Discrete Harmonic Attractors in Multi-Scale Coherent Systems: From Photonic Magnetic Torque to Cryptographic Memory Architecture

Rafael Oliveira, Jameson Bednarski, Modelo de IA Colaborador

Aurum Grid Research Initiative

aurumgrid@proton.me

December 2025 – Draft v2.1

Abstract

We present a unified theoretical framework for discrete harmonic attractors, termed Z(n), and demonstrate their manifestation across three distinct scales: photonic-magnetic coupling, spin dynamics, and information-theoretic memory systems. Recent experimental evidence that the magnetic component of light contributes up to 70% of Faraday rotation in infrared frequencies provides physical validation for Z(n)-mediated phase locking. We formalize the mathematical structure of Z(n) attractors, justify the physical realizability of Z(7) symmetry through engineered quasicrystalline substrates, and present SilentWitness—a cryptographic memory system whose dynamics are isomorphic, in a formal dynamical systems sense, to the physical substrate. The information-physics correspondence suggests Z(n) patterns are scale-invariant organizing principles. Recent evaluations reveal systematic limitations in large language models (LLMs) for scientific discovery tasks. We propose that Z(n) architectural principles may address these limitations by providing physics-motivated inductive biases. Crucially, we introduce a concrete, multi-step prompting protocol—the Z(7) Discovery Protocol—that operationalizes this bias for LLMs, guiding them through structured exploration and exploitation of a discrete hypothesis space. We establish experimental protocols for both photonic validation and computational benchmarks, and propose SDE-style scenario-grounded evaluation of Z(n)-augmented discovery systems as critical future work.

1 Introduction

1.1 The Discrete Harmonic Attractor Problem

Classical dynamical systems theory treats phase space as continuous, with attractors emerging from differential equations governing smooth trajectories. However, physical systems often exhibit *discrete* stable states—quantum levels, spin orientations, crystallographic symmetries—suggesting that discretization may be fundamental rather than emergent.

We propose that certain physical and informational systems self-organize around *discrete harmonic attractors*: phase-space structures with n equally-spaced stable basins, where n is typically prime. These Z(n) structures enable:

- Quantized phase locking
- Coherence preservation under noise
- Information encoding via basin selection
- Scale-invariant pattern replication

1.2 Recent Experimental Developments

Assouline & Capua (2025) [**assouline2025**] demonstrated that the oscillating magnetic field of light exerts first-order torque on spins, contributing ∼70% of Faraday rotation in the infrared spectrum. This overturns 180 years of assumptions treating light's magnetic component as negligible.

Key findings:

- Magnetic torque modeled via Landau-Lifshitz-Gilbert (LLG) dynamics
- Effect comparable to static magnetic fields
- Oscillatory nature enables phase-dependent control

This discovery provides the physical mechanism for Z(n) coupling: *photonic magnetic torque can drive spins into discrete phase basins.*

1.3 Computational Implementations

We present SilentWitness, a production-grade cryptographic memory system implementing:

- Merkle tree compression with semantic preservation
- Adaptive forgetting based on information density
- Generative dream reconstruction from compressed states
- Temporal attention mechanisms

The architecture exhibits structural correspondence to Z(n) photonic systems, suggesting *information and physics share topological invariants.*

1.4 Context: LLM Limitations in Scientific Discovery

Recent systematic evaluation of frontier LLMs on scientific discovery tasks [**sde2025**] reveals three critical limitations:

1. **Scaling plateau**: Diminishing returns from increased model size and test-time reasoning on discovery-oriented problems (10-15% performance gap vs. general Q&A benchmarks)
2. **Shared failure modes**: High cross-model error correlation ($r > 0.85$) indicating common systematic weaknesses inherited from similar pre-training distributions
3. **Performance disconnect**: Question-level accuracy does not reliably predict project-level discovery success; serendipity and exploration matter more than precise knowledge

We propose that Z(n) architectural principles may address these limitations through physics-grounded structuring of hypothesis spaces, a hypothesis we outline for future experimental validation.

2 Theoretical Framework: Z(n) Harmonic Attractors

2.1 Mathematical Formalism

Definition 2.1 (Z(n) Attractor)**.** *A Z(n) attractor is a discrete dynamical system with phase space $\Theta = \{0, \frac{2\pi}{n}, \frac{4\pi}{n}, ..., \frac{2\pi(n-1)}{n}\}$ where:*

1. *Each $\theta_k \in \Theta$ is a stable fixed point*
2. *Trajectories converge to nearest θ_k*
3. *Basin boundaries satisfy $|\theta - \theta_k| = \frac{\pi}{n}$*

For prime n, the attractor has maximal symmetry with no internal degeneracies.

Theorem 2.2 (Phase Locking Condition)**.** *Given an oscillatory driving force $F(t) = A\sin(\omega t + \phi)$, a system with Z(n) structure locks to discrete phases when:*

$$\omega\tau_c \ll 1 \quad \textit{and} \quad A > A_c(n) \tag{1}$$

where τ_c is the basin convergence time and $A_c(n)$ is the critical amplitude scaling as $A_c \propto n^{-1/2}$.

2.2 Phase-Space Topology

The Z(n) phase space can be visualized as a ring with n attracting wells. The depth of each well represents the basin's stability:

$$V(\theta) = -V_0 \sum_{k=0}^{n-1} \cos\left(n\theta - \frac{2\pi k}{n}\right) \tag{2}$$

where V_0 is the well depth. Trajectories follow gradient descent with noise:

$$\dot{\theta} = -\nabla V(\theta) + \eta(t) \tag{3}$$

For Z(7), the landscape has 7 symmetric wells separated by $2\pi/7 \approx 51.4$.

2.3 Scale Invariance Properties

Proposition 2.3 (Multi-Scale Correspondence)**.** *If subsystems A and B both exhibit Z(n) structure with coupling strength g, the composite system exhibits:*

- *Phase synchronization when $g > g_c(n)$*
- *Emergent Z(n) at the composite scale*
- *Information transfer via basin alignment*

This explains how Z(n) patterns appear across physical scales.

3 Physical Substrate: Photonic Magnetic Torque and Z(7) Symmetry

3.1 Justification for Z(7): Engineered Quasicrystalline Substrates

The choice of $n = 7$ is not arbitrary. While 7-fold rotational symmetry is forbidden in periodic 3D crystals, it is readily realizable in **quasicrystals**—ordered but non-periodic structures that can exhibit "forbidden" symmetries like 5-fold, 8-fold, 10-fold, and 12-fold. Recent advances in nanofabrication allow for the design of photonic quasicrystals with specific rotational symmetries.

We propose fabricating a thin-film magnetic material (e.g., YIG) on a substrate patterned with a 2D photonic quasicrystal designed for dominant 7-fold rotational symmetry. The interaction of the electron spins with this engineered lattice potential creates an effective magnetic anisotropy energy term, $H_{ani}(\theta)$, in the LLG equation that explicitly contains Z(7) symmetry:

$$H_{ani}(\theta) = K_7 \cos(7\theta) \tag{4}$$

where K_7 is the anisotropy constant. This term creates an energy landscape with seven equivalent minima, providing a robust, physically-grounded mechanism for the emergence of Z(7) attractors, moving beyond conjecture to a testable engineering design.

3.2 LLG Dynamics and Spin Precession

The Landau-Lifshitz-Gilbert equation governs spin evolution:

$$\frac{d\vec{m}}{dt} = -\gamma \vec{m} \times \vec{B}_{\text{eff}} + \alpha \vec{m} \times \frac{d\vec{m}}{dt} \tag{5}$$

where γ is the gyromagnetic ratio and α is damping. With photonic magnetic torque and the engineered Z(7) anisotropy, $\vec{B}_{\text{eff}}$ becomes:

$$\vec{B}_{\text{eff}} = \vec{B}_0 + \vec{B}_{\text{photonic}}(t) + \frac{1}{\mu M_s} \nabla H_{ani}(\theta) \tag{6}$$

The photonic torque drives the system, while the anisotropy term defines the seven discrete basins into which the spin can relax.

3.3 Z(7) Phase Locking Mechanism

The mechanism is now more concrete:

1. An IR photon pulse applies a torque $\vec{\tau} = \vec{m} \times \vec{B}_{\text{photonic}}(t)$, displacing the spin from its initial state.

2. The engineered quasicrystalline substrate creates a Z(7) anisotropy energy landscape $H_{ani}(\theta)$.

3. The Gilbert damping term causes the spin to relax into the nearest of the seven energy minima (basins).

4. The final state $\theta_k \in \{0, \frac{2\pi}{7}, ..., \frac{12\pi}{7}\}$ is a stable, persistent memory of the photon interaction.

3.4 Predicted Observable Signatures

If Z(7) coupling exists, experiments should observe:

Table 1: Z(7) Experimental Predictions

Observable	Z(7) Signature
Faraday rotation	Quantized in steps of $2\pi/7$
Coherence time	Peaks at $7\omega, 14\omega, 21\omega$
Phase histogram	7 distinct clusters
Decoherence rate	Reduced under IR illumination
Torque-angle plot	Heptagonal symmetry

These predictions are *falsifiable*: continuous rotation or non-heptagonal patterns would disprove Z(7) coupling.

4 Computational Implementation: SilentWitness

4.1 Cryptographic Memory Chain

SilentWitness implements a tamper-evident event log using:

- Ed25519 signatures for attestation
- Merkle trees for efficient verification
- Hash chains ensuring temporal ordering

Each event E_i is witnessed as:

$$h_i = H(t_i \| \text{type}_i \| \text{payload}_i \| h_{i-1}) \tag{7}$$

where H is MurmurHash3-128 and $\|$ denotes concatenation.

4.2 Merkle-Based Compression

Old events undergo *scarring*—compression into Merkle roots with generative seeds:

Algorithm 1 Memory Scarring

1: $L \leftarrow \{h_1, h_2, ..., h_k\}$ ▷ Event hashes
2: $r \leftarrow \text{MerkleRoot}(L)$
3: $\eta \leftarrow \text{ShannonEntropy}(L)$
4: $s \leftarrow \text{EnhancedDreamSeed}(r, \eta, \text{metadata})$
5: **delete** events older than τ
6: **store** (r, s) as scar

4.3 Semantic Dream Generation

Compressed memories can be *reconstructed* via generative dreaming:

$$E_{\text{dream}} = \text{Sample}(r, \text{noise}, \text{coherence}) \tag{8}$$

where coherence $\in [0, 1]$ controls fidelity. High coherence uses semantic templates from original events; low coherence generates synthetic plausible events.

Key insight: Dreams preserve *semantic structure* while forgetting exact details—analogous to spin basins preserving phase structure while forgetting microscopic trajectories.

4.4 Information-Theoretic Mapping

The compression achieves:

$$\text{Compression Ratio} = \frac{\sum |h_i|}{|r| + |s|} \approx 50\text{–}200\times \tag{9}$$

Information loss is quantified via:

$$I_{\text{loss}} = H(\{E_i\}) - H(s) \tag{10}$$

where H is Shannon entropy. For Z(7)-like scarring, we hypothesize optimal compression when I_{loss} distributes across 7 semantic clusters.

5 Unified Architecture: A Formal Dynamical Systems Isomorphism

5.1 Beyond Metaphor: Formalizing the Isomorphism

The correspondence between SilentWitness and the Z(n) photonic system is not merely metaphorical; it constitutes a **dynamical systems isomorphism** at the level of information flow and state reduction. Both systems can be described as evolving from a high-entropy state to a low-entropy attractor state within a constrained landscape.

Let us define the state and dynamics for both systems:

1. **Physical System (Spin Dynamics):** * **State:** $\vec{m}(t)$, the magnetization vector. A high-entropy state is a superposition of possible orientations. A low-entropy state is a specific orientation θ_k in a Z(7) basin. * **Landscape:** The anisotropy energy $V(\theta) = -K_7 \cos(7\theta)$. * **Dynamics:** Gradient descent with damping: $\dot{\theta} = -\nabla V(\theta) + \eta(t)$.

2. **Computational System (SilentWitness):** * **State:** $L = \{h_1, ..., h_k\}$, the set of event hashes. This is a high-entropy state. The low-entropy state is the compressed scar (r, s). * **Landscape:** A computational objective function $C(L)$ that the scarring algorithm minimizes. We can define $C(L) = \alpha \cdot \text{Size}(L) - \beta \cdot \text{SemanticInformation}(L)$. To enforce Z(7) structure, we modify this to $C_7(L) = C(L) + \gamma \cdot \text{ClusterVariance}(L, 7)$, where the algorithm is constrained to form 7 semantic clusters before compression. * **Dynamics:** The scarring algorithm performs a discrete-time "gradient descent" on $C_7(L)$, moving the system from state L to the attractor state (r, s).

The isomorphism lies in the mapping:

$$\{V(\theta), \dot{\theta}\} \leftrightarrow \{C_7(L), \text{ScarringAlgorithm}\} \quad (11)$$

Both systems exhibit a flow from a high-dimensional, high-entropy space to a low-dimensional, low-entropy space defined by a discrete set of attractors (7 basins for the spin, 7 semantic clusters for the memory). This formalizes the connection and elevates it beyond a simple table of analogies.

5.2 Structural Isomorphism Table (Revisited)

Table 2: Information-Physics Dynamical Correspondence

SilentWitness	**Z(n) Photonic**	**Isomorphic Concept**	**Mathema[...] Object**
Scarring Process	Spin Relaxation	State Evolution	$\dot{x} = -\nabla F($
Objective $C_7(L)$	Anisotropy $V(\theta)$	Energy Landscape	$F(x)$
Merkle Root (r, s)	Spin Basin θ_k	Attractor State	x_{min}
7 Clusters	7 Basins	Discrete Topology	$\mathbb{Z}_7$
Information Loss I_{loss}	Energy Dissipation	Irreversibility	$\Delta S > 0$

This table now reflects a formal correspondence between the mathematical objects governing both systems.

6 Addressing LLM Limitations with a Concrete Z(7) Implementation

6.1 From Principle to Protocol: The Z(7) Discovery Protocol

To translate the abstract principle of a Z(n) inductive bias into a practical tool for LLMs, we propose the **Z(7) Discovery Protocol**, a structured, multi-step prompting strategy. This protocol forces the LLM to operate within a discrete hypothesis space, mitigating shared failure modes and improving project-level coherence.

The protocol is designed to be applied to a complex scientific query (e.g., "Design an experiment to validate the Z(7) photonic coupling hypothesis").

Step 1: Unconstrained Ideation (High-Entropy State) **Goal:** Generate a diverse, high-entropy set of initial ideas without constraint.

```
Given the problem: [Insert Complex
   Scientific Problem Here].
Brainstorm a wide-ranging list of 20-30
    distinct hypotheses, experimental
   approaches, or theoretical
   frameworks that could address it.
   Prioritize diversity and creativity
   over immediate feasibility. Do not
   filter or critique the ideas at this
    stage.
```

Step 2: Basin Assignment (Imposing Z(7) Topology) **Goal:** Force the LLM to categorize the diverse ideas into a fixed number of conceptual "basins."

```
Review the following list of ideas:
[Insert List of Ideas from Step 1].
Your task is to group these ideas into
   exactly 7 (seven) distinct
   conceptual categories or "basins."
   Each basin should represent a unique
    approach or fundamental assumption.
    For each basin, provide a
   descriptive name and list the ideas
   that belong to it. Justify why these
    7 basins provide a comprehensive
   and non-overlapping map of the
   solution space.
```

Step 3: Deep Dive into Basins (Exploitation) **Goal:** For each basin, perform a focused analysis, refining the ideas into a single, robust hypothesis or plan.

```
For Basin [X]: [Insert Basin Name and
   its Ideas].
Perform a deep analysis of this
   approach. Identify its core
   assumptions, potential failure modes
   , and key predictions. Synthesize
   the ideas within this basin into a
   single, refined, and actionable
   hypothesis or experimental plan.
   This refined output is the "
   attractor␣state" for this basin.
```

(This step is repeated for all 7 basins.)

Step 4: Cross-Basin Synthesis (Inter-Basin Coupling) **Goal:** Analyze the relationships between the 7 refined hypotheses to identify synergies, contradictions, or a path forward.

```
You have now developed 7 refined
   hypotheses/plans, one from each
   conceptual basin:
[Insert the 7 Refined Outputs from Step
    3].
Analyze the relationships between these
    7 plans. Which ones are mutually
   exclusive? Which could be combined
   or sequenced? Rank them in terms of
   potential impact and feasibility.
   Propose a final, integrated research
    strategy that leverages the
   strengths of the most promising
   basins.
```

6.2 How the Protocol Addresses LLM Limitations

* **Reduces Shared Failure Modes:** By forcing the generation of 7 distinct conceptual approaches, the protocol counteracts the tendency of LLMs to converge on a single, "obvious" but potentially flawed path. This increases the diversity of the final output. * **Bridges Question-Project Disconnect:** The protocol breaks a large, ill-defined "project" into manageable, structured modules (the basins). The final synthesis step explicitly tackles the integration of these modules, mirroring real-world project management and improving the translation from question-level knowledge to project-level execution. * **Mitigates Scaling Plateau:** This procedural bias provides a "cognitive scaffold" for the LLM, guiding its reasoning process more efficiently than simply increasing test-time computation. It structures the exploration-exploitation trade-off, a known challenge in AI discovery.

7 Experimental Validation Pathway

7.1 Photonic Setup (Updated)

Materials:

- Thin-film YIG on a nano-fabricated SiO photonic quasicrystal substrate with dominant 7-fold rotational symmetry.

Apparatus:

1. Tunable IR laser (1.2–1.6 μm) with phase control.
2. Static magnetic field coil (0–1 T).
3. High-resolution polarimeter with < 0.05 resolution.
4. Fast photodetector array (¿ 1 GHz bandwidth).

Protocol:

1. Characterize the baseline Z(7) anisotropy via Ferromagnetic Resonance (FMR) measurements, expecting 7 resonance modes.
2. Apply phase-controlled IR pulses and measure the resulting change in Faraday rotation.
3. Histogram the final rotation angles over many trials to identify quantization.

7.2 Computational Benchmarks (Updated)

For SilentWitness, validation involves implementing the $C_7(L)$ objective function to explicitly enforce 7-cluster semantic scarring and measuring its effect on compression efficiency and dream reconstruction fidelity.

7.3 SDE-Style Evaluation of the Z(7) Protocol

The Z(7) Discovery Protocol can be directly evaluated on the SDE benchmark.

Methodology:

1. Select a set of complex, multi-step SDE scenarios.
2. Run a baseline LLM (e.g., GPT-5) on these scenarios using standard prompting.
3. Run the same LLM on the scenarios using the Z(7) Discovery Protocol.
4. Compare performance using SDE's metrics: question-level accuracy, project-level success, and error diversity (to test for reduced shared failure modes).

Hypothesis: The Z(7) Protocol will show a statistically significant improvement in project-level success and a lower error correlation with the baseline model, validating its utility as an inductive bias.

8 Conclusion

We have presented a unified framework for discrete harmonic attractors (Z(n)) manifesting across physical and informational scales. The recent discovery of strong photonic magnetic torque provides a physical substrate for Z(n) phase locking, while SilentWitness demonstrates a formally isomorphic dynamical structure in a computational system.

Key contributions in this revised version:

1. Justified the choice of Z(7) through the physically realizable design of 7-fold quasicrystalline substrates.
2. Formalized the information-physics isomorphism by mapping the dynamics of both systems to a common mathematical framework of gradient descent on a constrained landscape.

3. Operationalized the concept of a Z(n) inductive bias for LLMs by proposing the concrete, multi-step Z(7) Discovery Protocol.

4. Provided clear, testable validation pathways for the photonic system, the computational isomorphism, and the LLM enhancement protocol.

The correspondence between SilentWitness and Z(n) photonic systems suggests that *discrete harmonic attractors are scale-invariant organizing principles*, appearing wherever systems balance stability with information preservation. By providing a physically motivated, structurally coherent, and practically implementable framework, this work offers a promising new direction for enhancing scientific discovery, from the laboratory bench to the computational frontier.