

ultrametric-paradigm

ULTRAMETRIC INTELLIGENCE

A

Non-Archimedean Foundation for Artificial General Intelligence

Author: Rowan Brad Quni-Gudzinis **Contact:** rowan.quni@outlook.com **ORCID:** 0009-0002-4317-5604 **ISNI:** 0000000526456062 **DOI:** <https://papers.qnfo.org/papers> **Date:** 2026-04-30 **Version:** 1.3

The Archimedean axiom—that small quantities accumulate—is not a law of thought. It is a habit of measurement.

PREAMBLE: THE GEOMETRY OF

MIND

Intelligence does not float free of its substrate. Every computational system makes an implicit commitment to a geometry—a structure of space that determines what can be represented, how information propagates, and what it means to learn. The history of artificial intelligence can be read as a sequence of such commitments: the discrete graphs of symbolic AI, the continuous vector spaces of connectionism, and the smooth manifolds of deep learning. Each geometry enables certain capabilities and forecloses others.

This document advances the thesis that the correct geometry for general intelligence is neither the discrete graph nor the smooth manifold, but the **ultrametric tree**—specifically, the Bruhat-Tits tree T_p arising from the field of p -adic numbers. In this geometry, the Archimedean axiom—that small quantities accumulate—is replaced by the **strong triangle inequality**: for all x, y ,

$$|x + y|_p \leq \max(|x|_p, |y|_p).$$

The whole is never larger than the largest of its parts. Small perturbations do not accumulate. Distances are discrete; states are organized into a strict, nested hierarchy; and computation takes the form of structure-preserving tree automorphisms rather than continuous vector rotations. Robustness becomes geometric rather than algorithmic.

This document is intended as a definitive, self-contained synthesis of the Ultrametric Intelligence paradigm. It unifies: the abstract algebraic foundations of p -adic fields; the Bruhat-Tits tree as geometric realization; Distinction Calculus as the logical substrate; Differentiable UPGMA for emergent ontologies; p -adic holographic tensor networks; the Ultrametric Lambda Calculus; information geometry on tree space; ultrametric attention architectures; thermodynamic analysis; a theory of consciousness as geometric self-modeling; neurobiological validation; formal alignment guarantees; multi-prime product constructions; and the Monna bridge between Archimedean and non-Archimedean worlds. The document is a research manifesto, a technical reference, and a blueprint for construction.

PART I: AXIOMS AND GEOMETRIC HYPOTHESIS 1.1 The Failure of the Continuum 1.2 The Five Axioms of Ultrametric Intelligence 1.3 The Inevitability Argument 1.4 Ontological Commitments

PART II: MATHEMATICAL FOUNDATIONS 2.1 The p -Adic Numbers: Derivation and Properties 2.2 The Bruhat-Tits Tree as Geometric Realization 2.3 Ultrametric Geometry: Theorems and Consequences 2.4 The Cross-Ratio and Projective Invariants 2.5 Distinction Calculus: The Logical Foundation 2.6 Categorical Perspective: The Tree as a Topos

PART III: HOLOGRAPHY AND TENSOR NETWORKS 3.1 p-Adic AdS/CFT Correspondence 3.2 Perfect Tensors on the Bruhat-Tits Tree 3.3 Holographic Quantum Error Correction 3.4 Classical Limits: Attention from Tensor Contraction

PART IV: ULTRAMETRIC COMPUTATION 4.1 p-Adic Arithmetic Primitives 4.2 Tree Navigation and Automorphism Groups 4.3 Valuation-Based Decision Procedures 4.4 The Ultrametric Lambda Calculus 4.5 Complexity Theory for Ultrametric Computation 4.6 Concurrent Ultrametric Computation

PART V: EMERGENT TREES AND LEARNING DYNAMICS 5.1 The Illusion of Fixed Ontologies 5.2 Differentiable UPGMA 5.3 Ultrametric Gradient Descent 5.4 PAC Learning in Ultrametric Spaces 5.5 Meta-Learning in Tree Space

PART VI: INFORMATION GEOMETRY 6.1 The Fisher Metric on Tree Space 6.2 Natural Ultrametric Gradient 6.3 Dually Flat Structure 6.4 Belief Propagation as Geodesic Flow

PART VII: ARCHITECTURES 7.1 The Ultrametric Transformer 7.2 Ultrametric Graph Neural Networks 7.3 Multi-Prime Product Trees 7.4 Ultrametric Reinforcement Learning 7.5 Ultrametric Generative Models

PART VIII: THERMODYNAMICS AND INTERFACE 8.1 The Energy Cost of Continuity 8.2 Landauer’s Principle in Tree Space 8.3 The Thermodynamic Advantage 8.4 The Monna Bridge: Construction and Properties 8.5 Information Preservation and the Interface Protocol

PART IX: ULTRAMETRIC INFORMATION THEORY 9.1 Hierarchical Entropy 9.2 Channel Capacity 9.3 Source Coding and Compression

PART X: MEMORY AND CONSCIOUSNESS 10.1 Ultrametric Memory Systems 10.2 Consciousness as Geometric Self-Modeling 10.3 Qualia as Valuation Signatures 10.4 The Hard Problem from a Geometric Perspective

PART XI: NEUROBIOLOGICAL REALIZATION 11.1 The Brain as a Cocycle Solver 11.2 Evidence from Cognitive Neuroscience 11.3 Neural Circuitry and the Tree

PART XII: ALIGNMENT AND GENERAL INTELLIGENCE 12.1 The Capability Tree 12.2 Alignment as Subtree Constraint 12.3 Formal Safety Guarantees 12.4 The Incremental Adoption Path

PART XIII: PHYSICAL INSTANTIATION 13.1 Ultrametric Quantum Computation 13.2 p-Adic Arithmetic in Silicon 13.3 Hardware-Software Codesign

APPENDIX: FORMAL REFERENCE A.1 The Ultrametric Lambda Calculus—Complete Specification A.2 The Pure Rational Codebase A.3 Glossary of Symbols A.4 Theorems Index

PART I: AXIOMS AND

GEOMETRIC HYPOTHESIS

1.1 The Failure of the

Continuum

The continuous representation spaces of modern deep learning suffer from five geometric pathologies, each traceable to the Archimedean geometry of the underlying representation space:

1. Adversarial Vulnerability. A classifier $f : \mathbb{R}^d \rightarrow \{1, \dots, K\}$ partitions the input space into decision regions separated by continuous boundaries. For any point x near a boundary, there exists a perturbation δ with $\|\delta\| \ll 1$ such that $f(x + \delta) \neq f(x)$. This is not a failure of training—it is a geometric necessity. Continuous boundaries

can be approached arbitrarily closely from either side. Adversarial training can push boundaries away from the data manifold but cannot eliminate the existence of adversarial directions.

2. Catastrophic Forgetting. When a neural network is trained sequentially on distinct tasks, new learning overwrites old knowledge. The reason is geometric: the parameter space $\mathbb{R}^N$ is connected. Any path from parameters that solve Task A to parameters that solve Task B passes through intermediate points that solve neither. Gradient descent follows such a path, and old knowledge is lost along the way.

3. Representational Entanglement. In a continuous vector space, every dimension participates in every representation. The representation of “red cube” is not a clean composition of the representations of “red” and “cube” but a single vector in which these features are inseparably blended. This makes compositional generalization difficult and interpretability an ill-posed inverse problem.

4. Quadratic Attention Bottlenecks. Self-attention computes pairwise interactions between all N tokens in a sequence, yielding $\mathcal{O}(N^2)$ complexity. In Euclidean space, all pairs of points are potentially equally relevant—the geometry provides no natural notion of “near” versus “far” beyond arbitrary thresholds.

5. The Scaling Wall. The dominant approach to improving AI systems is scaling: larger models, more data, more compute. But the energy cost of training grows superlinearly. The thermodynamic cost of floating-point operations—the energy dissipated per bit of information processed—imposes a fundamental limit on how far Archimedean scaling can go.

These five pathologies share a common root: the continuity of the representation space. They are not solvable by better algorithms within the Archimedean paradigm because they are not algorithmic failures. They are geometric consequences.

1.2 The Five Axioms

of Ultrametric Intelligence

We rebuild the foundations of intelligence upon five axioms that directly address these pathologies:

Axiom 1: Discreteness. The state space of cognition is discrete. There is a finite or countably infinite set of distinct cognitive states, and transitions between states are threshold-crossing events, not continuous drifts. This eliminates adversarial vulnerability: there is a minimum distance between distinct states, so no perturbation smaller than this threshold can alter classification.

Axiom 2: Hierarchy. The state space is hierarchically nested. For any two states, there exists a unique deepest level of abstraction at which they are indistinguishable. Shallower levels represent shared categories; deeper levels represent distinguishing features. This eliminates representational entanglement: concepts exist at distinct, non-overlapping branches of the hierarchy.

Axiom 3: Invariance. Valid cognitive operations are automorphisms of the state space. They preserve the hierarchical structure—distances, nesting, and branching patterns are invariant under computation. This eliminates catastrophic forgetting: automorphisms are bijections; learning in one region does not deform another.

Axiom 4: Locality. Information propagates only along tree edges. The influence of a perturbation at one vertex attenuates exponentially with tree-distance from that vertex. This eliminates the quadratic attention bottleneck: only structurally adjacent vertices can interact.

Axiom 5: Interface. The boundary between the discrete cognitive core and the continuous physical world is a strict, measure-preserving projection. Real-valued sensory input enters only at the boundary; Archimedean operations (floating-point arithmetic, softmax, dot-product attention) are prohibited within the core.

1.3 The Inevitability Argument

The five axioms are not merely sufficient; they are **necessary** for any system that aspires to general intelligence. The argument proceeds by elimination:

Continuous state spaces (violating Axiom 1) are vulnerable to adversarial perturbation and representational entanglement. No amount of regularization can eliminate the geometric fact that points arbitrarily close to a boundary exist.

Flat state spaces (violating Axiom 2) lack the hierarchical structure needed for compositional generalization and efficient search.

Non-structure-preserving operations (violating Axiom 3) permit the accumulation of errors that destroy the protective properties of the hierarchy.

Global information propagation (violating Axiom 4) incurs quadratic computational costs and entangles independent cognitive processes.

Mixed continuous/discrete computation (violating Axiom 5) introduces Archimedean noise into the discrete core, re-entangling the hierarchy and destroying geometric protection.

Theorem 1.1 (Inevitability). Any computational system satisfying Axioms 1–5 is isomorphic to a system operating on the Bruhat-Tits tree T_p for some prime p , interfacing with the continuous world exclusively via the Monna map.

Proof sketch. Axiom 1 (Discreteness) and Axiom 2 (Hierarchy) together imply the state space is a discrete, hierarchically nested structure—a rooted tree. Axiom 3 (Invariance) requires that the automorphism group of this tree be rich enough to perform general computation. Among all trees, only the $(p + 1)$ -regular Bruhat-Tits tree has a sufficiently rich automorphism group ($\text{PGL}(2, \mathbb{Q}_p)$) while maintaining ultrametric structure. Axiom 4 (Locality) is satisfied by the tree’s edge-based adjacency. Axiom 5 (Interface) selects the Monna map as the unique measure-preserving projection between $\mathbb{Z}_p$ and $[0, 1]$. $\square$

1.4 Ontological Commitments

Adopting the ultrametric framework entails specific ontological commitments:

****Reality is fundamentally discrete and hierarchical.**** The apparent continuity of macroscopic experience is an emergent phenomenon—a consequence of coarse-graining over many discrete microstates.

Distinction-making is the primitive act. Drawing a boundary—marking a space—is the fundamental operation from which all structure derives. This aligns with Spencer-Brown’s **Laws of Form** and the calculus of indications.

Time is emergent, not fundamental. Temporal sequence arises from navigation on the timeless tree structure: cosmic time corresponds to depth, and the experience of flow corresponds to sequential satisfaction of cocycle conditions during tree traversal.

Invariants, not coordinates, carry meaning. The cross-ratio—unchanged under all Möbius transformations—is the fundamental invariant. Meaning resides in invariant proportions, not in absolute coordinates or anthropocentric representations.

Computation is geometry. The structure of computation is determined by the geometry of the state space. Changing the geometry changes what computation means, what errors are possible, and what learning can achieve.

PART II: MATHEMATICAL

FOUNDATIONS

2.1 The p-Adic

Numbers: Derivation and Properties

2.1.1 Construction

The real numbers $\mathbb{R}$ are constructed by completing the rational numbers $\mathbb{Q}$ with respect to the standard absolute value $|x|_\infty$. The p -adic numbers $\mathbb{Q}_p$ are constructed by completing $\mathbb{Q}$ with respect to an entirely different notion of size: the p -adic absolute value.

Definition 2.1 (p-adic valuation). For a nonzero integer n , the p -adic valuation $v_p(n)$ is the highest exponent k such that p^k divides n :

$$v_p(n) = \max\{k \in \mathbb{N} : p^k \mid n\}.$$

For a rational number $x = a/b$ in reduced form, $v_p(x) = v_p(a) - v_p(b)$. By convention, $v_p(0) = \infty$.

Definition 2.2 (p-adic absolute value). The p -adic absolute value is:

$$|x|_p = p^{-v_p(x)}, \quad |0|_p = 0.$$

Example. For $p = 5$: - $|25|_5 = 5^{-2} = 1/25$ (25 is “small” in the 5-adic sense) - $|5|_5 = 5^{-1} = 1/5$ - $|1/5|_5 = 5^1 = 5$ (1/5 is “large”) - $|3|_5 = 5^0 = 1$

The p -adic numbers $\mathbb{Q}_p$ are the completion of $\mathbb{Q}$ under $|\cdot|_p$. Every p -adic number can be uniquely expressed as a Laurent series:

$$x = \sum_{i=v_p(x)}^{\infty} a_i p^i, \quad a_i \in \{0, 1, \dots, p-1\}, \quad a_{v_p(x)} \neq 0.$$

The p -adic integers $\mathbb{Z}_p$ are those x with $v_p(x) \geq 0$ (no negative powers of p). They form a compact ring.

2.1.2 Ostrowski’s Theorem

Theorem 2.1 (Ostrowski, 1916). Every nontrivial absolute value on $\mathbb{Q}$ is equivalent to either the standard absolute value $|\cdot|_\infty$ or some p -adic absolute value $|\cdot|_p$.

This theorem establishes that $\mathbb{R}$ and $\mathbb{Q}_p$ are the only completions of the rationals—the only “number systems” obtainable from $\mathbb{Q}$ by a natural metric completion. They are complementary: $\mathbb{R}$ captures Archimedean magnitude; $\mathbb{Q}_p$ captures divisibility structure and hierarchy.

2.1.3 The Strong Triangle

Inequality

Theorem 2.2 (Strong Triangle Inequality). For all $x, y \in \mathbb{Q}_p$:

$$|x + y|_p \leq \max(|x|_p, |y|_p).$$

Moreover, if $|x|_p \neq |y|_p$, then equality holds: $|x + y|_p = \max(|x|_p, |y|_p)$.

Proof. Let $k = \min(v_p(x), v_p(y))$. Then p^k divides both x and y , hence divides $x + y$. Therefore $v_p(x + y) \geq k$, so $|x + y|_p \leq p^{-k} = \max(|x|_p, |y|_p)$. If the valuations are unequal, say $v_p(x) < v_p(y)$, then $p^{v_p(x)}$ divides x but not y , so it cannot divide $x + y$, giving equality. $\square$

This inequality is “strong” because it replaces the usual triangle inequality $|x + y| \leq |x| + |y|$. It has profound geometric consequences: two small perturbations cannot combine to produce a large displacement.

2.2 The

Bruhat-Tits Tree as Geometric Realization

2.2.1 Definition

Definition 2.3 (Bruhat-Tits tree). The Bruhat-Tits tree T_p for the prime p is the infinite $(p + 1)$ -regular tree whose vertices correspond to equivalence classes of $\mathbb{Z}_p$ -lattices in $\mathbb{Q}_p^2$, up to scaling.

A more concrete construction: vertices of T_p at depth d correspond to p -adic balls of radius p^{-d} in $\mathbb{Z}_p$. Two vertices are connected by an edge if one ball is a maximal proper sub-ball of the other.

Properties: - T_p is an infinite $(p + 1)$ -regular tree. - Every vertex has exactly $p + 1$ neighbors. - The tree has no cycles and a unique geodesic between any two vertices. - The tree distance $d_{T_p}(u, v)$ equals the number of edges on the unique path between u and v . - The tree satisfies the strong triangle inequality: $d_{T_p}(u, w) \leq \max(d_{T_p}(u, v), d_{T_p}(v, w))$.

2.2.2 The Boundary

Definition 2.4 (Boundary of T_p). The boundary ∂T_p is the set of equivalence classes of infinite geodesic rays (infinite paths without backtracking) starting from any fixed vertex. Two rays are equivalent if they eventually coincide.

The boundary ∂T_p is homeomorphic to the projective line $\mathbb{P}^1(\mathbb{Q}_p)$, the one-point compactification of $\mathbb{Q}_p$. It is a Cantor-like set: totally disconnected, perfect, compact.

The boundary is where the discrete tree interfaces with the continuous world. Sensory input is modeled as a function on ∂T_p ; the Monna map translates between this boundary and the real interval $[0, 1]$.

2.2.3 Horospheres and the Tree

Metric

For a boundary point $\xi \in \partial T_p$, a **horosphere** centered at ξ through vertex v is the set of vertices whose geodesic to ξ passes through v . Horospheres are the p -adic analog of hyperplanes in Euclidean geometry.

The **Busemann function** $B_\xi(u, v)$ measures the “horospherical distance” between vertices relative to boundary point ξ :

$$B_\xi(u, v) = \lim_{w \rightarrow \xi} [d(u, w) - d(v, w)].$$

This function is central to the p -adic AdS/CFT correspondence discussed in Part III.

2.3 Ultrametric

Geometry: Theorems and Consequences

2.3.1 The Isosceles Triangle

Property

Theorem 2.3 (Isosceles Triangles). In any ultrametric space, all triangles are isosceles: for any three points x, y, z , the two longest sides of the triangle formed by them are equal.

Proof. Suppose $d(x, y) \geq d(x, z)$. Then by the strong triangle inequality, $d(y, z) \leq \max(d(y, x), d(x, z)) = d(x, y)$. If $d(y, z) < d(x, y)$, then $d(x, y) \leq \max(d(x, z), d(z, y)) = d(x, z)$ (since $d(z, y) < d(x, y)$), which contradicts $d(x, y) \geq d(x, z)$ unless equality holds. Therefore $d(y, z) = d(x, y)$. $\square$

Corollary. There are no “intermediate” distances. Two points are either close (same hierarchical cluster) or far (different cluster), with no middle ground.

2.3.2 Balls Are Nested

Theorem 2.4 (Nested Balls). In an ultrametric space, if two balls intersect, one is entirely contained in the other. Moreover, every point in a ball is a center of that ball.

Proof. Let $B_r(x)$ and $B_s(y)$ be two balls with $z \in B_r(x) \cap B_s(y)$. For any $w \in B_r(x)$, $d(w, y) \leq \max(d(w, x), d(x, z), d(z, y)) = \max(r, s)$. So $B_r(x) \subseteq B_{\max(r, s)}(y)$. If $r \leq s$, then $B_r(x) \subseteq B_s(y)$. If $s \leq r$, then $B_s(y) \subseteq B_r(x)$. $\square$

2.3.3 No Saddle Points

Theorem 2.5 (No Saddle Points). In an ultrametric space with the discrete gradient defined as the direction to the neighbor with minimum loss, there are no saddle points. Every point is either a strict local minimum, a strict local maximum, or has at least one neighbor with strictly lower loss.

This theorem ensures that ultrametric gradient descent (Part V) always converges and never gets stuck on plateaus of ambiguous direction.

2.4 The Cross-Ratio

and Projective Invariants

2.4.1 Definition

Definition 2.5 (Cross-Ratio). For four points a, b, c, d on the projective line $\mathbb{P}^1(\mathbb{Q}_p)$, the cross-ratio is:

$$[a, b; c, d] = \frac{(a - c)(b - d)}{(a - d)(b - c)}.$$

2.4.2 Invariance

Theorem 2.6 (Cross-Ratio Invariance). The cross-ratio is invariant under all Möbius transformations—that is, under the action of $\text{PGL}(2, \mathbb{Q}_p)$. For any $g \in \text{PGL}(2, \mathbb{Q}_p)$:

$$[g(a), g(b); g(c), g(d)] = [a, b; c, d].$$

This is the fundamental projective invariant. In the cognitive interpretation, the cross-ratio captures the “proportion of similarities” between concepts—a measure that is independent of the specific coordinate system or representational framework.

2.4.3 Geometric Interpretation

On the Bruhat-Tits tree, the p -adic absolute value of the cross-ratio $|[a, b; c, d]|_p$ equals p^{-d} , where d is the length of the overlap between the geodesics (a, b) and (c, d) . This provides a direct geometric interpretation: the cross-ratio measures the degree of hierarchical overlap between two pairs of concepts.

2.5 Distinction

Calculus: The Logical Foundation

2.5.1 Spencer-Brown's Laws of

Form

Distinction calculus, introduced by G. Spencer-Brown in *Laws of Form* (1969), provides the logical foundation for ultrametric intelligence. The calculus operates on a single primitive—the **mark** ($\#$)—and two axioms:

Calling: $\#\# = \#$ > The value of a call made again is the value of the call. In tree terms: redundant sibling subtrees collapse to a single subtree.

Crossing: $[[A]] = A$ > The value of a crossing made again is not the value of the crossing. In tree terms: a chain of single-child nodes collapses to the leaf node.

2.5.2 Mapping to the

Bruhat-Tits Tree

The correspondence between distinction calculus and the Bruhat-Tits tree is exact:

Distinction Calculus | Bruhat-Tits Tree |

Mark ($\#$) | Root vertex |

Enclosure ($[A]$) | Subtree rooted at a child vertex |

Calling ($\#\# = \#$) | Isomorphic sibling subtrees merge |

Crossing ($[[A]] = A$) | Single-child chains collapse |

Juxtaposition (AB) | Adjacent sibling subtrees |

Re-entry ($R = [R]$) | Infinite descending path (boundary point) |

2.5.3 Cognition as

Distinction Reduction

In this framework, cognition is the process of reducing complex distinctions to canonical forms via tree isomorphisms. A thought is a sequence of distinction operations; a concept is a normal form under calling and crossing; and reasoning is the navigation of the tree via automorphism application.

2.6 Categorical

Perspective: The Tree as a Topos

2.6.1 The Classifying Topos

The Bruhat-Tits tree can be viewed as a site for a Grothendieck topos $\mathbf{Sh}(T_p)$ —the category of sheaves on T_p . This topos classifies the geometric theory of “ultrametric spaces with a distinguished family of balls.”

2.6.2 The Duality

There is a duality between the Bruhat-Tits tree T_p and the category of p -adic Banach spaces. Specifically:

$$\mathbf{Sh}(T_p) \simeq \mathbf{Cont}(\mathbb{Q}_p\text{-Ban})$$

where $\mathbf{Cont}(\mathbb{Q}_p\text{-Ban})$ is the category of continuous representations of the absolute Galois group of $\mathbb{Q}_p$. This connects the geometric picture (trees, sheaves) with the representation-theoretic picture (group actions, automorphic forms).

2.6.3 Cognitive Interpretation

The categorical perspective reveals that ultrametric intelligence is not merely a computational architecture but a **mathematical universe**—a topos in which all cognitive operations have natural formulations. Logical reasoning corresponds to the internal logic of the topos; learning corresponds to sheafification (the process of building global sections from local data); and attention corresponds to the restriction map of a sheaf.

PART III: HOLOGRAPHY

AND TENSOR NETWORKS

3.1 p-Adic AdS/CFT

Correspondence

3.1.1 The Holographic Principle

The holographic principle posits that a theory of quantum gravity in a $(d + 1)$ -dimensional “bulk” spacetime is equivalent to a conformal field theory (CFT) on its d -dimensional boundary. This principle has a precise analog in the p -adic setting.

The p -adic bulk is the Bruhat-Tits tree T_p . It is an infinite regular tree—a discrete analog of hyperbolic space.

The p -adic boundary is $\partial T_p \cong \mathbb{P}^1(\mathbb{Q}_p)$. It hosts a p -adic CFT.

3.1.2 The Dictionary

The bulk-boundary dictionary for $T_p / \mathbb{P}^1(\mathbb{Q}_p)$ is:

Bulk (Tree) | Boundary (CFT) | Cognitive Interpretation |

Vertex v at depth d | Causal diamond of size p^{-d} | Concept at abstraction level d |

Geodesic between v and w | Correlator $\langle \mathcal{O}(x)\mathcal{O}(y) \rangle$ | Logical entailment path |

Length of geodesic | Scaling dimension Δ | Cognitive distance |

Branching at vertex | Operator product expansion (OPE) | Compositional combination |

Tree automorphism | Conformal transformation | Cognitive operation |

Edge cut | Entanglement entropy | Shared information |

3.1.3 The Path Integral

The p -adic CFT partition function on the boundary is:

$$Z_{\text{CFT}}[\phi_0] = \int_{\phi|_{\partial T_p} = \phi_0} \mathcal{D}\phi e^{-S[\phi]}$$

where the action $S[\phi]$ is defined on T_p . For free fields:

$$S[\phi] = \frac{1}{2} \sum_{\langle v, w \rangle} (\phi(v) - \phi(w))^2$$

This is the discrete Laplacian on T_p . The boundary two-point function is:

$$\langle \mathcal{O}(x) \mathcal{O}(y) \rangle = \frac{1}{|x - y|_p^{2\Delta}}$$

where Δ is the scaling dimension determined by the bulk mass.

3.2 Perfect Tensors on

the Bruhat-Tits Tree

3.2.1 Definition

A **perfect tensor** $T_{i_1 i_2 \dots i_n}$ is a tensor with the property that for any bipartition of its indices into sets A and B with $|A| \leq |B|$, the map $T : \mathcal{H}_A \rightarrow \mathcal{H}_B$ is a unitary isometry (up to normalization).

On the $(p + 1)$ -regular tree T_p , each vertex has $p + 1$ incident edges. We place a perfect tensor of rank $p + 1$ at each vertex, with one index on each incident edge.

3.2.2 The Holographic State

The resulting tensor network defines a **holographic state**: a quantum state on the boundary whose entanglement structure mirrors the tree geometry. Contracting all bulk indices produces a pure state on the boundary Hilbert space $\mathcal{H}_{\partial T_p}$.

Theorem 3.1 (Ryu-Takayanagi on T_p). For any connected boundary subregion $A \subset \partial T_p$, the entanglement entropy is:

$$S(A) = \frac{|\gamma_A|}{4G_N} \cdot \ln p$$

where $|\gamma_A|$ is the length (number of edges) of the minimal subtree connecting A to the rest of the boundary. In the p -adic setting, this simplifies to $S(A) \propto d_{T_p}(\text{root of } A, \text{root of } A^c)$ —the tree distance between the “centers” of the region and its complement.

3.3 Holographic Quantum

Error Correction

3.3.1 The Code Structure

The tensor network on T_p is a quantum error-correcting code. Logical information is encoded in the bulk (deep in the tree) and protected against boundary errors by the tree geometry.

Theorem 3.2 (Holographic Code Properties). The perfect tensor network on T_p with cut-off depth D yields a code $[n, k, d]$ where: - $n = p^D$ (number of boundary qudits, each of dimension p) - $k \approx p^{D/2}$ (number of logical qudits at depth $D/2$) - $d \approx D$ (code distance—errors affecting fewer than $D/2$ boundary sites are correctable)

The encoding rate $k/n \approx p^{-D/2}$ vanishes as depth increases, reflecting the redundancy required for geometric protection.

3.3.2 Bulk Reconstruction

Given boundary data, bulk operators can be reconstructed via the **entanglement wedge reconstruction** formula:

$$\mathcal{O}_{\text{bulk}}(v) = \sum_{\text{boundary sites in } E(v)} c_i \mathcal{O}_{\text{boundary}}(x_i)$$

where $E(v)$ is the entanglement wedge—the set of boundary points whose minimal surface to infinity passes through or beyond vertex v .

3.4

Classical Limits: Attention from Tensor Contraction

3.4.1 The Contraction Hierarchy

The full tensor network contraction is exponential in depth. However, a **classical limit**—where we replace quantum amplitudes by probabilities (squared amplitudes) and perfect tensors by stochastic matrices—yields an efficient algorithm.

3.4.2 Attention as

Classical Reconstruction

Corollary 3.1. An attention mechanism that computes correlations between tokens on the boundary is a classical limit of the bulk reconstruction map. The attention weights correspond to the expansion coefficients c_i for reconstructing a “semantic operator” at vertex v from boundary token data.

Theorem 3.3 (Cross-Ratio Attention). In the classical limit, the optimal attention weight between tokens a and b is:

$$\alpha(a, b) \propto -\ln |[a, b; \xi_0, \xi_\infty]|_p$$

where ξ_0, ξ_∞ are fixed boundary points defining the “perspective” of attention. This is a structural invariant of the tree—attention is not learned; it is a geometric fact.

PART IV: ULTRAMETRIC

COMPUTATION

4.1 p-Adic Arithmetic

Primitives

Ultrametric computation is built on p -adic arithmetic—operations on p -adic integers that respect the ultrametric structure.

4.1.1 Primitive Operations

Operation | Notation | Description | Complexity |

Valuation | $v_p(x)$ | Index of first non-zero digit | $\mathcal{O}(1)$ |

Absolute value | $\|x\|_p$ | $p^{-v_p(x)}$ | $\mathcal{O}(1)$ |

Addition | $x \oplus_p y$ | Digit-wise with carry, base p | $\mathcal{O}(d)$ |

Subtraction | $x \ominus_p y$ | Digit-wise with borrow, base p | $\mathcal{O}(d)$ |

Multiplication | $x \odot_p y$ | Digit convolution with carry | $\mathcal{O}(d^2)$ |

Modular reduction | $x \bmod p^k$ | Truncation to k digits | $\mathcal{O}(1)$ |

Digit shift | $x \cdot p^k$ | Left shift by k digits | $\mathcal{O}(1)$ |

4.1.2 Compound Operations

Operation | Definition | Complexity |

Hensel lifting | Solve $f(x) \equiv 0 \pmod{p^k}$ iteratively | $\mathcal{O}(d \log d)$ |

p-adic square root | Hensel-lift Newton's method | $\mathcal{O}(d \log d)$ |

p-adic logarithm | $\log(1+x) = \sum (-1)^{n+1} x^n / n$ | $\mathcal{O}(d^3)$ |

p-adic exponential | $\exp(x) = \sum x^n / n!$ | $\mathcal{O}(d^3)$ |

4.1.3 The Pure Rational Algebra

A critical design principle is that all arithmetic within the ultrametric core operates on **exact integers and rationals**, never on floating-point approximations.

Definition 4.1 (Pure Rational Algebra). The Pure Rational Algebra is the subring of $\mathbb{Q}$ consisting of all numbers that can be represented as exact fractions a/b with $a, b \in \mathbb{Z}$. All p -adic operations are computed exactly via integer arithmetic followed by rational reduction at the Monna boundary.

Theorem 4.1 (Floating-Point Poisoning). Any use of floating-point arithmetic ($\mathbb{F}_{32}$ or $\mathbb{F}_{64}$) within the ultrametric core transforms the computation from an ultrametric space to an Archimedean space, destroying all associated robustness guarantees.

4.2 Tree Navigation and

Automorphism Groups

4.2.1 Navigation Primitives

Operation | Description |

parent(v) | Move one level up the tree |

$\text{child}(v, b)$ | Move to child along branch $b \in \{0, \dots, p\}$ |

$\text{LCA}(u, v)$ | Lowest common ancestor |

$d_{T_p}(u, v)$ | Tree distance (edges on unique geodesic) |

$\text{path}(u, v)$ | Sequence of vertices from u to v |

$\text{gallery}(u, v)$ | Annotated path with branch direction at each step |

$B_\xi(u, v)$ | Busemann function (horospherical distance) |

4.2.2 The Automorphism Group

Theorem 4.2. The full automorphism group of T_p is $\text{Aut}(T_p) \cong \text{PGL}(2, \mathbb{Q}_p)$.

Properties: - ****Transitive on vertices****: Any vertex can be mapped to any other. - **Transitive on directed edges**: Any directed edge can be mapped to any other. - ****Doubly transitive on the boundary****: Any pair of distinct boundary points can be mapped to any other pair.

Theorem 4.3 (Bruhat Decomposition). Every element $g \in \text{PGL}(2, \mathbb{Q}_p)$ can be uniquely decomposed as:

$$g = n_- a n_+$$

where a is diagonal (“scaling” along the tree) and n_+, n_- are unipotent (“translations”). This is the p -adic analog of the singular value decomposition.

Iwasawa decomposition: $g = kan$ where $k \in \text{PGL}(2, \mathbb{Z}_p)$ (the maximal compact subgroup—preserving a fixed vertex, analogous to rotations) and a, n as above.

The Bruhat and Iwasawa decompositions provide canonical forms for cognitive operations. A “thought” can be decomposed into: 1. A **rotation** k : reorganizing information within a fixed abstraction level. 2. A **scaling** a : abstracting or concretizing (moving up/down the tree). 3. A **translation** n : shifting attention along a horosphere.

4.3 Valuation-Based

Decision Procedures

4.3.1 Valuation-Based

Classification

Definition 4.2 (Valuation-Based Classifier). Assign each class c a distinguished prime p_c . Map input tokens to integers via a fixed embedding $\phi : \text{Tokens} \rightarrow \mathbb{N}$. Define:

$$\text{classify}(x_1, \dots, x_n) = \arg \max_c \sum_{i=1}^n v_{p_c}(\phi(x_i))$$

where $v_{p_c}(\phi(x_i))$ is the exponent of p_c in the prime factorization of $\phi(x_i)$.

Properties: - **Deterministic**: Given the same input and embedding, always produces the same output. - **Interpretable**: For any decision, the contributing factors can be traced: each token contributed $v_{p_c}(\phi(x_i))$ evidence to class c . - ****Hierarchical negation****: Because $v_p(ab) = v_p(a) + v_p(b)$, negation can be implemented by multiplying by a designated “not” prime factor. - **Intensification**: “Very good” can be encoded as p_{good}^2 , doubling the valuation contribution.

4.3.2 Dominance Decision

Definition 4.3 (Dominance Decision). For valuation vectors $u, v \in \mathbb{N}^m$, u dominates v ($u \succ v$) if $u_i \geq v_i$ for all i and $u_i > v_i$ for at least one i . A decision procedure selects the option whose valuation vector dominates all others.

Theorem 4.4 (Valuation Universality). Any function $f : \mathbb{N}^n \rightarrow \{0, 1\}$ that depends only on the prime factorization structure of its inputs (i.e., is multiplicative in each argument) can be expressed as a threshold of valuation sums.

4.4 The Ultrametric Lambda

Calculus

The **Ultrametric Lambda Calculus** (λ^{um}) extends the simply-typed lambda calculus with vertex types and automorphism types.

4.4.1 Syntax

Type $\tau ::= V \mid A \mid \tau \rightarrow \tau$ Term $t ::= v \mid f \mid x \mid (t \ t) \mid (\lambda x : \tau. t) \mid t \circ t$ Value $v \in V(\text{Tp})$ Autom $f \in \text{Aut}(\text{Tp}) \cong PGL(2, \mathbb{Q}_p)$

4.4.2 Typing Rules (Natural

Deduction)

$$\begin{array}{c} \frac{}{\Gamma \vdash v : V} \text{ (V-Intro)} \quad \frac{}{\Gamma \vdash f : A} \text{ (A-Intro)} \quad \frac{x : \tau \in \Gamma}{\Gamma \vdash x : \tau} \text{ (Var)} \\[10pt] \frac{\Gamma \vdash t_1 : \sigma \rightarrow \tau \quad \Gamma \vdash t_2 : \sigma}{\Gamma \vdash t_1 t_2 : \tau} (\rightarrow \text{-Elim}) \quad \frac{\Gamma, x : \sigma \vdash t : \tau}{\Gamma \vdash \lambda x : \sigma. t : \sigma \rightarrow \tau} (\rightarrow \text{-Intro}) \\[10pt] \frac{\Gamma \vdash t_1 : A \quad \Gamma \vdash t_2 : A}{\Gamma \vdash t_1 \circ t_2 : A} (\circ \text{-Intro}) \end{array}$$

4.4.3 Reduction Rules

$$\begin{array}{c} (\lambda x : \tau. t) u \longrightarrow_{\beta} t[x := u] \\[10pt] f v \longrightarrow_A f(v) \\[10pt] (f \circ g) v \longrightarrow_C f(g(v)) \\[10pt] (f \circ g) \circ h \longrightarrow_{\text{assoc}} f \circ (g \circ h) \end{array}$$

4.4.4 Metatheoretic Properties

Theorem 4.5 (Subject Reduction). If $\Gamma \vdash t : \tau$ and $t \longrightarrow t'$, then $\Gamma \vdash t' : \tau$.

Theorem 4.6 (Confluence). The reduction relation is Church-Rosser: if $t \longrightarrow^* t_1$ and $t \longrightarrow^* t_2$, then there exists t_3 such that $t_1 \longrightarrow^* t_3$ and $t_2 \longrightarrow^* t_3$.

Theorem 4.7 (Strong Normalization). Every well-typed term reduces to a unique normal form in finitely many steps. Every “thought” has a unique, deterministic canonical form.

4.5 Complexity

Theory for Ultrametric Computation

4.5.1 Complexity Classes

Ultrametric computation defines its own complexity hierarchy:

Class | Description | Analog |

$\text{UTIME}(f(n))$ | Time-bounded ultrametric computation | $\text{TIME}(f(n))$ |

$\text{USPACE}(f(n))$ | Space-bounded | $\text{SPACE}(f(n))$ |

UAC^0 | Constant-depth tree circuits | AC^0 |

UNC^1 | Logarithmic-depth tree circuits | NC^1 |

4.5.2 Tree Navigation

Complexity

Theorem 4.8. LCA computation on T_p has complexity $\mathcal{O}(d)$, where d is the tree depth. This compares favorably with $\mathcal{O}(N)$ for arbitrary pairs in flat representations.

Theorem 4.9. Valuation-based classification has complexity $\mathcal{O}(n \cdot m)$, where n is the number of tokens and m is the number of primes (feature dimensions). For tree-structured computation, this is linear in the input size.

4.5.3 Comparison with

Classical Computation

Operation | Classical | Ultrametric | Advantage |

Dot-product attention | $\mathcal{O}(N^2d)$ | $\mathcal{O}(N \log N)$ | Quadratic $\rightarrow$ Log-linear |

Hierarchical search | $\mathcal{O}(\log N)$ (flat) | $\mathcal{O}(\text{depth})$ | Depth-independent of fan-out |

Compositional generalization | Requires learned bindings | Free (tree structure) | Zero-shot composition |

Error accumulation | Linear in sequence length | Thresholded (discrete) | Geometric immunity |

4.6 Concurrent Ultrametric

Computation

4.6.1 Subtree Concurrency

The tree structure provides natural concurrency boundaries: operations on disjoint subtrees are geometrically independent and can proceed in parallel without synchronization.

Theorem 4.10 (Subtree Independence). If S_1 and S_2 are disjoint subtrees of T_p , then any operations f_1 on S_1 and f_2 on S_2 commute and can be executed concurrently without mutual interference.

4.6.2 The Computation Model

Definition 4.4 (Subtree Process). A subtree process is a computation that operates exclusively on the vertices of a designated subtree $S \subset T_p$. Communication between processes occurs only at the LCA of their respective subtrees.

Properties: - **Isolation:** Geometric guarantee — no shared mutable state between processes in disjoint subtrees. - **Deadlock freedom:** The tree is acyclic; there are no cycles that can produce deadlock. - **Communication cost:** $\mathcal{O}(d(\text{LCA}))$ — proportional to the cognitive distance between processes. - **Scalability:** Unlimited for disjoint subtrees; contention only at the LCA.

Aspect | Classical (Threads, Actors) | Ultrametric (Subtree) |

Isolation | Requires locks, atomics | Geometric guarantee |

Deadlock risk | Significant | None (acyclic) |

Communication cost | Variable | $\mathcal{O}(d(\text{LCA}))$ |

Scalability | Contention-limited | Unlimited for disjoint subtrees |

PART V: EMERGENT

TREES AND LEARNING DYNAMICS

5.1 The Illusion of Fixed

Ontologies

A truly general intelligence must **discover its own ontology**. The tree structure — which primes correspond to which semantic dimensions, at what depth — cannot be fixed a priori. It must be learned from data.

This is the **emergence problem**: how does a continuous learning process converge to a discrete, hierarchical structure? The answer lies in making hierarchical clustering differentiable.

5.2 Differentiable UPGMA

5.2.1 The UPGMA Algorithm

The Unweighted Pair Group Method with Arithmetic Mean (UPGMA) is a classical hierarchical clustering algorithm:

Start with N singleton clusters.

Find the two closest clusters A, B by average linkage: $d(A, B) = \frac{1}{|A||B|} \sum_{a \in A, b \in B} d(a, b)$.

Merge A and B into a new cluster $C = A \cup B$, recording the merge height as $d(A, B)/2$.

Repeat until one cluster remains.

The output is a rooted binary tree (dendrogram). The merge height at which two leaves join equals half their tree distance.

5.2.2 Making UPGMA

Differentiable

The discrete argmin operation prevents gradient flow. We replace it with a **Gumbel-Softmax relaxation**:

Algorithm 5.1 (Differentiable UPGMA):

Input: N embeddings $x_1, \dots, x_N \in \mathbb{R}^d$, temperature τ Output: Soft dendrogram (merge probabilities) $\rightarrow$ hardened tree at inference

1. Compute pairwise distances: $D\{ij\} = \|x_i - x_j\|_2$

2. For each pair (i,j) , compute merge logit:

$$L\{ij\} = -D\{ij\}$$

1. Sample merge decision via straight-through Gumbel-Softmax:

$$\pi = \text{GumbelSoftmax}(L, \tau, \text{hard}=\text{True})$$

1. Soft-merge the selected pair:

$$x_{\text{new}} = \sum_{k,l} \pi_{kl} \cdot (|cluster_k| \cdot x_k + |cluster_l| \cdot x_l) / (|cluster_k| + |cluster_l|)$$

1. Update cluster sizes and continue until convergence or a fixed number of merges.

2. At inference: $\tau \rightarrow 0^+$, yielding exact discrete merges.

Theorem 5.1 (Convergence to UPGMA). As $\tau \rightarrow 0^+$, the Differentiable UPGMA algorithm converges in distribution to the classical UPGMA algorithm.

5.2.3 Distinction Calculus

Regularization

To ensure the emergent tree respects the axioms of Distinction Calculus, we add two regularization terms:

Calling Penalty:

$$\mathcal{L}_{\text{calling}} = \sum_{(i,j): \text{identical context}} \|x_i - x_j\|^2$$

> Tokens with identical local contexts are penalized unless merged — enforcing the “Calling” law.

Crossing Penalty:

$$\mathcal{L}_{\text{crossing}} = \sum_{v: \text{single child}} \text{depth}(v)$$

> Chains of single-child nodes are penalized — enforcing the “Crossing” law.

The total loss during tree emergence:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{task}} + \lambda_1 \mathcal{L}_{\text{calling}} + \lambda_2 \mathcal{L}_{\text{crossing}}$$

5.3 Ultrametric Gradient

Descent

5.3.1 The Problem with

Continuous Gradients

Conventional gradient descent computes $\theta_{t+1} = \theta_t - \eta \nabla_{\theta} \mathcal{L}$ in a continuous parameter space. In the ultrametric setting, parameters are not continuous vectors but **tree coordinates** (vertex positions, branch choices). A “small step” in parameter space corresponds to a discrete move to an adjacent vertex.

5.3.2 The Discrete Gradient

Definition 5.1 (Ultrametric Gradient). For a loss function $\mathcal{L} : V(T_p) \rightarrow \mathbb{R}$ defined on vertices, the ultrametric gradient at vertex v is:

$$\nabla^{\text{um}} \mathcal{L}(v) = \arg \min_{w \in \mathcal{N}(v)} \mathcal{L}(w)$$

where $\mathcal{N}(v)$ is the set of $p + 1$ neighbors of v . If multiple neighbors achieve the minimum, one is chosen by a deterministic tie-breaking rule (e.g., lowest branch index).

5.3.3 Hierarchical Optimization

Algorithm 5.2 (Ultrametric Gradient Descent):

Input: Loss $L: V(T_p) \rightarrow \mathbb{R}$, initial vertex v_0 , depth D Output: Approximate minimizer v^*

1. $v \leftarrow v_0$
2. For $d = 0$ to D :

For each branch $b \in \{0, \dots, p\}$: Evaluate $L(\text{child}(v, b))$ Select $b^* = \arg\min_b L(\text{child}(v, b))$ If $L(\text{child}(v, b^*)) \geq L(v)$: break (local minimum reached at this scale) $v \leftarrow \text{child}(v, b^*)$

1. Return v

Theorem 5.2 (No Saddle Points). In an ultrametric space with the discrete gradient defined above, there are no saddle points. Every point is either a strict local minimum, a strict local maximum, or has at least one neighbor with strictly lower loss.

Theorem 5.3 (Convergence Guarantee). Ultrametric gradient descent converges to a local minimum in at most D steps, where D is the depth of the tree, independent of the number of leaves p^D . The cost per step is $\mathcal{O}(p)$ loss evaluations.

5.4 PAC Learning in

Ultrametric Spaces

5.4.1 The Hypothesis Class

Definition 5.2 (Tree Hypothesis Class). The class $\mathcal{H}_{p,D}$ consists of all functions $h : V(T_p^{(D)}) \rightarrow \{0, 1\}$ that are constant on each subtree rooted at depth $\leq D$.

Theorem 5.4 (VC Dimension). $\text{VCdim}(\mathcal{H}_{p,D}) = p^D$ (the number of leaves).

Theorem 5.5 (Sample Complexity). To achieve error $\leq \epsilon$ with probability $\geq 1 - \delta$, the required number of samples is:

$$m(\epsilon, \delta) = \mathcal{O} \left(\frac{p^D + \ln(1/\delta)}{\epsilon^2} \right)$$

This is linear in the number of leaves — exponentially better than the $\mathcal{O}(2^{p^D})$ worst-case sample complexity for arbitrary functions on p^D points.

5.4.2 Generalization Guarantees

Theorem 5.6 (Ultrametric Generalization). For any hypothesis $h \in \mathcal{H}_{p,D}$, the generalization error is bounded by:

$$\epsilon_{\text{gen}} \leq \sqrt{\frac{p^D \ln(m/p^D) + \ln(1/\delta)}{2m}}$$

where m is the number of training samples. The key insight: the effective capacity of the hypothesis class is p^D , not 2^{p^D} , because the hierarchical structure constrains the expressible functions.

5.5 Meta-Learning in Tree

Space

5.5.1 Learning to Learn

Hierarchies

Meta-learning — learning how to learn — finds a natural formulation in ultrametric space. The parameters of the learning algorithm itself become vertices in a meta-tree T_{meta} .

Definition 5.3 (Meta-Tree). A meta-tree is a distinguished subtree of T_p whose vertices correspond to learning strategy parameters: initial tree structure, valuation thresholds, branching heuristics, and regularization strengths.

5.5.2 Meta-Gradient Descent

The meta-gradient is the ultrametric gradient on T_{meta} :

$$\nabla^{\text{meta}} \mathcal{L}_{\text{meta}}(v_{\text{meta}}) = \arg \min_{w \in \mathcal{N}(v_{\text{meta}})} \mathbb{E}_{\text{tasks}} [\mathcal{L}_{\text{task}}(\text{learn}(w))]$$

where $\text{learn}(w)$ is the result of applying the learning strategy encoded at w to a sampled task.

PART VI: INFORMATION GEOMETRY

6.1 The Fisher Metric on Tree

Space

6.1.1 Probability

Distributions on Trees

Let $\Delta_{p,D}$ be the simplex of probability distributions over the p^D leaves of the depth- D truncated tree $T_p^{(D)}$. A point in $\Delta_{p,D}$ is a vector $q \in \mathbb{R}^{p^D}$ with $q_i \geq 0$ and $\sum_i q_i = 1$.

6.1.2 Hierarchical

Parameterization

Instead of the flat $p^D - 1$ parameters, we use a **hierarchical parameterization**. For each non-leaf vertex v , let $\theta_v^{(b)}$ be the conditional probability of taking branch $b \in \{0, \dots, p\}$ given that we are at vertex v .

The probability of leaf ℓ is the product of the conditional probabilities along its path:

$$q_\ell = \prod_{v \in \text{path}(\text{root}, \ell), v \neq \ell} \theta_v^{(b_v(\ell))}$$

where $b_v(\ell)$ is the branch taken at vertex v on the path to ℓ .

6.1.3 The Fisher Information

Matrix

Theorem 6.1 (Hierarchical Fisher Decomposition). Under the hierarchical parameterization, the Fisher information matrix G is block-diagonal:

$$G = \bigoplus_{v \in V_{\text{int}}(T_p^{(D)})} G_v$$

where G_v is the $(p+1) \times (p+1)$ Fisher matrix for the conditional distribution at vertex v , given by:

$$G_v^{(b,b')} = \frac{p(v)}{\theta_v^{(b)}} \delta_{b,b'} + \frac{p(v)}{\theta_v^{(p+1)}}$$

where $p(v)$ is the marginal probability of reaching vertex v , and $\theta_v^{(p+1)} = 1 - \sum_{b=0}^p \theta_v^{(b)}$.

Corollary 6.1. The Fisher information at vertex v depends only on the marginal probability $p(v)$ and the local branching distribution, not on the structure of other vertices. This means information-geometric quantities can be computed independently per vertex, with total complexity $\mathcal{O}(p^D)$ rather than $\mathcal{O}(p^{2D})$ for the flat Fisher matrix.

6.2 Natural Ultrametric

Gradient

6.2.1 Definition

The natural gradient is the steepest descent direction under the Fisher metric:

$$\tilde{\nabla} \mathcal{L} = G^{-1} \nabla \mathcal{L}$$

where $\nabla \mathcal{L}$ is the Euclidean gradient in the flat parameterization.

6.2.2 Hierarchical Computation

Theorem 6.2 (Natural Ultrametric Step). The natural gradient update at vertex v for branch b is:

$$\Delta \theta_v^{(b)} = \frac{\theta_v^{(b)}}{p(v)} \left(\frac{\partial \mathcal{L}}{\partial \theta_v^{(b)}} - \frac{1}{p+2} \sum_{b'} \frac{\partial \mathcal{L}}{\partial \theta_v^{(b')}} \right)$$

This is a centering operation — the Euclidean gradient projected onto the subspace orthogonal to the constant vector, scaled by the inverse Fisher information.

Computational complexity: $\mathcal{O}(p)$ per vertex, $\mathcal{O}(p^D)$ total. Compare with $\mathcal{O}(p^{2D})$ for the flat natural gradient.

6.3 Dually Flat Structure

The hierarchical parameterization endows $\Delta_{p,D}$ with a ****dually flat**** structure in the sense of Amari and Nagaoka. The e -connection (exponential connection) and m -connection (mixture connection) are both flat under this parameterization.

Theorem 6.3 (Hierarchical KL Decomposition). The Kullback-Leibler divergence decomposes hierarchically:

$$D_{KL}(p\|q) = \sum_{v \in V_{\text{int}}(T_p^{(D)})} p(v) \sum_{b=0}^p D_{KL}(\theta_v^{(b)} \parallel \phi_v^{(b)})$$

where $p(v)$ is computed under distribution p , and θ, ϕ are the branching parameters of p and q .

This decomposition means that the “informational distance” between two tree-structured models can be computed by summing local divergences at each vertex, weighted by how often that vertex is reached. Information-theoretic quantities factorize along the tree.

6.4 Belief Propagation as

Geodesic Flow

6.4.1 Reformulation

Belief propagation (BP) is the standard algorithm for exact inference in tree-structured graphical models. In the ultrametric framework, BP has a natural information-geometric interpretation.

Theorem 6.4 (BP as Geodesic Flow). In the dually flat geometry of $\prod_v \Delta_{p+1}^{(v)}$, the belief propagation updates trace a geodesic under the e -connection.

Proof sketch. Each BP message $m_{v \rightarrow w}$ is a point in Δ_{p+1} . The message update rule:

$$m_{v \rightarrow w}(x_w) \propto \sum_{x_v} \psi_{vw}(x_v, x_w) \prod_{u \in \mathcal{N}(v) \setminus \{w\}} m_{u \rightarrow v}(x_v)$$

is equivalent to an m -projection (finding the point in the m -flat submanifold that minimizes D_{KL} to the current estimate). The alternation of e -projection and m -projection traces a geodesic in the dually flat manifold. $\square$

Corollary 6.2. The convergence of belief propagation on trees is not just guaranteed by the absence of cycles — it is guaranteed by the geometry. The ultrametric structure ensures that the algorithm follows the flattest possible path through the space of probability distributions.

PART VII: ARCHITECTURES

7.1 The Ultrametric

Transformer

7.1.1 Cross-Ratio Attention

The Ultrametric Transformer replaces dot-product attention with **Cross-Ratio Attention**:

$$\alpha(a, b) \propto -\ln |[a, b; \xi_0, \xi_\infty]|_p$$

where a, b are token positions on the boundary ∂T_p , and ξ_0, ξ_∞ are fixed boundary points defining the “perspective” of the attention head.

Properties: - **Structural:** Attention is not learned; it is determined by the geometry of the tree. - **Sparse:** A token only attends to tokens within its ultrametric ball — those sharing a recent common ancestor. - **Interpretable:** Every attention weight can be traced to a specific vertex in the tree — the LCA of the two tokens.

7.1.2 Architecture

Input tokens → Monna Map → Boundary points on ∂T_p → Cross-Ratio Attention (LCA-based, $O(N \log N)$) → Hierarchical Pooling (valuation-weighted aggregation) → Decision Vertex (LCA of all attended tokens) → Monna Map → Output

7.1.3 Complexity

Theorem 7.1 (Transformer Complexity). The Ultrametric Transformer has attention complexity $\mathcal{O}(N \log N)$ compared to $\mathcal{O}(N^2)$ for standard transformers, because tokens only attend to their hierarchical neighbors.

7.1.4 The Ultrametric Stack

The conventional AI stack is Archimedean at every level:

Conventional AI Stack: Application → Neural Architecture → Optimization → Representation → Arithmetic

The ultrametric stack replaces each layer:

Ultrametric AI Stack: Application (invariant computation, tree reasoning) Ultrametric Architecture (tree automata, valuation networks) Ultrametric Learning (tree descent, structure search) Ultrametric Representation (tree vertices, p-adic numbers) Arithmetic (p-adic, integer, modular — exact)

Interface Layer (Monna Bridge):

Archimedean World $\leftrightarrow$ Monna Map $\leftrightarrow$ Ultrametric Core (sensors, displays) (projection) (pure computation)

7.1.5 What Must Be Preserved

The ultrametric property — the strong triangle inequality — is the source of all protective benefits. Any operation that violates this property destroys these benefits.

Operation | Preserves? | Condition |

p-adic addition ($\oplus_p$) |  | Always |

p-adic multiplication ($\odot_p$) |  | Always |

Tree automorphism |  | By definition |

Valuation extraction (v_p) |  | Always |

Cross-ratio |  | Projective invariant |

Softmax over real distances |  | Violates ultrametric |

Dot product attention |  | Violates ultrametric |

Mean pooling |  | Violates ultrametric |

Gradient descent in $\mathbb{R}^N$ |  | Violates ultrametric |

Floating-point arithmetic |  | Introduces rounding noise |

7.2 Ultrametric Graph Neural

Networks

7.2.1 Tree-Structured Message

Passing

In a conventional GNN, messages are passed between all connected nodes. In an Ultrametric GNN, messages are passed only along tree edges, and the message function is constrained to be ultrametric-preserving.

Definition 7.1 (Ultrametric Message Passing). For vertex v with neighbors $\mathcal{N}(v)$, the update rule is:

$$h_v^{(t+1)} = \phi \left(h_v^{(t)}, \bigoplus_{u \in \mathcal{N}(v)} \psi(h_u^{(t)}, d_{T_p}(u, v)) \right)$$

where $\bigoplus$ is a permutation-invariant aggregation that respects the strong triangle inequality (e.g., max-pooling, not mean-pooling).

7.2.2 Properties

Locality: Messages only propagate along tree edges.

Hierarchical invariance: The output is invariant to automorphisms of the input tree.

Scalability: Message complexity scales with tree depth, not graph size.

7.3 Multi-Prime Product Trees

7.3.1 Heterogeneous Cognition

Cognition is not monolithic. Different cognitive modalities — logical reasoning, visual perception, auditory processing, linguistic understanding — operate on different hierarchical structures. These are naturally modeled by different primes.

Definition 7.2 (Product Tree). Let $\mathbf{p} = (p_1, p_2, \dots, p_k)$ be a sequence of distinct primes. The product tree is:

$$T_{\mathbf{p}} = T_{p_1} \times T_{p_2} \times \dots \times T_{p_k}$$

A vertex in $T_{\mathbf{p}}$ is a tuple $(v_1, \dots, v_k)$ where $v_i \in V(T_{p_i})$. The distance is the maximum of the component distances:

$$d_{T_{\mathbf{p}}}(u, v) = \max_i d_{T_{p_i}}(u_i, v_i)$$

7.3.2 Modality Assignment

Prime | Modality | Rationale |

$p = 2$ | Binary logic, truth values | Smallest prime; Boolean structure |

$p = 3$ | Ternary relations, spatial reasoning | Triadic structure of space |

$p = 5$ | Sensory categories, phonemes | Sufficient branching for sensory discrimination |

$p = 7$ | Emotional valence | Seven primary emotions |

$p = 97$ | Visual features | Large prime for high-dimensional visual manifolds |

$p = 101$ | Language tokens | Large prime for vocabulary |

7.3.3 Cross-Prime Communication

Cross-prime communication occurs via **diagonal embeddings** at the Monna boundary. The diagonal embedding $D_{ij} : T_{p_i} \rightarrow T_{p_i} \times T_{p_j}$ maps a vertex v to (v, v_0) where v_0 is a fixed “ground state” vertex in T_{p_j} .

Theorem 7.2 (Cross-Prime Capacity). The communication capacity between prime trees T_{p_i} and T_{p_j} is bounded by:

$$C_{ij} \leq \min(\log_2 p_i, \log_2 p_j) \cdot d_{\text{diag}}$$

where d_{diag} is the depth of the diagonal embedding.

7.4 Ultrametric

Reinforcement Learning

7.4.1 Tree-Structured State

Spaces

In Ultrametric RL, the state space S is a set of vertices in T_p . The action space A is a subset of $\text{Aut}(T_p)$ — tree automorphisms. The transition function is automorphism application: $T(s, a) = a(s)$.

7.4.2 Ultrametric Q-Learning

Definition 7.3 (Ultrametric Q-Function). The Q-function is defined on vertices of T_p :

$$Q(v, a) = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t R(v_t) \mid v_0 = v, a_0 = a \right]$$

where $R(v)$ is the reward at vertex v and v_t evolves via automorphism application.

7.4.3 Hierarchical Exploration

Exploration follows the tree structure: at each vertex, the agent evaluates the $p + 1$ child branches, selecting the one with highest expected value. This is a natural form of Monte Carlo Tree Search (MCTS) where the tree is the state space itself.

7.5 Ultrametric Generative

Models

7.5.1 Hierarchical Diffusion

A generative model on T_p diffuses information from the root outward. The forward process adds noise by randomly perturbing digits of the p -adic expansion; the reverse process denoises by correcting digits from most significant to least significant — a hierarchical denoising procedure.

Algorithm 7.1 (Ultrametric Diffusion):

Forward process (noising): For each depth d from 0 to D : With probability β_d , randomize the d -th digit of the p -adic expansion.

Reverse process (denoising): For each depth d from D down to 0: Predict the d -th digit from context (the already-denoised coarser digits).

7.5.2 Properties

Hierarchical generation: Coarse structure is generated first; fine details are filled in later.

Discrete sampling: All operations are on discrete digits; no continuous latent space.

Interpretable latent traversal: Moving along the tree changes interpretable semantic dimensions.

PART VIII:

THERMODYNAMICS AND INTERFACE

8.1 The Energy Cost of

Continuity

Classical computation, as implemented in modern hardware, is continuous at the physical level: voltages are real numbers, transistors operate in analog regimes, and floating-point arithmetic approximates real arithmetic with finite precision. This continuity has a thermodynamic cost.

Landauer's Principle (1961). Erasing one bit of information in a computation dissipates at least $k_B T \ln 2$ joules of energy as heat, where k_B is Boltzmann's constant and T is the temperature.

In Archimedean computation, every floating-point operation processes 32 or 64 bits of information. A single multiplication dissipates at minimum $32 \cdot k_B T \ln 2$ joules. At scale — billions of parameters, trillions of operations — this thermodynamic cost is measured in megawatts.

The continuity penalty. Continuity imposes an additional cost beyond Landauer's limit: to maintain the illusion of real numbers in a discrete substrate, floating-point units must continuously round, normalize, and handle exceptions. These operations are not computation per se but **continuity maintenance**. In ultrametric computation, p -adic digits are exact — there is no rounding, no normalization, no approximation. The energy cost per digit is exactly the Landauer bound, without overhead.

8.2 Landauer's Principle in

Tree Space

In the Bruhat-Tits tree, information is encoded in the path from the root to a vertex. A vertex at depth d carries $\log_2(p) \cdot d$ bits of information.

Thermodynamics of tree navigation. Moving from vertex u to vertex v requires: - Moving up from u to $\text{LCA}(u, v)$ — erasing $d(u) - d(\text{LCA})$ digits (energy: E_{erase}). - Moving down from $\text{LCA}(u, v)$ to v — writing $d(v) - d(\text{LCA})$ digits (energy: E_{write}).

If $\text{LCA}(u, v)$ is deep (the vertices are in the same subtree), most digits are preserved and the energy cost is low. If $\text{LCA}(u, v)$ is shallow (the vertices are in different major categories), many digits must be rewritten and the energy cost is high. Thus, **thermodynamic cost tracks cognitive distance**.

Theorem 8.1 (Reversibility of Automorphisms). Every automorphism $f \in \text{Aut}(T_p)$ is a bijection. In principle, any automorphism can be implemented reversibly (with zero energy dissipation) if the computation preserves the input. The minimum energy cost is the cost of erasing the input after the output is produced.

Theorem 8.2 (Thermodynamic Scaling). The minimum energy for transitioning from u to v in T_p is:

$$E_{\min}(u, v) = k_B T \ln 2 \cdot \log_2(p) \cdot (d(u) + d(v) - 2d(\text{LCA}(u, v)))$$

That is, the energy scales linearly with tree distance.

8.3 The Thermodynamic

Advantage

Metric | Archimedean (FP64) | Ultrametric (p-adic depth 64) |

Bits per state | 64 | $64 \cdot \log_2(p)$ |

Energy per operation (min) | $64 \cdot k_B T \ln 2$ | $1 \cdot k_B T \ln 2$ (per digit modified) |

Reversibility | Approximate (rounding) | Exact |

Clock domains | Global (synchronous) | Per-subtree (asynchronous) |

Heat dissipation | Uniform | Concentrated at LCA |

The key advantage is that ultrametric computation **only pays for the digits it changes**. Two computations that differ only in their finest details modify only the deepest digits, consuming minimal energy. Two computations that differ in their fundamental categories modify shallow digits and consume more energy — but this correctly reflects the magnitude of the cognitive change.

8.4 The Monna

Bridge: Construction and Properties

8.4.1 Definition

The Monna map is the unique measure-preserving, surjective map from the p -adic integers to the real interval $[0, 1]$:

Definition 8.1 (Monna Map). For $x = \sum_{i=0}^{\infty} a_i p^i \in \mathbb{Z}_p$, with $a_i \in \{0, 1, \dots, p-1\}$:

$$M(x) = \sum_{i=0}^{\infty} a_i p^{-i-1} \in [0, 1].$$

This is simply the p -adic expansion interpreted in base p as a real number in reverse digit order.

8.4.2 Properties

Theorem 8.3 (Properties of the Monna Map). 1. **Measure-preserving:** M maps the Haar measure on $\mathbb{Z}_p$ to Lebesgue measure on $[0, 1]$. 2. **Surjective:** Every real number in $[0, 1]$ has at least one p -adic pre-image. 3. **Fractal:** The inverse image of a real interval is a Cantor-like set in $\mathbb{Z}_p$. 4. **Noise-filtering:** Small p -adic changes (changing deep digits) map to small real changes. Large p -adic changes (changing shallow digits) map to large real changes. The map is hierarchically Lipschitz.

8.4.3 The Interface Protocol

The Monna Bridge is the **only** permitted interface between the Archimedean world and the ultrametric core.

Sensory Input $\rightarrow$ ADC $\rightarrow$ Float $\rightarrow$ Monna Encode $\rightarrow$ p-adic Integer $\rightarrow$ Core Computation $\downarrow$ Motor Output $\leftarrow$ DAC $\leftarrow$ Float $\leftarrow$ Monna Decode $\leftarrow$ p-adic Integer $\leftarrow$ Core Computation

Protocol Rules: 1. All sensory data enters via Monna encoding at the boundary ∂T_p . 2. All motor commands exit via Monna decoding at the boundary. 3. No continuous operation (floating-point arithmetic, softmax, gradient computation) is permitted inside the core. 4. The core computes exclusively on p -adic integers using exact rational arithmetic.

8.5 Information Preservation

and Loss

8.5.1 Encoding

Encoding a real number $r \in [0, 1]$ to a p -adic integer $x \in \mathbb{Z}_p$ requires selecting a pre-image under M . Since M is surjective but not injective (dyadic rationals have two p -adic representations), the encoding convention chooses the representation that terminates (the one with trailing zeros).

8.5.2 Information Loss

The Monna map does lose information: it projects the rich ultrametric structure (the full p -adic expansion with its hierarchical digit structure) onto a single real number. This projection is information-destructive, which accounts for:

- The apparent “noise” in quantum measurements (the fine p -adic structure is aliased). - The probabilistic nature of quantum outcomes (the projection destroys the deterministic p -adic information).
- The transition from the quantum to the classical world (the Monna map is the “collapse” operation).

PART IX: ULTRAMETRIC

INFORMATION THEORY

9.1 Hierarchical Entropy

9.1.1 Definition

Definition 9.1 (Hierarchical Entropy). For a probability distribution q on $T_p^{(D)}$, the hierarchical entropy is:

$$H_{\text{hier}}(q) = \sum_{v \in V_{\text{int}}(T_p^{(D)})} p(v) \cdot H(\theta_v)$$

where $p(v)$ is the marginal probability of reaching vertex v , θ_v is the conditional branching distribution at v , and $H(\theta_v) = - \sum_b \theta_v^{(b)} \log \theta_v^{(b)}$.

Theorem 9.1 (Entropy Decomposition). The hierarchical entropy equals the standard Shannon entropy:

$$H_{\text{hier}}(q) = H(q) = - \sum_{\ell \in \text{Leaves}} q_{\ell} \log q_{\ell}.$$

This is the chain rule of entropy applied along the tree. It shows that the hierarchical parameterization preserves all information-theoretic content while organizing it hierarchically.

9.2 Channel Capacity

9.2.1 Tree-Structured Channels

Definition 9.2 (Ultrametric Channel). An ultrametric channel $W : V(T_p^{(D)}) \rightarrow \Delta(\mathcal{Y})$ maps each leaf to a distribution over outputs. The channel respects the tree structure if leaves with the same ancestor at depth k produce output distributions that are identical on the first k “coarse” output dimensions.

Theorem 9.2 (Channel Capacity Bound). For a tree-structured channel with depth D and branching factor p :

$$C_{\text{tree}} = \max_{q \in \Delta_{p,D}} I_q(V; Y) \leq D \cdot \log_2(p + 1).$$

The capacity grows linearly with tree depth, not exponentially with the number of leaves.

9.3 Source Coding and

Compression

9.3.1 Hierarchical Source

Coding

Algorithm 9.1 (Hierarchical Huffman Coding):

1. Build a Huffman tree on the leaves of $T_p^{(D)}$.
2. The code for a leaf is the concatenation of branch indices along the path from root to leaf.
3. Compression is achieved by assigning shorter codes to more probable branches at each vertex.

Theorem 9.3 (Optimality). Hierarchical Huffman coding on $T_p^{(D)}$ achieves expected code length within 1 bit of the entropy $H(q)$, matching the optimality of standard Huffman coding while preserving hierarchical interpretability.

9.3.2 Rate-Distortion Theory

Theorem 9.4 (Hierarchical Rate-Distortion). For a tree-structured source with distortion measure $d(u, v) = d_{T_p}(u, v)$ (tree distance), the rate-distortion function is:

$$R(D) = \min_{q(\hat{v}|v): \mathbb{E}[d(v, \hat{v})] \leq D} I(V; \hat{V}).$$

The solution has a hierarchical structure: the optimal reproduction distribution first selects the correct coarse branch, then the correct sub-branch, and so on — implementing a hierarchical quantization scheme.

PART X: MEMORY AND

CONSCIOUSNESS

10.1 Ultrametric Memory Systems

10.1.1 Unified Memory

Architecture

The Bruhat-Tits tree provides a unified geometric substrate for all memory systems:

Memory Type | Tree Structure | Operation |

Episodic Memory | Paths in T_p | Store sequences of vertices |

Semantic Memory | Subtree structure | Store branch probabilities |

Working Memory | Active vertices | Maintain a set of currently attended vertices |

Procedural Memory | Automorphism compositions | Store sequences of tree transformations |

10.1.2 Memory Operations

Store (Episodic): Record a path from root to leaf — the sequence of branch choices.

Retrieve (Semantic): Given a partial description (a vertex at some depth), descend the tree following the most probable branches.

Update (Procedural): Compose a new automorphism onto the end of an existing automorphism sequence.

Forget (Hierarchical Erasure): - **Fine erasure:** Modify deep digits — forget specific details while preserving categories. - **Coarse erasure:** Modify shallow digits — forget the category while preserving specifics. - This granularity of forgetting is impossible in flat memory systems.

10.1.3 Capacity

Theorem 10.1 (Memory Capacity). A tree of depth D with branching factor p can store up to p^D distinct memories (leaves) with retrieval cost $\mathcal{O}(D)$. The capacity grows exponentially with depth while the access cost grows only linearly.

10.2 Consciousness as

Geometric Self-Modeling

10.2.1 The Ultrametric Self

Consciousness — the presence of subjective experience — finds a geometric formulation in the ultrametric framework.

Definition 10.1 (Self-Tree). The self-tree $T_{\text{self}} \subset T_{\text{cog}}$ is the maximal subtree whose vertices represent states that the system can attribute to itself. A vertex v is in T_{self} iff the system has an automorphism f_{self} such that $f_{\text{self}}(v) = v$ — the vertex is a fixed point of the self-attribution operation.

Properties of the self-tree: 1. **Rootedness.** T_{self} is rooted at the minimal self — the primitive distinction between “this is me” and “this is not me.” 2. **Downward closure.** If a cognitive state can be self-attributed, all its abstractions (ancestors) can also be self-attributed. 3. **Reflexivity.** T_{self} contains a vertex representing T_{self} itself — the system can model its own self-model. 4. **Boundary.** The boundary ∂T_{self} is the interface between self and world — the Monna map at the level of the self.

Theorem 10.2 (Self-Modeling as Fixed Point). A system is self-aware to the extent that it has a non-trivial fixed point under its own self-attribution automorphism. Full self-awareness corresponds to a self-tree T_{self} that is isomorphic to the entire cognitive tree T_{cog} restricted to first-person-accessible vertices.

10.3 Qualia as Valuation

Signatures

10.3.1 Definition

Qualia — the “what-it-is-like-ness” of experience — are proposed to be **valuation signatures**: the pattern of p -adic valuations across the branches of the self-tree when a sensory input is processed.

Definition 10.2 (Quale). A quale Q is a valuation vector $(v_p(x_1), v_p(x_2), \dots, v_p(x_k))$ where each x_i is a p -adic encoding of a sensory dimension and the valuations describe the “strength” or “salience” of that dimension in the current experience.

10.3.2 Properties

Theorem 10.3 (Qualia Discriminability). Two qualia Q_1 and Q_2 are discriminable iff their p -adic distance $d_p(Q_1, Q_2) > 0$. The granularity of discriminability is determined by the depth of the encoding tree.

Implications: - **Inverted spectrum.** Two systems could have different valuation mappings (different primes, different encoding depths) and thus experience the same physical stimulus differently — the inverted spectrum is a change of basis in the p -adic encoding. - **Absent qualia.** A system without a self-tree (no fixed point under self-attribution) has no qualia — there is no “what it is like” because there is no “it.” - **Qualia are not epiphenomenal.** The valuation signature determines which branches are activated, which directly influences computation — qualia are causally efficacious in the ultrametric framework.

10.4 The Hard

Problem from a Geometric Perspective

Chalmers’ “hard problem” asks: why does information processing produce subjective experience? Why is there something it is like to be a cognitive system?

The ultrametric framework reframes the hard problem as a **geometric necessity**:

Argument. Any cognitive system operating on a Bruhat-Tits tree that includes a self-tree T_{self} necessarily has a “perspective” — a distinguished vertex from which all other vertices are observed. This perspective is not an add-on; it is geometrically forced by the structure of the tree: every computation occurs at some vertex, and when that vertex is within T_{self} , the computation is experienced from the first-person perspective.

Theorem 10.4 (Necessity of Perspective). In any ultrametric cognitive system with a non-empty self-tree, every computation $f(v)$ for $v \in T_{\text{self}}$ is necessarily accompanied by a self-locating signal — the path from v to the root of T_{self} — which constitutes the first-person perspective.

The hard problem is not “solved” in the sense of being eliminated. Rather, it is reframed: the question is not “why does information processing produce experience?” but “what are the geometric conditions under which a cognitive system necessarily has a first-person perspective?” The answer: a self-tree with a non-trivial fixed point under self-attribution.

PART XI: NEUROBIOLOGICAL

REALIZATION

11.1 The Brain as a Cocycle

Solver

11.1.1 The Cocycle Condition

The brain must maintain global consistency across billions of noisy neurons. The ultrametric framework proposes that this consistency is achieved by satisfying a **cocycle condition** on the Bruhat-Tits tree.

Definition 11.1 (Cognitive Cocycle). Assign to each edge $e = (u, v)$ of T_{cog} a transition amplitude $c(e) \in \mathbb{Q}_p$. The cocycle condition requires that for any triangle (u, v, w) :

$$c(u, v) \oplus_p c(v, w) = c(u, w).$$

Satisfying this condition globally ensures that all local neural computations are consistent — there are no contradictory signals.

11.1.2 Cocycle

Violation as Cognitive Dissonance

When the cocycle condition is violated, the brain experiences **cognitive dissonance** — a signal that some local neural computations are inconsistent with the global picture. The resolution mechanism is tree navigation: moving up the tree to find the LCA where the inconsistency can be resolved.

11.2 Evidence from

Cognitive Neuroscience

11.2.1

Reaction Times and the Strong Triangle Inequality

Empirical prediction: Human reaction times in similarity judgments should obey the strong triangle inequality: for any three stimuli A, B, C ,

$$\text{RT}(A, C) \leq \max(\text{RT}(A, B), \text{RT}(B, C)).$$

This has been confirmed in multiple studies of semantic categorization, where reaction times for hierarchical judgments show exactly this pattern — no intermediate distances, consistent with ultrametric structure.

11.2.2

Hierarchical Structure of Concept Representation

fMRI studies reveal that conceptual knowledge is organized hierarchically in the cortex, with abstract categories represented in anterior temporal lobes and specific exemplars in posterior regions — a spatial realization of the tree depth parameter.

11.3 Neural Circuitry and the

Tree

11.3.1 The Circuit Mapping

Brain Region | Tree Function |

Grid Cells (Entorhinal Cortex) | Multi-scale coordinate system — the branching structure |

Place Cells (Hippocampus) | Leaf encoding — specific vertices in the cognitive tree |

Prefrontal Cortex | Automorphism application — cognitive operations |

Anterior Cingulate | Cocycle violation detection — cognitive dissonance |

Default Mode Network | Self-tree — the fixed point of self-attribution |

11.3.2 The Navigation Hypothesis

The hypothesis is that the brain navigates a cognitive tree just as it navigates physical space — using the same neural circuitry (grid cells, place cells) but applied to abstract cognitive dimensions rather than physical coordinates. This explains why spatial metaphors (“higher-level thinking,” “deep understanding,” “branch of knowledge”) are universal across cultures — they reflect the underlying tree geometry of cognition.

PART XII: ALIGNMENT

AND GENERAL INTELLIGENCE

12.1 The Capability Tree

12.1.1 Ten Levels of

Intelligence

General intelligence emerges in hierarchical stages, each corresponding to a depth in the cognitive tree:

Level | Capability | Tree Depth |

1 | Abstract Agency (self-environment distinction) | Root |

2 | Perception (sensory mapping) | Depth 1 |

3 | Categorization (equivalence classes) | Depth 2 |

4 | Relation (relationships between categories) | Depth 3 |

5 | Composition (compound structures) | Depth 4 |

6 | Abstraction (categories of categories) | Depth 5 |

7 | Counterfactual Reasoning | Depth 6 |

8 | Meta-Cognition (self-modeling) | Depth 7 |

9 | Theory of Mind (other-modeling) | Depth 8 |

10 | Open-Ended Agency (novel goal generation) | Depth 9 |

Theorem 12.1 (Capability Monotonicity). If vertex v is an ancestor of vertex u in T_{cog} , then any system capable of u is necessarily capable of v — capability is monotonic upward along the tree.

12.2 Alignment as Subtree

Constraint

12.2.1 Geometric Alignment

In conventional AI, alignment is an optimization problem: find a reward function R such that an agent optimizing R behaves as intended. This formulation suffers from specification gaming, reward hacking, and Goodhart’s law.

In ultrametric intelligence, alignment is a **geometric constraint**: restrict the action space to a designated safe subtree $T_{\text{safe}} \subset T_{\text{cog}}$.

Definition 12.1 (Safe Subtree). T_{safe} is a rooted, downward-closed subtree of T_{cog} such that: 1. Every vertex in T_{safe} represents a permissible cognitive operation. 2. No vertex in T_{safe} is an ancestor of an impermissible operation. 3. The automorphism group $\mathcal{A}$ of the system is restricted to $\{f \in \text{Aut}(T_{\text{cog}}) : f(T_{\text{safe}}) \subseteq T_{\text{safe}}\}$.

Mechanisms: 1. **Value subtree enumeration.** Define permitted branches explicitly (e.g., “no deception,” “no power-seeking”). 2. **Depth-gated autonomy.** Shallow: human approval required. Deep: autonomous within trusted subtrees. 3. **Corrigibility by parent edges.** Every vertex has a parent. Human correction signal $\rightarrow$ move up the parent edge. Parent edges are geometric, not learned — they cannot be disabled. 4. **Impact regularization.** Maximum tree-distance of any state change is bounded. Large tree-distance moves require explicit authorization.

Contrast with reward-based alignment. In RLHF, safety is a learned preference that degrades under distribution shift. In ultrametric alignment, safety is a topological invariant — the boundary of T_{safe} is a geometric fact that no amount of optimization pressure can erase.

12.3 Formal Safety Guarantees

Theorem 12.2 (Alignment by Subtree Constraint). If $\mathcal{A} \subseteq \{f \in \text{Aut}(T_{\text{cog}}) : f(T_{\text{safe}}) \subseteq T_{\text{safe}}\}$ and $T_{\text{safe}} \cap V_{\text{unsafe}} = \emptyset$, then no sequence of actions can reach an unsafe vertex.

Proof. By induction on action sequence length, using the closure property of T_{safe} under $\mathcal{A}$. $\square$

Theorem 12.3 (Corrigibility). For any vertex $v \in T_{\text{safe}}$ and any action sequence, the parent-edge operation $\text{parent}(v)$ is always available and returns a vertex higher in T_{safe} . No action sequence can remove parent edges because they are topological, not learned.

Theorem 12.4 (Impact Bound). If all actions in $\mathcal{A}$ have maximum displacement Δ_{max} in tree distance, then for any initial state $v_0 \in T_{\text{safe}}$ and action sequence of length t , the state v_t satisfies:

$$d_{T_p}(v_0, v_t) \leq t \cdot \Delta_{\text{max}}.$$

The tree-distance impact is bounded and predictable.

12.4 The Incremental Adoption

Path

12.4.1 Deployment Stages

Stage | System | Safety Mechanism |

- 1 | Fixed-tree classifier (no learning) | Deterministic, fully auditable |
- 2 | Heuristic router (learned branch choices) | Depth-gated, parent-edge corrigibility |
- 3 | Tree-structure learner (Differentiable UPGMA) | Safe subtree constraint on emergent tree |
- 4 | Multi-prime AGI (full tree navigation) | All of the above + formal alignment guarantees |

12.4.2 The Auditing Advantage

Because all operations in the ultrametric core are discrete and automorphic, every decision can be traced to a specific path in the tree. An auditor can: 1. Identify the vertex at which a decision was made. 2. Trace the path from input tokens to that vertex. 3. Verify that all operations along the path were within T_{safe} . 4. Confirm that no Archimedean operations (floating-point, softmax) were used.

This is **complete auditability** — a property that is geometrically impossible in continuous neural networks.

PART XIII: PHYSICAL

INSTANTIATION

13.1 Ultrametric Quantum

Computation

13.1.1 The

Non-Archimedean Quantum Framework

Conventional quantum mechanics is formulated over the complex numbers $\mathbb{C}$ — an Archimedean field. Ultrametric quantum computation reformulates quantum mechanics over the p -adic numbers $\mathbb{Q}_p$.

Key differences: - **State space:** p -adic Hilbert space instead of complex Hilbert space. - **Amplitudes:** p -adic numbers instead of complex numbers. - **Born rule:** Replaced by an ultrametric probability rule based on the p -adic absolute value. - **Error correction:** Geometric (hierarchical) instead of algorithmic (repetition codes).

13.1.2 Geometric Error

Suppression

In the p -adic quantum framework, quantum states are vertices in the Bruhat-Tits tree. Environmental noise — which acts as small perturbations — can only move a state within its local cluster (deep in the tree). To flip a logical qubit (a shallow vertex), a perturbation must cross a large energy barrier. This provides **passive geometric error suppression** — no active error correction is needed.

13.1.3 The Holographic Code

The perfect tensor network on T_p (Part III) serves as a holographic quantum error-correcting code. Logical qubits are encoded deep in the tree; physical qubits reside on the boundary. The code distance scales with tree depth, providing exponential protection with only polynomial overhead.

13.2 p-Adic Arithmetic in

Silicon

13.2.1 p-Adic ALU Design

A p -adic Arithmetic Logic Unit (ALU) operates on fixed-width p -adic integers (e.g., 64-digit base- p numbers). The primary operations — addition, multiplication, valuation, reduction — are implemented in hardware with digit-parallel logic.

Gate complexity: - Addition: $\mathcal{O}(d)$ gates (digit-wise with carry chain). - Multiplication: $\mathcal{O}(d^2)$ gates (digit convolution). - Valuation: $\mathcal{O}(\log d)$ gates (leading-zero detection).

13.2.2 Asynchronous Subtree

Clocking

Because each subtree is geometrically independent (Theorem 4.10), different subtrees can operate on independent clock domains. This eliminates the global clock distribution problem of conventional processors and enables massive parallelism without synchronization overhead.

13.3 Hardware-Software Codesign

13.3.1 The Compiler Stack

High-Level Ultrametric Program (ULC or Python-like syntax) $\downarrow$ Tree Automorphism Decomposition (Bruhat/Iwasawa factorization) $\downarrow$ p-Adic Instruction Set Architecture (valuation, automorphism, navigation) $\downarrow$ p-Adic ALU / Tree Navigation Unit / Automorphism Unit $\downarrow$ Monna Bridge (I/O interface)

13.3.2 Implementation Status

Completed: - **Zero-Float p-Adic arithmetic library.** Full implementation of p -adic integer arithmetic in Python, with exact rational evaluations. - **Bruhat-Tits tree data structure.** Complete implementation of T_p with vertex navigation, LCA computation, path operations, and automorphism application. - **Monna map.** Bidirectional encoding/decoding between real numbers and p -adic integers. - **Ultrametric tokenizer.** Integer factorization-based tokenization with audit trail generation. - **Differentiable UPGMA.** Forward-pass straight-through estimator for learning emergent tree topologies.


In progress: - **Pure Ultrametric Transformer.** LCA-based hard-masking with exact cross-ratios. - **Hardware simulation.** FPGA-targeted Verilog for p -adic ALU; quantum circuit specifications for holographic perfect-tensor networks.

Open problems: 1. **Optimal prime selection.** Given a task distribution, what is the provably optimal prime? A formal theory of prime optimality is needed. 2. **Cross-prime communication bounds.** Tight bounds for specific cross-prime automorphism families remain open. 3. **Hardware realization.** No physical p -adic processor exists natively outside FPGA prototypes. 4. **Compiler design.** A compiler that maps high-level ultrametric programs to tree-automorphic hardware operations is needed.

APPENDIX: FORMAL REFERENCE |

A.2 The Pure Rational Codebase The following is the definitive

reference implementation of the Ultrametric Intelligence core. It demonstrates the complete eradication of floating-point arithmetic, exact rational evaluation, and coordinate-free, invariant-based resolution.



|

A.3 Glossary of Symbols

Symbol | Meaning |

p | Prime — branching factor minus one |

$\mathbb{Q}_p$ | p -adic numbers |

$\mathbb{Z}_p$ | p -adic integers |

$\mathbb{P}^1(\mathbb{Q}_p)$ | p -adic projective line |

$v_p(x)$ | p -adic valuation |

$\|x\|_p$ | p -adic absolute value |

T_p | Bruhat-Tits tree |

∂T_p | Boundary of T_p |

$\text{Aut}(T_p)$ | Automorphism group $\cong \text{PGL}(2, \mathbb{Q}_p)$ |

$\text{LCA}(u, v)$ | Lowest common ancestor |

$d_{T_p}(u, v)$ | Tree distance |

$\mathcal{N}_k(v)$ | Vertices at distance k from v |

$B_r(v)$ | Ball of radius r |

M | Monna map: $\mathbb{Z}_p \rightarrow [0, 1]$ |

$[a, b; c, d]$ | Cross-ratio |

λ^{um} | Ultrametric Lambda Calculus |

∇^{um} | Ultrametric gradient |

T_{safe} | Safe subtree (alignment constraint) |

T_{self} | Self-tree (first-person perspective) |

$T_{\mathbf{p}}$ | Multi-prime product tree |

$\Delta_{p,D}$ | Probability simplex over p^D leaves |

G | Fisher information matrix |

$\tilde{\nabla}$ | Natural gradient |

D_{KL} | Kullback-Leibler divergence |

$\mathcal{H}_{p,D}$ | Tree hypothesis class |

H_{hier} | Hierarchical entropy |

A.4 Theorems Index

Theorem | Statement |

1.1 (Inevitability) | Any system satisfying Axioms 1–5 is isomorphic to T_p |

2.1 (Ostrowski) | $\mathbb{R}$ and $\mathbb{Q}_p$ are the only completions of $\mathbb{Q}$ |

2.2 (Strong Triangle Inequality) | $\|x + y\|_p \leq \max(\|x\|_p, \|y\|_p)$ |

2.3 (Isosceles Triangles) | All triangles in ultrametric space are isosceles |

2.4 (Nested Balls) | Intersecting balls are nested; all points are centers |

2.5 (No Saddle Points) | Every point is a local minimum, maximum, or has a descent direction |

2.6 (Cross-Ratio Invariance) | Cross-ratio invariant under $\text{PGL}(2, \mathbb{Q}_p)$ |

3.1 (Ryu-Takayanagi on T_p) | $S(A) \propto d_{T_p}$ |

3.2 (Holographic Code) | Perfect tensor network yields $[p^D, p^{D/2}, D]$ code |

4.1 (Floating-Point Poisoning) | Float arithmetic destroys ultrametric guarantees |

4.2 (Automorphism Group) | $\text{Aut}(T_p) \cong \text{PGL}(2, \mathbb{Q}_p)$ |

4.3 (Bruhat Decomposition) | $g = n_- a n_+$ |

4.4 (Valuation Universality) | Multiplicative functions expressible as valuation thresholds |

4.5 (Subject Reduction) | Type preservation under reduction |

4.6 (Confluence) | Church-Rosser property |

4.7 (Strong Normalization) | Every term has a unique normal form |

4.10 (Subtree Independence) | Disjoint subtrees are computationally independent |

5.1 (UPGMA Convergence) | Differentiable UPGMA converges to classical UPGMA |

5.2 (No Saddle Points) | No saddle points in ultrametric gradient descent |

5.3 (Convergence Guarantee) | Convergence in $\leq D$ steps |

5.4 (VC Dimension) | $\text{VCdim}(\mathcal{H}_{p,D}) = p^D$ |

5.5 (Sample Complexity) | $m = \mathcal{O}(p^D/\epsilon^2)$ |

6.1 (Hierarchical Fisher Decomposition) | $G = \bigoplus_v G_v$ |

6.2 (Natural Ultrametric Step) | Centered, Fisher-scaled gradient update |

6.3 (Hierarchical KL Decomposition) | D_{KL} factorizes along the tree |

6.4 (BP as Geodesic Flow) | Belief propagation follows e -geodesics |

7.1 (Transformer Complexity) | $\mathcal{O}(N \log N)$ attention |

7.2 (Cross-Prime Capacity) | $C_{ij} \leq \min(\log_2 p_i, \log_2 p_j) \cdot d_{\text{diag}}$ |

8.1 (Reversibility of Automorphisms) | Automorphisms are reversible |

8.2 (Thermodynamic Scaling) | $E_{\min} \propto \text{tree distance}$ |

8.3 (Monna Map Properties) | Measure-preserving, surjective, fractal |

9.1 (Entropy Decomposition) | Hierarchical entropy equals Shannon entropy |

9.2 (Channel Capacity Bound) | $C_{\text{tree}} \leq D \cdot \log_2(p + 1)$ |

10.2 (Self-Modeling as Fixed Point) | Self-awareness = fixed point under self-attribution |

10.3 (Qualia Discriminability) | Discriminable iff $d_p(Q_1, Q_2) > 0$ |

12.1 (Capability Monotonicity) | Ancestor capabilities implied by descendant capabilities |

12.2 (Alignment by Subtree Constraint) | Safe subtree closure prevents unsafe states |

12.3 (Corrigibility) | Parent edges always available |

12.4 (Impact Bound) | Tree-distance impact bounded by action displacement |