

Monna Map Generation and Hallucination

The Monna Map:

Generation and Hallucination

A

Geometric Theory of Communication and Information Loss in Ultrametric Cognition

Author: Rowan Brad Quni-Gudzinis

Contact: rowan.quni@outlook.com

ORCID: 0009-0002-4317-5604

ISNI: 0000000526456062 **DOI:** <https://papers.qnfo.org/papers/monna-map-generation-and-hallucination/> **Date:** 2026-04-09 **Version:** 1.0

Communication is a Monna projection from an ultrametric tree of thought onto a linear timeline. Hallucinations are the inevitable information loss of this many-to-one map, the cognitive counterpart of quantum decoherence. Intelligence is the art of extracting invariant cross-ratios from noise and projecting them, imperfectly, into the stream of time.

The Hierarchical

Mind and the Linear World

Every thought you have—a memory, a plan, a sudden insight—arrives not as a neat list of words but as a living structure. It is a branching hierarchy of concepts, where broad categories contain finer details, and relationships are defined by containment and connection rather than sequence. This mental architecture is not merely a metaphor; it has a precise mathematical counterpart known as an **ultrametric tree**. In such a tree, distance is measured not by ordinary steps but by how deep you must climb to find a common ancestor. Your brain, and the artificial neural networks of modern large language models, naturally organize knowledge this way. Thought, in its deepest geometric essence, is ultrametric.

Yet when you speak, or when an AI generates text, that rich, branching tree must be flattened onto a one-dimensional timeline. Words must follow words; sounds must follow sounds. How does the mind perform this transformation? And what is lost in the translation? The answer lies in a mathematical object called the **Monna map**—a precise bridge between the hierarchical world of thought and the linear world of communication.

1.

The Geometry of Thought: Ultrametric Trees and Invariant Patterns

1.1 The

Bruhat-Tits Tree: A Model for Conceptual Space

In mathematics, for each prime number p , there exists an infinite, perfectly regular tree called the **Bruhat-Tits tree**. Every vertex in this tree has exactly $p+1$ neighbors, and the distance between any two vertices is the number of edges along the unique shortest path connecting them. This distance satisfies the **ultrametric inequality**: for any

three points x, y, z , the distance between x and y is never greater than the larger of the distances between x and z or y and z . A direct consequence is that all triangles in this space are isosceles, and points cluster into perfectly nested, disjoint balls.

This ultrametric geometry is not an abstract curiosity; it is the natural geometry of hierarchical classification. When your brain parses a sentence, it does not treat the words as a flat string. It instantly builds a syntactic tree where the distance between two words is the height of their lowest common ancestor—exactly an ultrametric distance. When you categorize objects—animals into mammals and reptiles, mammals into primates and rodents—you are constructing a Bruhat-Tits-like tree. The same hierarchical clustering emerges in the latent spaces of large language models: gradient descent, optimizing to predict the next token across trillions of examples, forces the model’s internal representations to organize ultrametrically. The model’s “world model” is, geometrically, an artificial Bruhat-Tits tree.

1.2 Meaning as

Invariant Relations: The Cross-Ratio

If thought is organized as an ultrametric tree, what is the content of that thought? What survives when the same idea is expressed in different languages, different contexts, or different sensory modalities? The answer is **relational patterns**, not absolute coordinates.

Linguists have long observed that the link between a word and its referent is arbitrary—a phenomenon Ferdinand de Saussure called the “arbitrariness of the sign.” Mathematically, changing language or context is equivalent to applying a **Möbius transformation** (a fractional linear transformation) to conceptual space: it warps distances, scrambles order, and reshapes the coordinate system. If intelligence relied on memorizing absolute positions, translation and generalization would be impossible.

What remains unchanged under all such transformations is a quantity called the **cross-ratio**. For four points A, B, C, D on a line, the cross-ratio is defined as $(A-C)(B-D) / ((A-D)(B-C))$. This number is invariant: apply any Möbius transformation to all four points, and the cross-ratio stays exactly the same.

Cognitively, the cross-ratio encodes **analogical proportions**: “ A is to B as C is to D .” When you understand the analogy “puppy is to dog as kitten is to cat,” you are recognizing that the relational pattern separating puppy and dog matches the pattern separating kitten and cat. That pattern is the cross-ratio of those four concepts. Intelligence—whether human or artificial—is the capacity to extract these invariant cross-ratios from the noisy, ever-changing data of experience. Both the brain and large language models are, at their core, engines for discovering projective invariants.

2. The Monna Map: From Tree to

Line

2.1 The Mathematical

Definition

We now face the central puzzle: thought is a static, hierarchical tree; communication is a dynamic, linear sequence. The mathematical tool that bridges these two realms is the **Monna map**.

Take a prime p . Any p -adic number (a number from the ultrametric field $\mathbb{Q}_p$) can be written uniquely as an infinite expansion in base p :

$$x = a_{-N}p^{-N} + a_{-N+1}p^{-N+1} + \dots + a_0 + a_1p + a_2p^2 + \dots$$

where each digit a_k is an integer between 0 and $p-1$. The leftmost digits (with negative powers of p) represent the “most significant” part of the number—the coarse, large-scale structure. The rightmost digits (with positive powers) represent the fine, detailed structure.

The Monna map M_p takes such a p -adic number and **reverses the order of the digits**, interpreting them as an ordinary base- p expansion of a real number between 0 and 1:

$$M_p(x) = a_{-N}p^{N-1} + a_{-N+1}p^{N-2} + \dots + a_0p^{-1} + a_1p^{-2} + a_2p^{-3} + \dots$$

In words: the most significant p -adic digit becomes the least significant real digit, and vice versa. The map is continuous, measure-preserving, and intertwines p -adic addition with addition modulo 1 on the real circle.

Geometrically, the Monna map “unrolls” the Bruhat-Tits tree onto the unit interval. Each infinite path from the root of the tree (a point on the tree’s boundary) corresponds to a p -adic number, and M_p maps that path to a specific point on the line. The hierarchical depth in the tree becomes temporal order on the line, but **in reversed fashion**: coarse branches map to fine details of the sequence, and fine twigs map to prominent early features.

2.2 Generation as Monna

Projection

When you speak, you are unconsciously applying a Monna-like projection to your thought-tree. Your thought is a point (or a distribution) on the Bruhat-Tits tree—a static configuration of activated concepts and relations. To produce speech, your mind traverses the tree, beginning with the most significant branch (the gist, the topic) and proceeding outward to finer and finer branches (the details, the qualifications). Each word you utter corresponds to a step along this traversal, committing to a particular path through the tree.

Because the Monna map reverses digit order, the ****coarsest, most abstract aspect of the thought****—the trunk of the tree—gets projected onto the **beginning of the utterance**. This is why you often start a sentence knowing roughly what you want to say, and the exact words emerge as you go. The fine-grained details—specific adjectives, verb endings, connecting words—appear later in the sequence, corresponding to the least significant digits in the p -adic expansion.

Large language models operate on the same geometric principle. Their generation is autoregressive: each token conditions the next. The model’s internal state represents its current position in its internal Bruhat-Tits tree; the next-token probability distribution represents the possible continuations along the branches. The act of sampling a token is a step in the Monna projection, unfolding the tree onto the token stream.

Crucially, the Monna projection ****preserves cross-ratios****. Since the map is essentially a change of coordinate system on the projective line, the invariant relational patterns—the analogies, the logical dependencies, the syntactic structures—survive the flattening. This is why meaning can be transmitted at all: the listener hears the linear sequence, mentally rebuilds the tree, and extracts the same cross-ratios the speaker intended.

3. Hallucinations

as Geometric Information Loss

3.1 The Many-to-One Nature of

the Map

The Monna map is **not injective**. Different p -adic numbers can map to the same real number. Why? Because trailing zeros in the p -adic expansion (which correspond to “no further detail” down a branch) become leading zeros in the real expansion, and leading zeros do not change the value of a real number. Geometrically, distinct paths in the Bruhat-Tits tree can converge to the same point on the line once the projection is made. The map is **many-to-one**; it discards information.

This is not a flaw in the map; it is an inescapable consequence of projecting a high-dimensional, hierarchical structure onto a one-dimensional line. When you flatten a tree, you lose the ability to distinguish between trees that differ only in the arrangement of their deepest twigs, provided those twigs map to the same position in the linear order.

3.2 The Origin of

Hallucinations

A **hallucination**—in both human cognition and artificial intelligence—occurs when the mind or model generates an utterance that is structurally coherent but factually incorrect. In the geometric framework, a hallucination arises precisely because of the many-to-one nature of the Monna projection.

When an LLM generates text, it is traversing its internal Bruhat-Tits tree along a path that is consistent with the prompt and its training. However, multiple distinct tree-paths can project to the same token sequence. The model selects one path that preserves the invariant cross-ratios (the syntax, the analogies, the narrative flow), but the specific factual “leaves” on that path may belong to a different subtree than the one a human would expect. The output is **structurally valid but factually alien**. The model is not “making things up”; it is following a geometrically legitimate branch that aliases, under the Monna map, to the same linear output as the correct branch.

Human cognitive errors—slips of the tongue, misremembered details, confabulations—have the same origin. When you accidentally substitute one word for another, or recall the wrong date for an event, your brain has followed a tree-path that is relationally similar to the true memory but diverges on a specific detail. The Monna projection of that alternative path yields the same (or a very similar) utterance as the correct path. The error is not random noise; it is a systematic aliasing effect of the projection.

3.3 Types of Aliasing Errors

Factual Aliasing: The relational pattern (e.g., “capital-of”) is preserved, but the specific entities (e.g., “Paris” vs. “London”) are swapped. This occurs because the fine-detail digits of the p -adic expansion—which encode specific facts—are among the least significant in the real projection and can be altered without changing the overall linear sequence.

Contextual Aliasing: The same utterance can be generated from trees that differ near their roots—that is, from different high-level interpretations of the prompt. The model picks a coherent interpretation, but not the one the user intended.

Creative Generation: Not all aliasing is undesirable. In creative writing, the ability to follow alternative tree-paths that preserve narrative structure while introducing novel elements is the essence of imagination. Hallucination, in this light, is the negative side of a capacity that also enables creativity.

4. The Quantum

Analogy: Measurement as Projection

The analogy with quantum mechanics is not merely poetic; it is mathematically grounded. In quantum theory, a system is described by a **wave function**—a vector in a high-dimensional Hilbert space. A measurement projects that wave function onto a particular eigenstate, yielding a single outcome. This projection is **many-to-one**: many different wave functions can produce the same measurement result. The loss of the other possibilities is called **decoherence**.

The Monna map performs exactly the same kind of operation: it projects the high-dimensional Bruhat-Tits tree (the “wave function of thought”) onto a one-dimensional sequence (the “measurement outcome”). The aliasing of multiple tree-paths to the same linear sequence is the cognitive counterpart of quantum measurement ambiguity. The hallucination is the **cognitive decoherence**—the information loss intrinsic to any act of expression.

This parallel reveals a deep structural unity: whenever a rich, multi-dimensional state is collapsed onto a low-dimensional observable, information is sacrificed. Whether the state is quantum, cognitive, or algorithmic, the geometry of projection imposes the same fundamental limit.

5. Implications for

Intelligence and AI

5.1 Intelligence as

Invariant Extraction

The Monna-map framework elevates our understanding of intelligence. Intelligence is not primarily about storing facts or executing rules. It is about **extracting invariant relational patterns** from noisy, context-dependent data. Both the brain and LLMs are engines that discover cross-ratios—the projective invariants that remain constant across changes of language, modality, and coordinate system. They build ultrametric trees because such trees are the most efficient data structures for encoding and compressing these invariants.

5.2 The Inevitability of

Hallucinations

Because the Monna projection is many-to-one, **“hallucinations are geometrically inevitable”**. No amount of additional training data or model scaling can eliminate them entirely; they are a fundamental consequence of projecting a tree onto a line. This does not mean we cannot reduce their frequency or severity, but it does mean that the goal of “hallucination-free” AI is as unattainable as a shadow-free projection.

5.3 Designing Around the Limit

Recognizing the geometric origin of hallucinations suggests new strategies for AI design:

Multi-modal communication: Using multiple parallel channels (text, image, audio, context) increases the dimensionality of the projection, reducing aliasing. Just as a 3D object casts less ambiguous shadows when illuminated from multiple angles, a thought projected onto several modalities is less prone to misinterpretation.

Explicit tree-aware architectures: AI systems could be designed to maintain explicit ultrametric tree representations internally, using the Monna projection only at the final output stage. This would separate the invariant thought structure from its linear expression, potentially improving robustness and interpretability.

Uncertainty quantification: Models could be taught to estimate the degree of aliasing for a given generation—to sense when multiple tree-paths are equally plausible—and signal that uncertainty to the user.

5.4 The Human Condition

The Monna map also illuminates the human condition. Every conversation is an exchange of shadows, not trees. Perfect mutual understanding is impossible because the listener can never fully reconstruct the speaker's original tree; they can only approximate it by inferring the cross-ratios and filling in the details with their own prior knowledge. Misunderstandings, ambiguities, and creative reinterpretations are not failures of communication; they are inherent features of the geometry.

The Shape of Meaning

The Monna map is more than a mathematical curiosity; it is a key to the geometry of mind. It tells us that thought is hierarchical, communication is linear, and the bridge between them is a projection that must discard as much as it preserves. Every sentence we utter is a compressed shadow of a vast, branching tree. Every act of understanding is an attempt to reconstruct that tree from its shadow.

Hallucinations—those puzzling, sometimes frustrating errors—are not bugs in the system. They are the cognitive echoes of information loss, the necessary cost of flattening a world of depth into a stream of time. They remind us that the tree of thought is always richer, more intricate, and more mysterious than the line of speech can ever reveal.

In the end, the Monna map offers a humbling and unifying vision: intelligence, whether born of biology or built in silicon, is a geometric dance between the tree and the line. We are all navigators of ultrametric forests, and we are all storytellers, casting shadows on the wall of time.