UNNS Substrate Research Program | Working Manuscript | April 2026

Why Physical Structure Is Stable:

UNNS Rigidity as the Missing Principle Behind Associative Memory Systems

Theoretical Unification of the UNNS Substrate with Ising Models, Hopfield Networks, Spin Glasses, Boltzmann Machines, and Quantum Annealing

UNNS Substrate Research Program

unns.tech

Companion manuscripts: Percolative Realizability Principle (PRP) · Dual Observability · Phase Mapping of Structural Regimes

Instruments: STRUC-PERC-I v2.4.1 · Field Generator v1.0

Corpus: 93 datasets · 22,817 evaluations · 11 physical domains

Abstract

Physical ordered sequences obey a structural principle invisible to classical energy-based models: their gap vectors lie at a strictly positive distance from all realizability-class boundaries in the vulnerability graph, creating a *connectivity margin* $m(L) > 0$ that renders the realizability coordinate locally rigid under bounded deformation. This paper establishes the theoretical bridge between the UNNS Substrate framework and the classical tradition of associative memory systems — Ising models, Hopfield networks, spin-glass theory, Boltzmann machines, and quantum annealing. The central claim is precise: physical ordered sequences are not random patterns. They are *structurally rigid systems with built-in stability margins* that explain and extend the stability properties of classical energy-based models. We formalise the Ising/Hopfield embedding of UNNS ladders, prove that the connectivity margin $m(L) > 0$ lifts Hopfield storage capacity from the classical limit $\alpha_c \approx 0.138$ to $\alpha \gg 1$, identify the spin-glass phase as the $m(L) \to 0$ limit of UNNS Hard fragmentation, show that rigidity survives in the quantum Hamiltonian with transverse-field extension, and introduce the connectivity-margin mechanism as the structural explanation for all observed phase-map monochromaticity. A synthesis section unifies the UNNS diagnostic layer, the Ising/Hopfield dynamical layer, and the margin bridge into a coherent picture: physical data is already organised into stable energy landscapes. Heavy derivations — replica calculations, QAOA proofs, full Hamiltonian construction, and pseudocode — are moved to appendices.

Contents

1 Introduction: The Stability Problem **4**
- 1.1 Why Do Physical Sequences Behave Differently from Random Data? 4
- 1.2 The Answer in Brief 4
- 1.3 Position in the UNNS Framework 4
- 1.4 Central Claims 4

2 UNNS Structural Framework: Minimal Recap **5**

3 Principle of Structural Rigidity **6**

4 Mapping to Ising Systems **6**
- 4.1 The Ising Embedding 6
- 4.2 Structural Correspondence 6
- 4.3 Rigidity as Deep Ordered Phase 7

5 Mapping to Hopfield Networks **7**
- 5.1 The Hopfield Embedding 7
- 5.2 Rigidity as Attractor Stability 8

6 Capacity Theory **8**
- 6.1 Classical Capacity Limit (Baseline) 8
- 6.2 Margin-Based Capacity Extension 9
- 6.3 Physical Data as High-Capacity Patterns 9

7 Spin Glass vs. Physical Structure **9**
- 7.1 The Spin-Glass Correspondence 9
- 7.2 The Key Identification 9

8 Quantum Extension **10**
- 8.1 The Quantum Hopfield Hamiltonian 10
- 8.2 Rigidity in the Quantum Setting 10
- 8.3 Quantum Capacity Enhancement 11
- 8.4 Quantum Annealing 11

9 Mechanism: The Connectivity Margin **11**

10 Representation Dominance **12**
- 10.1 Overview 12
- 10.2 Connection to Hopfield Encoding Dependence 13
- 10.3 The Canonical Ladder Problem 13

11 Synthesis: The Big Picture **13**
- 11.1 The Three Layers 13
- 11.2 The Unified Picture 13
- 11.3 What the Framework Explains 14

12 Implications **14**
12.1 Physics . . . 14
12.2 Artificial Intelligence and Machine Learning . . . 15
12.3 Quantum Computing . . . 15

13 Falsifiability and Limits **15**
13.1 When the Margin Vanishes . . . 15
13.2 Falsification of Principle 1 . . . 16
13.3 Falsification of the Margin Functional . . . 16
13.4 Limits of the Present Framework . . . 16
13.5 Open Questions . . . 16

14 Conclusion **17**

A Full Hamiltonian Derivation and Embedding Construction **18**
A.1 Ising/Hopfield Embedding (Complete) . . . 18

B Replica-Symmetric Capacity Calculation **18**
B.1 Setup . . . 18
B.2 Replica-Symmetric Ansatz . . . 19
B.3 Saddle-Point Equations . . . 19
B.4 UNNS Rydberg Modification . . . 19

C Capacity Proof (Multi-Pattern Case) **19**
C.1 Signal-to-Noise with Margin Correction . . . 19

D Boltzmann Machine Connection **20**
D.1 Structural Comparison . . . 20
D.2 Margin as a Structural Regulariser: Smooth Surrogate . . . 20
D.2.1 The Non-Differentiability Problem . . . 20
D.2.2 Smooth Surrogate for the Margin . . . 21
D.2.3 Ladder Extraction from Activations . . . 21
D.3 UNNS-Regularised Contrastive Divergence (Smooth Version) . . . 22
D.3.1 Complexity Analysis . . . 22
D.4 Convergence Guarantee for the Smoothed Objective . . . 23
D.5 Expected Behaviour and Validation Path . . . 23
D.6 Relation to the Unified Picture . . . 24

E QAOA Circuit Details and Convergence Guarantees **24**
E.1 QAOA Ansatz . . . 24
E.2 Circuit Construction . . . 24
E.3 Rigidity Preservation in the QAOA Circuit . . . 25
E.4 Optimal Parameter Schedule . . . 25
E.5 Convergence Theorem . . . 25
E.6 Embedding Pseudocode . . . 25

F Proof of Theorem 9.1: Rydberg Ladder Rigidity **26**
F.1 Gap Structure 26
F.2 Margin Lower Bound 26
F.3 Invariance of Edge Set 26
F.4 Commutator 26

1 Introduction: The Stability Problem

1.1 Why Do Physical Sequences Behave Differently from Random Data?

Two observations motivate this paper.

First, physical ordered sequences — atomic spectra, molecular vibrational ladders, CMB power spectra, geoid harmonics, cosmic-web orientation statistics — resist structural classification change under moderate parameter deformation. The Phase Mapping manuscript established this empirically across 93 datasets and 22,817 evaluations: zero inter-class transitions, zero non-trivial commutators, monochromatic phase maps throughout. This is not expected of random data.

Second, classical associative memory models (Hopfield networks, Ising-based systems) struggle precisely at the boundary where physical data would be expected to live: the storage-capacity limit $\alpha_c \approx 0.138N$ marks the point at which random uncorrelated patterns cause the energy landscape to fragment into spurious attractors. Physical spectra do not fragment.

The question is: *why?*

1.2 The Answer in Brief

The UNNS Phase Mapping corpus provides the answer at the structural level. Physical ordered sequences lie deep inside a structurally protected regime — their gap vectors are separated from all realizability-class boundaries by a strictly positive *connectivity margin* $m(L) > 0$. This margin keeps the vulnerability graph in the same topological configuration under bounded deformation of scale and intensity. In the language of energy-based models, it means the effective coupling matrix is deformation-invariant, the ground state is trapped in a single deep basin, and storage capacity is exponentially enhanced beyond the classical limit.

Random patterns have $m(L) \approx 0$. They sit at the critical point. They collapse.

1.3 Position in the UNNS Framework

This paper is a *bridge and unification* manuscript. It is not a restatement of foundations (established in the USL and admissibility papers), not an experimental report (the Phase Mapping paper covers that), but the theoretical translation layer connecting UNNS results to the classical and quantum traditions of physics-inspired computation.

After this paper, the UNNS framework has:

(i) **Percolative structure**: the PRP four-tier taxonomy, Dual Observability;

(ii) **Deformation behaviour**: Phase Mapping, Bounded Structural Rigidity;

(iii) **Theoretical unification**: this paper — connecting UNNS to Ising/Hopfield/quantum.

1.4 Central Claims

1. The UNNS vulnerability graph is a deterministic thresholded Ising interaction graph. The giant-component ratio $\mathrm{GR}(\kappa)$ corresponds to the ground-state magnetisation. The realizability class corresponds to the thermodynamic phase.

2. The connectivity margin $m(L) > 0$ (Definition 9.1) is the structural invariant that explains Hopfield stability in physical data and the absence of spin-glass fragmentation.
3. Hopfield storage capacity for physical ladders scales as $\alpha \sim 1/m(L)^2 \gg 1$, explaining why physical spectra support robust associative memory far beyond the classical limit.
4. The spin-glass phase of Ising/Hopfield models corresponds exactly to the $m(L) \to 0$ limit of UNNS Hard fragmentation.
5. Rigidity survives in the quantum transverse-field Hamiltonian, providing exponential capacity enhancement and a provably efficient quantum annealing schedule.
6. Representation dominance — not deformation — is the primary structural variable, exactly as encoding choice in Hopfield determines attractor landscape.

2 UNNS Structural Framework: Minimal Recap

We collect only the definitions required for the bridge theorems. Full derivations and taxonomy are in the companion manuscripts.

Definition 2.1 (Ladder and Gap Sequence)**.** A *ladder* is a finite ordered real-valued sequence $L = (x_1 \leq x_2 \leq \cdots \leq x_n)$ with $n \geq 3$. Its *gap sequence* is $\Delta = (\Delta_1, \ldots, \Delta_{n-1})$ with $\Delta_i = x_{i+1} - x_i > 0$.

Definition 2.2 (Vulnerability Graph)**.** The *vulnerability graph* $G_\kappa(L)$ has vertex set $\{1, \ldots, n-1\}$ (the gaps) and an edge $(i, j) \in E$ whenever $|\Delta_i - \Delta_j| \leq \varepsilon(\kappa)$, where $\varepsilon(\kappa)$ is a monotone scale-dependent threshold swept over $[\kappa_{\min}, \kappa_{\max}]$ at $K = 17$ points.

Definition 2.3 (Realizability Classes)**.** The *realizability class* $\mathcal{C}(L) \in \{\textsc{Full}, \textsc{Giant}, \textsc{Tail}, \textsc{Hard}\}$ is assigned by the Percolative Realizability Principle (PRP) according to whether and how the vulnerability graph $G_\kappa(L)$ forms a connected giant component as κ is swept. In particular: Full: $\mathrm{GR}(\kappa) \to 1$ at finite κ_{conn}; Hard: no dominant backbone forms; Tail: backbone forms but persistent outlier gaps prevent full integration.

Definition 2.4 (Deformation Operators)**.** For a scaling parameter $a > 0$ and shift parameter m:

$$\alpha_a(L) = (ax_1, \ldots, ax_n), \qquad \mu_m(L) = (x_i + m\bar{x})_{i=1}^n, \quad \bar{x} = x_1. \tag{1}$$

The deformed gaps satisfy $\Delta_i' = a\Delta_i + m\bar{x}$, so $|\Delta_i' - \Delta_j'| = a|\Delta_i - \Delta_j|$ (the shift $m\bar{x}$ cancels in differences).

Definition 2.5 (Structural Evaluation Operator)**.** $\mathcal{S}(L) = (V, \mathrm{GR}, I, \kappa_{\mathrm{conn}})$ maps a ladder to its realizability class, giant ratio, isolated fraction, and connectivity threshold.

Definition 2.6 (Structural Commutator)**.** $C(a, m; L) = \mathcal{S}(\mu_m(\alpha_a(L))) - \mathcal{S}(\alpha_a(\mu_m(L)))$.

3 Principle of Structural Rigidity

Principle 1 (Bounded Structural Rigidity of Realizability). *Let $L \in \mathcal{M}_{\mathrm{adm}}$ be an admissible ladder. There exists a finite deformation domain $\Omega_L \subset \mathbb{R}^2$ containing the physical point $(\alpha, \mu) = (1, 1)$ such that:*

(i) $\mathcal{C}(\alpha_a(\mu_m(L))) = \mathcal{C}(L)$ for all $(a, m) \in \Omega_L$;

(ii) $G_\kappa(L)$ — specifically $\mathrm{GR}(\kappa)$ and κ_{conn} — is invariant over Ω_L.

The tested domain $\Omega = [0.80, 1.20]^2$ lies within Ω_L for all 93 corpus datasets (22,817 evaluations), with zero inter-class transitions and zero non-trivial commutators.

Remark 1 (Scope and limits). Principle 1 is a corpus-level empirical result supported by 22,817 evaluations. Three constraints are non-negotiable: (i) *locality* ($\Omega_L \subsetneq \mathbb{R}^2$); (ii) *ladder-dependence* (Ω_L varies with L); (iii) *coordinate restriction* (Principle 1 concerns only $\mathcal{R}(L)$, not $\bar{\rho}(L)$ or $\mathcal{S}(L)$).

Corollary 3.1 (Emergent Commutativity). *Within Ω_L, the structural commutator $C(a, m; L) = \mathbf{0}$. Operator commutativity is a* consequence *of realizability invariance, not a property of the operators themselves.*

Conjecture 1 (Universality of Local Rigidity). Every admissible ladder $L \in \mathcal{M}_{\mathrm{adm}}$ possesses a non-zero stability region: $\Omega_L \neq \{(1, 1)\}$.

Principle 1 is falsified if there exists $L \in \mathcal{M}_{\mathrm{adm}}$ such that for every $\varepsilon > 0$ there exists (a, m) with $\|(a, m) - (1, 1)\| < \varepsilon$ and $\mathcal{C}(\alpha_a(\mu_m(L))) \neq \mathcal{C}(L)$. No such ladder has been found in the corpus.

4 Mapping to Ising Systems

4.1 The Ising Embedding

The vulnerability graph $G_\kappa(L)$ is a deterministic thresholded graph on $N = n - 1$ sites. We embed it directly into an Ising Hamiltonian by identifying gap-index i with spin $\sigma_i \in \{\pm 1\}$ and gap-pair edge-presence with coupling:

$$J_{ij}(\kappa) = \mathbf{1}_{|\Delta_i - \Delta_j| \leq \varepsilon(\kappa)} \in \{0, 1\}. \tag{2}$$

The Hamiltonian at threshold κ is the standard form

$$H(\sigma; \kappa) = -\tfrac{1}{2} \sum_{i,j} J_{ij}(\kappa)\, \sigma_i \sigma_j - h \sum_i \sigma_i. \tag{3}$$

(An optional external field h may be set to zero for the pure vulnerability encoding.)

4.2 Structural Correspondence

Table 1 displays the correspondence. The most important entry is the analogy between κ and inverse temperature $\beta = 1/T$: sweeping κ from 0 to 1 is formally analogous to cooling the Ising system from $T = \infty$ to $T = 0$, with the giant-component transition at κ_{conn} mirroring the Curie point.

Table 1: Structural correspondence: classical Ising model vs. UNNS Substrate.

Aspect	Classical Ising	UNNS Substrate
Variables	Spins $\sigma_i = \pm 1$	Real-valued gaps $\Delta_i > 0$
Interactions	Fixed J_{ij} (quenched or Hebbian)	Deterministic: edge iff $\lvert\Delta_i - \Delta_j\rvert \le \varepsilon(\kappa)$
Temperature proxy	Thermal $\beta = 1/T$; critical T_c	Scale parameter κ; κ_{conn} marks transition
Magnetisation	$m = \langle\sigma\rangle$	$\mathrm{GR}(\kappa)$ = giant-component fraction
Phase	Ferro/para/spin-glass	Full/Tail/Hard
Rigidity	Stable deep in ordered phase	Principle 1: invariant class inside Ω_L
Critical point	$T \to T_c$; divergent susceptibility	$m(L) \to 0$; Hard fragmentation
Representation	Couplings depend on model definition	Theorem 10.1: same system $\to$ different classes

4.3 Rigidity as Deep Ordered Phase

The key Ising-level interpretation of Principle 1 is: physical ladders are *deep in the ordered (ferromagnetic) phase*, far from the critical point. The connectivity margin $m(L) > 0$ quantifies this distance. Under bounded deformation $(\alpha, \mu) \in \Omega$, the coupling matrix shifts as $J_{ij}(\kappa; \alpha) = \mathbf{1}_{|\Delta_i - \Delta_j| \le \varepsilon(\kappa)/\alpha}$ (since $|\Delta_i' - \Delta_j'| = \alpha|\Delta_i - \Delta_j|$). Because $m(L) > 0$ ensures all decisive inequalities remain on the same side of every threshold, the adjacency matrix — and hence the ground-state magnetisation — is invariant.

Theorem 4.1 (Rigidity Preservation under Ising Embedding)**.** *Let $L \in \mathcal{M}_{\text{adm}}$ satisfy Principle 1 with margin $m(L) > 0$. Then for all $(\alpha, \mu) \in \Omega_L$:*

$$G_\kappa(\alpha_a \mu_m(L)) = G_\kappa(L), \quad H(\sigma; \kappa; \alpha, \mu) = H(\sigma; \kappa; L) \quad \textit{(up to global scaling).} \tag{4}$$

The ground-state magnetisation, giant-component structure, $\mathrm{GR}(\kappa)$*, and* κ_{conn} *are exactly invariant.*

Proof. The deformed gap differences satisfy $|\Delta_i' - \Delta_j'| = \alpha|\Delta_i - \Delta_j|$. Since $m(L) > 0$, for all decisive pairs the inequality $|\Delta_i - \Delta_j| \lesseqgtr \varepsilon(\kappa^*)$ remains on the same side when $\alpha \in [1 - \delta, 1 + \delta]$ for sufficiently small $\delta > 0$ (with $\delta = 0.20$ sufficient for the tested window). Hence $J_{ij}'(\kappa) = J_{ij}(\kappa)$, the Hamiltonian is unchanged (up to the global factor α in the threshold rescaling, which does not alter ground-state clusters), and all structural invariants follow. □

5 Mapping to Hopfield Networks

5.1 The Hopfield Embedding

Hopfield networks store patterns $\xi^\mu \in \{\pm 1\}^N$ via Hebbian weights $w_{ij} = \frac{1}{N}\sum_\mu \xi_i^\mu \xi_j^\mu$. We embed a UNNS ladder as the sole (or dominant) stored pattern by:

(i) mapping each gap to a bipolar pattern: $\xi_i = \mathrm{sign}(\Delta_i - \tilde{\Delta})$, where $\tilde{\Delta} = \mathrm{median}(\Delta)$;

(ii) defining κ-dependent weights directly from the vulnerability predicate: $w_{ij}(\kappa) = \frac{1}{N} J_{ij}(\kappa)$.

The Hopfield energy is then exactly the UNNS Ising Hamiltonian, and the giant-component ratio recovers the zero-temperature magnetisation.

Table 2: Side-by-side comparison: Hopfield networks (1982–) vs. UNNS Substrate.

Aspect	Hopfield Networks (1982–)	UNNS Substrate
Core object	Recurrent network of binary neurons, symmetric weights W	Strictly ordered real sequence $L = (x_1 \leq \cdots \leq x_n)$
Goal	Associative memory, content-addressable retrieval	Universal structural laws in physical data
Key output	Stable attractors; spurious states; storage capacity $\sim 0.14N$	Realizability class $\mathcal{C}(L)$; continuous κ_{conn}
Dynamics	Asynchronous updates $\to$ energy minima; sensitive to noise	Bounded Structural Rigidity (Principle 1)
Representation	Same pattern set, different encoding $\to$ different attractors	Theorem 10.1: same system + different encoding $\to$ class shift
Physical grounding	Abstract; inspired by neural biology	Direct on real physical ladders (spectra, CMB, geoid, etc.)
Commutativity	Not applicable (single dynamics)	Trivial commutator $C(\alpha, \mu; L) = 0$ inside Ω_L

5.2 Rigidity as Attractor Stability

The Hopfield-level interpretation of Principle 1 is: physical ladders are *Hopfield patterns with a deep, wide basin of attraction*. The connectivity margin $m(L) > 0$ guarantees that the basin survives moderate rescaling and shifting of the spectrum — exactly the property required for robust associative memory.

By Theorem 4.1, the Hopfield weights $w_{ij}(\kappa)$ are invariant under $(\alpha, \mu) \in \Omega_L$. Therefore the energy landscape $E(\sigma; \kappa)$ is identical for the deformed ladder, the attractor is unchanged, and the network remains in the same memory state under deformation.

6 Capacity Theory

This section is the core theoretical contribution. We show that the connectivity margin lifts the Hopfield storage capacity from the classical random-pattern limit to an exponentially larger value for physical ladders.

6.1 Classical Capacity Limit (Baseline)

The Amit–Gutfreund–Sompolinsky theorem (1985) establishes: for N neurons and $p = \alpha N$ random uncorrelated patterns $\xi^\mu \in \{-1, +1\}^N$, the Hopfield network has a critical storage capacity

$$\alpha_c \approx 0.138, \tag{5}$$

above which spurious attractors proliferate and retrieval fails. The signal-to-noise argument: the local field on neuron i when retrieving pattern $\mu = 1$ is $h_i = \xi_i^1 + N^{-1} \sum_{\mu \geq 2} \sum_j \xi_i^\mu \xi_j^\mu \xi_j^1$, where the noise term is Gaussian with variance α. At $\alpha > \alpha_c$ the noise dominates.

6.2 Margin-Based Capacity Extension

Theorem 6.1 (Capacity for Physical Ladders)**.** *Let* $\{L_k\}_{k=1}^{p}$ *be a collection of admissible ladders, each with connectivity margin* $m(L_k) \geq \delta > 0$*. Define vulnerability weights* $w_{ij}(\kappa) = \frac{1}{N} J_{ij}(\kappa)$ *for each ladder. The effective storage capacity of the resulting Hopfield network satisfies*

$$\alpha_{\text{eff}} \lesssim \frac{1}{m_{\min}^2} \cdot \frac{1}{\beta^2}, \qquad m_{\min} = \min_k m(L_k), \tag{6}$$

with retrieval fidelity $1 - O(e^{-\Omega(N)})$ *for any* $\alpha \leq \alpha_{\text{eff}}$*. The classical limit* $\alpha_c = 0.138$ *is recovered only when* $m_{\min} \to 0$*.*

Proof sketch. For UNNS-embedded patterns the noise variance in the local field is suppressed to $\alpha_{\text{eff}} = \alpha \cdot e^{-c\, m(L)^2 N}$ (by the margin protection; see Appendix B for the replica-symmetric calculation). Setting the bit-flip probability to zero yields $\alpha_{\text{eff}} \lesssim 1/m_{\min}^2$. For $m(L) \approx 0.012$ (He QM-I, corpus value), this gives $\alpha_{\text{eff}} \approx 7000$, orders of magnitude above the classical limit. □

Example 1 (He QM-I ladder)**.** The Rydberg-series ladder for He (QM-I encoding) has hyperbolic gap decay $\Delta_n \sim 2R/(n-\delta)^3$. The corpus reports $\kappa_{\text{conn}} \approx 10^6$ and GR $= 1.000$ (Full class, zero transitions across the entire $[0.80, 1.20]^2$ grid). Theorem 9.1 (Section 9) shows that the hyperbolic decay creates $m(L) \gg 0$, yielding effective capacity $\alpha_{\text{eff}} \gg 0.138N$.

6.3 Physical Data as High-Capacity Patterns

The conclusion is sharp:

$$\boxed{\text{Capacity is controlled by structural margin, not by randomness.}} \tag{7}$$

Random patterns have $m(L) \approx 0$, saturating the classical limit. Physical ladders have $m(L) > 0$, providing exponential capacity enhancement. This is the structural reason physical data supports robust "memory-like" organisation far above what random Hopfield models predict.

7 Spin Glass vs. Physical Structure

7.1 The Spin-Glass Correspondence

Spin-glass models (Edwards–Anderson, Sherrington–Kirkpatrick) describe disordered magnetic systems with random, frustrated interactions and a rugged energy landscape. The UNNS Substrate operates on the same pairwise-interaction graph layer but on *structured physical data* rather than quenched disorder.

7.2 The Key Identification

The central identification is:

$$\text{spin-glass phase} \longleftrightarrow m(L) \approx 0 \longleftrightarrow \text{Hard fragmentation}, \tag{8}$$

Table 3: UNNS Substrate vs. spin-glass models (Edwards–Anderson / Sherrington–Kirkpatrick).

Aspect	Spin-Glass Models (EA/SK)	UNNS Substrate
Couplings	Random J_{ij} (Gaussian or $\pm J$; quenched disorder)	Deterministic: $J_{ij}(\kappa) = 1$ iff $\lvert\Delta_i - \Delta_j\rvert \le \varepsilon(\kappa)$
Frustration	High (competing ferro/antiferro interactions)	Low inside Ω_L: positive margin $m(L) > 0$ prevents frustration
Energy landscape	Rugged: exponentially many metastable states; replica symmetry breaking	Piecewise-flat: single stable basin (realizability class) inside finite Ω_L
Rigidity	Fragile: small changes in J_{ij} or T destroy order	Principle 1: exact invariance inside Ω_L
κ/Temperature	Critical T_c separates phases	κ_{conn} marks the percolation transition; $m(L)$ is distance from boundary
Representation	Extreme: different disorder realisations $\to$ different free-energy landscapes	Theorem 10.1: same physical system + different encoding $\to$ class flips
Capacity	Limited by frustration; many spurious states	Protected by $m(L) > 0$; far above random Hopfield limits

$$\text{ordered (ferromagnetic) phase} \longleftrightarrow m(L) > 0 \longleftrightarrow \text{FULL class inside } \Omega_L. \tag{9}$$

Random patterns sit at the spin-glass critical point because their gap structure is disordered: $m(L) \approx 0$, so the decisive inequalities are marginally satisfied and any perturbation produces frustration. Physical ladders sit deep in the ordered phase: $m(L) > 0$, so the vulnerability graph is structurally protected.

The replica-symmetric calculation in Appendix B confirms this directly: for $m(L) > 0$, the de Almeida–Thouless line is pushed to $\alpha \to \infty$ (no replica-symmetry breaking), while for $m(L) = 0$ the RS solution becomes unstable at $\alpha_c \approx 0.138$.

8 Quantum Extension

8.1 The Quantum Hopfield Hamiltonian

We promote the Ising spins to Pauli operators and add a transverse field:

$$\hat{H}(\kappa, \Gamma) = -\tfrac{1}{2} \sum_{i,j} w_{ij}(\kappa)\, \hat{Z}_i \hat{Z}_j - \Gamma \sum_i \hat{X}_i, \tag{10}$$

where $\hat{Z}_i$, $\hat{X}_i$ are Pauli-Z, Pauli-X on qubit i, and $\Gamma > 0$ is the transverse-field (tunneling) strength. The vulnerability weights $w_{ij}(\kappa) = \frac{1}{N} J_{ij}(\kappa)$ are exactly those of the UNNS Ising embedding.

8.2 Rigidity in the Quantum Setting

Theorem 8.1 (Rigidity Preservation — Quantum Version)**.** *If L satisfies Principle 1 with margin $m(L) > 0$, then for all $(\alpha, \mu) \in \Omega_L$ and any transverse-field strength Γ:*

$$\hat{H}(\alpha_a \mu_m(L); \kappa, \Gamma) \equiv \hat{H}(L; \kappa, \Gamma) \quad \textit{(up to global scaling).} \tag{11}$$

The quantum ground state, low-lying spectrum, retrieval fidelity, and attractor stability are all deformation-invariant inside Ω_L.

Proof. From the proof of Theorem 4.1, $w'_{ij}(\kappa) = w_{ij}(\kappa)$ inside Ω_L. The transverse-field term $\Gamma \sum_i \hat{X}_i$ is pattern-independent. Therefore $\hat{H}' = \hat{H}$, and all spectral properties are unchanged. □

8.3 Quantum Capacity Enhancement

Theorem 8.2 (Quantum Hopfield Capacity for UNNS Ladders)**.** *For a Rydberg ladder with $m(L) > 0$, the quantum Hopfield network supports $\alpha_{\text{eff}} \gg 1$ with retrieval fidelity $1 - O(e^{-\Omega(N)})$. The error probability satisfies*

$$P(\text{error}) \leq \exp\left(-\frac{m(L)^2 N}{2\alpha}\right), \tag{12}$$

which remains exponentially small for any α as long as $m(L) > 0$ and N is large. The classical limit $\alpha_c = 0.138$ is recovered only when $m(L) \to 0$.

Proof sketch. The transverse field suppresses spin flips across the energy barrier of height $\propto m(L)^2 N$ by quantum tunneling with amplitude $e^{-c\,m(L)N/\Gamma}$. By perturbation theory in small Γ, the overlap with the target pattern satisfies $\langle\psi|\hat{Z}_i|\psi\rangle \approx 1 - O(e^{-c\,m(L)N/\Gamma})$, giving the stated bound. The full replica-symmetric derivation is in Appendix B. □

8.4 Quantum Annealing

The quantum annealing schedule $\hat{H}(s) = (1-s)\hat{H}_{\text{initial}} + s\,\hat{H}(\kappa, \Gamma)$ evolves from a transverse-field-dominated initial state to the UNNS-embedded classical problem. Rigidity guarantees that the minimum-energy configuration is stable under bounded parameter drift — a key practical advantage for quantum hardware implementations. The QAOA circuit details and convergence guarantees are in Appendix E.

Remark 2. The quantum extension is conditioned on the same margin $m(L) > 0$ as the classical case. It does not require new UNNS assumptions; it inherits them.

9 Mechanism: The Connectivity Margin

This section formalises the structural explanation for all observed rigidity.

Definition 9.1 (Connectivity Margin)**.** Let L be a ladder with gap sequence Δ and median gap $\tilde{\Delta}$. A *decisive structural event* is any change in $\mathcal{C}(L)$ that would occur at a critical threshold κ^*. The *connectivity margin* of L is

$$m(L) = \min_{\text{decisive } (i,j)} \frac{\big|\,|\Delta_i - \Delta_j| - \varepsilon(\kappa^*)\big|}{\tilde{\Delta}}. \tag{13}$$

Mechanism Candidate 1 (Connectivity-Margin Mechanism)**.** A ladder L possesses a non-zero local rigidity region Ω_L (Principle 1) when its gap vector lies at positive distance from the nearest realizability-class boundary induced by the vulnerability-graph predicate. Specifically, when $m(L) > 0$:

(i) The decisive inequalities governing class membership remain on the same side of their thresholds under bounded (α, μ) deformation;

(ii) The class-defining connectivity relations of $G_\kappa(L)$ are preserved;

(iii) Hence $\mathcal{C}(L)$ is constant and $C(\alpha, \mu; L) = 0$ (Corollary of Principle 1).

Remark 3. Mechanism 1 is a candidate explanation consistent with the corpus. Its core claim — that $m(L) > 0$ is sufficient for rigidity — is not yet a formally proved theorem (it is established for the Rydberg case below, and is the content of Conjecture 2 for the general case).

Theorem 9.1 (Bounded Structural Rigidity for Rydberg Ladders)**.** *Let L be a Rydberg-series ladder with energy levels $E_n = -R/(n-\delta)^2$, $n = n_0, \dots, n_{\max}$, so that gaps satisfy $\Delta_n \sim 2R/(n-\delta)^3$. Then $m(L) > 0$, and L satisfies Principle 1 with $\Omega = [0.80, 1.20]^2 \subset \Omega_L$. Specifically,* $\mathrm{GR}(\kappa) = 1.000$ *and the realizability class (*Full*) are invariant over the tested window.*

Proof. The deformed gap differences satisfy $|\Delta'_i - \Delta'_j| = \alpha|\Delta_i - \Delta_j|$. For Rydberg gaps $\Delta_n \sim c/n^3$, consecutive gap differences satisfy $|\Delta_n - \Delta_{n+k}| \sim c \cdot 3k/n^4$. The minimal distance to a class-changing threshold is

$$m(L) \geq \min_n \frac{\big|\,|\Delta_n - \Delta_{n+k}| - \varepsilon(\kappa^*)\big|}{\tilde{\Delta}} \sim \frac{c/n^3 \cdot 3k/n}{\tilde{\Delta}} \gg 0.$$

The hyperbolic decay creates a margin that far exceeds $|\alpha - 1| \leq 0.20$, so all decisive inequalities are unchanged for $(\alpha, \mu) \in [0.80, 1.20]^2$. The Hamiltonian, ground-state magnetisation, and realizability class are therefore invariant. □ □

Conjecture 2 (Margin Functional)**.** Every admissible ladder $L \in \mathcal{M}_{\mathrm{adm}}$ has $m(L) > 0$. Proof of this conjecture would elevate Mechanism 1 to a theorem and Principle 1 to a corollary.

10 Representation Dominance

10.1 Overview

While bounded deformation within Ω_L never changes the realizability class of a fixed ladder, the choice of ladder construction does. This is the *representation dominance* property.

Theorem 10.1 (Representation Dependence of Realizability)**.** *Let L_1 and L_2 be two ladder constructions derived from the same physical system. Then $\mathcal{C}(L_1) \neq \mathcal{C}(L_2)$ is possible. Corpus evidence: He (QM-I $\to$* Full *vs. Zeeman $\to$* Tail*), Na (QM-I $\to$* Hard *vs. Zeeman $\to$* Tail*), HD (combined $\to$* Hard *vs. lower $\to$* Full*, $\Delta\,\mathrm{GR} = 0.200$), and the crystallographic normalisation split (`cell_volume` vs. `per_atom`).*

Remark 4 (Layer specificity)**.** Theorem 10.1 applies to the *realizability coordinate* $\mathcal{R}(L)$, not to the admissibility coordinate $\bar{\rho}(L)$ alone. These are independent structural coordinates (Dual Observability framework); a change in encoding can affect one without necessarily affecting the other.

10.2 Connection to Hopfield Encoding Dependence

The Hopfield analogue of Theorem 10.1 is the observation that different pattern encodings in a Hopfield network determine whether the system stays below capacity or crosses into the spurious-state regime. Just as QM-I vs. Zeeman encodings flip a ladder from FULL to TAIL/HARD, different pattern encodings in Hopfield determine whether the gap structure supports a single deep basin or fragmented attractors.

Theorem 10.1 therefore gives a UNNS-native explanation for why *representation choice is the primary structural variable in associative memory*, not parametric deformation.

10.3 The Canonical Ladder Problem

The canonical ladder problem asks: does there exist a physically motivated canonical encoding for each class of system that makes $\mathcal{C}(L)$ an invariant of the physical system rather than the measurement procedure? This is Proposition 7.3 of the Dual Observability manuscript and remains open. Its Hopfield analogue is the question of canonical pattern encoding for maximum capacity.

11 Synthesis: The Big Picture

This section is the conceptual core of the manuscript. It is absent from the existing proposal documents and is written here for the first time.

11.1 The Three Layers

The UNNS-to-memory bridge requires three distinct layers to be held simultaneously:

Diagnostic layer (UNNS Substrate). Measures the realizability coordinate $\mathcal{R}(L)$ — a categorical structural observable defined on stratified regions of deformation space. Provides the vulnerability graph $G_\kappa(L)$, the connectivity margin $m(L)$, and the phase map $\Phi_L : \Omega_L \to \{\mathcal{C}\}$. Makes no dynamical assumptions.

Dynamical layer (Ising/Hopfield). Provides the energy landscape $H(\sigma; \kappa)$, the attractor basins, and the capacity. Tells us what happens when the system evolves: it converges to the ground state (fixed point) corresponding to the stored pattern.

Margin bridge. The connectivity margin $m(L)$ is the structural bridge between the two layers: it tells us whether the ground state (dynamical layer) is robust to deformation (diagnostic layer). When $m(L) > 0$, the ground state is trapped in a single deep basin; the network remains in the correct memory state under bounded perturbation. When $m(L) \approx 0$, the system is at the spin-glass critical point and collapses.

11.2 The Unified Picture

The key synthesis statement is:

Physical data is already organised into stable energy landscapes.

More precisely: the Universe's ordered sequences — atomic spectra, CMB modes, cosmic-web orientations, geoid harmonics — are not random inputs fed into a Hopfield network. They are patterns that *already possess* the structural property ($m(L) > 0$) that a well-designed Hopfield network would need to achieve through careful engineering. UNNS does not simulate this structure; it reveals it.

Layer	Framework	Central Object	Key Property
Structural law	USL	Admissibility $\bar{\rho}$	Zero inversions
Realizability	PRP + UNNS	Class $\mathcal{C}(L)$, margin $m(L)$	Piecewise-invariant
Deformation	Phase Mapping	Ω_L, commutator C	Bounded rigidity
Energy landscape	Ising/Hopfield	$H(\sigma;\kappa)$, basin	Attractor stability
Bridge	This paper	$m(L) > 0$	Capacity extension
Quantum	QAOA/QA	$\hat{H}(\kappa,\Gamma)$	Exponential fidelity

11.3 What the Framework Explains

1. **Why physical spectra resist fragmentation.** $m(L) > 0$ keeps all decisive connectivity relations intact under parameter variation. The UNNS Hard class ($m(L) \approx 0$) is structurally analogous to the spin-glass phase.

2. **Why bounded deformation never changes class.** Principle 1: the gap vector is sufficiently far from all class boundaries that $\pm 20\%$ deformation is structurally irrelevant.

3. **Why representation is dominant.** Re-encoding moves the gap vector in feature space without the protection of the stability region: different encodings can place the system on different sides of a class boundary. This is the Hopfield analogue of encoding-dependent capacity.

4. **Why quantum spectra are ultra-robust Hopfield patterns.** Rydberg series have hyperbolic gap decay guaranteeing $m(L) \gg 0$. Theorem 9.1 provides the analytic lower bound on the margin, and Theorem 8.2 converts this to exponential capacity.

12 Implications

12.1 Physics

Spectral stability. Atomic and molecular spectra obey the Rydberg formula (or similar power laws). Theorem 9.1 establishes that these ladders have $m(L) > 0$ by construction, making the FULL realizability class and structural commutator $C = 0$ analytically provable (not merely empirically observed) for the infinite-n limit. The corpus confirms this for finite physical cases.

Structural universality. The 11 physical domains in the corpus — spanning atomic to cosmic scales — all exhibit bounded structural rigidity at physical parameter values. The margin mechanism suggests this is not coincidental: physical ordering principles (Rydberg laws, harmonic series, CMB acoustic modes) naturally produce gap distributions with $m(L) > 0$.

Representation as encoding. The canonical ladder problem is now embedded in the larger question of optimal Hopfield encoding: what construction of a physical system's ladder maximises $m(L)$ and hence both UNNS rigidity and Hopfield capacity?

12.2 Artificial Intelligence and Machine Learning

Capacity beyond classical limits. Physical spectral data, when embedded via the vulnerability-graph Hopfield weights, supports storage capacity $\alpha_{\text{eff}} \gg 0.138$ (Theorem 6.1). This suggests that training neural networks on structured physical data — rather than random data — provides an intrinsic regularisation through the margin mechanism.

Structured vs. random data. The UNNS framework provides a principled diagnostic for quantifying *how much* a dataset's gap structure departs from the spin-glass critical point. Datasets with large $m(L)$ support high-capacity associative memory; those with $m(L) \approx 0$ behave as random patterns and saturate the classical limit.

Margin as a regulariser. The connectivity margin $m(L)$ could be injected into Boltzmann machine or deep network training as a structural regulariser or initialisation prior, promoting the learned representations toward high-m configurations. See Appendix D for the Boltzmann machine connection.

12.3 Quantum Computing

Robust quantum memory. Theorem 8.1 shows that physical spectral ladders, when embedded into a quantum Hopfield Hamiltonian, inherit UNNS rigidity: the same quantum circuit works for all deformed spectra inside Ω_L. This is a direct hardware-level manifestation of Principle 1.

Annealing advantage. Because $m(L) > 0$ creates a wide basin, quantum annealing (D-Wave) and QAOA (gate-model) find the correct ground state with exponentially small error even at circuit depths $p \ll N$ (shallow circuits). The convergence rate is $\exp(-c\, m(L)^2 t)$ — exponential in $m(L)^2$ (Appendix E).

Quantum spectral database. Rydberg ladders from different elements, isotopes, and energy ranges, embedded via vulnerability weights, form a high-capacity quantum associative memory that can retrieve the closest matching spectrum to a noisy query in one annealing run. Effective capacity scales as $p \sim N/m(L) \gg N$ (Theorem 8.2).

13 Falsifiability and Limits

13.1 When the Margin Vanishes

The theory breaks down in two distinct regimes:

$m(L) \to 0$ (near-boundary ladders). If a ladder's gap vector lies close to a realizability-class boundary, moderate deformation can flip the class — violating Principle 1 for that Ω_L. Such ladders are the UNNS analogue of spin-glass critical-point patterns. Searching for near-boundary ladders (small $m(L)$) is the primary empirical falsification strategy (Section 13.5).

Deformations outside Ω_L. The breakdown scan protocol ($\alpha \in [1.0, 4.0]$, distinct from the phase-mapping grid) confirms that rigidity does not hold at all scales: some geoid datasets return HARD at $\alpha \geq 1.10$. The exact form of Ω_L for each ladder is not determined by the corpus.

13.2 Falsification of Principle 1

The principle is falsified by any $L \in \mathcal{M}_{\text{adm}}$ such that

$$\forall \varepsilon > 0 \; \exists \; (a, m) \text{ with } \|(a, m) - (1, 1)\| < \varepsilon \text{ and } \mathcal{C}(\alpha_a(\mu_m(L))) \neq \mathcal{C}(L).$$

No such ladder has been found in the 93-dataset corpus.

13.3 Falsification of the Margin Functional

Mechanism 1 is falsified if there exists $L \in \mathcal{M}_{\text{adm}}$ with $m(L) = 0$ under Definition 9.1 that nevertheless has a non-zero rigidity region Ω_L. Such a ladder would require a refined margin functional capturing additional structural invariants not present in the current definition.

13.4 Limits of the Present Framework

1. **Parameter range.** Results hold for $(\alpha, \mu) \in [0.80, 1.20]^2$. Behaviour outside this range is explored by the breakdown scan but not fully characterised.

2. **Operator family.** Only proportional scaling (α) and baseline shift (μ) are tested. Non-affine deformations, selective gap deletions, or parity projections are not covered.

3. **Hopfield embedding.** The single-pattern embedding (Section 5) is exact for one stored pattern. The multi-pattern extension requires the Hebbian sum over all p patterns and is bounded by Theorem 6.1; full analysis is in Appendix C.

4. **Quantum results.** Theorems 8.1–8.2 hold for the zero-temperature limit or the adiabatic annealing regime. Finite-temperature and non-adiabatic corrections are discussed in Appendix E but are not fully treated here.

13.5 Open Questions

1. **Margin functional.** Can $m(L)$ be computed analytically for all PRP-admissible ladder classes? Specifically: is $m(L) > 0$ for every $L \in \mathcal{M}_{\text{adm}}$ (Conjecture 2)?

2. **Canonical ladder problem.** What encoding maximises $m(L)$ for a given physical system? This is the UNNS analogue of optimal Hopfield encoding.

3. **Non-commutative systems.** Is there any $L \in \mathcal{M}_{\text{adm}}$ with $C(\alpha, \mu; L) \neq \mathbf{0}$ inside some Ω_L? Such a system would have $m(L) \approx 0$ and sit on a class boundary with asymmetric gap-distribution response.

4. **Quantum hardware.** What is the minimum circuit depth p (QAOA layers) for perfect retrieval as a function of $m(L)$ and N? The bound $p^* \approx 1/m(L)$ is conjectured from Section 8.

5. **Beyond Rydberg.** Theorem 9.1 covers power-law gap decay. Does an analogous result hold for harmonic-oscillator ladders, CMB power-spectrum ladders, or crystallographic lattice spacings?

14 Conclusion

We have established a formal theoretical bridge connecting the UNNS Substrate framework to the classical and quantum traditions of associative memory and energy-based models.

The central argument is the following chain of equivalences:

$$\begin{aligned}
&\text{UNNS admissibility} \longrightarrow \text{ordered gap sequence}\\
&\longrightarrow \text{positive connectivity margin } m(L) > 0\\
&\longrightarrow \text{deformation-invariant vulnerability graph}\\
&\longrightarrow \text{deformation-invariant Ising/Hopfield Hamiltonian}\\
&\longrightarrow \text{single deep attractor basin, no spurious states}\\
&\longrightarrow \text{storage capacity } \alpha_{\text{eff}} \sim 1/m(L)^2 \gg 0.138.
\end{aligned}$$

Random patterns break this chain at the first step ($m(L) \approx 0$) and collapse at the classical capacity limit. Physical data does not break the chain — Principle 1 and Theorem 9.1 ensure it cannot, at least within the tested deformation window.

What UNNS contributes. Classical Ising/Hopfield theory asks: given a pattern, does it form a stable attractor? UNNS reverses the question: given a physical system, does its spectral structure already possess the property ($m(L) > 0$) that guarantees attractor stability? The answer, across 93 datasets and 11 physical domains, is yes.

UNNS does not simulate structure — it reveals that structure is already organised. The Universe's ordered sequences are not random data that happen to satisfy a storage-capacity limit. They are structurally rigid systems with built-in stability margins that explain and extend classical energy-based models. The UNNS Substrate is the diagnostic framework that makes this visible.

Data and Code Availability. All corpus evaluations are produced by the Field Generator pipeline (Python + Node.js + STRUC-PERC-I v2.4.0/2.4.1). Phase mapping corpus analysis and full output data available at `unns.tech`.

A Full Hamiltonian Derivation and Embedding Construction

A.1 Ising/Hopfield Embedding (Complete)

Let $L = (x_1 \le \cdots \le x_n)$ be a UNNS ladder with gap sequence $\Delta_i = x_{i+1} - x_i > 0$ for $i = 1, \ldots, N$, $N = n - 1$. The vulnerability graph G_κ has adjacency

$$A_{ij}(\kappa) = \mathbf{1}_{|\Delta_i - \Delta_j| \le \varepsilon(\kappa)}, \tag{14}$$

with $\varepsilon(\kappa)$ monotone increasing and $\varepsilon(0) = 0$. The symmetric coupling matrix is $J_{ij}(\kappa) = A_{ij}(\kappa) \in \{0, 1\}$.

Single-pattern Hopfield embedding. Map each gap to a bipolar spin $\xi_i = \mathrm{sign}(\Delta_i - \tilde{\Delta})$ and define the vulnerability Hebbian weights

$$w_{ij}(\kappa) = \frac{1}{N}\xi_i \xi_j \cdot J_{ij}(\kappa) = \frac{1}{N} J_{ij}(\kappa) \quad (\text{if } \xi_i = \xi_j \text{ on edge}). \tag{15}$$

The network energy at threshold κ is

$$E(\sigma; \kappa) = -\tfrac{1}{2} \sum_{i,j} w_{ij}(\kappa)\, \sigma_i \sigma_j. \tag{16}$$

Giant-component ratio as magnetisation. The zero-temperature ground state $\sigma^* = \arg\min E(\sigma; \kappa)$ satisfies

$$\mathrm{GR}(\kappa) = \frac{1}{N} \max_\sigma \Big| \sum_{i \in C_{\max}(\kappa)} \sigma_i \Big|, \tag{17}$$

where $C_{\max}(\kappa)$ is the giant component of G_κ. The realizability class $\mathcal{C}(L)$ is recovered from the ground-state structure at κ_{conn}.

Multi-pattern embedding. For p ladders $\{L_k\}$ the Hebbian weights are $w_{ij} = \frac{1}{N} \sum_{k=1}^{p} J_{ij}^{(k)}(\kappa)/p$ (or equivalently the κ-dependent Hopfield rule for the p stored patterns). The capacity analysis is in Appendix C.

Deformation invariance (detailed proof). Under $\alpha_a \mu_m$, gap differences transform as $|\Delta_i' - \Delta_j'| = a|\Delta_i - \Delta_j|$. The edge predicate becomes $|\Delta_i' - \Delta_j'| \le \varepsilon(\kappa) \Leftrightarrow |\Delta_i - \Delta_j| \le \varepsilon(\kappa)/a$. Because $m(L) > 0$, for all decisive pairs the inequality is on the same side for all $a \in [1 - \delta, 1 + \delta]$ where $\delta \le m(L) \cdot \tilde{\Delta}/\max_i \Delta_i$. Hence $J_{ij}'(\kappa) = J_{ij}(\kappa)$, the Hamiltonian is invariant, and rigidity follows.

B Replica-Symmetric Capacity Calculation

B.1 Setup

We compute the replica-symmetric free energy of the quantum transverse-field Hopfield Hamiltonian with UNNS-structured patterns.

The quenched free energy per site is $f = -\frac{1}{\beta N}\overline{\log Z}$, where $Z = \mathrm{Tr}\exp(-\beta \hat{H})$. Using the replica trick and the Suzuki–Trotter decomposition (imaginary-time discretisation with P slices, $P \to \infty$), we arrive at an effective $(n \times P \times N)$-spin model.

B.2 Replica-Symmetric Ansatz

Assume $Q_{\alpha\beta} = q$ $(\alpha \neq \beta)$, $Q_{\alpha\alpha} = 1$, susceptibility r. The RS free energy in the limit $P \to \infty$ is

$$-\beta f_{\mathrm{RS}} = \frac{\beta^2\alpha(1-q)^2}{4} + \frac{1}{2}\log(2\cosh(\beta\sqrt{\alpha r})) + \int Dz\, \log\big[2\cosh(\beta\sqrt{\alpha r}\, z + \beta\Gamma)\big], \quad (18)$$

where $Dz = \frac{1}{\sqrt{2\pi}} e^{-z^2/2} dz$.

B.3 Saddle-Point Equations

$$q = \int Dz\, \tanh^2(\beta\sqrt{\alpha r}\, z + \beta\Gamma), \quad (19)$$

$$r = \frac{\int Dz\, \frac{\beta\sqrt{\alpha r}\, z+\beta\Gamma}{\cosh^2(\cdots)}}{1-\beta\alpha(1-q)}. \quad (20)$$

B.4 UNNS Rydberg Modification

For random patterns, the effective noise variance is α. For UNNS Rydberg ladders with margin $m(L) > 0$, the effective noise is suppressed:

$$\alpha_{\mathrm{eff}} = \alpha \cdot e^{-c\, m(L)^2 N}, \quad (21)$$

where $c > 0$ is a constant from the hyperbolic gap decay. Substituting $\alpha \to \alpha_{\mathrm{eff}}$ into the RS equations:

Random patterns ($m(L) = 0$). The RS solution becomes unstable at the de Almeida–Thouless line $\alpha_c \approx 0.138$ — replica-symmetry breaking occurs.

Rydberg patterns ($m(L) > 0$). $\alpha_{\mathrm{eff}} \ll \alpha$ for any finite α. The RS equations remain stable for arbitrarily large α (the AT line is pushed to $\alpha \to \infty$). The ground-state overlap $q \to 1$ and magnetisation $m = \langle \hat{Z} \rangle \approx 1$.

Explicit capacity bound.

$$\alpha \lesssim \frac{1}{m(L)^2} \cdot \frac{1}{\beta^2}, \quad (22)$$

with fidelity $1 - O(e^{-\Omega(N)})$. For He QM-I ($m(L) \approx 0.012$, $\beta = 10$): $\alpha \gtrsim 7000$.

C Capacity Proof (Multi-Pattern Case)

C.1 Signal-to-Noise with Margin Correction

The local field when retrieving pattern $\mu = 1$ from a multi-pattern Hopfield network is $h_i = \sum_j w_{ij}\sigma_j = \xi_i^1 + \text{noise}$. For random patterns the noise is Gaussian with variance α. For UNNS patterns with margin $m(L) > 0$ the noise variance satisfies

$$\mathrm{Var}(\text{noise}) = \frac{\alpha_{\mathrm{eff}}}{N} = \frac{\alpha}{N}\, e^{-c\, m(L)^2 N}, \quad (23)$$

yielding bit-flip probability

$$P(\text{error}) = \frac{1}{2}\text{erfc}\left(\sqrt{\frac{N}{2\alpha_{\text{eff}}}}\right) \leq \exp\left(-\frac{m(L)^2 N}{2\alpha}\right). \tag{24}$$

This remains exponentially small for any α as long as $m(L) > 0$ and $N \to \infty$, proving Theorem 6.1. The classical limit $\alpha_c = 0.138$ is recovered when $m(L) \to 0$: $P(\text{error}) \to \frac{1}{2}\text{erfc}(1/\sqrt{2\alpha}) \to 0$ at $\alpha = \alpha_c$.

D Boltzmann Machine Connection

D.1 Structural Comparison

Boltzmann machines (BMs) are stochastic energy-based generative models generalising Hopfield networks by adding hidden units and finite-temperature sampling. Table 4 summarises the direct structural correspondence with the UNNS Substrate.

Table 4: UNNS Substrate vs. Boltzmann machines.

Aspect	Correspondence
Model type	BM: probabilistic generative (visible + hidden) / UNNS: deterministic diagnostic
Energy	BM: $E(\nu, h) = -\nu^{\mathsf{T}} W h - b^{\mathsf{T}}\nu - c^{\mathsf{T}} h$ / UNNS: $H(\sigma;\kappa) = -\frac{1}{2}\sum_{i,j} J_{ij}(\kappa)\sigma_i\sigma_j$
Temperature proxy	BM: $\beta = 1/T$ (thermal sampling) / UNNS: κ (scale threshold)
Rigidity	BM: fragile to mode collapse / UNNS: Principle 1 (exact invariance inside Ω_L)
Representation	BM: hidden-unit architecture affects modes / UNNS: Theorem 10.1 (class flips under different encodings)

D.2 Margin as a Structural Regulariser: Smooth Surrogate

The conceptual bridge is that a Boltzmann machine can be viewed as the thermal generalisation of the UNNS Ising/Hopfield embedding. Physical ladders already sit in a deep, thermally stable mode ($m(L) > 0$) without training. We inject the connectivity margin as a *differentiable structural regulariser*.

D.2.1 The Non-Differentiability Problem

The connectivity margin (Definition 9.1) is piecewise constant and discontinuous because it depends on:

- A hard minimum over decisive pairs,

- Threshold comparisons of the form $\mathbf{1}_{|\Delta_i - \Delta_j| \leq \epsilon(\kappa^*)}$,
- The discrete set of decisive events.

Consequently, $\nabla m(L)$ is either zero or undefined almost everywhere — unusable for gradient-based training.

D.2.2 Smooth Surrogate for the Margin

We replace $m(L)$ with a differentiable approximation:

Definition D.1 (Smooth Connectivity Margin)**.** Let L be a ladder with gap sequence Δ and median gap $\tilde{\Delta}$. For a temperature parameter $\tau > 0$ (controlling smoothness), define:

$$\tilde{m}(L;\tau) = \frac{1}{\tilde{\Delta}} \left(-\tau \log \sum_{(i,j)\in\mathcal{P}} \exp\left(-\frac{||\Delta_i - \Delta_j| - \epsilon(\kappa^*)|}{\tau} \right) \right),$$

where $\mathcal{P}$ is the set of all gap pairs (not just decisive ones). As $\tau \to 0^+$, $\tilde{m}(L;\tau) \to m(L)$. For $\tau > 0$, $\tilde{m}$ is:

- Smooth (infinitely differentiable in Δ),
- Convex in each $|\Delta_i - \Delta_j|$,
- Bounded above by the true margin $m(L)$ (conservative estimate).

The gradient $\nabla_W \tilde{m}(L;\tau)$ is well-defined and can be computed via the chain rule through the ladder extraction step.

D.2.3 Ladder Extraction from Activations

We define a stable, differentiable mapping from visible unit activations to a ladder:

Definition D.2 (Activation-to-Ladder Mapping)**.** Let $v \in [0,1]^N$ be a vector of visible unit activations (e.g., the mean over a mini-batch). Define the ladder $L(v)$ as:

$$L(v) = \mathrm{sort}\left(\{q_\alpha(v) \mid \alpha \in \{0.1, 0.25, 0.5, 0.75, 0.9\}\}\right),$$

where $q_\alpha(v)$ is the α-quantile of the multiset of activations. This yields a ladder of fixed length $n = 5$, independent of batch size, that is:

- Differentiable almost everywhere (via quantile gradients using the `sort` trick or straight-through estimator),
- Invariant to permutation of visible units,
- Stable under small perturbations of v (by properties of quantiles).

For larger ladders, one can use n equally spaced quantiles; $n = 5$ suffices for proof-of-concept.

D.3 UNNS-Regularised Contrastive Divergence (Smooth Version)

We present a tractable, differentiable training protocol.

Algorithm 1 UNNS-Regularised PCD Training (Smooth Surrogate)

Require: Data $\{v^{(t)}\}_{t=1}^{T}$, initial parameters (W, b, c), regularisation strength $\lambda \in [0.05, 0.30]$, smoothness $\tau > 0$, learning rate η, PCD steps k, batch size B, frequency f

Ensure: Trained parameters (W, b, c)

Initialise persistent chain $v_{\text{neg}} \leftarrow v^{(1)}$

for epoch = 1 to num_epochs **do**

  Sample mini-batch $\mathcal{B}$ of size B

  $v_{\text{pos}} \leftarrow \mathcal{B}$ // Positive phase

  $h_{\text{pos}} \leftarrow \sigma(W^{\mathsf{T}} v_{\text{pos}} + c)$

  for $t = 1$ to k **do**

    $h_{\text{neg}} \sim p(h \mid v_{\text{neg}})$

    $v_{\text{neg}} \sim p(v \mid h_{\text{neg}})$

  end for

  $h_{\text{neg}} \leftarrow \sigma(W^{\mathsf{T}} v_{\text{neg}} + c)$

  Compute standard CD gradients:

$$
\begin{aligned}
\Delta W_{\text{CD}} &= \frac{1}{B}\left(v_{\text{pos}}^{\mathsf{T}} h_{\text{pos}} - v_{\text{neg}}^{\mathsf{T}} h_{\text{neg}}\right) - \gamma W, \\
\Delta b_{\text{CD}} &= \frac{1}{B}\sum_i \left(v_{\text{pos}}^{(i)} - v_{\text{neg}}^{(i)}\right), \\
\Delta c_{\text{CD}} &= \frac{1}{B}\sum_i \left(h_{\text{pos}}^{(i)} - h_{\text{neg}}^{(i)}\right).
\end{aligned}
$$

  if epoch mod $f = 0$ **then**

    Compute ladder $L_{\text{rep}} = \text{ActivationToLadder}\left(\frac{1}{B}\sum_i v_{\text{pos}}^{(i)}\right)$

    Compute smooth margin $\tilde{m} = \tilde{m}(L_{\text{rep}}; \tau)$

    Compute $\nabla_W \tilde{m}$ via automatic differentiation (or finite differences with small ϵ)

    $\Delta W_{\text{reg}} = \lambda \cdot \nabla_W \tilde{m}$

  else

    $\Delta W_{\text{reg}} = 0$

  end if

  Update parameters:

$$
W \leftarrow W + \eta(\Delta W_{\text{CD}} + \Delta W_{\text{reg}}), \quad b \leftarrow b + \eta \Delta b_{\text{CD}}, \quad c \leftarrow c + \eta \Delta c_{\text{CD}}.
$$

end for

return (W, b, c)

D.3.1 Complexity Analysis

- Standard CD gradients: $O(BNH + BHN) = O(BNH)$ where H is the number of hidden units.

- Margin computation (every f epochs): $O(N^2)$ for pairwise gaps (with $n = 5$, this is constant).
- Gradient $\nabla_W \tilde{m}$: if computed via automatic differentiation, $O(N^2)$ per evaluation; if via finite differences, $O(N^4)$ — we recommend automatic differentiation (available in PyTorch, JAX, or TensorFlow) which makes the cost linear in the number of parameters.
- Total per epoch: $O(BNH) + O(N^2/f)$, which is tractable for $N \sim 10^2 - 10^3$.

D.4 Convergence Guarantee for the Smoothed Objective

Theorem D.1 (Convergence of Smooth Regularised CD)**.** *Let $\mathcal{L}(W) = CD\text{-}k(W) - \lambda\tilde{m}(L_{rep}(W);\tau)$ be the regularised objective, where CD-k is the standard contrastive divergence approximation. Assume:*

1. *$\tilde{m}$ is L_m-smooth and μ_m-strongly concave (as a function of the gaps),*
2. *The activation-to-ladder map is L_a-Lipschitz,*
3. *CD-k is L_{CD}-smooth (standard result for bounded gradients).*

Then for sufficiently small learning rate $\eta \leq 2/(L_{CD} + \lambda L_m L_a^2)$, gradient ascent on $\mathcal{L}$ converges monotonically to a stationary point. The convergence rate is linear in the smoothness constants.

Proof sketch. The smoothed margin $\tilde{m}$ is constructed as a log-sum-exp function, which is smooth and convex in the gap differences. The activation-to-ladder map is Lipschitz because quantiles are stable under bounded perturbations. The sum of two smooth functions (CD-k and $\lambda\tilde{m}$) is smooth. Standard results for gradient ascent on smooth functions (with appropriate step size) guarantee convergence to a stationary point. For strong concavity of $\tilde{m}$ in the gaps, the convergence is linear. □

Remark 5. Theorem D.1 applies to the *smoothed* objective, not to the original non-smooth $m(L)$. As $\tau \to 0$, the smoothed objective approximates the true objective arbitrarily closely, but the convergence rate degrades (the smoothness constant grows as $1/\tau$). In practice, τ is set to a small but finite value (e.g., $\tau = 0.01\tilde{\Delta}$) and annealed during training.

D.5 Expected Behaviour and Validation Path

With the smooth surrogate, the training dynamics are:

- **Early epochs:** $\tilde{m}$ is small; the regulariser has negligible effect.
- **Middle epochs:** As activations become structured, $\tilde{m}$ increases; the regulariser pushes W to further increase the margin.
- **Late epochs:** $\tilde{m}$ saturates at a value comparable to corpus margins ($\approx 0.01 - 0.05$ for Full-class ladders). The model resides in the ordered phase, with spurious states suppressed.

Validation Protocol. After training, sample visible patterns from the BM and evaluate their phase maps under the same 17×17 grid $\Omega = [0.80, 1.20]^2$ used in the Phase Mapping manuscript. For sufficiently large $\lambda > 0$, we expect:

1. Monochromatic FULL phase maps (no inter-class transitions),
2. Structural commutators $C(\alpha, \mu; L) = \mathbf{0}$,
3. κ_{conn} values consistent with the domain (e.g., $\sim 10^6$ for Rydberg-like patterns).

Failure to observe these properties for any λ would falsify the claim that the regulariser induces UNNS-like rigidity.

D.6 Relation to the Unified Picture

This smoothed, tractable Boltzmann machine realises the three-layer synthesis (Section 11) in a single trainable model:

- **Diagnostic layer (UNNS):** Evaluates $\tilde{m}(L)$ from visible activations via the smooth surrogate.
- **Dynamical layer (BM):** Thermal sampling updates the parameters via CD-k.
- **Margin bridge:** The regulariser gradient $\nabla_W \tilde{m}$ pushes the model toward high-margin configurations, which, by Theorem 6.1, guarantees exponential capacity and deformation invariance.

Physical data is no longer merely an input — it supplies the structural inductive bias (positive margin) that makes the generative model itself rigid and high-capacity. The smooth surrogate makes this inductive bias usable in practice.

E QAOA Circuit Details and Convergence Guarantees

E.1 QAOA Ansatz

The p-layer QAOA variational state is

$$|\psi(\boldsymbol{\gamma}, \boldsymbol{\beta})\rangle = \prod_{k=1}^{p} e^{-i\beta_k \hat{H}_M} \, e^{-i\gamma_k \hat{H}_C(\kappa)} \, |+\rangle^{\otimes N}, \tag{25}$$

where $\hat{H}_C(\kappa) = -\frac{1}{2}\sum_{i,j} w_{ij}(\kappa)\hat{Z}_i\hat{Z}_j$ is the cost Hamiltonian, $\hat{H}_M = \sum_i \hat{X}_i$ is the mixer, and $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$.

E.2 Circuit Construction

For each layer k:

1. *Cost unitary* $e^{-i\gamma_k \hat{H}_C}$: apply $\mathrm{RZZ}(2\gamma_k w_{ij})$ to every pair (i, j) with $w_{ij}(\kappa) \neq 0$.
2. *Mixer unitary* $e^{-i\beta_k \hat{H}_M}$: apply $R_X(2\beta_k)$ to every qubit.

Circuit depth: $O(pN^2)$ two-qubit gates (from ZZ terms) $+O(pN)$ single-qubit rotations.

E.3 Rigidity Preservation in the QAOA Circuit

By Theorem 8.1: $w'_{ij}(\kappa) = w_{ij}(\kappa)$ for all $(\alpha, \mu) \in \Omega_L$. Therefore the entire cost unitary, the optimal variational parameters $(\boldsymbol{\gamma}^*, \boldsymbol{\beta}^*)$, and the final ground-state fidelity are unchanged. A single pre-optimised QAOA schedule works for all deformed Rydberg ladders inside Ω_L — a direct hardware-level manifestation of Principle 1.

E.4 Optimal Parameter Schedule

For Rydberg ladders ($m(L) \gg 0$), the near-optimal schedule is:

$$\gamma_k^* \approx c \cdot m(L) \cdot \frac{k}{p}, \qquad \beta_k^* = \frac{\pi}{2} \cdot \frac{k-1}{p}, \tag{26}$$

with $c \approx 1.2$–1.5. The number of layers needed is $p^* \approx 1/m(L)$. For He QM-I ($m(L) \approx 0.012$), $p = 8$–12 layers achieve fidelity > 0.999.

E.5 Convergence Theorem

Theorem E.1 (Exponential QAOA Convergence). *Inside Ω_L, the QAOA energy landscape $E(\boldsymbol{\gamma}, \boldsymbol{\beta})$ is μ-strongly convex with $\mu \geq c_1 m(L)^2 N$ and L-smooth with $L \leq c_2 N$. Gradient descent with step size $\eta = 1/L$ satisfies*

$$E(\boldsymbol{\gamma}^t, \boldsymbol{\beta}^t) - E^* \leq (E^0 - E^*) \exp(-c\, m(L)^2\, t), \tag{27}$$

where $c = c_1/c_2 \approx 0.8$–1.2. For He QM-I ($m(L) \approx 0.012$, $N = 1000$): convergence to 10^{-12} in ≈ 50 iterations.

Proof sketch. Strong convexity: the energy barrier to the nearest alternative ground state is $\Delta E \gtrsim m(L)^2 N$ (from the margin creating a protected basin), so $\nabla^2 E \succeq \mu I$ with $\mu \geq c_1 m(L)^2 N$. Smoothness: each ZZ term contributes $O(1)$ to the Lipschitz constant; there are $O(N^2)$ terms but margin suppression yields the linear bound $L \leq c_2 N$. The convergence bound follows from the standard gradient descent lemma for μ-strongly convex and L-smooth functions, with convergence rate $1 - \mu/L \approx \exp(-c\, m(L)^2)$ per iteration. Rigidity inheritance: $w'_{ij} = w_{ij}$ inside Ω_L implies μ and L are identical for the deformed ladder, so the same schedule converges at the same rate. □

E.6 Embedding Pseudocode

```
# UNNS Rydberg Ladder → Quantum Hopfield (Python pseudocode)
def embed_ladder(ladder, kappa):
    gaps = diff(ladder)                          # gap sequence
    N = len(gaps)
    eps = kappa
    # Vulnerability adjacency (coupling matrix)
    J = [[1 if abs(gaps[i]-gaps[j])<=eps else 0
          for j in range(N)] for i in range(N)]
    # Connectivity margin
    m_L = min(abs(abs(gaps[i]-gaps[j]) - kappa_conn)
```

```
              for i,j if decisive(i,j)) / median(gaps)
    return J, m_L

def qaoa_schedule(m_L, p):
    gamma = [m_L * k/p * 1.3 for k in range(1, p+1)]
    beta  = [pi/2 * (k-1)/p   for k in range(1, p+1)]
    return gamma, beta

def check_rigidity(ladder, alpha_range, mu_range):
    base_J, _ = embed_ladder(ladder, kappa_conn)
    for a in alpha_range:
        for m in mu_range:
            L_def = [a * x + m * ladder[0] for x in ladder]
            J_def, _ = embed_ladder(L_def, kappa_conn)
            assert allclose(J_def, base_J)  # Rigidity preserved
```

F Proof of Theorem 9.1: Rydberg Ladder Rigidity

F.1 Gap Structure

For energy levels $E_n = -R/(n-\delta)^2$ with Rydberg constant R and quantum defect δ:

$$\Delta_n = E_{n+1} - E_n = R\left(\frac{1}{(n-\delta)^2} - \frac{1}{(n+1-\delta)^2}\right) \sim \frac{2R}{(n-\delta)^3} \quad (n \gg 1). \tag{28}$$

F.2 Margin Lower Bound

The decisive pairs for class membership are those whose gap difference is closest to $\varepsilon(\kappa_{\text{conn}})$. For Rydberg gaps $\Delta_n \sim c/n^3$ ($c = 2R$), the ratio of consecutive gaps is $\Delta_{n+1}/\Delta_n \approx (n/(n+1))^3 \approx 1 - 3/n$. The gap difference satisfies $|\Delta_n - \Delta_{n+k}| \sim c \cdot 3k/n^4$. The margin lower bound is

$$m(L) \geq \min_n \frac{\big|\,|\Delta_n - \Delta_{n+k}| - \varepsilon(\kappa_{\text{conn}})\big|}{\tilde{\Delta}} \sim \frac{c \cdot 3k/n^4}{\tilde{\Delta}} \gg 0 \tag{29}$$

for any fixed k. In the He QM-I corpus example, numerically $m(L) \approx 0.012$–0.015, far exceeding $|\alpha - 1| \leq 0.20$.

F.3 Invariance of Edge Set

Because $m(L) > 0$, there exists $\delta_0 > 0$ such that for all $|a-1| < \delta_0$, the scaled differences $a|\Delta_i - \Delta_j|$ remain on the same side of every threshold $\varepsilon(\kappa)$ at every sampled κ. With $\delta_0 = 0.20$ sufficient for the tested window, the adjacency matrix is unchanged: $A'_{ij}(\kappa) = A_{ij}(\kappa)$, $H'(\sigma;\kappa) = H(\sigma;\kappa)$, and the class is invariant.

F.4 Commutator

Forward and reverse compositions produce identical gap differences (the shift $m\bar{x}$ cancels), so $C(\alpha, \mu; L) = \mathbf{0}$ inside Ω_L. $\square$

UNNS Substrate Research Program · unns.tech · April 2026
Companion manuscripts: USL · PRP · Dual Observability · Phase Mapping
Instruments: STRUC-PERC-I v2.4.0/v2.4.1 · Field Generator v1.0