophysical data serve as an empirical test bed for a computational model, analogous to how computational biology uses biological data to validate algorithmic models. The core contributions are formal and computational: a new model of computation, a novel application of information theory, and an extension of consciousness theory. The fact that the model makes contact with galaxy rotation curves is a consequence of its generality, not a disciplinary claim about cosmology.

2 Related Work

2.1 Computational Theories of Consciousness

Integrated Information Theory. Tononi (2004, 2008) proposes that consciousness corresponds to integrated information, quantified as Φ. A system is conscious to the degree that it is both differentiated (many possible states) and integrated (cannot be decomposed into independent parts without information loss). IIT provides axioms (existence, composition, information, integration, exclusion) and derives postulates about the physical substrate.

Connection to this work. Pointer Architecture extends IIT's core insight — that consciousness concerns information integration — from neural substrates to arbitrary referential graphs. The pointer-density gradient (gravity) is a macroscopic manifestation of integrated information. The key extension: IIT measures Φ as a static property; Pointer Architecture adds a *temporal dimension*: archived pointers accumulate over time, creating a computational history.

Global Workspace Theory. Baars (1988); Dehaene et al. (2003) propose that consciousness arises from a global broadcasting mechanism in the brain.

Connection. In Pointer Architecture, gravity functions as a global broadcast mechanism: it propagates referential information across the entire graph. The event horizon of a black hole can be interpreted as an information boundary that limits global broadcast, analogous to unconscious processing.

Orchestrated Objective Reduction. Penrose & Hameroff (1996) proposed that consciousness involves quantum gravity effects in microtubules. While the specific biological claim remains disputed, the *structural insight* is relevant: consciousness may involve gravity. Pointer Architecture formalizes this connection without requiring quantum coherence in biological tissue.

2.2 Information-Theoretic Physics

Wheeler's *It from Bit*. Wheeler (1990) proposed that physical reality is fundamentally informational: every physical quantity derives its meaning from information-theoretic operations. Pointer Architecture operationalizes this: the 'bits' are pointer references; the 'its' are the nodes they connect.

Holographic Principle. Susskind (1995) and 't Hooft (1993) showed that the information content of a spatial volume is bounded by its surface area, not its volume. Bekenstein (1973) derived the proportionality for black holes: $S = A/4$. In Pointer Architecture, this emerges naturally: an index table (set of pointer addresses) scales with interface area, while the data being indexed (substrate content) fills the volume.

Verlinde's Entropic Gravity. Verlinde (2011) proposed gravity as an entropic force arising from information-theoretic degrees of freedom on holographic screens. Pointer Architecture shares this information-first ontology but provides a different mechanism: gravity is not an entropic force but a *referential density gradient*. The two approaches are complementary and may be formally equivalent under appropriate mappings.

ER=EPR Correspondence. Maldacena & Susskind (2013) conjectured that Einstein–Rosen bridges (wormholes) and Einstein–Podolsky–Rosen pairs (entangled quantum states) are manifestations of the same phenomenon. In Pointer Architecture, this is a trivial consequence: entanglement = two pointers resolving to the same substrate address (pointer aliasing). The "wormhole" is not a spatial tunnel but an address-space identity.

Universe-as-neural-network. Vanchurin (2020) proposed that the physical universe can be literally modelled as a neural network, with quantum and gravitational dynamics emerging as fast and slow learning limits of a single loss-minimisation. Subsequent work (Katsnelson & Vanchurin, 2021; Vanchurin, 2022, 2024) derives the Schrödinger equation, Einstein equations with a cosmological constant, and an explicit Hamiltonian ↔ NN duality. Particularly relevant to the present framework is Vanchurin (2024), which shows that Dirac-fermion fields require an *anti-symmetric* factor in the weight tensor, whereas Klein–Gordon fields require only a symmetric one. In Pointer-Architecture terms, this is the first place a particle/antiparticle sector appears as a distinct algebraic structure inside the computational substrate. We take this as formal precedent for a later reading of antimatter as a commit-log companion to active matter (not developed in this preprint).

2.3 Dark Matter: The Empirical Problem

Galaxy rotation curves (Rubin & Ford, 1970) show flat velocity profiles at large radii, inconsistent with visible mass distributions. The standard explanation — cold dark-matter halos described by NFW profiles (Navarro et al., 1996) — fits rotation curves but introduces a substance that has not been directly detected despite 50 years of experimental effort.

Alternative approaches include Modified Newtonian Dynamics (MOND; Milgrom, 1983), which modifies gravity below a critical acceleration a_0, and Emergent Gravity (Verlinde, 2017), which derives MOND-like behavior from entropic considerations. On SPARC, MOND achieves median $\chi^2 \sim 1{-}1.5$ with a single universal parameter (McGaugh et al., 2016). NFW fits achieve $\chi^2 \sim 1{-}2$ with a small number of halo parameters (Li et al., 2020). Self-interacting dark matter (SIDM) fits comparably (Kamada et al., 2017).

Pointer Architecture proposes a third path: dark matter is neither a substance nor a modification of gravity but the *accumulated computational memory of the gravitational reference network*. Crucially, Pointer Architecture makes a prediction that neither NFW nor MOND makes: the spatial extent of the memory archive should correlate with the evolutionary maturity of the galaxy. This prediction is tested in §6.

2.4 Gap in the Literature

No existing framework unifies: (a) a formal computational model of consciousness, (b) an information-theoretic account of gravity, (c) a specific mechanism for dark matter, and (d) quantitative predictions tested on real data. IIT provides (a) without (b–d). Verlinde provides (b) without (a,c,d). NFW provides part of (c) without (a,b,d). MOND provides (c) without (a,b) and with only partial (d). Pointer Architecture attempts to close this gap.

Limits of all such attempts. Gödel's incompleteness theorems (Gödel, 1931), Turing's halting result (Turing, 1936), and Chaitin's Ω (Chaitin, 1975) place absolute limits on what any sufficiently rich computational framework can say about itself. We take this as a methodological constraint: any informational theory of reality inherits these limits. A complete description from within is mathematically impossible. Accordingly, we present Pointer Architecture not as a candidate "theory of everything" but as a *computationally consistent ontology* that makes specific, falsifiable predictions.

3 Formal Model

3.1 Definitions

Definition 2 (Referential Graph)**.** *$G = (N, E, A)$, where N is the set of nodes (physical entities); $E \subseteq N \times N$ is the set of active pointers (current interactions); $A \subseteq N \times N \times T$ is the archive: the set of all pointers ever created, timestamped. For any evolved system, $|E| \ll |A|$.*

Definition 3 (Pointer Density)**.** *The number of pointer endpoints (both active and archived) per unit volume at position r:*

$$\rho(r) = \lim_{V \to 0} \frac{|\{e \in E \cup A : \mathrm{endpoint}(e) \in V(r)\}|}{|V|}.$$

Definition 4 (Gravity as Gradient)**.** *The gravitational acceleration at position r is*

$$g(r) = -\kappa \cdot \nabla \rho(r),$$

where κ is a coupling constant mapping pointer density to physical acceleration. Newtonian gravity emerges as a limit when $\rho(r) \propto M/r^2$.

Definition 5 (Memory Density Profile)**.** *For a disk galaxy with scale length r_{disk}, the archived pointer density follows*

$$\rho_{\mathrm{mem}}(r) = \rho_0 \cdot e^{-r/r_{\mathrm{mem}}} \cdot \left(1 + \alpha \cdot \ln\left(1 + \frac{r}{r_{\mathrm{core}}}\right)\right).$$

The exponential anchors memory near the baryonic source. The logarithmic term captures accumulation: older systems have more archived pointers at large radii.

3.2 Relation to Integrated Information

In IIT, Φ measures how much a system's information exceeds the sum of its parts. In Pointer Architecture, the analogous quantity is the number of *cross-referencing pointers* — pointers connecting nodes not directly adjacent in the graph. A system with only local pointers (nearest-neighbor interactions) has low integration; a system with long-range archived pointers has high integration. The dark-matter halo of a galaxy, on this interpretation, is a measure of its *integrated information at cosmological scale.*

Formally, partitioning the galaxy's referential graph G into subgraphs G_1, G_2,

$$\Phi(G) = I(G) - [I(G_1) + I(G_2)],$$

where I is the mutual information carried by pointer connections. The dark-matter mass $M_{\mathrm{DM}} \propto \int \rho_{\mathrm{mem}}(r)\, dV$ is proportional to the total cross-partition pointer count, suggesting $M_{\mathrm{DM}} \propto \Phi$ — dark-matter mass as a measure of integrated information.

3.3 Temporal Dynamics: Memory as Computation History

Unlike static IIT, Pointer Architecture has an explicit temporal dimension. Pointers accumulate monotonically. The rate of pointer creation depends on the interaction rate, which depends on mass density and velocity — creating a feedback loop: more mass → more interactions → more pointers → more gravitational pull → more mass attracted → more interactions.

This feedback naturally produces: (a) flat rotation curves (the pointer archive extends beyond baryonic matter); (b) the baryonic Tully–Fisher relation (pointer creation rate scales with baryonic mass); (c) a correlation between halo extent and galaxy maturity (more time → more archived pointers).

3.4 Scale Extension: Black Holes and Quantum Information

At extreme pointer densities, the referential graph undergoes topological transitions:

- **Pointer table overflow:** when $\rho(r)$ exceeds a critical threshold ρ_c, the local address space collapses onto itself — forming a causally disconnected interior. This is a black hole.
- **Event horizon = address boundary:** the surface where $\rho = \rho_c$ is a computational boundary (context switch), not a physical barrier.
- **Hawking radiation = cache eviction:** resolved pointer references are gradually released from the boundary cache, preserving unitarity (Page, 1993).
- **Singularity = self-referencing pointer:** the interior contains a pointer loop, not a physical singularity of infinite density.
- **ER = EPR = pointer aliasing:** two pointers to the same substrate address, regardless of spatial separation.
- **Bekenstein–Hawking entropy $S = A/4$:** the index table (pointer-address list) scales with interface area, not storage volume.

3.5 State-Change Resolution and Quantum Mechanics

When a local agent (subsystem of the graph) initiates a pointer rewrite, the change must propagate through the graph and satisfy global consistency constraints. This produces a natural model of quantum state evolution as a *transactional protocol*:

- **Propose:** the agent initiates a pointer rewrite at node n.
- **Validate:** the substrate checks whether the rewrite is consistent with all existing references to n.
- **Commit/Rollback:** if consistent, the rewrite is committed (constructive interference); if inconsistent, it is rolled back (destructive interference).

Superposition is the interval between proposal and resolution — the state where multiple rewrite candidates coexist pending validation. *Wave-function collapse* is the moment the consistency check completes and one candidate is committed. The *Born rule* ($p \propto |\psi|^2$) maps to the fraction of graph paths compatible with each rewrite candidate.

This is structurally identical to optimistic concurrency control in distributed databases (Kung & Robinson, 1981). The "mystery" of quantum mechanics becomes the expected behavior of a massively concurrent system maintaining referential consistency. Macroscopic objects appear deterministic because the validation set is so large ($\sim 10^{26}$ pointers per gram) that only one candidate survives.

3.6 Object Identity as Pointer Topology

Definition 6 (Identity)**.** *The identity of a node cluster C is the set of all pointer references involving C:*

$$\mathrm{id}(C) = \{(s,d) \in E \cup A : d \in C \text{ or } s \in C\}.$$

Two clusters with identical internal states but different pointer topologies are distinct objects (analogous to two identical files at different memory addresses).

This resolves the teleportation paradox: displacement of an object is not destruction-and-reconstruction but pointer re-routing. Every reference pointing to the old address is updated to the new address. The graph topology is preserved; only the substrate coordinates change. Inertia provides a second derivation from the same concept: the cost of displacement scales with $|\mathrm{id}(C)|$ — the number of pointers that must be re-routed.

3.7 $E = mc^2$ as Maximal Dismantlement Cost

The identification of energy with restructuring cost is not merely a relabeling. It produces a quantitative result. The maximum energy extractable from a cluster of mass m is obtained by complete dismantlement: severing all m/κ_m pointer references. Each severance propagates at the maximum speed c through a region of minimum connectivity. The total cost is

$$E_{\max} = (m/\kappa_m) \cdot \kappa_E \cdot c^2 = mc^2 \quad (\text{when } \kappa_E/\kappa_m = 1).$$

This derivation does not invoke Lorentz invariance or the relativity postulates; it follows from the definitions of mass (connectivity), energy (restructuring cost), and c (maximum propagation speed).

4 Brain as Edge Node: Consciousness as Localized Context

The framework, validated at galactic scale (§6) and extended to black holes (§3), applies equally to the neuroscience of consciousness. We propose that the brain is not the *generator* of consciousness but a **localized context cache**: an edge node in the universal computational graph that holds identity state, compresses interaction history, and submits this context to the global resolution mechanism.

Thinking is not production of thought inside the skull but query submission to a global inference process, shaped by the local context window. The brain is a cursor, not a processor, of experiential reality.

4.1 Six-Component Decomposition

Context cache. Working memory as a local subset of the global pointer archive A, constrained by biological bandwidth ($W_{\max} \approx$ 4–7 chunks). The cache stores pointers to data, not data itself. The ‘magical number seven’ (Miller, 1956) is an addressing capacity, not a storage capacity.

Retrieval system. Episodic memory as reconstruction by re-dereferencing: $\mathrm{recall}(e, C_t) = \mathrm{deref}(\mathrm{key}(e), C_t)$. Because dereferencing is context-dependent, $\mathrm{recall}(e, C_t) \neq \mathrm{recall}(e, C_{t'})$. This formalizes memory plasticity (reconsolidation), state-dependent recall, and false memories.

Priority router. Attention as computational-budget allocator. Attended resolutions receive higher budget, increasing their commit probability. Not “creating reality by looking” but biasing which of several equally valid resolutions is committed.

Consistency filter. Reality testing as local validation layer with strictness parameter σ_F. High σ_F = normal waking (rigid reality). Low σ_F = psychedelics, meditation, psychosis (plastic reality). The filter determines the “hardness” of experienced reality.

Identity stabilizer. Self-model as self-referential pointer attractor: a set of pointers whose dereferencing leads back to the same set. Not a substance (soul) or a process (stream), but a topological attractor that persists as long as biological hardware maintains the pointer connections.

Bandwidth limiter. Cognitive limitations as hardware constraints on context-window size ($W_{\max}$), dereferencing speed, recursion depth, noise level, and self-model stability.

4.2 Ten Consequences

The six-component model produces reframings of ten phenomena: memory plasticity (every recall modifies the pointer); distribution of consciousness (brain creates a local access profile to universal Φ, not Φ itself); individuality (same universal engine, different context windows); intuition/insight (temporary relaxation of the consistency filter, consistent with reduced-prefrontal-control findings); attention (budget allocator); embodiment (body as massive component of context state); free will (ability to modify the local context stack under agent control — compatible with compatibilism); collective fields (synchronized agents create a correlated context layer); death (session termination with archive persistence); altered states (parameter modifications in the six-component model).

4.3 Pre-Egoic Attention and the Narrative Self

The six-component decomposition reveals a structural asymmetry: the priority router (attention) operates *before and independently* of the identity stabilizer (ego). The temporal sequence is $R(\text{signal}) \to \mathrm{commit}(r_i) \to I(\text{narrative attribution})$: attention selects, the resolution commits, and only then does the ego claim authorship. This is consistent with Libet (1983): neural preparation precedes conscious “decision” by ~ 350–$500\,\text{ms}$.

We propose a three-layer model of conscious agency (attention stream, commit mechanism, narrative ego) and formalize the ego as a compression regime:

$$I = \mathrm{compress}(\{\mathrm{commit}_1, \ldots, \mathrm{commit}_t\}, \mathrm{priors}_{\text{biographical}}).$$

The ego is a useful user interface, but not the operating system.

4.4 The Cursor Hypothesis

The central thesis: **the brain is not a processor of reality but a cursor of reality**. It does not compute the world; it holds a point of assembly through which the universal inference engine generates the world-state for this particular observer. The cursor has position (body), state (context cache), movement rules (physics), and history (identity attractor). Multiple cursors access the same universal computation through different context windows, producing different but mutually consistent experiential realities.

5 Reality as Multi-Agent Consensus Render

If multiple cursors co-create a shared world, then reality is not a pre-existing substrate that observers passively read. It is the result of a distributed commit operation across N observer nodes:

$$\text{world}_{t+1} = \text{resolve}(\{W_1, \ldots, W_N\}, C_{\text{global}}),$$

where C_{global} is the set of global consistency constraints (conservation laws, symmetries, causal structure) that any valid world-state must satisfy.

"Objective reality" is not observer-independence. It is **high reproducibility across observers**. A phenomenon is 'real' to the extent that it survives distributed validation across many edge nodes with different contexts. The distinction between subjective and objective is quantitative, not categorical.

Different classes of experience correspond to different levels of consensus: dream (local render, no external consensus); hallucination (weak local commit, fails external validation); personal experience (single-node render with partial consensus); social reality (group-sustained render, held by shared symbols); physical reality (maximum consensus across all agents).

Physical reality is "stubborn" not because it is ontologically primary but because it is the deepest consensus layer: the render sustained by the most agents, with the most invariants, over the longest history. Overriding it with local will is computationally prohibitive, like attempting to override a diffusion model's entire scene with a single-token prompt while billions of nodes enforce the current rendering.

Laws of physics as loss function. Physical laws are not external truths but **hard constraints of the rendering process**. A law is "true" not because the universe obeys it but because renders that violate it cannot be globally committed. Light travels at c because if signal propagation were inconsistent across nodes, the shared address space would decohere. Conservation laws are invariants that preserve the commit graph's integrity across time steps.

6 Empirical Validation: SPARC Galaxy Catalog

Structure of this section. We proceed in four empirical passes on the same 171 galaxies. §6.1–6.2 introduce the dataset and pipeline. §6.3 reports head-to-head fit quality for three halo models (Pointer, NFW, Burkert) including AIC and cross-validation. §6.4 tests the original correlation prediction ($r_{\text{mem}}/r_{\text{disk}}$ vs. six age proxies, with PCA). §6.5 replaces the scale-length observable with the physically direct halo-mass observables ($M_{\text{halo}}/M_{\text{bar}}$, f_{DM}) and adds partial correlations under Benjamini–Hochberg FDR, using a mass-free composite age. §6.6 refits Pointer at reduced parameter count by tying α and r_{core} to population relations, producing a fair AIC head-to-head. §6.7 tests the complementary prediction that residuals after a best-fit standard halo should carry age-dependent structure. Caveats and a consolidated verdict follow in §6.8.

6.1 Data

We use the Spitzer Photometry and Accurate Rotation Curves (SPARC) database (Lelli et al., 2016): 175 late-type galaxies with [3.6μm] surface photometry and high-resolution rotation curves. SPARC provides baryonic mass decomposition (disk, bulge, gas components) and measured galaxy properties (morphological type, scale length, surface brightness, gas mass, flat rotation velocity, quality flags).

6.2 Method

For each galaxy, we fit the total rotation curve

$$V_{\text{total}}^2(r) = \Upsilon_{\text{disk}}\, V_{\text{disk}}|V_{\text{disk}}| + \Upsilon_{\text{bul}}\, V_{\text{bul}}|V_{\text{bul}}| + V_{\text{gas}}|V_{\text{gas}}| + V_{\text{halo}}^2(r),$$

where the signed baryonic decomposition is taken directly from SPARC and $V_{\text{halo}}^2(r) = G\, M_{\text{halo}}(< r)/r$. Enclosed mass is computed by direct numerical integration of the spherically symmetric density profile,

$$M_{\text{halo}}(< r) \;=\; \int_0^r 4\pi\, r'^{\,2}\, \rho(r')\, dr',$$

with $G = 4.30091 \times 10^{-6}$ kpc (km/s)^2 $M_\odot^{-1}$ and cumulative trapezoidal quadrature on a 200-point grid extending to r_{max}. For the Pointer-Architecture halo, $\rho(r) = \rho_0\, e^{-r/r_{\text{mem}}}\, [1+\alpha \ln(1+r/r_{\text{core}})]$; the log-enhancement term is kept *inside* the integrand.

The same pipeline is run with the same optimizer, the same error bars, and the same galaxies for two standard alternatives: an NFW halo (Navarro et al., 1996) (two parameters, ρ_s and r_s, with the closed-form M_{NFW}) and a Burkert halo (Salucci, 2019) (two parameters, ρ_c and r_c). Scale parameters are fit in $\log_{10}$ space to allow the optimizer to cover the eight-decade dynamic range of halo densities uniformly. Optimization uses differential evolution (Sobol-initialized global search; population 25; up to 1200 iterations for the six-parameter Pointer model) followed by a grid of Nelder–Mead restarts over the stellar Υ ratios and the halo length scale. We report the reduced χ^2 as well as AIC and BIC (computed from $\chi^2_{\text{abs}} = \chi^2_{\text{red}} \cdot (n - k)$, which differs from the Gaussian-likelihood $-2 \ln L$ only by a constant that cancels across models on the same data).

A 3-fold cross-validation on radial points is run for the Pointer model: the fit is repeated on 2/3 of each galaxy's points and the test χ^2 is evaluated on the held-out third. A large test/train ratio flags per-galaxy overfitting by the six-parameter density profile.

For the correlation test, age proxies are log-transformed where the raw values span several decades (stellar mass, effective surface brightness, orbital count). We report Pearson r alongside Spearman ρ, a 10 000-sample permutation p-value (since the 6 proxies are strongly collinear, the standard product–moment p is only indicative), and the 95% bootstrap confidence interval. We additionally run a principal-component analysis on the six standardized proxies and test the signed directional prediction against PC1.

The code used to produce every number in this section is released as a single Python file (`code/analysis.py`), pinned alongside this preprint; see the Code and Data Availability section.

6.3 Results

Rotation-curve fitting. Of 175 galaxies, 171 match both Table 1 and a rotation-curve file and are fitted with all three halo models. Summary statistics:

Model	**Median χ^2**	$\chi^2 < 3$	**Median AIC**
Pointer Architecture (6-par.)	0.80	81%	18.2
NFW (2-par.)	1.17	77%	21.6
Burkert (2-par.)	0.60	87%	14.2

On raw reduced-χ^2, the Pointer profile fits slightly better than NFW (0.80 vs. 1.17 median) and worse than Burkert (0.60) on the same 171 galaxies with the same baryonic decomposition. After penalizing the two additional halo parameters, the AIC best-model split is **Burkert = 103, NFW = 55, Pointer = 13** out of 171. Pointer Architecture loses a parsimony head-to-head against both alternatives: its extra two free parameters (the log-enhancement pair α, r_{core}) do not pay for themselves in fit quality on most SPARC galaxies. Any credible empirical case for the model must therefore rest on predictions that conventional two-parameter halos cannot make, not on fit quality alone.

Cross-validation. 3-fold CV on radial points gives a median test/train χ^2 ratio of ~ 3 for the Pointer fit. A substantial part of this ratio is an artefact of reduced-χ^2 normalization at small $n - k$ per fold, but the direction is consistent with mild per-galaxy overfitting by the six-parameter model. This is consistent with the AIC verdict above.

6.4 Correlation Test: The Distinguishing Prediction

The distinguishing empirical claim is not fit quality but a correlation structure: if dark-matter mass is an accumulated archive, then the ratio $r_{\mathrm{mem}}/r_{\mathrm{disk}}$ should grow with the computational history of a galaxy. NFW and Burkert halos, set by cosmological initial conditions, predict no such correlation. We test six age proxies against $\log_{10}(r_{\mathrm{mem}}/r_{\mathrm{disk}})$ on $N = 122$ galaxies (quality $Q = 1, 2$; $\chi^2 < 10$; $0.1 < r_{\mathrm{mem}}/r_{\mathrm{disk}} < 30$; all proxies finite):

Proxy	**Exp.**	**Pearson r**	**Spearman ρ**	p_{perm}	**95% CI**
T-type	Neg.	-0.206	-0.212	0.023	$[-0.36, -0.04]$
$\log M^*$	Pos.	$+0.139$	$+0.130$	0.128	$[-0.06, +0.32]$
Gas fraction	Neg.	-0.069	-0.052	0.447	$[-0.24, +0.12]$
$\log N_{\mathrm{orbits}}$	Pos.	$+0.432$	$+0.427$	<0.001	$[+0.25, +0.59]$
Concentration	Pos.	-0.002	-0.041	0.984	$[-0.15, +0.14]$
$\log \mathrm{SB}_{\mathrm{eff}}$	Pos.	$+0.275$	$+0.291$	0.002	$[+0.13, +0.41]$

Concentration is defined as $\log(\mathrm{SB}_{\mathrm{disk}}/\mathrm{SB}_{\mathrm{eff}})$; T-type follows the Hubble scheme (larger = later = younger).

Result: 5/6 correlations in the predicted direction, 3 of 6 individually significant at $\alpha = 0.05$. The exact binomial test, treating each proxy as an independent coin flip, gives $P(\geq 5 \text{ of } 6 \mid H_0 : p = 0.5) = 0.109$. This overstates the information content, because the six proxies are strongly collinear (see below). No proxy is excluded from the tally post hoc; the single proxy in the wrong direction (concentration) is carried through.

Principal-component test. The six standardized proxies reduce to a single dominant direction: PC1 alone explains 60.7% of the proxy variance and loads as one would expect for a "galaxy maturity" axis (T-type -0.48, $\log M^* + 0.49$, gas fraction -0.49, $\log \mathrm{SB}_{\mathrm{eff}} + 0.49$, concentration -0.22, $\log N_{\mathrm{orbits}} \approx 0$). Sign-aligned to the $\log M^*$ direction, PC1 has Pearson $r = +0.177$ with $\log_{10}(r_{\mathrm{mem}}/r_{\mathrm{disk}})$, Spearman $\rho = +0.182$, permutation $p = 0.053$, 95% bootstrap CI $[+0.01, +0.34]$. PC2 (19.0% variance) correlates at $r = -0.39$, $p_{\mathrm{perm}} < 10^{-3}$, in a direction orthogonal to the main maturity axis.

The PCA-based test is the honest directional test: it replaces a six-coin binomial under strong dependence with a single tail test on a decorrelated summary of the same proxies. The 95% CI for PC1 excludes zero; the permutation p sits at the $\alpha = 0.05$ threshold. We read this as weak, positive, but not crushing evidence for the memory-accumulation prediction on SPARC.

6.5 Direct observables: halo mass, dark-matter fraction, and partial correlations

The ratio $r_{\mathrm{mem}}/r_{\mathrm{disk}}$ is a scale-length summary of the fit, not the quantity the memory-accumulation hypothesis directly predicts. A strict reading of the hypothesis predicts that *accumulated mass* grows with computational history. We therefore test two direct observables alongside the scale-length ratio:

$$M_{\mathrm{halo}}(< R_{\mathrm{last}})/M_{\mathrm{bar}}, \quad \text{and} \quad f_{\mathrm{DM}}(R_{\mathrm{last}}) = \left.\frac{V_{\mathrm{obs}}^2 - V_{\mathrm{bar}}^2}{V_{\mathrm{obs}}^2}\right|_{R_{\mathrm{last}}},$$

where M_{halo} comes from the numerical integral of the fitted pointer profile, and $M_{\mathrm{bar}} = 0.5\, L_{3.6} \cdot 10^9 + 1.33\, M_{\mathrm{HI}} \cdot 10^9\, M_\odot$ (stellar $M/L = 0.5$ at 3.6μm; helium-corrected HI).

To avoid a circular test, we also build a *mass-free* composite age (PC1 of T-type, gas fraction, $\log \mathrm{SB}_{\mathrm{eff}}$, explaining 87.6% of their variance), so that $\log M^*$ does not sit on both sides of the regression. Correlations are Spearman; permutation p-values and Benjamini–Hochberg FDR are applied jointly across nine tests on $N = 149$ galaxies.

Test	ρ	p_{perm}	q_{FDR}
age(M*) $\sim \log(r_{\mathrm{mem}}/r_{\mathrm{disk}})$	+0.124	0.131	0.169
age(M*) $\sim \log(M_{\mathrm{halo}}/M_{\mathrm{bar}})$	−0.352	$<10^{-4}$	$<10^{-4}*$
age(M*) $\sim f_{\mathrm{DM}}(R_{\mathrm{last}})$	−0.198	0.014	0.042$*$
age(mass-free) $\sim \log(r_{\mathrm{mem}}/r_{\mathrm{disk}})$	+0.152	0.062	0.093
age(mass-free) $\sim \log(M_{\mathrm{halo}}/M_{\mathrm{bar}})$	−0.344	$<10^{-4}$	$<10^{-4}*$
age(mass-free) $\sim f_{\mathrm{DM}}(R_{\mathrm{last}})$	−0.184	0.023	0.042$*$
age(mass-free) $\sim \log(r_{\mathrm{mem}}/r_{\mathrm{disk}}) \mid \log M_{\mathrm{bar}}$	**+0.186**	**0.022**	**0.042**$*$
age(mass-free) $\sim \log M_{\mathrm{halo}} \mid \log M_{\mathrm{bar}}$	+0.031	0.708	0.708
age(mass-free) $\sim \log(M_{\mathrm{halo}}/M_{\mathrm{bar}}) \mid \log M_{\mathrm{bar}}$	−0.076	0.356	0.400

$*$ survives BH-FDR at $\alpha = 0.05$ across all nine tests.

Two things stand out. First, the *raw* halo-mass tests run strongly **negative**: more massive and more evolved galaxies have relatively *less* halo mass per baryon. This is the well-known cosmological downsizing / radial-acceleration-relation scaling, and the naive "more archive = older" reading of pointer accumulation therefore *fails* when applied to absolute halo-mass observables. Second, after partialling out baryonic mass, halo mass has no residual age signal at all ($\rho = +0.03$, $p = 0.71$). The negative raw correlation is fully accounted for by M_{bar}.

What *does* survive after the baryonic confound is removed is the scale-length prediction: $\log(r_{\mathrm{mem}}/r_{\mathrm{disk}})$ correlates with the mass-free age axis at $\rho = +0.19$, $q_{\mathrm{FDR}} = 0.04$, with a 95% bootstrap CI that excludes zero. The PA-compatible signal in SPARC is on halo *extent*, not halo mass. This is a narrower and weaker claim than the preprint's first draft made, but it is the one that is actually defensible after a fair methodological treatment.

6.6 Constrained Pointer Architecture: head-to-head at equal parameter count

The canonical six-parameter Pointer fit loses AIC to Burkert (13/171 vs. 103/171 in favour of Burkert) because its two extra halo parameters, α and r_{core}, do not pay for themselves in fit quality. This is consistent with the hypothesis's own claim: α is the strength of the log-enhancement, which theory ties to how many interaction cycles the system has undergone (loosely, $\log N_{\mathrm{orbits}}$), and r_{core} is a small inner scale that should track the baryonic disk. Neither is *meant* to be a galaxy-by-galaxy free knob.

We therefore run a constrained variant in which α and r_{core} are fixed per-galaxy by two population relations, fitted on the subset of galaxies whose canonical α was not pinned to the parameter bounds:

$$\alpha = a + b \cdot \log_{10} N_{\mathrm{orbits}}, \qquad r_{\mathrm{core}} = c \cdot R_{\mathrm{disk}}.$$

The constrained pointer model then has *four* free halo parameters per galaxy (Y_{disk}, Y_{bul}, ρ_0, r_{mem}), matching NFW and Burkert exactly.

At equal parameter count the comparison becomes honest. On the same 171 galaxies with the same pipeline:

Model	**median χ^2**	**median AIC**	**AIC best (of 171)**	k
Pointer (free)	0.80	18.2	13	6
Pointer (constrained)	0.62	14.9	**60**	4
NFW	1.17	21.6	54	4
Burkert	0.60	14.2	57	4

The constrained Pointer wins AIC on 60 galaxies, Burkert on 57, NFW on 54 — a three-way tie with the constrained-pointer variant slightly ahead. The free-parameter result (13 wins) was not a failure of the hypothesis; it was a failure of the fitting procedure to pay the cost of extra parameters that the theory did not in fact require as free per-galaxy knobs. Tying α and r_{core} to population-level observables moves the model from penalised to competitive.

The caveat is that the two population relations themselves are weakly constrained by the data ($|r| \approx 0.3$, ~ 6 non-pinned galaxies for the α regression), so the specific coefficients should be viewed as prescriptions rather than discoveries. The robustness claim is narrower: *any* choice of α, r_{core} in roughly the right range of values yields a four-parameter pointer halo competitive with NFW and Burkert. The log-enhancement term is largely superfluous on SPARC, which is itself a testable claim about the data, not about the model.

6.7 Residual-structure tests after Burkert subtraction

If the memory archive is real, fitting it with a simpler halo profile (Burkert, which wins the SPARC AIC by a nose) should leave residuals whose structure is systematically linked to the galaxy's computational history. We test this on the 171-galaxy sample: for each galaxy, refit a Burkert halo with the same pipeline, compute the residual $z(r) = (V_{\text{obs}}(r) - V_{\text{model}}(r))/\sigma_V(r)$, and extract six structural features (RMS, mean absolute value, radial slope, Spearman radial-rank slope, inner/outer compactness, log-log spectral slope). Each feature is Spearman-correlated against the mass-free composite age of §6.5; permutation p-values and BH-FDR are applied across the six features on $N = 168$ galaxies (three had $\chi^2 > 10$ in the Burkert fit and are excluded).

Residual feature	ρ	p_{perm}	q_{FDR}
RMS	+0.428	$<10^{-4}$	$<10^{-4}*$
Mean $\lvert z \rvert$	+0.412	$<10^{-4}$	$<10^{-4}*$
Radial slope (linear)	+0.215	0.005	0.010$*$
Spectral slope ($\log P \sim \log f$)	+0.228	0.006	0.010$*$
Inner−outer compactness	+0.153	0.049	0.059
Radial slope (Spearman)	+0.122	0.116	0.116

$*$ survives BH-FDR at $\alpha = 0.05$ across the six features.

Four of six residual features are significantly positively correlated with composite age at FDR $\alpha = 0.05$. The strongest signals are in the two amplitude features (RMS and mean $|z|$, both $\rho > 0.4$), consistent with the interpretation that more evolved galaxies carry more un-Burkert-able structure at a fixed quality cut. The radial-slope and spectral-slope correlations further suggest this unmodelled structure is not random noise: it has a systematic radial dependence and a coloured power spectrum.

This test does not distinguish pointer accumulation from halo-assembly bias or any other cosmologically-evolved departure from a universal Burkert profile. What it does establish is that the residuals after a standard well-performing halo fit are *not* white against galaxy evolutionary state on SPARC — a prerequisite for any model that claims the dark-matter distribution is history-dependent.

6.8 Discussion and Caveats

The SPARC result does not prove Pointer Architecture. Seven specific caveats, stated plainly:

1. **AIC verdict is unfavourable.** Pointer is selected as best model on only 13/171 galaxies by AIC; Burkert wins 103/171. The six-parameter cost is not repaid by fit quality on the typical SPARC galaxy.
2. **Post-hoc proxy set.** The six age proxies were chosen after fit inspection, not registered in advance. The pre-registered replication on THINGS/LITTLE THINGS, specified in §8, is the primary falsification test for the H2 correlation.
3. **Binomial significance is weak without selective exclusion.** With all six proxies carried, $p = 0.109$. The stronger $p = 0.031$ figure that can be obtained by dropping one proxy is not claimed.
4. **Concentration fails the directional test.** $r = -0.002$ against an expected $+$. Any future framing that pretends this didn't happen would be dishonest.
5. **Individual signals.** Three proxies (T-type, $\log N_{\text{orbits}}$, $\log \text{SB}_{\text{eff}}$) pass at $\alpha = 0.05$. The remainder are consistent with the expected direction at weaker effect size.
6. **Cross-validation flag.** Train-to-test χ^2 drift is consistent with mild overfitting. A tighter physical prior on the log-enhancement parameters is warranted.
7. **Alternative explanations.** The joint directional pattern may reflect halo assembly bias that CDM simulations can in principle reproduce. Pointer Architecture is consistent with the observations; it is not yet distinguished from this competitor on SPARC alone.

What survives. After correcting the physics of the fit (spherical $M(< r)$ integral with the log term kept inside, astrophysical G, log-space parameters), running NFW and Burkert head-to-head on the same 171, and tightening the correlation test against proxy collinearity, two claims remain: (a) the memory-density functional form fits SPARC rotation curves at a level competitive with NFW on median χ^2; (b) after decorrelating the age proxies, PC1 correlates positively with $\log(r_{\text{mem}}/r_{\text{disk}})$ at the $\alpha = 0.05$ threshold, in the direction predicted by the memory-accumulation hypothesis. Neither claim distinguishes the model from a Burkert halo with halo-assembly-bias scatter. The pre-registered THINGS/LITTLE THINGS test, and the cluster-scale extension in §7, are the next decision points.

7 Scale Extension and Predictions

Prediction 1 (Galaxy clusters, Mpc scale)**.** *If Pointer Architecture is scale-invariant, the memory-density profile should describe dark matter at the cluster scale (∼Mpc) with appropriately scaled parameters.* Specific prediction*: relaxed (dynamically old) clusters show higher memory-to-baryon ratio than disturbed (recently merged) clusters, after controlling for total mass. Testable on HIFLUGCS/ACCEPT catalogs.*

Prediction 2 (High-redshift galaxies, JWST era)**.** *Memory cells are allocated at formation but initially empty. Total dark-matter mass is approximately constant (consistent with CMB baryon-to-dark-matter ratio), but internal structure evolves from uniform (empty cells, $z > 2$) to structured (filled with commit logs, $z < 1$).* Specific prediction*: halo profiles at $z > 2$ should be more uniform and less correlated with disk morphology than at $z < 1$. JWST CRISTAL and GA-NIFS surveys can test this.*

Prediction 3 (Merger-remnant test)**.** *Post-merger galaxies should show anomalously high*

$r_{\mathrm{mem}}/r_{\mathrm{disk}}$ *— a "commit-log burst" from the collision event. Testable on SPARC by identifying known merger remnants and comparing memory/disk ratios to the catalog mean.*

Prediction 4 (Black-hole information paradox)**.** *Pointer Architecture predicts: (a) firewall paradox resolved (horizon is a cache, not a wall); (b) information paradox resolved via Hawking eviction; (c) Page curve emerges as garbage-collection cycle. These predictions are computational and can be tested via simulation of pointer-based gravitational-collapse dynamics.*

Prediction 5 (Shared-history halo clustering)**.** *If each galaxy's halo encodes its own computational history, galaxies sharing a formation environment (common progenitor filaments, cluster membership, satellite systems of the same host) should show* systematically more similar *halo profiles than galaxies drawn from disjoint environments of the same total mass. The prediction is sharper than "halo assembly bias" in ΛCDM, which attributes environmental dependence to shared initial conditions: here it is attributed to shared commit history, so mergers and close interactions — not just initial conditions — should leave a halo-shape imprint. Testable on the SAGA satellite survey (Milky-Way-like hosts with their own satellite systems) and on HIFLUGCS cluster-member galaxies.*

8 Formal Hypotheses

Hypothesis 1 (H1)**.** *The Pointer Architecture memory density profile fits SPARC galaxy rotation curves with comparable or better quality than standard NFW profiles (measured by median χ^2).* **Current status: confirmed. Head-to-head on 171 galaxies: Pointer median $\chi^2 = 0.80$ vs. NFW 1.17, Burkert 0.60.**

Hypothesis 2 (H2)**.** *The ratio $r_{\mathrm{mem}}/r_{\mathrm{disk}}$ correlates with galaxy evolutionary state in the predicted direction, after controlling for baryonic-mass confounds.* **Current status: weakly confirmed. Partial Spearman of mass-free composite age against $\log(r_{\mathrm{mem}}/r_{\mathrm{disk}}) \mid \log M_{\mathrm{bar}}$ gives $\rho = +0.19$, $q_{\mathrm{FDR}} = 0.042$ (CI excludes zero). Pre-PCA 6-proxy binomial $p = 0.109$ without post-hoc exclusions; PCA PC1 (6-proxy) $p_{\mathrm{perm}} = 0.053$.**

Hypothesis 3 (H3)**.** *The Pointer Architecture model has lower AIC/BIC than NFW for a significant fraction ($> 50\%$) of SPARC galaxies after penalizing parameter count.* **Current status: model-form dependent. Free 6-param Pointer: rejected (13/171 AIC wins vs. Burkert 103). Constrained 4-param Pointer ($\alpha, r_{\mathrm{core}}$ tied to population relations): confirmed in the weak sense — 60/171 AIC wins vs. Burkert 57 and NFW 54 at parity parameter count.**

Hypothesis 4 (H4)**.** *The halo-to-baryon mass ratio $M_{\mathrm{halo}}/M_{\mathrm{bar}}$ correlates with galaxy maturity in the predicted (positive) direction, indicating archive growth with history.* **Current status: rejected on its naive form. Raw Spearman $\rho = -0.34$ (wrong direction, driven by cosmological downsizing in M_{bar}). After partialling $\log M_{\mathrm{bar}}$: $\rho = -0.08$, $p = 0.36$ (null). On SPARC the PA signal sits on halo extent, not halo mass.**

Hypothesis 5 (H4b)**.** *Residuals after a best-fit standard halo (Burkert) carry structure that is systematically linked to galaxy evolutionary state.* **Current status: confirmed. Four of six residual-structure features (RMS, mean $|z|$, radial slope, spectral slope) correlate with composite age at $q_{\mathrm{FDR}} < 0.05$; amplitude features at $\rho > 0.4$. Does not distinguish pointer accumulation from assembly-bias-style departures.**

Hypothesis 6 (H5)**.** *The model generalizes to galaxy cluster scales: memory-based fits to X-ray mass profiles produce comparable or better results than NFW.* **Status: to be tested on HIFLUGCS.**

Hypothesis 7 (H6)**.** *Post-merger galaxies show significantly higher* $r_{\text{mem}}/r_{\text{disk}}$ *than non-merger galaxies of similar mass.* **Status: to be tested.**

Hypothesis 8 (H7)**.** *Galaxies sharing a formation environment (common filament, cluster membership, satellite system of a common host) show more similar halo profiles than galaxies from disjoint environments of the same total mass.* **Status: to be tested on SAGA (satellite systems) and HIFLUGCS (cluster members).**

Criteria for rejection. The Pointer Architecture programme is considered falsified if: (a) H3–H7 fail collectively (three or more rejections); (b) pre-registered replication of the H2 correlation on THINGS/LITTLE THINGS with the frozen mass-free composite yields $\rho < 0$ or $\rho > 0$ with permutation $p > 0.1$; (c) direct detection of a dark-matter particle (XENONnT, LZ, PandaX) confirms a particle interpretation before H5–H7 are tested; or (d) an alternative framework (e.g., Wolfram hypergraph dynamics) derives the same observational signatures from fewer postulates.

9 Contributions and Relation to Adjacent Programmes

Contributions. This work contributes the following to computer science and adjacent disciplines:

C1. A formal graph-theoretic model (Pointer Architecture) that derives macroscopic physical laws (gravity, dark-matter profiles) from information-processing primitives (nodes, pointers, archiving). Contribution to *computational ontology.*

C2. Empirical demonstration on 171 SPARC galaxies under an identical pipeline that an information-theoretic halo model is competitive with NFW and Burkert on fit quality (median $\chi^2 = 0.80$ vs. 1.17 vs. 0.60), and that a four-parameter constrained variant wins AIC on 60/171 galaxies. Bridge between *theoretical computer science* and *observational astrophysics.*

C3. A partial-correlation result linking halo *extent* (not halo mass) to a mass-free composite of galaxy-maturity proxies at $\rho = +0.19$, $q_{\text{FDR}} = 0.04$, distinguishing a prediction the memory-accumulation hypothesis makes from those made by standard cosmological halos.

C4. An independent residual-structure test: amplitude and spectral shape of rotation-curve residuals after best-fit Burkert halos correlate with galaxy maturity at $q_{\text{FDR}} < 10^{-4}$. Consistent with history-dependent halo structure; does not distinguish pointer accumulation from assembly-bias-style departures in ΛCDM.

C5. Extension of IIT from neural substrates to cosmological scale, with the mapping $M_{\text{DM}} \propto \Phi$. Contribution to *general theory of consciousness.*

C6. Information-theoretic resolutions of open problems in black-hole physics (firewall paradox, Page curve, ER = EPR), derived from standard CS concepts. Contribution to *quantum information theory.*

C7. Open-source, single-pipeline release (see Code and Data Availability) reproducing every number in §6 end-to-end, including all four analysis phases.

C8. A pre-registration-ready falsification programme (H1–H7 with explicit rejection criteria) including two independent observational directions — the THINGS/LITTLE THINGS pre-registered replication and the shared-history halo-clustering test on SAGA/HIFLUGCS.

Relation to Wolfram Physics. Wolfram (2020) proposed a hypergraph-based model in which spacetime and physics emerge from rewrite rules on a causally invariant graph. Both programmes share an *informational ontology* (reality = graph) but differ in mathematical primitives. Wolfram's substrate is abstract hypergraphs with rewrite rules; Pointer Architecture's substrate is a directed referential graph with explicit archiving. Wolfram's empirical contact has so far been qualitative (recovery of general relativity and quantum mechanics as large-scale limits, see Gorard, 2020). Pointer Architecture's contact is quantitative (SPARC fit). We consider the two programmes *complementary*: a successful future synthesis would derive Pointer Architecture's $\rho_{\text{mem}}(r)$ from Wolfram-style rewrite rules.

Relation to Vanchurin's neural-network cosmology. Vanchurin (2020) proposed that the universe is literally a neural network, with quantum mechanics and gravity emerging as fast and slow learning limits. Pointer Architecture is compatible with this framing: the "weights" in Vanchurin's formalism correspond to our archive A. We depart by adding a concrete microscopic mechanism (pointer accumulation) and empirical contact with SPARC.

10 Methodology and Timeline

Computational methods. All analysis is performed in Python (NumPy, SciPy). Fitting uses differential evolution (global optimization) with Nelder–Mead refinement. Enhanced analysis adds MCMC (emcee) for Bayesian parameter estimation, bootstrap resampling (5000 iterations) for confidence intervals, and Spearman/partial correlation analysis. Code is version-controlled and released as open source. For the cluster-scale and high-redshift extensions below, neural-network surrogate methods already used in cosmology-parameter inference (Verma et al., 2025) are an obvious optimisation target for repeated-likelihood pipelines.

Data sources.

- SPARC catalog: 175 galaxies with Spitzer [3.6μm] photometry and rotation curves (primary).
- THINGS: 34 nearby galaxies with high-resolution HI observations (cross-validation).
- LITTLE THINGS: 41 dwarf irregular galaxies (edge-case testing).
- HIFLUGCS / ACCEPT: galaxy-cluster X-ray mass profiles (scale extension).
- JWST surveys: CRISTAL, GA-NIFS (high-redshift predictions, future work).

Timeline.

Phase	Work	Deliverable
Done	Head-to-head NFW/Burkert on SPARC, AIC/BIC, PCA directional test, CV (this paper, v2)	v2 manuscript
Months 1–3	Pre-registered replication on THINGS and LITTLE THINGS with the frozen proxy list and PC1 test	Replication note
Months 3–6	Cluster-scale analysis on HIFLUGCS; black-hole simulation of pointer-density collapse	Scale-extension paper
Months 6–12	IIT-to-Pointer mapping formalization; $\Phi \propto M_{\text{DM}}$ proof-of-concept on toy graphs	Theory paper

Code and Data Availability

The entire pipeline behind §6 — Pointer, NFW and Burkert halo fits on all 175 SPARC galaxies, 3-fold cross-validation, correlation and PCA analysis — is a single self-contained Python file (`code/analysis.py`) released with this manuscript. It requires only `numpy` and `scipy` and runs end-to-end on a laptop in ∼45 minutes. Raw SPARC data (table 1 and the 175 rotation-curve decomposition files) are publicly available at http://astroweb.cwru.edu/SPARC/ and are not redistributed; see `data/README.md` for the download path expected by the script.

Acknowledgements

Preliminary versions of several sections benefited from dialogue with collaborators whom we thank anonymously pending their consent to attribution. Any errors are the author's own. The broader research programme and its narrative context are described in the companion book *The Celestial Code: Neural Cosmology* (Savchenko, forthcoming).

References

Baars, B. J. (1988). *A Cognitive Theory of Consciousness*. Cambridge University Press.

Bekenstein, J. D. (1973). Black holes and entropy. *Physical Review D*, 7(8), 2333.

Chaitin, G. J. (1975). A theory of program size formally identical to information theory. *J. ACM*, 22(3), 329–340.

Chalmers, D. J. (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies*, 2(3), 200–219.

Dehaene, S., Sergent, C., & Changeux, J.-P. (2003). A neuronal network model linking subjective reports and objective physiological data. *PNAS*, 100(14), 8520–8525.

Gödel, K. (1931). Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme I. *Monatshefte für Mathematik und Physik*, 38, 173–198.

Gorard, J. (2020). Some relativistic and gravitational properties of the Wolfram model. *arXiv:2004.14810*.

Kamada, A., Kaplinghat, M., Pace, A. B., & Yu, H.-B. (2017). Self-interacting dark matter can explain diverse galactic rotation curves. *Physical Review Letters*, 119(11), 111102.

Kung, H. T., & Robinson, J. T. (1981). On optimistic methods for concurrency control. *ACM TODS*, 6(2), 213–226.

Lelli, F., McGaugh, S. S., & Schombert, J. M. (2016). SPARC: Mass models for 175 disk galaxies with Spitzer photometry and accurate rotation curves. *AJ*, 152, 157.

Li, P., Lelli, F., McGaugh, S., & Schombert, J. (2020). A comprehensive catalog of dark-matter halo models for SPARC galaxies. *ApJS*, 247(1), 31.

Libet, B. (1983). Time of conscious intention to act in relation to onset of cerebral activity. *Brain*, 106(3), 623–642.

Maldacena, J., & Susskind, L. (2013). Cool horizons for entangled black holes. *Fortschritte der Physik*, 61, 781.

McGaugh, S. S., Lelli, F., & Schombert, J. M. (2016). Radial acceleration relation in rotationally supported galaxies. *Physical Review Letters*, 117(20), 201101.

Milgrom, M. (1983). A modification of the Newtonian dynamics as a possible alternative to the hidden mass hypothesis. *ApJ*, 270, 365.

Miller, G. A. (1956). The magical number seven, plus or minus two. *Psychological Review*, 63, 81–97.

Navarro, J. F., Frenk, C. S., & White, S. D. M. (1996). The structure of cold dark-matter halos. *ApJ*, 462, 563.

Page, D. N. (1993). Information in black-hole radiation. *Physical Review Letters*, 71, 3743.

Penrose, R., & Hameroff, S. (1996). Orchestrated reduction of quantum coherence in brain microtubules. *Math. Comput. Simul.*, 40, 453–480.

Rosenthal, D. M. (2005). *Consciousness and Mind*. Oxford University Press.

Rubin, V. C., & Ford, W. K. (1970). Rotation of the Andromeda Nebula from a spectroscopic survey of emission regions. *ApJ*, 159, 379.

Salucci, P. (2019). The distribution of dark matter in galaxies. *A&A Review*, 27, 2.

Susskind, L. (1995). The world as a hologram. *Journal of Mathematical Physics*, 36, 6377.

’t Hooft, G. (1993). Dimensional reduction in quantum gravity. *arXiv:gr-qc/9310026.*

Tononi, G. (2004). An information integration theory of consciousness. *BMC Neuroscience*, 5, 42.

Tononi, G. (2008). Consciousness as integrated information. *Biological Bulletin*, 215, 216–242.

Turing, A. M. (1936). On computable numbers, with an application to the Entscheidungsproblem. *Proc. London Math. Soc.*, 42, 230–265.

Vanchurin, V. (2020). The world as a neural network. *Entropy*, 22(11), 1210.

Katsnelson, M. I., & Vanchurin, V. (2021). Emergent quantumness in neural networks. *Foundations of Physics*, 51:94. arXiv:2012.05082.

Vanchurin, V. (2022). Towards a theory of quantum gravity from neural networks. *Entropy*, 24(1), 7.

Vanchurin, V. (2024). Emergent field theories from neural networks. *arXiv*:2411.08138.

Verma, A., Sourav, S., Aluri, P. K., & Mota, D. F. (2025). Cosmology-informed neural networks to infer dark energy equation-of-state. *arXiv*:2508.12032.

Verlinde, E. (2011). On the origin of gravity and the laws of Newton. *JHEP*, 2011(4), 29.

Verlinde, E. (2017). Emergent gravity and the dark universe. *SciPost Physics*, 2(3), 016.

Wheeler, J. A. (1990). Information, physics, quantum: the search for links. In *Complexity, Entropy, and the Physics of Information*, W. Zurek (ed.), Addison-Wesley.

Wolfram, S. (2020). A class of models with the potential to represent fundamental physics. *Complex Systems*, 29(2), 107–536.