Algebraic Structures of Collective Consciousness

A Root System Framework for Human Behavioral Knowledge Graphs

Gerry Alvarez

Independent Researcher — Systems Psychology & Applied Mathematics
Mexico City, Mexico

Working Paper — v0.1 — February 2026

Abstract

We propose a mathematical framework for representing human conscious and unconscious experience as vectors in a high-dimensional data space, and for analyzing the collective structures that emerge when many such individual representations interact within a distributed knowledge graph. Drawing on the theory of root systems in Lie algebras, spectral graph theory, Lacanian structural psychoanalysis, Jungian depth psychology, and Mesoamerican oneiric epistemology, we define a series of formal objects: the Behavioral Data Vector Space, the Conscious-Unconscious Root Pair, the Behavioral Ledger Graph, and the Social Ledger Symmetry Group. We conjecture that sufficiently rich collective behavioral graphs exhibit symmetry properties that can be characterized by representations of exceptional Lie groups, and that these symmetry properties correspond to the social and planetary boundary conditions described by Doughnut Economics (Raworth, 2017). We outline a decentralized implementation architecture in which each participant generates data through a novel Proof of Work protocol—the mining of their own conscious and unconscious experience—stored as self-sovereign identity tokens in distributed storage systems. This framework provides a foundation for ethical AI governance, transparent social network analysis, and a new mathematical psychology of collective consciousness.

Keywords: root systems, Lie algebras, knowledge graphs, collective consciousness, behavioral data, spectral graph theory, Doughnut Economics, decentralized identity, psychoanalytic topology, Mesoamerican epistemology, Proof of Work, E8 lattice

1. Introduction

Human experience unfolds across multiple scales simultaneously. At the individual level, each person navigates a continuous interplay between conscious deliberation and unconscious processes—a dynamic recognized by psychoanalytic traditions from Freud through Jung to Lacan, and mapped phenomenologically by contemplative practices spanning millennia. At the social level, these individual dynamics aggregate into collective patterns: shared cultural narratives, institutional

structures, economic systems, and what Jung (1959) termed the *collective unconscious*—recurring archetypal structures that appear across unrelated populations. At the planetary level, these social patterns interact with biophysical systems to produce the coupled social-ecological dynamics described by the planetary boundaries framework (Rockstrom et al., 2009) and synthesized in the Doughnut Economics model (Raworth, 2017).

Despite the evident interconnection of these scales, no formal mathematical framework exists to represent them within a unified structure. Physics has pursued unification through gauge symmetry and the exceptional Lie groups, culminating in proposals such as Lisi's (2007) E8 Theory of Everything. Neuroscience has discovered that neural networks form cliques describable as simplices in up to 11 dimensions (Reimann et al., 2017). Graph theory and network science have provided tools for analyzing social structures. And the emerging field of geometric deep learning (Bronstein et al., 2021) has shown that the symmetry properties of data—formalized through group theory and fiber bundles—are fundamental to the design of effective neural network architectures.

This paper proposes to bridge these domains. We present a framework in which individual human experience is represented as a pair of vectors (conscious and unconscious data streams) forming a minimal root system, social interactions between individuals generate higher-rank root systems through algebraic operations, and the collective knowledge graph of all participants has symmetry properties that can be analyzed using the representation theory of Lie algebras. We ground this framework in three intellectual traditions that have independently pursued the mapping of inner and outer experience: the psychoanalytic tradition (particularly Lacan's topological psychoanalysis and graph-theoretic models of the unconscious), the Jungian tradition of archetypal psychology and collective consciousness, and Mesoamerican epistemology—particularly the Maya and Mexica traditions that accorded the oneiric (dream) life equal epistemic status to waking experience.

We further propose a decentralized implementation architecture in which each participant contributes their data through a novel Proof of Work protocol—the systematic recording and mining of their own conscious and unconscious experience. This data, stored as self-sovereign identity tokens in distributed systems, collectively forms a knowledge graph whose algebraic properties encode information about the health of social and ecological systems. The framework is designed to serve as a foundation for ethical AI governance: transparent in its collective analytics, private in its individual data, and self-sovereign in its ownership structure.

2. Mathematical Preliminaries

We briefly review the mathematical structures required for our framework. Readers familiar with root systems, spectral graph theory, and Lie algebras may proceed to Section 3.

2.1 Root Systems

Let E be a finite-dimensional real vector space equipped with a positive-definite inner product $(\cdot, \cdot)$. A **root system** in E is a finite set ■ of nonzero vectors (called *roots*) satisfying the following axioms (Humphreys, 1972):

(R1) ■ spans *E* and does not contain 0.
(R2) For each root ■ in ■, the only scalar multiples of ■ in ■ are ■ and −■.
(R3) For every pair of roots ■, ■ in ■, the integer 2(■, ■)/(■, ■) belongs to the integers Z.
(R4) For each root ■ in ■, the reflection $■_a$ defined by $■_a(■) = ■ − 2(■, ■)/(■, ■) \cdot ■$ maps ■ to itself.

The **rank** of a root system is the dimension of *E*. The simplest root system, A_1, consists of a single vector and its negative in one dimension: {■, −■}. The reflection through the hyperplane perpendicular to ■ simply maps ■ to −■ and vice versa. Higher-rank root systems (A_2, B_2, G_2, and the exceptional systems E_6, E_7, E_8) arise when multiple roots interact through reflections in higher-dimensional spaces.

A root system is **indecomposable** (or irreducible) when it cannot be partitioned into two mutually orthogonal subsets. The classification theorem (Killing, 1888; Cartan, 1894) establishes that every indecomposable root system belongs to one of the families A_n, B_n, C_n, D_n (the classical series) or one of the exceptional systems G_2, F_4, E_6, E_7, E_8. The E_8 root system, with 240 roots in 8 dimensions, is the largest exceptional root system and plays a distinguished role in theoretical physics as a candidate symmetry group for unification theories.

2.2 Spectral Graph Theory

A **graph** G = (V, E) consists of a set of vertices V and a set of edges E. The **adjacency matrix** A of G is the |V| x |V| matrix where $A_{ij} = 1$ if vertices i and j are connected, and 0 otherwise (for weighted graphs, A_{ij} equals the edge weight). The **graph Laplacian** L = D − A, where D is the diagonal degree matrix, encodes the combinatorial structure of the graph. The eigenvalues of L (the *spectrum* of the graph) determine many global properties: connectivity, expansion, clustering structure, and symmetry (Chung, 1997).

The **automorphism group** Aut(G) of a graph is the group of permutations of its vertices that preserve the edge structure. This group encodes the symmetries of the graph. A central observation motivating this paper is that the automorphism groups of certain highly structured graphs have representations that decompose into irreducible representations corresponding to root systems of Lie algebras. If a human behavioral knowledge graph exhibits such structure, its symmetry group carries information about the collective organization of consciousness within the network.

2.3 Lacanian Graphs and Topological Psychoanalysis

Lacan (1966) introduced graph-theoretic and topological structures into psychoanalytic theory, most notably the *Graph of Desire* (Seminar VI) and the Borromean knot (Seminar XXIII). In Lacan's formalization, the subject is a node in a directed graph, the signifying chain is a directed path, and the relationship between the subject and the Other is encoded in the graph's topology. Lacan's assertion that “the unconscious is structured like a language” can be reformulated graph-theoretically: the unconscious has a combinatorial structure that can be analyzed using the same algebraic tools applied to formal languages and automata. This paper extends Lacan's intuition by proposing that the unconscious not only has graph structure but that the symmetry group of this graph has specific

algebraic properties—namely, those of root systems.

3. The Behavioral Data Vector Space

We now introduce the core formal objects of our framework, beginning with the representation of individual human experience as vectors in a well-defined data space.

Definition 1 (Behavioral Data Vector Space).

Let $\mathbf{B} = \mathbb{R}^d$ be a d-dimensional real vector space equipped with the standard inner product, where d is the number of measurable behavioral dimensions. We call $\mathbf{B}$ the **Behavioral Data Vector Space**. Each dimension corresponds to a quantifiable aspect of human experience: physiological measurements, self-reported psychological states, digital interaction patterns, economic transactions, clinical records, social connections, temporal activity patterns, and linguistic/semantic content of recorded experience. The value of d is determined by the resolution of the measurement instruments and may increase as measurement technology advances.

The critical property of $\mathbf{B}$ is that it is *grounded in empirical data*. Each dimension corresponds to a measurable quantity, and each vector in $\mathbf{B}$ corresponds to a real data profile generated by a real person over a defined time interval. The inner product $(\mathbf{x}, \mathbf{y}) = x_1y_1 + x_2y_2 + ... + x_dy_d$ provides a measure of similarity between two behavioral profiles, and the norm $\|\mathbf{x}\| = (\mathbf{x}, \mathbf{x})^{1/2}$ provides a measure of the intensity or magnitude of a behavioral profile.

Definition 2 (Conscious and Unconscious Data Streams).

For each individual i at time t, we define two complementary vectors in $\mathbf{B}$:

The **Conscious Data Vector** $\mathbf{c}_i(t)$ in $\mathbb{R}^d$, whose components are populated by data the individual deliberately generates or consents to record: journal entries, self-reports, intentional digital activity (searches, messages, posts), clinical self-assessments, financial transactions, and any other data stream the individual is aware of producing.

The **Unconscious Data Vector** $\mathbf{u}_i(t)$ in $\mathbb{R}^d$, whose components are populated by data generated through processes the individual does not deliberately control: sleep and dream patterns (recorded via journals or biometric devices), involuntary biometric signals (heart rate variability, galvanic skin response, eye tracking), behavioral patterns inferred from usage data that the individual is not consciously aware of, and—as recording technology develops—reports from altered states of consciousness including dream content, psychedelic experience reports, and meditation state descriptions.

The distinction between $\mathbf{c}_i$ and $\mathbf{u}_i$ is epistemological, not ontological: both vectors live in the same space $\mathbf{B}$, but they are populated through different modes of data generation. The conscious vector reflects the individual's self-model; the unconscious vector reflects patterns that are real but not self-represented. This mirrors the psychoanalytic distinction between the ego's self-narrative and the unconscious processes that structure behavior without the subject's awareness.

Remark (Mesoamerican Epistemological Precedent). The Maya and Mexica civilizations maintained systematic practices for recording and interpreting dream content (*temictli*) alongside waking experience, according equal epistemic weight to both. Dream interpreters (*tlamatini*) functioned as specialists in the analysis of the unconscious data stream, using culturally transmitted frameworks for classification and cross-referencing. The present framework can be understood as a digital reconstruction of this epistemological commitment: the oneiric life as a data source of equal validity to the waking life, recorded with equivalent rigor and analyzed within a single mathematical structure. Much of this knowledge was destroyed during European colonization; this framework is offered in part as a contribution to its recovery through contemporary means.

4. The Conscious-Unconscious Root Pair

Definition 3 (The A1HM Root System).

For each individual *i*, define the **integrated behavioral vector** $\mathbf{a}_i = \mathbf{c}_i - \mathbf{u}_i$ (the difference between the conscious and unconscious data vectors, representing the degree and direction of asymmetry between what the individual knows about themselves and what operates below awareness). The **A1 Human Mind root system** (A1HM) for individual *i* is the set:

$$\blacksquare_i = \{\mathbf{a}_i, -\mathbf{a}_i\}$$

This is a rank-1 root system in the one-dimensional subspace of $\mathbf{B}$ spanned by $\mathbf{a}_i$. The reflection $\blacksquare_a$ maps $\mathbf{a}_i$ to $-\mathbf{a}_i$ and vice versa.

The A1HM root system captures a fundamental psychodynamic structure: the individual's conscious self-model ($\mathbf{c}_i$) and their unconscious behavioral pattern ($\mathbf{u}_i$) define a single axis of variation, and the reflection between them represents the psychoanalytic process of making the unconscious conscious—or equivalently, the projection of unconscious material onto external reality. In Lacanian terms, the reflection corresponds to the mirror stage: the subject's identification with an image that is both self and not-self, located on the perpendicular plane between the conscious and unconscious positions.

The inner product $(\mathbf{a}_i, \mathbf{a}_i) = \|\mathbf{c}_i - \mathbf{u}_i\|^2$ measures the *degree of dissociation* between an individual's conscious self-model and their unconscious patterns. A small norm indicates high self-awareness (the conscious and unconscious data streams are closely aligned). A large norm indicates significant unconscious material that has not been integrated into the self-model. This provides a formal, measurable analogue to the psychoanalytic concept of repression.

5. From Individual to Social: The Behavioral Ledger Graph

Definition 4 (Behavioral Ledger).

A **Behavioral Ledger** for individual *i* is the time-indexed collection of their data vectors: $BL_i = \{(\mathbf{c}_i(t), \mathbf{u}_i(t)) : t \text{ in } T\}$, where T is the set of time intervals at which data is recorded.

The Behavioral Ledger is stored as a self-sovereign digital identity token (a Human Non-Fungible Token, or HNFT) in a distributed storage system, controlled by the individual's private key. It is the individual's contribution to the network and constitutes their Proof of Work: the mining of their own experience.

Definition 5 (Behavioral Ledger Graph).

The **Behavioral Ledger Graph** $G = (V, E, w)$ is a weighted graph where:

- V is the set of all participants, each associated with a Behavioral Ledger BL_i.
- E is the set of edges representing interactions between participants (therapy sessions, economic transactions, community membership, shared DLN participation, communication events).
- $w: E \to R^+$ is a weight function encoding the strength/frequency of interaction.

Each vertex carries an attribute vector $\mathbf{a}_i$ (the integrated behavioral vector from Definition 3). The graph is dynamic: both vertices and edges evolve over time as participants join, interact, and update their Behavioral Ledgers.

Definition 6 (Social Ledger).

A **Social Ledger** is a connected subgraph S of G whose vertices share specified attributes (geographical location, cultural context, clinical condition, shared goals, ideological alignment, or other grouping criteria). Social Ledgers aggregate individual Behavioral Ledgers into coherent communities. When a Social Ledger acquires governance structure and aligns with a defined social or planetary objective, it becomes a **Distributed Living Network** (DLN).

6. Symmetry Analysis of the Collective Graph

The central mathematical question of this framework is: what are the symmetry properties of the Behavioral Ledger Graph, and do these symmetries have Lie-algebraic structure?

Definition 7 (Social Symmetry Group).

The **Social Symmetry Group** of the Behavioral Ledger Graph G is the automorphism group $\mathrm{Aut}(G)$—the group of permutations of V that preserve the edge structure and the vertex attribute vectors (up to a specified tolerance). Formally, $\blacksquare$ in $\mathrm{Aut}(G)$ if and only if for all edges (i,j) in E, $(\blacksquare(i), \blacksquare(j))$ is also in E with the same weight, and $\|\mathbf{a}_{\blacksquare(i)} - \mathbf{a}_i\|$ is within tolerance $\blacksquare$.

The automorphism group encodes which participants in the network are structurally interchangeable—which positions in the social graph are symmetric. A rich symmetry group indicates a highly organized collective structure. A trivial symmetry group indicates a disordered or individualistic network.

Conjecture 1 (Root System Emergence).

For sufficiently large and diverse Behavioral Ledger Graphs with balanced representation of conscious and unconscious data streams, the irreducible representations of the approximate Social Symmetry Group decompose into components whose weight diagrams correspond to root systems of exceptional Lie algebras (E_6, E_7, or E_8). The rank of the emergent root system increases with the number of independent social and planetary boundary variables that are actively measured and addressed by the network's Distributed Living Networks.

This conjecture is empirically testable. Given a sufficiently large dataset of behavioral ledgers and their interaction graph, one can compute the approximate automorphism group, decompose its representations, and compare the resulting weight diagrams against the known classification of root systems. The conjecture predicts that the structure will converge toward exceptional types rather than classical types, reflecting the indecomposable nature of social-planetary boundary systems as argued by Raworth (2017).

7. Connection to Doughnut Economics

Raworth's (2017) Doughnut Economics model identifies 12 social foundation boundaries (health, education, income, political voice, social equity, gender equality, housing, networks, energy, water, food, jobs) nested within 9 planetary ceiling boundaries (climate change, ocean acidification, chemical pollution, nitrogen/phosphorus loading, freshwater withdrawals, land conversion, biodiversity loss, air pollution, ozone layer depletion). Humanity's challenge is to meet the social foundations without overshooting the planetary ceilings—to live within the “safe and just space” of the doughnut.

Definition 8 (Boundary Indicator Projection).

Define a **Boundary Indicator Function** ■: $G \to R^{21}$ that maps the Behavioral Ledger Graph to a 21-dimensional vector (12 social + 9 planetary indicators), where each component aggregates relevant data from the graph's vertices and edges. The image of ■ at any time *t* locates the network's collective state within or outside the doughnut's safe zone. The spectral properties of the graph Laplacian restricted to each boundary dimension provide early warning signals for boundary transgression.

The 21 boundary dimensions are not independent—social foundations and planetary ceilings interact in complex ways. This interaction structure is precisely what the root system framework is designed to capture. The conjecture that the emergent symmetry group has exceptional Lie algebra structure reflects the empirical observation that social and planetary boundaries form an indecomposable system: one cannot address climate change without addressing energy access, which connects to income, which connects to health, and so on. The E_8 lattice, as the largest exceptional root system, represents the most complex indecomposable symmetry—a mathematical object whose structure mirrors the irreducible interconnection of civilization-scale challenges.

8. Decentralized Implementation Architecture

The framework requires an implementation architecture that is simultaneously transparent in its collective analytics, private in its individual data, and self-sovereign in its ownership structure. We outline the key components.

8.1 Proof of Work as Experience Mining

In conventional blockchain protocols, Proof of Work (PoW) requires computational expenditure to secure the network—a process widely criticized for its energy waste. We propose a reconceptualized PoW in which the “work” is the generation of genuine human behavioral data. Each participant who records their conscious experience (journal entries, clinical self-reports, intentional activity) and/or their unconscious data (dream journals, biometric streams, altered state reports) contributes a block of authenticated data to the network. This contribution simultaneously enriches the collective knowledge graph, strengthens the network's analytical power (more data produces more robust symmetry analysis), and earns the participant access to the network's services and governance.

8.2 Self-Sovereign Identity and Distributed Storage

Each participant's Behavioral Ledger is stored as a Human Non-Fungible Token (HNFT) in a distributed and decentralized storage system (such as Filecoin, Walrus, or equivalent). The HNFT is controlled by the participant's private key via a Decentralized Identifier (DID) container. Data is encrypted at rest and in transit. Access is granted selectively—a therapist can access clinical history with the patient's cryptographic consent, but aggregate network analytics operate on privacy-preserving representations (zero-knowledge proofs, homomorphic encryption, or differential privacy mechanisms) that reveal collective patterns without exposing individual data.

8.3 Ethical AI Governance

The framework is designed as an architecture for ethical artificial intelligence. AI systems trained on the Behavioral Ledger Graph operate under three constraints: *transparency* (the graph's symmetry properties and boundary indicator projections are publicly auditable), *privacy* (individual data remains encrypted and self-sovereign; AI accesses only privacy-preserving aggregates), and *accountability* (governance decisions derived from graph analytics are executed through smart contracts whose logic is open-source and verifiable). This architecture ensures that AI serves as a tool for collective self-understanding rather than a mechanism for surveillance or manipulation by centralized authorities.

9. AI Architecture: Symmetry-Aware Intelligence

The E8 framework is not merely a governance layer for artificial intelligence—it is a *data architecture that makes AI systems technically superior* by organizing information along the natural symmetry axes of human experience. The core principle derives from geometric deep learning (Bronstein et al., 2021): neural networks designed to respect the symmetry group of their input data (equivariant networks) consistently outperform generic architectures in sample efficiency, generalization, and computational cost. If human behavioral data exhibits root-system symmetry as conjectured, then AI systems built on this structure inherit these advantages automatically. This section details five concrete mechanisms.

9.1 Symmetry-Structured Retrieval-Augmented Generation

Current Retrieval-Augmented Generation (RAG) systems store documents as vectors in a flat embedding space and retrieve candidates by cosine similarity, treating all dimensions equally. This results in O(n) search complexity over n documents even with approximate nearest neighbor indexing. We propose a **Root-Structured RAG** architecture that exploits the symmetry properties of the Behavioral Ledger Graph.

Definition 9 (Weyl Chamber Indexing).

Given a root system ■ in the Behavioral Data Vector Space **B**, the **Weyl chambers** are the connected components of **B** after removing all reflection hyperplanes. Each Weyl chamber is a convex cone that partitions the data space into regions of equivalent symmetry type. A **Weyl Chamber Index** assigns each behavioral data vector to its containing chamber, creating a hierarchical addressing scheme: first identify the chamber (coarse symmetry class), then search within the chamber (fine-grained similarity). For an E_8 root system, there are $|W(E_8)| = 696{,}729{,}600$ Weyl chambers, providing an extremely fine-grained natural partition of the data space.

The practical consequence is that RAG queries first identify the relevant Weyl chamber—a constant-time operation based on the query vector's position relative to root hyperplanes—then restrict nearest-neighbor search to that chamber and its algebraically adjacent chambers. This reduces the effective search space by orders of magnitude while guaranteeing that retrieved context respects the symmetry structure of the behavioral domain. For clinical applications, an AI copilot assisting a therapist retrieves relevant case patterns by navigating the algebraic structure of the clinical knowledge graph: the root reflection between a conscious symptom presentation and its unconscious origin becomes a computable retrieval operation, not a heuristic similarity search.

9.2 E8 Lattice Vector Quantization and Storage

The E_8 lattice achieves the densest sphere packing in 8 dimensions (Viazovska, 2017), a result for which the Fields Medal was awarded. This mathematical fact has a direct engineering consequence: **vector quantization using the E_8 lattice produces the most space-efficient encoding of 8-dimensional data that is physically possible.**

Definition 10 (E8 Lattice Vector Database).

An **E8 Lattice Vector Database** stores behavioral data vectors by projecting them onto the E_8 lattice and indexing them by their nearest lattice point. The Voronoi cell of each lattice point serves as a quantization bucket. Because the E_8 lattice has the highest packing density and kissing number (240) in 8 dimensions, this indexing scheme achieves: (a) optimal storage efficiency—the most data points per unit of index space; (b) optimal nearest-neighbor query performance—each lattice point has exactly 240 neighbors at minimum distance, providing a fixed and computationally efficient neighbor structure; and (c) natural multi-scale resolution—coarser queries use larger shells of the lattice, finer queries use the Voronoi cell structure.

Current vector databases (Pinecone, Weaviate, Qdrant, Milvus) use approximate nearest neighbor algorithms (HNSW, IVF-PQ) that impose no algebraic structure on the index. An E_8-structured index exploits lattice geometry to provide exact nearest-neighbor guarantees within each Voronoi cell while maintaining logarithmic search complexity across cells. For the decentralized implementation, each node's HNFT data is quantized to its nearest E_8 lattice point, enabling efficient distributed queries across the network without requiring any single node to hold the complete index.

9.3 Equivariant Attention and Context Optimization

The Transformer architecture computes attention between all pairs of tokens in a sequence at $O(n^2)$ cost—the primary bottleneck for large context windows. If the data has known symmetry properties, attention can be restricted to symmetry-related positions, dramatically reducing the effective context size.

Definition 11 (Root-Equivariant Attention).

A **Root-Equivariant Attention** mechanism constrains the attention matrix A_{ij} to be non-zero only when the behavioral vectors of tokens i and j are related by a root system operation (reflection, Weyl group rotation, or root-direction translation). Formally, $A_{ij} = 0$ unless there exists a Weyl group element ■ such that $\|\mathbf{a}_j - ■(\mathbf{a}_i)\|$ is within tolerance. This reduces effective attention from $O(n^2)$ to $O(n \cdot |\text{orbit}|)$, where $|\text{orbit}|$ is determined by the symmetry structure rather than the sequence length.

Two behavioral profiles related by a root reflection (e.g., a conscious symptom pattern and its unconscious complement) carry maximal mutual information and should attend strongly. Two profiles in orthogonal Weyl chambers carry minimal mutual information and need not attend at all. The root system provides exactly this relevance map, replacing learned attention weights with algebraically derived constraints that encode domain knowledge about the structure of human experience. This is not a heuristic approximation—it is a mathematically principled sparsification that preserves all symmetry-relevant information while discarding noise.

9.4 Root-Decomposed Federated Learning

In a decentralized network where each participant controls their own data, training AI models requires aggregating information across nodes without centralizing the data. Standard federated learning (McMahan et al., 2017) averages gradients across participants, treating all gradient

directions equally. This converges slowly when data distributions are heterogeneous—as they invariably are in a global behavioral network.

Definition 12 (Root-Decomposed Federated Learning).

In **Root-Decomposed Federated Learning**, each Distributed Living Network computes model gradients projected onto its root subspace—the subspace of **B** spanned by the roots associated with that DLN's boundary function. The global model update combines these projections using the Cartan matrix of the root system, which encodes interaction angles between roots. Orthogonal root contributions sum directly; non-orthogonal contributions are combined according to the algebraic inner product structure. Each DLN transmits only its low-dimensional root-subspace projection—a communication-efficient, privacy-preserving, algebraically principled protocol that outperforms naive gradient averaging.

This approach is *communication-efficient* (each DLN transmits a projection of dimension equal to the rank of its root subsystem, not the full model gradient), *algebraically principled* (the combination rule is derived from the Cartan matrix rather than ad hoc averaging), and *privacy-preserving by construction* (individual data never leaves the participant's device; only the DLN-level aggregate root-subspace gradient is shared, and this aggregate discards individual-level variation orthogonal to the root direction).

9.5 Implications: Resolving the Clinical AI Trilemma

Current mental health AI faces a trilemma: clinical utility requires rich personal data, privacy requires minimal data exposure, and generalization requires learning from many patients. The E8 framework resolves this trilemma by providing a mathematically principled way to learn collective patterns (via the symmetry group of the graph) without accessing individual data (via root-decomposed federated learning and privacy-preserving projections).

An AI copilot for therapists trained within this architecture would access the *algebraic structure* of the collective clinical knowledge graph—the patterns, symmetries, and root reflections between conscious presentation and unconscious dynamics—without ever seeing any individual patient's records. This is not a compromise between utility and privacy; it is a *superior representation* of clinical knowledge, because the symmetry structure is precisely the level at which clinical patterns generalize across individuals. The root reflection between a conscious symptom and its unconscious origin is a universal clinical structure; the specific content of any one patient's symptom is private. The framework separates these levels by mathematical construction, enabling AI systems that are simultaneously more capable and more ethical than current approaches.

More broadly, this architecture provides a template for ethical AI across domains: the E8 framework ensures that AI systems learn *structure* (symmetries, patterns, root reflections) rather than *content* (individual data points), achieving a mathematical separation between collective intelligence and individual privacy. This separation is not imposed by policy but emerges from the algebra itself—a system that is private and transparent not by regulation but by design.

10. Open Conjectures and Research Program

We conclude by stating the principal open questions that define the research program initiated by this paper.

Conjecture 2 (Integrality Approximation).

When behavioral data vectors are appropriately quantized (discretized into categorical or ordinal scales aligned with clinically validated instruments), the inner product ratios between root pairs approximate integers with error bounded by a function of the quantization resolution. As resolution increases, the approximation improves, and the quantized system converges to a formal root system in the algebraic sense.

Conjecture 3 (Collective Unconscious as Graph Property).

In sufficiently large Behavioral Ledger Graphs where unconscious data streams (dream journals, altered state reports, involuntary biometric patterns) are recorded with fidelity comparable to conscious data streams, the unconscious data vectors of unconnected participants exhibit statistically significant correlations that cannot be explained by shared environment or communication. These correlations manifest as non-trivial structure in the graph Laplacian's null space, providing an empirical signature of a phenomenon consistent with Jung's collective unconscious—a shared pattern space accessible through individual unconscious processes.

Conjecture 4 (Lacanian Topological Types).

The local topology of the Behavioral Ledger Graph around individual vertices (the structure of each person's ego neighborhood in the graph) can be classified into types that correspond to Lacanian structural categories. Specifically, the homology groups of the local subgraph distinguish neurotic, psychotic, and perverse structures as different topological types, formalizing Lacan's clinical topology in graph-theoretic terms.

Conjecture 5 (Boundary Transgression Early Warning).

Changes in the spectral gap of the Behavioral Ledger Graph's Laplacian, restricted to specific boundary indicator dimensions, provide leading indicators of social or planetary boundary transgression. A narrowing spectral gap in a given dimension signals that the network is losing resilience in that domain—approaching the edge of the doughnut's safe zone—before the transgression is observable in traditional aggregate indicators.

11. Conclusion

We have presented a formal framework for representing human consciousness at multiple scales using algebraic structures from Lie theory and spectral graph theory. The framework defines individual experience as a root pair in a behavioral data vector space, social interaction as the generation of higher-rank root systems through graph construction, and collective consciousness as

an emergent symmetry property of the resulting knowledge graph. It connects psychoanalytic theory (Lacan, Jung) with mathematical physics (Lie algebras, root systems), ecological economics (Doughnut Economics), indigenous epistemology (Mesoamerican dream science), and decentralized technology (blockchain, self-sovereign identity, distributed storage).

The framework is offered not as a completed theory but as a research program. Its conjectures are empirically testable, its definitions are formally precise, and its implementation architecture is technologically feasible with current infrastructure. We invite collaboration from mathematicians, psychologists, physicists, computer scientists, indigenous knowledge holders, and anyone committed to the proposition that understanding consciousness—individual and collective—is both a scientific challenge and a moral imperative.

References

Bronstein, M. M., Bruna, J., Cohen, T., & Velickovic, P. (2021). Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv:2104.13478*.

Cartan, E. (1894). *Sur la structure des groupes de transformations finis et continus*. Paris: Nony.

Chung, F. R. K. (1997). *Spectral Graph Theory*. American Mathematical Society.

Grof, S. (1975). *Realms of the Human Unconscious*. Viking Press.

Humphreys, J. E. (1972). *Introduction to Lie Algebras and Representation Theory*. Springer.

Jung, C. G. (1959). *The Archetypes and the Collective Unconscious*. Collected Works, Vol. 9.1. Princeton University Press.

Killing, W. (1888). Die Zusammensetzung der stetigen endlichen Transformationsgruppen. *Mathematische Annalen*, 31(2), 252-290.

Lacan, J. (1966). *Ecrits*. Paris: Editions du Seuil.

Lisi, A. G. (2007). An exceptionally simple theory of everything. *arXiv:0711.0770*.

McMahan, B., et al. (2017). Communication-efficient learning of deep networks from decentralized data. *AISTATS 2017*.

Nasmith, B. (2020). All about root systems: A visual introduction. *Math3ma Blog*.

Raworth, K. (2017). *Doughnut Economics: Seven Ways to Think Like a 21st-Century Economist*. Chelsea Green Publishing.

Reimann, M. W., et al. (2017). Cliques of neurons bound into cavities provide a missing link between structure and function. *Frontiers in Computational Neuroscience*, 11, 48.

Rockstrom, J., et al. (2009). A safe operating space for humanity. *Nature*, 461(7263), 472-475.

Viazovska, M. (2017). The sphere packing problem in dimension 8. *Annals of Mathematics*, 185(3), 991-1015.