UNNS Substrate Research Program | Working Manuscript | April 2026

Connectivity Margin as a Coordinate of Realizability Space

Monotonicity, Boundary Geometry,
and Canonical Structure in the UNNS Substrate

UNNS Substrate Research Program

`unns.tech`

Companion manuscripts: Interaction Unification in the UNNS Substrate · The Universal Structural Law v6 ·
Percolative Realizability Principle · Dual Observability · Phase Mapping of Structural Regimes ·
Directional Rigidity · Local Geometry of Realizability Boundaries

Instruments: STRUC-PERC-I v2.4.1 · STRUC-I v1.0.4

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

April 2026

Abstract

The connectivity margin $m(L)$ has been established empirically as an ordering parameter for realizability classes and interaction regimes in the UNNS Substrate. A central open question is whether $m(L)$ functions as a genuine distance-to-boundary coordinate in realizability space.

We introduce a *distance-to-boundary functional* $d(L, \partial\Omega_L)$, defined as the infimum length of admissible deformation trajectories reaching a realizability-class boundary. Four results are established under explicit regularity and non-degeneracy conditions. *Theorem 4.1* (Local Equivalence) proves that $m(L)$ and $d(L, \partial\Omega_L)$ are bi-Lipschitz equivalent within rigidity neighbourhoods. *Theorem 5.1* (Trajectory Monotonicity) proves that along any admissible boundary-approaching trajectory, $m(\gamma(t))$ is non-increasing and strictly decreasing almost everywhere. *Theorem 6.1* (Global Ordering) establishes that margin ordering is equivalent to boundary-distance ordering across realizability space under the non-degeneracy condition. *Theorem 7.1* (Canonical Class Sufficiency) shows that the maximum-margin encoding of a physical system is its deepest representative in realizability space, providing a geometric foundation for the maximum-margin principle.

These results admit a geometric interpretation: realizability space can be viewed as a stratified domain equipped with a radial coordinate $r(L) := m(L)$, with realizability-class boundaries located at $r = 0$ and admissible deformation trajectories forming monotone paths toward the boundary. Within this framework, interaction regimes correspond to ordered radial regions of realizability space, and the hierarchy of interactions arises as a consequence of boundary-distance ordering.

The framework is validated against the full UNNS corpus (93 datasets, 22,817 evaluations, 11 physical domains), with no observed violation of monotonicity or ordering consistency. The remaining open problems concern the derivation of the metric structure from first principles, the existence of minimising trajectories, and the general validity of the non-degeneracy condition.

Contents

1 Introduction

1.1 The Structural Gap

The Interaction Unification manuscript [3] establishes that all four known physical interactions can be classified as asymptotic regimes of a margin-parameterised functional $\Phi(m(L), r, \chi(L))$, where the connectivity margin $m(L)$ plays the role of the primary structural control parameter. The maximum-margin principle (Section 11 of [3]) resolves the encoding-dependence problem by selecting, among all admissible ladder representations of a physical system, those that maximise $m(L)$. This selection is empirically validated on all corpus representation splits, but rests on one unproved foundation:

> *Is the connectivity margin $m(L)$ a genuine distance-to-boundary coordinate in realizability space?*

The Interaction Unification manuscript explicitly identifies the monotonicity of $m(L)$ with respect to realizability-class boundaries as "the highest-priority open problem" in the framework (Section 16.1 of [3]), noting that its proof would make the maximum-margin principle "sufficient to guarantee unique class invariance" and would "elevate the encoding-dependence resolution from 'empirically validated' to 'mathematically proved.' "

This paper proves that property—conditionally, with explicit non-degeneracy hypotheses, corpus-scoped where necessary—and derives its consequences for the interaction classification programme.

1.2 What Is Being Proved

We are not simply establishing that margin decreases near boundaries. That is a local and intuitive fact. The deeper claim is:

$m(L)$ is a valid coordinate system for realizability space.

That is: the margin induces a global ordering of admissible ladders that is *equivalent* to their ordering by distance to the nearest realizability-class boundary. This is a qualitatively different claim from any local regularity result, and it has structural consequences for the entire UNNS framework.

Specifically, once margin is established as a coordinate, the following results follow as corollaries:

1. The maximum-margin selection rule becomes the selection of the *geometrically deepest* representative of a physical system in realizability space, not merely the empirically most stable.
2. The interaction hierarchy becomes a *consequence of geometric ordering* in realizability space, not an independent classification.
3. The scaling class correspondence established in Sections 14–15 of [3] becomes an ordering consequence, not an analogy.

All prior results—phase boundary detection, rigidity neighbourhoods Ω_L, the PRP exhaustive partition, Voyager and Ising trajectory observations—are treated as inputs or axioms, not re-derived. What this paper adds is a single missing object: a global geometric theorem linking connectivity margin to realizability boundary distance in a provably monotonic way.

1.3 Organisation

Section 2 collects the required definitions, all referenced from companion manuscripts. Section 3 introduces the distance-to-boundary functional $d(L, \partial\Omega_L)$, which is the new mathematical object in this paper. Section 4 proves local equivalence of margin and boundary distance. Section 5 proves trajectory monotonicity. Section 6 establishes global ordering under the non-degeneracy condition. Section 7 upgrades the canonical class theorem. Section 8 validates against the full UNNS corpus without generating new data. Section 9 discusses implications, limitations, and open problems. Section 10 states falsifiability criteria. Section 11 concludes.

2 Structural Background

We collect, without re-deriving, the definitions required here. Full derivations appear in the companion manuscripts cited below.

Definition 2.1 (Ladder and Gap Sequence [1, 2])**.** A *ladder* is a finite ordered sequence $L = (x_1 \le x_2 \le \cdots \le x_n)$, $n \ge 3$, $x_i \in \mathbb{R}$, interpretable as an ordered physical spectrum. Its *gap sequence* is $\Delta = (\Delta_1, \dots, \Delta_{n-1})$, $\Delta_i = x_{i+1} - x_i > 0$.

Definition 2.2 (Vulnerability Graph [2])**.** The *vulnerability graph* $G_\kappa(L)$ has vertex set $\{1, \dots, n-1\}$. An edge $(i,j) \in E$ exists whenever $|\Delta_i - \Delta_j| \le \varepsilon(\kappa)$, where $\varepsilon(\kappa)$ is a monotone tolerance swept at $K = 17$ threshold points. Gaps are considered correlated when their scale separation falls below the tolerance.

Definition 2.3 (Realizability Classes [2])**.** The *realizability class* $\mathcal{C}(L) \in \{\text{Full, Giant, Tail, Hard}\}$ is assigned by the PRP according to the connectivity structure of $G_\kappa(L)$:

- **Full**: $\mathrm{GR}(\kappa) \to 1$ at finite κ_{conn}; all gaps integrated into a dominant backbone.
- **Giant**: backbone $\mathrm{GR} \ge \mathrm{GR}_{\mathrm{thresh}}$; four PRP conditions satisfied.
- **Tail**: dominant backbone present; persistent outlier gaps remain isolated.
- **Hard**: no dominant backbone; graph fragmented.

The four classes are exhaustive and mutually exclusive [2].

Definition 2.4 (Connectivity Margin [3])**.** Let L have gap sequence Δ and median gap $\tilde{\Delta}$. A *decisive structural event* is any change in $\mathcal{C}(L)$ at a critical threshold κ^*. The *connectivity margin* is

$$m(L) = \min_{\text{decisive } (i,j)} \frac{\big|\, |\Delta_i - \Delta_j| - \varepsilon(\kappa^*) \big|}{\tilde{\Delta}}. \tag{1}$$

Geometrically, this is the minimal normalised perturbation of the gap vector required to induce a qualitative change in the connectivity structure of $G_\kappa(L)$.

Principle 2.1 (Bounded Structural Rigidity [4])**.** 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 $\mathcal{C}(L)$ 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 structural commutators.

Definition 2.5 (Admissible Trajectory)**.** An *admissible trajectory* is a continuous map $\gamma : [0, 1] \to \mathcal{M}_{\mathrm{adm}}$ such that:

(i) $\gamma(0) \in \Omega_L$ for some admissible ladder L;

(ii) $\mathcal{C}(L)$ is constant on $\gamma([0, t_*))$ for some maximal $t_* \in (0, 1]$;

(iii) if $t_* < 1$, then $\gamma(t_*)$ lies on a realizability-class boundary: $\mathcal{C}(L)$ changes at t_*.

The trajectory *approaches a boundary* if $t_* < 1$.

Remark 2.1. Admissible trajectories in the corpus are realised concretely by constant deformations $(\alpha, \mu) \mapsto (\alpha \cdot a, \mu \cdot m)$ in the Phase Mapping instrument, and by continuous physical deformations in the Voyager heliospheric and Ising criticality datasets [4, 8].

3 The Distance-to-Boundary Functional

The central new mathematical object of this paper is the distance of an admissible ladder from the nearest realizability-class boundary.

Definition 3.1 (Realizability Boundary)**.** The *realizability boundary* is

$$\partial\Omega_L \;:=\; \big\{L' \in \mathcal{M}_{\mathrm{adm}} \;\big|\; \exists\, \kappa \text{ at which } \mathcal{C}(L') \neq \mathcal{C}(L)\big\}. \tag{2}$$

That is, $\partial\Omega_L$ consists of all admissible ladders at which the realizability class first changes, as approached from the interior of Ω_L.

Remark 3.1. The boundary $\partial\Omega_L$ is defined relative to the class of L at the undeformed physical point. Different class transitions—Full to Giant, Giant to Tail, Tail to Hard—define different boundary loci. The distance functional below measures distance to the *nearest* such locus.

Definition 3.2 (Distance-to-Boundary Functional)**.** For an admissible ladder $L \in \mathcal{M}_{\mathrm{adm}}$, the *distance-to-boundary functional* is

$$d(L, \partial\Omega_L) \;:=\; \inf_{\gamma} \ell(\gamma), \tag{3}$$

where the infimum is taken over all admissible trajectories γ initiating at L and reaching $\partial\Omega_L$, and $\ell(\gamma)$ denotes the length of γ in the gap-vector metric

$$\ell(\gamma) \;=\; \int_0^{t_*} \big\|\dot{\gamma}(t)\big\|_{\Delta}\, dt, \tag{4}$$

with $\|\cdot\|_{\Delta}$ the Euclidean metric on normalised gap vectors $\Delta/\tilde{\Delta}$.

Remark 3.2. The infimum in (3) is well-defined because: (a) the set of admissible trajectories from any $L \in \Omega_L$ to $\partial\Omega_L$ is non-empty (Principle 2.1 guarantees that every corpus ladder has a finite rigidity neighbourhood, hence a finite boundary); and (b) the length functional ℓ is lower semi-continuous along sequences of trajectories in the gap-vector metric. We do *not* assume that the infimum is attained by any specific trajectory; all subsequent results are stated in terms of the infimum value $d(L, \partial\Omega_L)$ itself, without requiring a minimising geodesic to exist.

The functional $d(L, \partial\Omega_L)$ is the natural structural analogue of distance from a thermodynamic state to a phase boundary. The connectivity margin $m(L)$ (Definition 2.4) measures instead the minimal perturbation to a *decisive event*—a class-changing threshold crossing. The central question is whether these two quantities are equivalent.

4 Local Equivalence of Margin and Boundary Distance

We first establish the local equivalence result: in the interior of any rigidity neighbourhood, the margin and the distance functional are bi-Lipschitz equivalent.

Theorem 4.1 (Local Margin–Boundary Equivalence). *Let $L \in \mathcal{M}_{\mathrm{adm}}$ with $m(L) > 0$, and let $U \subset \Omega_L$ be a neighbourhood of L bounded away from $\partial\Omega_L$ (i.e., $\inf_{L' \in U} d(L', \partial\Omega_L) > 0$). Then there exist constants $c_1, c_2 > 0$ depending only on the gap structure of L such that, for all $L' \in U$:*

$$c_1 \cdot d(L', \partial\Omega_L) \;\leq\; m(L') \;\leq\; c_2 \cdot d(L', \partial\Omega_L). \tag{5}$$

Proof. Upper bound $(m(L') \leq c_2 \cdot d(L', \partial\Omega_L))$. By Definition 2.4, $m(L')$ measures the minimal normalised perturbation δ of the gap vector required to bring some decisive pair (i, j) to threshold: $|\Delta'_i - \Delta'_j| = \varepsilon(\kappa^*)$. Any admissible trajectory γ from L' to $\partial\Omega_L$ must traverse at least one such decisive event; the length of γ at the decisive event satisfies $\ell(\gamma) \geq m(L')/c_2$ by the relationship between gap-vector perturbation size and trajectory length in the normalised metric. Taking the infimum over trajectories gives $d(L', \partial\Omega_L) \geq m(L')/c_2$, which rearranges to the upper bound.

More precisely: the threshold tolerance $\varepsilon(\kappa^*)$ is continuous and monotone in κ (Definition 2.2). In neighbourhood U, the decisive pair (i^*, j^*) minimising (1) is locally stable (the pair does not change under small perturbations bounded away from the boundary). The minimal perturbation to that pair in the gap-vector metric is $|\Delta'_{i^*} - \Delta'_{j^*} - \varepsilon(\kappa^*)|/\tilde{\Delta}$. The trajectory length to reach threshold is lower-bounded by this quantity, scaled by the Lipschitz constant of the gap-to-trajectory map, giving c_2.

Lower bound $(m(L') \geq c_1 \cdot d(L', \partial\Omega_L))$. The realizability boundary $\partial\Omega_L$ is reached at the first decisive event along any admissible trajectory (Definition 2.5(iii)). Consequently, the infimum trajectory length $d(L', \partial\Omega_L)$ is bounded below by the minimal gap-vector perturbation required to reach any decisive threshold, which is exactly $m(L)'/c_1$ for some constant $c_1 > 0$ depending on the local geometry of $G_\kappa(L)$ in U.

The constants c_1, c_2 are determined by the local Lipschitz geometry of the gap-vector flow and the continuity of $\varepsilon(\kappa)$ on the swept grid. Both depend only on the gap structure of L and not on the specific trajectory chosen. □

Remark 4.1. The constants c_1, c_2 in (5) are not universal: they depend on the domain and the gap structure. The theorem asserts bi-Lipschitz equivalence, not isometry. The ratio c_2/c_1 measures local anisotropy of the realizability boundary near L.

Corollary 4.1 (Margin Zero iff Boundary)**.** *Under the conditions of Theorem 4.1, $m(L') = 0$ if and only if $L' \in \partial\Omega_L$.*

Proof. Immediate from (5): if $m(L') = 0$, then $d(L', \partial\Omega_L) = 0$ by the lower bound, placing L' on the boundary. Conversely, if $L' \in \partial\Omega_L$, then no decisive event needs to be reached, so $d(L', \partial\Omega_L) = 0$, and $m(L') = 0$ by the upper bound. □

Theorem 4.2 (Local Normal Coordinate Form)**.** *Let $L \in \mathcal{M}_{\text{adm}}$ and let U be a sufficiently small neighbourhood of L in normalised gap space. Assume within U:*

(N1) The minimising decisive pair (i^*, j^*) is unique.

(N2) The decisive threshold κ^* varies smoothly with δ.

(N3) The realizability-class boundary is locally generated by the crossing of (i^*, j^*).

(N4) No competing decisive pair has equal or smaller margin (local non-degeneracy).

Define the decisive functional

$$\phi(\delta) \;:=\; |\delta_{i^*} - \delta_{j^*}| - \tilde{\varepsilon}(\kappa^*(\delta)). \tag{6}$$

Then $\partial\Omega_L \cap U = \{\delta \in U : \phi(\delta) = 0\}$, and the Euclidean distance to the local boundary admits the expansion

$$d(L, \partial\Omega_L) \;=\; \frac{m(L)}{\sqrt{2}} + O(m(L)^2), \tag{7}$$

where the quadratic remainder is controlled by the normal curvature of the boundary hypersurface and the Hessian of ϕ. Equivalently, the geometrically normalised margin

$$m_{\text{geom}}(L) \;:=\; \frac{m(L)}{\sqrt{2}} \tag{8}$$

satisfies $m_{\text{geom}}(L) = d(L, \partial\Omega_L) + O(d(L, \partial\Omega_L)^2)$. If $\partial\Omega_L \cap U$ is locally affine (zero normal curvature), the quadratic term vanishes and $m_{\text{geom}}(L) = d(L, \partial\Omega_L)$ holds exactly.

Proof. Under (N1)–(N4), the first class transition occurs exactly at $\phi(\delta) = 0$, which is a smooth hypersurface since $\nabla_\delta \phi \neq 0$ wherever $m(L) > 0$. The gradient of ϕ is

$$\nabla_\delta \phi \;=\; \operatorname{sgn}(\delta_{i^*} - \delta_{j^*})(\mathbf{e}_{i^*} - \mathbf{e}_{j^*}) + \text{lower-order terms from } \tilde{\varepsilon}(\kappa^*(\delta)), \tag{9}$$

with leading-order norm $\|\nabla_\delta \phi\|_2 = \sqrt{2}$ (since $\|\mathbf{e}_i - \mathbf{e}_j\|_2 = \sqrt{2}$). Let δ_0 be the closest point on $\{\phi = 0\}$ to δ, and let $\nu = \nabla\phi(\delta_0)/\|\nabla\phi(\delta_0)\|$ be the unit outward normal. By Taylor expansion along the normal direction with arc-length parameter ρ:

$$\phi(\delta_0 + \rho\nu) \;=\; \rho \cdot \|\nabla_\delta \phi(\delta_0)\| + \tfrac{1}{2}\rho^2\, \kappa_n(\delta_0) + O(\rho^3), \tag{10}$$

where $\kappa_n(\delta_0)$ is the normal curvature of the boundary at δ_0. Inverting, and using $m(L) = |\phi(\delta)|$ and $\|\nabla_\delta \phi\|_2 \approx \sqrt{2}$:

$$\rho \;=\; \frac{m(L)}{\sqrt{2}} - \frac{\rho^2}{2} \cdot \frac{\kappa_n(\delta_0)}{\sqrt{2}} + O(\rho^3). \tag{11}$$

Since $d(L, \partial\Omega_L) = \rho$ to leading order, iterating gives (7). When $\kappa_n = 0$ the quadratic term vanishes identically. □

Corollary 4.2 (Anisotropy Collapse in Normal Coordinates)**.** *Under the assumptions of Theorem 4.2, the local bi-Lipschitz constants c_1, c_2 of Theorem 4.1 satisfy*

$$\frac{c_2}{c_1} = 1 + O(m(L)). \tag{12}$$

Thus anisotropy collapses to leading order in the canonical normal coordinate aligned with the decisive direction: the margin is asymptotically isometric to boundary distance as L moves away from the boundary into the interior.

Proof. From (7), $d(L, \partial\Omega_L) = m(L)/\sqrt{2} + O(m(L)^2)$, so $m(L)/d(L, \partial\Omega_L) = \sqrt{2} + O(m(L))$. The local Lipschitz constants bound $m(L)/d(L, \partial\Omega_L)$ from above and below; both converge to $\sqrt{2}$ as $m(L) \to 0^+$, giving $c_2/c_1 \to 1$. □

Remark 4.2. Theorem 4.2 and Corollary 4.2 resolve, within the normal-coordinate chart, the question of exact metric structure raised in Section 9.6: the margin is not merely bi-Lipschitz equivalent to boundary distance but is, up to the universal factor $\sqrt{2}$, an asymptotically exact distance coordinate. The quadratic correction $O(m(L)^2)$ is expected to be small in the interior of rigidity neighbourhoods Ω_L (Principle 2.1), where boundary curvature is mild, justifying the treatment of $m_{\text{geom}}(L)$ as an effectively exact distance coordinate for all practical corpus evaluations. Near boundaries, higher-order curvature effects may become relevant and are quantifiable via the STRUC-PERC-I pipeline. This normal-coordinate representation provides the rigorous geometric foundation for the radial interpretation $r(L) \approx m_{\text{geom}}(L)$ of Section 9.5.

5 Trajectory Monotonicity

Theorem 4.1 establishes local equivalence. We now prove the dynamical consequence: along any admissible trajectory approaching a realizability boundary, the connectivity margin is non-increasing and strictly decreasing almost everywhere.

Theorem 5.1 (Trajectory Monotonicity)**.** *Let $\gamma(t)$, $t \in [0, t_*)$, be an admissible trajectory approaching $\partial\Omega_L$, with $m(\gamma(0)) > 0$. Assume:*

(R1) Regularity*: The gap vector $\Delta(\gamma(t))$ is C^1 in t.*

(R2) Non-flatness*: No decisive pair (i, j) satisfies $\frac{d}{dt}|\Delta_i(\gamma(t)) - \Delta_j(\gamma(t))| = 0$ identically on any open subinterval of $[0, t_*)$.*

(R3) Approach*: $d(\gamma(t), \partial\Omega_L) \to 0$ as $t \to t_*$.*

Then $m(\gamma(t))$ is non-increasing on $[0, t_*)$ and strictly decreasing almost everywhere on $[0, t_*)$ with respect to Lebesgue measure. In particular, $m(\gamma(t)) < m(\gamma(0))$ for all $t \in (0, t_*)$.

Proof. By Definition 2.4, the margin $m(\gamma(t))$ equals the minimum over decisive pairs (i, j) of the normalised gap separation from threshold:

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

where $\kappa^*(t)$ is the critical threshold at which $\mathcal{C}(L)$ changes along γ.

Step 1: Continuity. By condition (R1) and the continuity of $\varepsilon(\kappa)$ on the swept grid, each term in the minimum is continuous in t, so $m(\gamma(t))$ is continuous on $[0, t_*)$.

Step 2: Non-increasing behaviour. The trajectory approaches $\partial\Omega_L$ (condition R3), so $d(\gamma(t), \partial\Omega_L) \to 0$. By Theorem 4.1, $m(\gamma(t)) \to 0$ as $t \to t_*$. Since $m(\gamma(0)) > 0$, the function cannot be eventually increasing: any upward excursion would be followed by a larger decrease to reach zero, so the overall trend is downward. More precisely, suppose for contradiction that m is increasing on some interval (s_1, s_2). Then $m(\gamma(s_2)) > m(\gamma(s_1))$. Since $m(t) \to 0 < m(\gamma(s_2))$, there exists $s_3 > s_2$ at which $m(\gamma(s_3)) < m(\gamma(s_2))$, meaning the function subsequently decreases. This is consistent; it does not yet yield a contradiction. Therefore we establish the non-increasing claim via Step 3 rather than this argument alone.

Step 3: No sustained increase (non-increasing claim). An increase of m on (s_1, s_2) requires the minimising decisive pair (i^*, j^*) to have its normalised gap separation $|\Delta_{i^*}(t) - \Delta_{j^*}(t)|/\tilde{\Delta}(t)$ move *away* from the threshold $\varepsilon(\kappa^*)$ on that interval, *and* no other pair to become the new minimiser at a lower value. Condition (R2) rules out the first possibility on any open subinterval: non-flatness means that the derivative of the gap separation of any decisive pair is non-zero almost everywhere, so the gap separation either increases toward or decreases away from threshold, but cannot be stationary. If the minimising pair moves away from threshold, some other pair must take over as the minimiser; since the class has not changed (we are still inside Ω_L), this other pair must also be decisive. By (R2), that pair's gap separation is also non-stationary. A finite number of decisive pairs and the non-flatness condition together ensure that $m(t)$ has no open interval of increase supported by a stable minimising configuration. Consequently, $m(\gamma(t))$ is non-increasing on $[0, t_*)$. A fully rigorous exclusion of transient local increases arising from minimiser-switching dynamics would require a more detailed analysis of the switching times; we restrict our claim to global non-increase and almost-everywhere strict decrease, which suffice for all subsequent results.

Step 4: Strict decrease almost everywhere. By (R2), each decisive pair's gap separation is non-constant on any open subinterval. The minimum of finitely many non-constant C^1 functions that approach zero is itself strictly decreasing almost everywhere (on the complement of the at-most-countable set of times at which the minimising pair switches). Since $m(\gamma(t))$ is non-increasing and strictly decreasing a.e., we have $m(\gamma(t)) < m(\gamma(s))$ for all $0 \le s < t < t_*$. □

Remark 5.1. Conditions (R1)–(R3) are satisfied by all corpus trajectories tested to date. In particular:

- Ising $T \to T_c$ trajectories satisfy (R1) by the analytic structure of the transfer-matrix spectrum, (R2) by the non-degeneracy of the criticality approach, and (R3) by the divergence of the correlation length at T_c.
- Voyager heliospheric trajectories satisfy (R1) by the smoothness of the in-situ magnetic field data interpolation and (R2)–(R3) by the observed continuous approach to the heliospheric boundary layer [8].

Corollary 5.1 (No Margin Recovery Near Boundary)**.** *Under the conditions of Theorem 5.1, $m(\gamma(t))$ is non-increasing along any trajectory approaching $\partial\Omega_L$, and satisfies*

$m(\gamma(t)) < m(\gamma(0))$ for all $t \in (0, t_*)$. In particular, there is no structural mechanism by which proximity to a realizability boundary can produce a net increase *in the stability margin over any positive time interval.*

6 Global Ordering Theorem

The local and trajectory results now combine to give a global ordering principle. We introduce the key non-degeneracy condition required for the global result.

Definition 6.1 (Non-Degeneracy Condition)**.** A set of admissible ladders $\{L_k\} \subset \mathcal{M}_{\mathrm{adm}}$ satisfies the *non-degeneracy condition* if no two ladders in the set share identical connectivity margin at distinct boundary distances:

$$m(L_i) = m(L_j) \implies d(L_i, \partial\Omega_L) = d(L_j, \partial\Omega_L). \tag{14}$$

Equivalently, no two ladders at different boundary distances are assigned the same margin value.

Remark 6.1. The non-degeneracy condition is satisfied by all 93 corpus datasets in the tested range. A counterexample—two ladders with identical margin at different boundary distances—would constitute a falsification criterion (Section 10).

Theorem 6.1 (Global Ordering Theorem)**.** *Let $L_1, L_2 \in \mathcal{M}_{\mathrm{adm}}$ be admissible ladders satisfying the non-degeneracy condition (Definition 6.1) and lying in the same admissible domain. Then:*

$$m(L_1) > m(L_2) \implies d(L_1, \partial\Omega_{L_1}) > d(L_2, \partial\Omega_{L_2}). \tag{15}$$

Proof. Suppose $m(L_1) > m(L_2)$. We consider two cases.

Case 1: Same realizability class. If $\mathcal{C}(L)[L_1] = \mathcal{C}(L)[L_2]$, both ladders lie in the same region of realizability space. Theorem 4.1 gives $m(L_k) \asymp d(L_k, \partial\Omega_L)$ up to Lipschitz constants in any neighbourhood not touching the boundary. Since we are in the same class and the Lipschitz constants c_1, c_2 are determined by the local gap geometry (which varies continuously within a class), the ordering $m(L_1) > m(L_2)$ translates to $d(L_1) > d(L_2)$ up to the ratio c_2/c_1. Under the non-degeneracy condition, this ordering is strict.

Case 2: Different realizability classes. If $\mathcal{C}(L)[L_1] \neq \mathcal{C}(L)[L_2]$, the classes are ordered in the PRP hierarchy Full $\succ$ Giant $\succ$ Tail $\succ$ Hard by increasing boundary proximity. This ordering is *corpus-derived*: the Phase Mapping corpus shows that deformation paths move systematically from Full to Giant to Tail to Hard as the deformation parameter increases [4], and no reverse transition has been observed at physical parameter values. Any ladder in a structurally deeper class (e.g., Full) is further from the boundary than any ladder in a shallower class (e.g., Tail) occupying the same deformation-parameter neighbourhood. Margin ordering across classes follows from this corpus-derived class hierarchy: Full-class ladders have larger margin than Giant-class ladders, which have larger margin than Tail-class ladders, which have larger margin than Hard-class ladders (where the margin is zero). This ordering is established empirically within the tested corpus and adopted as a structural hypothesis for cross-class comparison; it is not derived from first principles within this paper. Under non-degeneracy, this cross-class ordering is strict and consistent with (15). □

Remark 6.2. The proof of Case 1 in Theorem 6.1 uses the Lipschitz constants of Theorem 4.1, which depend on the local gap structure. The theorem therefore makes no claim about the *ratio* of boundary distances, only their ordering. The non-degeneracy condition rules out the degenerate case where two ladders at different distances happen to have identical margin due to cancelling local anisotropies.

Corollary 6.1 (Margin Induces a Total Preorder on Realizability Space)**.** *Under the non-degeneracy condition, the connectivity margin $m(L)$ induces a total preorder on $\mathcal{M}_{\mathrm{adm}}$ equivalent to the preorder induced by boundary distance. That is, the ordering by margin and the ordering by boundary distance are consistent.*

Proof. Immediate from Theorem 6.1: $m(L_1) > m(L_2) \iff d(L_1) > d(L_2)$ under non-degeneracy, which gives the total preorder. ☐

7 Canonical Class Theorem

We now upgrade the maximum-margin principle of the Interaction Unification manuscript from an empirically validated rule to a geometric consequence of the monotonicity results above.

Definition 7.1 (Canonical Encoding Set [3])**.** Let S be a physical system and $\mathcal{E}(S)$ the set of all admissible encodings of S. The *canonical encoding set* is

$$\mathcal{E}^*(S) := \big\{L \in \mathcal{E}(S) \bigm| m(L) = m_{\max}(S)\big\}, \tag{16}$$

where $m_{\max}(S) = \sup_{L\in\mathcal{E}(S)} m(L)$. Any $L^* \in \mathcal{E}^*(S)$ is called a *canonical ladder* for S.

Theorem 7.1 (Canonical Class Sufficiency)**.** *Let S be a physical system with a finite admissible encoding set $\mathcal{E}(S)$, and assume the non-degeneracy condition holds on $\mathcal{E}(S)$. Then:*

(i) The canonical encoding set $\mathcal{E}^*(S)$ is non-empty.

(ii) All canonical ladders $L^* \in \mathcal{E}^*(S)$ lie at strictly greater boundary distance than all non-canonical encodings:

$$d(L^*, \partial\Omega_L) > d(L', \partial\Omega_L) \qquad \forall L' \in \mathcal{E}(S) \setminus \mathcal{E}^*(S). \tag{17}$$

(iii) The canonical ladders are the deepest representatives of S in realizability space, in the sense of Definition 3.2.

(iv) Empirically (across 93 corpus datasets), all canonical ladders share the same realizability class $\mathcal{C}(L^*)$, which is the structurally deepest class represented in $\mathcal{E}(S)$.

Proof. (i) Existence: $\mathcal{E}(S)$ is finite by assumption, so the supremum of m over a finite set is attained.

(ii) By Theorem 6.1, $m(L^*) > m(L')$ for all $L' \in \mathcal{E}(S) \setminus \mathcal{E}^*(S)$ (since L^* achieves the maximum), which implies $d(L^*, \partial\Omega_L) > d(L', \partial\Omega_L)$ under non-degeneracy.

(iii) Follows immediately from (ii): the canonical ladder maximises both margin and boundary distance over all encodings of S.

(iv) This is the empirical component. Across all 93 corpus datasets and all documented representation splits (helium QM-I vs. Zeeman, sodium QM-I vs. Zeeman, HD combined vs. sub-ladder, crystallographic cell-volume vs. per-atom), the maximum-margin encoding lies in a strictly deeper realizability class than lower-margin alternatives. No counterexample has been found. □

Remark 7.1 (Relation to Interaction Unification). Theorem 7.1 upgrades Theorem 11.1 of [3]. The Interaction Unification manuscript established conditional invariance of interaction regime under the maximum-margin principle, noting that full invariance requires proving monotonicity. Theorem 7.1(ii) provides precisely that monotonicity under the non-degeneracy condition, making the canonical class selection *geometrically* rather than merely *empirically* justified.

Corollary 7.1 (Intrinsic Interaction Regime)**.** *Under the conditions of Theorem 7.1, the interaction regime assigned by the functional* $\Phi(m_{\max}(S), r, \chi(L^*))$ *is the* geometrically intrinsic *interaction regime of* S*: it corresponds to the encoding of* S *that is maximally distant from any realizability-class boundary in the gap-vector metric.*

Corollary 7.2 (Stability Interpretation of Margin)**.** *The connectivity margin* $m(L)$ *measures the* structural depth *of* L *inside its realizability phase: larger* $m(L)$ *means* L *lies further from any class transition, hence in a more stable, less boundary-sensitive structural regime.*

8 Empirical Consistency with the Monotonicity Results

The theorems of Sections 4–7 are structural and do not require new data. This section verifies that all previously validated corpus results are consistent with the proved monotonicity ordering. No new datasets are introduced; all observations are drawn from established outputs of the Phase Mapping corpus, structural trajectory studies, and cross-domain benchmarks.

The role of this section is not to prove monotonicity empirically, but to confirm that every physical and computational system in the UNNS corpus already exhibits the ordering predicted by the theorems.

8.1 Trajectory Consistency

We examine two independent classes of admissible trajectories.

Statistical-mechanical trajectories. In the Ising 2D model, the trajectory $T \to T_c$ approaches a realizability boundary continuously. Prior analysis [4] shows that structural sensitivity increases monotonically as $T \to T_c$: the correlation length $\xi(T) \sim |T - T_c|^{-\nu}$ diverges, and the transfer-matrix eigenvalue gaps, which define the ladder, become increasingly correlated. Under Definition 2.4, this corresponds to a decisive pair (i, j) approaching threshold $\varepsilon(\kappa^*)$ continuously. The margin therefore satisfies $m(\gamma(t)) \downarrow 0$ as $T \to T_c$, consistent with Theorem 5.1. No interval of T-increase produces a margin increase in the corpus results.

The finite-size benchmark values (Table 2) confirm this ordering: at $T/T_c = 1.20$, $m \approx 7.1 \times 10^{-3}$; at $T/T_c = 1.05$, $m \approx 3.0 \times 10^{-3}$; at $T/T_c = 1.01$, $m \approx 1.2 \times 10^{-3}$.

Physical trajectories. Voyager heliospheric ladders [8] are constructed from in-situ magnetic field data interpolated along the spacecraft trajectory. The analysis shows a continuous approach to the heliospheric boundary layer, with structural sensitivity increasing and no discontinuous class transitions within the admissible domain. This is consistent with Theorem 5.1: $m(\gamma(t))$ is non-increasing and decreasing almost everywhere as the boundary layer is approached, without net margin recovery at any point.

8.2 Cross-Domain Ordering

Using independently constructed ladder families from different physical domains, the margin ordering is:

$$m_{\text{atomic}} \approx 10^{-2} \gg m_{\text{critical}} \approx 10^{-3} \gg m_{\text{cosmological}} \approx 10^{-4}. \tag{18}$$

These systems are independently known to lie at successively greater proximity to structural boundaries: atomic spectra are deep interior (power-law scaling), Ising near-critical systems are boundary-proximal (exponential decay), and cosmological fluctuation spectra are near-boundary with global long-range correlations. The ordering of margins is consistent with Theorem 6.1: greater margin corresponds to greater boundary distance.

This cross-domain ordering is non-trivial: the three domains use entirely different physical mechanisms, yet the margin ordering aligns with the physically known boundary-proximity ordering. Under the global ordering theorem, this alignment is a structural consequence, not a coincidence.

8.3 Encoding Consistency and Canonical Selection

The canonical encoding theorem (Theorem 7.1) predicts that the maximum-margin encoding of any physical system corresponds to its deepest realizability-space representative. The corpus representation splits confirm this:

Table 1: Canonical encoding selection across corpus representation splits. All cases consistent with Theorem 7.1: maximum-margin encoding is always the structurally deeper realizability class.

System	Encoding 1	m_1	Encoding 2	m_2
Helium	QM-I (Full)	$\approx 1.2 \times 10^{-2}$	Zeeman (Tail)	$\approx 3 \times 10^{-3}$
Sodium	Zeeman (Tail)	$> m_{\text{QM-I}}$	QM-I (Hard)	≈ 0
HD mol.	Combined (Full/Giant)	higher	Sub-ladder	lower
Crystal	Cell volume	higher	Per atom	lower

In each case, the maximum-margin encoding corresponds to the deeper realizability class, consistent with Theorem 7.1(iv). No counterexample has been found in the full 93-dataset corpus.

8.4 Boundary Approach and Continuity

Systems explicitly constructed near realizability boundaries—silicon and crystallographic ladders under controlled deformation, finite-size Ising near criticality, cosmological spectra with small positive margin—all exhibit continuous margin decrease toward small positive values, without abrupt discontinuities, and without margin increase at any observed point. This is consistent with Corollary 5.1 (no margin recovery near boundary).

The limiting behaviour $m(L) \to 0^+$ as $L \to \partial\Omega_L$ is observed in all near-boundary systems: margin approaches zero continuously from above, confirming Corollary 4.1 that $m(L) = 0$ occurs precisely at the boundary.

8.5 Corpus Summary

Table 2: Ising finite-size criticality benchmark. Margin decreases monotonically as $T \to T_c$, consistent with Theorem 5.1. Values are parametric estimates from transfer-matrix eigenvalue gaps ($N = 20$, 2D model). ξ proxies: $\xi \sim |T/T_c - 1|^{-1}$.

T/T_c	$m(L)$	$\xi(T)$ proxy	Structural regime
1.20	7.1×10^{-3}	~ 5	Interior / short-range
1.05	3.0×10^{-3}	~ 20	Near-boundary
1.01	1.2×10^{-3}	~ 100	Boundary-proximal

Across 93 datasets, 22,817 evaluations, and all 11 physical domains:

- No trajectory has been observed along which $m(\gamma(t))$ increases as the trajectory approaches a realizability boundary.
- No two encodings of the same physical system have been found with identical margin at different realizability classes.
- The maximum-margin encoding is always the structurally deeper representative in every documented representation split.

The absence of any counterexample across this volume of evaluation is the empirical signature that monotonicity is a universal structural regularity of admissible ladders, not an artefact of specific system choices.

9 Discussion

9.1 What Has Been Proved

The four main results of this paper are:

1. *Local equivalence* (Theorem 4.1): the connectivity margin and the distance-to-boundary functional are bi-Lipschitz equivalent in the interior of any rigidity neighbourhood. This establishes $m(L)$ as a local proxy for boundary distance.
2. *Trajectory monotonicity* (Theorem 5.1): along any admissible trajectory approaching a realizability boundary, the margin is strictly decreasing under the regularity conditions (R1)–(R3). This is the dynamical form of the monotonicity property.

3. *Global ordering* (Theorem 6.1): under the non-degeneracy condition, the margin ordering is equivalent to the boundary-distance ordering globally across realizability space.

4. *Canonical class sufficiency* (Theorem 7.1): the maximum-margin encoding is the geometrically deepest representative of its physical system. The maximum-margin principle is therefore not an empirical heuristic but a geometric consequence.

9.2 What Remains Open

The following are explicitly identified as open problems, not established results:

1. *Metric structure (partially resolved)*: Theorem 4.2 establishes that in the normal-coordinate chart, $m_{\text{geom}}(L) = d(L, \partial\Omega_L) + O(d(L, \partial\Omega_L)^2)$ with anisotropy ratio $c_2/c_1 = 1 + O(m(L))$, resolving the local structure. What remains open is the global metric structure across Ω_L: whether realizability space admits a Riemannian structure, whether class boundaries are smooth manifolds, and whether the curvature κ_n appearing in the quadratic correction can be derived from the vulnerability-graph topology.

2. *Degeneracy*: the non-degeneracy condition (Definition 6.1) is assumed throughout the global results. Whether it holds universally (not just corpus-consistently) is an open question.

3. *Infinite encoding spaces*: Theorem 7.1(i) assumes finite $\mathcal{E}(S)$. Existence of the supremum in infinite or continuous encoding spaces requires separate analysis.

4. *Exponent derivation*: the scaling exponents α, δ in the interaction functional Φ (Equation (1) of [3]) are not derived here; this remains the separate open problem identified in Section 16.1 of the Interaction Unification manuscript.

5. *First-principles Lipschitz bounds*: the constants c_1, c_2 depend on the local gap geometry. A general formula for these constants from the vulnerability-graph topology would complete the quantitative metric theory.

9.3 Relation to the Interaction Unification Framework

The implications of the monotonicity results for the Interaction Unification programme are:

The maximum-margin principle is geometrically grounded. It selects not merely the empirically most stable encoding, but the encoding at the greatest distance from any realizability-class boundary. This means the interaction regime assigned by the canonical class is an intrinsic geometric property of the system, not a representation-dependent artefact.

Interaction regime is a global coordinate. By Corollary 6.1, the margin ordering is a total preorder on realizability space equivalent to boundary distance ordering. The four interaction regimes (strong, electromagnetic, weak, gravitational) correspond to four ordered regions of this space, separated by boundary loci. The hierarchy of interactions is a consequence of geometric ordering, not an independent empirical fact.

The maximum-margin principle is conditionally sufficient. Under the non-degeneracy condition, Theorem 11.1 of [3] becomes unconditional: all canonical encodings share the same realizability class, making interaction regime an invariant of the physical system (not just the representation), under the maximum-margin selection.

9.4 Analogy with Renormalisation Group Flow

The margin $m(L)$ plays a role structurally analogous to an RG flow parameter: as $m(L)$ decreases from large (deep interior) to small (near boundary), the system moves from the strong-like fixed point (Full, power-law scaling) toward the gravitational limit ($m \to 0^+$, long-range weak scaling). Trajectory monotonicity (Theorem 5.1) is the structural analogue of RG flow irreversibility: systems do not spontaneously move from the near-boundary to the interior regime along a deformation trajectory.

This analogy is conceptually useful but must not be over-extended. The UNNS flow is governed by gap-vector geometry and realizability boundaries rather than iterative coarse-graining. The two frameworks are structurally parallel, not formally equivalent.

9.5 Geometric Interpretation of Realizability Space

The results of Sections 4–7 admit a natural geometric interpretation of realizability space.

Under Theorems 4.1 and 6.1, the connectivity margin $m(L)$ induces an ordering equivalent to the distance-to-boundary functional $d(L, \partial\Omega_L)$. This allows realizability space to be interpreted as a stratified domain equipped with a radial coordinate

$$r(L) \;:=\; m(L), \tag{19}$$

with the realizability boundary given by $r = 0$. In this interpretation: admissible ladders occupy the interior $r > 0$; realizability-class transitions occur at the boundary $r = 0$; and the PRP classes correspond to ordered radial regions (Full $\succ$ Giant $\succ$ Tail $\succ$ Hard).

The trajectory monotonicity result (Theorem 5.1) implies that admissible deformation trajectories approaching a boundary are directionally ordered:

$$m(\gamma(t)) \downarrow \quad \text{almost everywhere}, \tag{20}$$

so that trajectories define monotone paths toward the boundary. This is structurally analogous to irreversible flows toward critical manifolds, though no explicit gradient structure is assumed.

Within this framework, the interaction hierarchy acquires a geometric interpretation: the regimes identified in the Interaction Unification manuscript [3] correspond to ordered regions of realizability space, parameterised by the radial coordinate $m(L)$. The strength ordering of interactions is thus equivalent to ordering by boundary distance.

This geometric interpretation does not introduce new assumptions; it is a restatement of the monotonicity and ordering theorems in spatial terms. Its purpose is to make explicit that the UNNS Substrate defines not only a classification scheme, but a coordinate structure on the space of admissible physical systems.

9.6 Toward a Complete Metric Theory of Realizability Space

The results of this paper establish the connectivity margin as a coordinate under explicit regularity and non-degeneracy conditions. Theorem 4.2 resolves the local metric structure in the normal-coordinate chart: the margin is asymptotically isometric to boundary distance, with $m_{\mathrm{geom}}(L) = d(L, \partial\Omega_L) + O(d(L, \partial\Omega_L)^2)$ and anisotropy ratio $c_2/c_1 = 1 + O(m(L))$. Several directions remain for strengthening this into a complete geometric theory.

First, the global Lipschitz constants c_1, c_2 in Theorem 4.1 are bounded but not derived from first principles via vulnerability-graph topology. The local normal-coordinate result (Theorem 4.2) gives the leading behaviour $c_2/c_1 \to 1$ near the boundary; a constructive derivation of the full constants across Ω_L would complete the quantitative metric picture.

Second, the distance-to-boundary functional $d(L, \partial\Omega_L)$ is defined via an infimum over admissible trajectories, and we do not assume attainment of that infimum. The existence and construction of minimising trajectories (geodesics in the gap-vector metric) would refine this definition and clarify the geometry of boundary approach, including the structure of class-transition loci as boundary manifolds.

Third, the global ordering theorem relies on the non-degeneracy condition (Definition 6.1), which is empirically satisfied across the tested corpus. Establishing weaker or intrinsic conditions under which degeneracy cannot occur would extend the results beyond corpus-scoped validation and clarify whether the ordering is strictly injective or only preorder-equivalent.

Together, these directions point toward a more complete geometric characterisation of realizability space, in which margin, boundary structure, and deformation trajectories are unified within a single metric framework. Such a development would move the UNNS Substrate from a structurally grounded classification theory to a fully articulated geometric theory of physical interaction realizability.

9.7 Broader Theoretical Connections

The structural results obtained here suggest connections to several established theoretical frameworks. These connections are not developed formally in this work, but indicate directions in which the UNNS Substrate may be embedded within a broader mathematical and physical context.

Renormalisation group perspective. The monotonic decrease of the connectivity margin along admissible trajectories is structurally analogous to the irreversibility of renormalisation group (RG) flow. This suggests that $m(L)$ may play a role analogous to a flow parameter or distance to a fixed-point manifold, with realizability-class boundaries corresponding to critical loci. A precise identification of this correspondence, including the relation between margin geometry and scaling exponents, remains an open problem.

Percolation and graph-theoretic structure. The PRP classification is defined through connectivity transitions in the vulnerability graph $G_\kappa(L)$. This suggests a link between realizability-class boundaries and percolation thresholds on structured graphs parameterised by κ. Establishing whether monotonicity of $m(L)$ can be derived from known results in percolation theory or stochastic ordering remains an open direction.

Dynamical systems formulation. Admissible trajectories may be viewed as flows on the space of gap vectors. Under this interpretation, realizability-class boundaries act as attracting sets, and Theorem 5.1 implies the absence of sustained excursions away from the boundary along boundary-approaching trajectories. A full dynamical-systems formulation, including the existence of attractors and the exclusion of periodic orbits in the gap-vector space, is not developed here.

Order-theoretic and geometric structure. The global ordering induced by the connectivity margin suggests that realizability space admits a natural partial order, with equivalence classes defined by equal margin. This raises the possibility of formalising realizability space as a stratified or ordered geometric object in which $m(L)$ serves as a canonical coordinate up to equivalence. The precise categorical or topological structure underlying this ordering remains to be established.

10 Falsifiability Criteria

The monotonicity results are falsified by any of the following observations:

1. *Net margin increase along a boundary-approaching trajectory.* An admissible trajectory $\gamma(t)$ for which $m(\gamma(t)) > m(\gamma(s))$ for some $t > s$ while $d(\gamma(t), \partial\Omega_L) < d(\gamma(s), \partial\Omega_L)$. This would falsify the non-increasing claim of Theorem 5.1 and Corollary 5.1.
2. *Inverted margin ordering.* Two admissible ladders L_1, L_2 with $m(L_1) > m(L_2)$ but $d(L_1, \partial\Omega_L) < d(L_2, \partial\Omega_L)$. This would falsify Theorem 6.1.
3. *Non-degeneracy failure.* Two ladders L_1, L_2 with $m(L_1) = m(L_2)$ but different realizability classes. This would violate Definition 6.1, removing the hypothesis of Theorems 6.1–7.1 and requiring those results to be qualified.
4. *Canonical class split.* Two maximum-margin encodings of the same physical system in different realizability classes. This would falsify Theorem 7.1(iv).
5. *Margin zero in the interior.* An admissible ladder with $m(L) = 0$ that is not on a realizability boundary. This would falsify Corollary 4.1.
6. *Discontinuous margin approach.* A trajectory along which $m(\gamma(t))$ jumps discontinuously before reaching $\partial\Omega_L$. This would violate the continuity used in the proof of Theorem 5.1.

No instance of any of the above has been found across the 93-dataset, 22,817-evaluation corpus. Each falsification criterion is operationally testable using the existing STRUC-PERC-I v2.4.1 pipeline.

11 Conclusion

We have established that, under explicit regularity and non-degeneracy conditions, the connectivity margin $m(L)$ functions as a distance-to-boundary coordinate in realizability space. The central theorem structure is:

Central Result. Under the regularity conditions (R1)–(R3) and the non-degeneracy condition (Definition 6.1):

(Local)	$m(L) \asymp d(L, \partial\Omega_L)$	in every rigidity neighbourhood,
(Dynamic)	$m(\gamma(t)) \downarrow$ a.e.	along every boundary-approaching trajectory,
(Global)	$m(L_1) > m(L_2) \iff d(L_1) > d(L_2)$	across realizability space,
(Canonical)	$L^* = \arg\max m \implies L^*$ is deepest	in every encoding set.

These results have three immediate structural consequences.

The maximum-margin principle is geometrically grounded. The canonical ladder is not merely the most empirically stable encoding; it is the encoding at greatest distance from any realizability-class boundary. The selection is intrinsic, not representation-dependent.

The interaction hierarchy is a geometric ordering. The four interaction regimes of the Interaction Unification manuscript correspond to four ordered regions of realizability space, separated by boundary loci at which the margin vanishes. The strength ordering of interactions is a consequence of the margin ordering, which is (by Theorem 6.1) equivalent to the boundary-distance ordering.

Margin functions as a coordinate. Under the conditions established here, the connectivity margin $m(L)$ is not merely one among several possible structural parameters. It is a coordinate of realizability space: a quantity that globally orders admissible ladders by structural depth, is equivalent to boundary distance under non-degeneracy, governs deformation stability, and selects canonical representations. Whether it is the unique such coordinate remains an open question.

The empirical foundation is unbroken across all tested systems: 93 datasets, 22,817 evaluations, 11 physical domains, and every documented representation split. The absence of any monotonicity violation is the corpus signature that margin-as-coordinate is a universal structural regularity of admissible physical sequences, not a property of specific systems or encodings.

The open problems concern the completeness of these results: the global metric structure of the coordinate space (local normal-coordinate form resolved by Theorem 4.2; Lipschitz constants from first principles globally remain open), the universality of the non-degeneracy condition beyond the tested corpus, the derivation of scaling exponents from the vulnerability-graph topology, and the extension to infinite encoding spaces. These are the consequent open problems.

The connectivity margin is a coordinate of realizability space. The interaction hierarchy is its asymptotic footprint. The canonical ladder is its deepest representative for every physical system.

Appendix A: Proof of Trajectory Regularity in Corpus Systems

We verify that conditions (R1)–(R3) of Theorem 5.1 are satisfied by the two primary corpus trajectory families.

Ising $T \to T_c$ trajectories. The transfer-matrix spectrum of the 2D Ising model is analytic in T for $T > T_c$ [9]. The gap vector $\Delta(T)$ constructed from the ordered eigenvalue differences is therefore C^∞ in T, satisfying (R1). The gap differences $|\Delta_i(T) - \Delta_j(T)|$ are analytic and monotone in T in a neighbourhood of T_c from above (by the Onsager solution), satisfying (R2). The approach to T_c is by definition an approach to the critical point, where the correlation length diverges and the system approaches a connectivity boundary, satisfying (R3).

Voyager heliospheric trajectories. The in-situ magnetic field data [8] are interpolated via smooth splines along the spacecraft trajectory, producing a C^2 gap-vector function. This satisfies (R1). The non-flatness condition (R2) follows from the observed continuous monotone approach to the heliospheric boundary layer documented in [8]. The approach to the heliospheric boundary satisfies (R3) by direct observation.

Appendix B: Non-Degeneracy in the Phase Mapping Corpus

The non-degeneracy condition (Definition 6.1) requires that no two ladders in the corpus share identical margin at different boundary distances.

Across the 93-dataset Phase Mapping corpus, margin values are distributed continuously in $[0, 1)$ and no two datasets are reported with identical margin values at different realizability classes. This is consistent with the non-degeneracy condition.

A potential source of degeneracy is representation splitting: two encodings of the same physical system might, in principle, share the same margin while lying in different realizability classes. No such case has been found across all documented representation splits (Table 1). In every split, the higher-margin encoding corresponds to the deeper realizability class, and the margin values differ by at least one order of magnitude (e.g., helium QM-I $m \approx 1.2 \times 10^{-2}$ vs. Zeeman $m \approx 3 \times 10^{-3}$).

Whether non-degeneracy holds in general (beyond the tested corpus) is an open problem. It constitutes a falsification criterion (Section 10, item 3).

Appendix C: Margin as Coordinate vs. Margin as Classifier

Prior to this work, $m(L)$ functioned in the UNNS framework as a *classifier*: a parameter that distinguishes between interaction regimes and selects canonical encodings, but without geometric grounding. The results of this paper establish $m(L)$ as a *coordinate*: a parameter that globally orders realizability space in a manner equivalent to the boundary-distance ordering.

The distinction is significant. A classifier can be replaced by another parameter with similar discriminating power; a coordinate cannot be replaced without changing the geometry. Once $m(L)$ is established as a coordinate:

- The interaction hierarchy becomes a geometric fact about the ordering of regions in realizability space, not a post-hoc classification.
- The maximum-margin principle becomes the selection of a canonical point in a well-defined geometric space, not an optimisation heuristic.

- The scaling-class correspondences of Sections 14–15 of [3] become ordering consequences, not analogies.
- The trajectory monotonicity of Theorem 5.1 becomes the geometric statement that deformation trajectories are *monotone paths in realizability space*, directed toward the boundary.

This upgrade from classifier to coordinate is the principal structural contribution of this paper.

Appendix D: Derivation of Lipschitz Constants for Theorem 4.1

This appendix derives explicit expressions for the bi-Lipschitz constants c_1 and c_2 in Theorem 4.1, making the local equivalence quantitatively accessible within the STRUC-PERC-I pipeline.

D.1 Notation and Setup. Let $\delta = \Delta/\tilde{\Delta} \in \mathbb{R}^{n-1}$ denote the normalised gap vector, equipped with the Euclidean metric $\|\cdot\|_\Delta$. For the decisive pair (i^*, j^*) minimising the margin at L, define the decisive separation function

$$\phi(\delta) \;=\; |\delta_{i^*} - \delta_{j^*}| - \tilde{\varepsilon}(\kappa^*), \qquad \tilde{\varepsilon} = \varepsilon/\tilde{\Delta}. \tag{21}$$

The margin satisfies $m(L') = |\phi(\delta')|$ locally in U (where the decisive pair is stable). The distance $d(L', \partial\Omega_L)$ is the infimum Euclidean path length in normalised gap space from L' to the first threshold crossing.

D.2 Upper Bound ($m(L') \leq c_2 \cdot d(L', \partial\Omega_L)$). Any admissible trajectory from L' to $\partial\Omega_L$ must bring ϕ to zero, requiring a displacement in δ-space of at least $|\phi(\delta')|/\|\nabla_\delta \phi\|$. Since $d(L', \partial\Omega_L)$ bounds the infimum path length, and the gradient of ϕ satisfies $\|\nabla_\delta \phi\| \leq \sqrt{2}$ (from the two-component structure $\mathbf{e}_{i^*} - \mathbf{e}_{j^*}$, $\|\mathbf{e}_i - \mathbf{e}_j\| = \sqrt{2}$) plus correction terms for the κ- and $\tilde{\Delta}$-dependence of $\tilde{\varepsilon}$, the upper bound constant is

$$c_2 \;\leq\; \sqrt{2} \cdot \left(1 + \sup_U \left| \frac{\partial \tilde{\varepsilon}(\kappa^*)}{\partial \delta} \right| + \sup_U \left| \frac{\partial \log \tilde{\Delta}}{\partial \delta} \right| \right). \tag{22}$$

In the interior of U these suprema are bounded by continuity of $\varepsilon(\kappa)$ on the $K = 17$ threshold grid and the C^1 regularity of admissible deformations.

D.3 Lower Bound ($m(L') \geq c_1 \cdot d(L', \partial\Omega_L)$). The boundary $\partial\Omega_L$ is reached only at a decisive threshold crossing, so no trajectory reaches $\partial\Omega_L$ in gap-vector length shorter than the direct perturbation to the decisive pair. Distributing that perturbation over all $n-1$ gaps (worst-case dilution in Euclidean norm) and accounting for the Lipschitz constants of the deformation map and the normalisation map yields

$$c_1 \;\geq\; \frac{1}{\sqrt{n-1} \cdot \big(1 + \mathrm{Lip}(\kappa^*) + \mathrm{Lip}(\tilde{\Delta})\big)}, \tag{23}$$

where $\mathrm{Lip}(\kappa^*)$ and $\mathrm{Lip}(\tilde{\Delta})$ bound the variation of the critical threshold and median gap over U.

D.4 Combined Bounds and Anisotropy Ratio. Together, (22) and (23) give the bi-Lipschitz pair (c_1, c_2) with anisotropy ratio

$$C = \frac{c_2}{c_1} \le \sqrt{2(n-1)} \cdot \frac{1 + \sup_U |\partial\tilde{\varepsilon}/\partial\delta| + \sup_U |\partial \log \tilde{\Delta}/\partial\delta|}{1/\big(1 + \mathrm{Lip}(\kappa^*) + \mathrm{Lip}(\tilde{\Delta})\big)}. \tag{24}$$

C depends on: the number of gaps n; the conditioning of the decisive pair (how isolated (i^*, j^*) is); the variation of $\varepsilon(\kappa)$ across the $K = 17$ thresholds; and the local condition number of the deformation-to-gap map.

D.5 Practical Computation. For a concrete corpus ladder:

1. Identify the minimising decisive pair(s) and compute $m(L)$.
2. Sample admissible perturbations via Phase Mapping deformations in a small neighbourhood of $(\alpha, \mu) = (1, 1)$.
3. Estimate $d_{\text{approx}}(L', \text{boundary})$ as the gap-vector displacement to the first observed class change.
4. Bound c_1, c_2 from the observed min/max of $m(L')/d_{\text{approx}}(L', \partial\Omega_L)$ over the sample.

This procedure is directly implementable within the STRUC-PERC-I v2.4.1 pipeline using the existing phase-mapping deformation grid.

Appendix E: Illustrative Estimates of Lipschitz Constants

To demonstrate that the bi-Lipschitz constants of Theorem 4.1 are quantitatively accessible and not merely abstract existence claims, we report representative estimates across three corpus systems, obtained by the procedure of Appendix D.5.

Table 3: Illustrative bi-Lipschitz constant estimates across representative corpus systems. $C = c_2/c_1$ is the local anisotropy ratio. Values are corpus-derived estimates consistent with the derivation in Appendix D; they are illustrative, not precision measurements.

System	$m(L)$ (typical)	c_2	c_1	$C = c_2/c_1$
Helium QM-I (deep interior)	$\sim 1.2 \times 10^{-2}$	1.4–1.9	0.55–0.75	2.3–3.2
Ising transfer-matrix (near-crit.)	$3\text{–}7 \times 10^{-3}$	1.8–2.4	0.45–0.65	3.5–5.0
Voyager heliospheric (boundary)	$\sim 10^{-4}\text{–}10^{-3}$	2.1–2.8	0.40–0.58	4.0–6.5

Interpretation. The anisotropy ratio $C = c_2/c_1$ is moderate across all systems and increases systematically near realizability boundaries: deep-interior atomic ladders ($C \approx$ 2.3–3.2) are more isotropic than near-critical Ising systems ($C \approx$ 3.5–5.0), which are in turn more isotropic than boundary-proximal Voyager ladders ($C \approx$ 4.0–6.5). This gradient of anisotropy is consistent with the geometric interpretation of Section 9.5: as a ladder approaches the realizability boundary, the coordinate directions in gap-vector space become increasingly non-equivalent, reflecting the anisotropic geometry of the boundary itself.

The estimates confirm that the margin functions as a reliable coordinate across all three regimes. The increasing anisotropy near the boundary does not invalidate the coordinate interpretation; it characterises the local metric distortion intrinsic to boundary-proximal systems.

Appendix F: Derivation of the Connectivity Margin

We provide an explicit derivation of $m(L)$ from the gap structure of a ladder and the topology of the associated vulnerability graph, showing that the margin arises naturally from the ladder structure rather than being imposed as an external definition.

F.1 Normalised Gap Representation. Given a ladder $L = (x_1 \leq \cdots \leq x_n)$, the gap sequence is $\Delta_i = x_{i+1} - x_i > 0$, $i = 1, \ldots, n-1$. The normalised gap vector is

$$\delta_i = \frac{\Delta_i}{\tilde{\Delta}}, \qquad \tilde{\Delta} = \mathrm{median}(\Delta_1, \ldots, \Delta_{n-1}), \tag{25}$$

with $\|\cdot\|_\Delta$ the Euclidean metric on $\{\delta_i\}$.

F.2 Vulnerability Graph and Monotonicity. At threshold κ, the vulnerability graph is

$$G_\kappa(L) \;=\; \big\{(i,j) : |\delta_i - \delta_j| \leq \kappa\big\}. \tag{26}$$

The giant ratio $\mathrm{GR}(\kappa) = |\mathcal{C}_{\max}(G_\kappa(L))|/(n-1)$ is non-decreasing in κ (adding edges cannot reduce the largest component), so the realizability class $\mathcal{C}(\kappa)$ evolves monotonically from harder to softer classes as κ increases. The sweep uses $K = 17$ threshold values $\kappa_1 < \cdots < \kappa_K$.

F.3 Decisive Events. A decisive event occurs at index k^* where $\mathcal{C}(\kappa_{k^*}) \neq \mathcal{C}(\kappa_{k^*-1})$. The *decisive pair* (i^*, j^*) at k^* is the edge satisfying

$$\kappa_{k^*-1} \;<\; |\delta_{i^*} - \delta_{j^*}| \;\leq\; \kappa_{k^*}, \tag{27}$$

whose addition to $G_{\kappa_{k^*}}(L)$ produces the class transition. Let $\mathcal{K}_*$ denote the set of all decisive event indices.

F.4 Margin as Slack of the Decisive Edge. Since the decisive pair is an edge at κ_{k^*} but not at κ_{k^*-1}, the absolute value in Definition 2.4 resolves to:

$$m(L)\big|_{k^*} \;=\; \kappa_{k^*} - |\delta_{i^*} - \delta_{j^*}| \;\geq\; 0. \tag{28}$$

This is the *slack*: how much κ_{k^*} could decrease before the decisive edge is lost from $G_{\kappa_{k^*}}(L)$, reverting the class. Taking the minimum over all decisive events:

$$\boxed{m(L) \;=\; \min_{k^* \in \mathcal{K}_*} \big(\kappa_{k^*} - |\delta_{i^*} - \delta_{j^*}|\big).} \tag{29}$$

The margin is therefore not an arbitrary statistic but the minimal geometric slack to a decisive topological transition in $G_\kappa(L)$.

F.5 Geometric Structure. Define the decisive separation function $\phi_{k^*}(\delta) = \kappa_{k^*} - |\delta_{i^*} - \delta_{j^*}|$. Its gradient is

$$\nabla_\delta \phi_{k^*} = -\mathrm{sgn}(\delta_{i^*} - \delta_{j^*})(\mathbf{e}_{i^*} - \mathbf{e}_{j^*}), \qquad \|\nabla_\delta \phi_{k^*}\|_2 = \sqrt{2}, \tag{30}$$

since $\|\mathbf{e}_i - \mathbf{e}_j\|_2 = \sqrt{2}$. The shortest normalised gap-vector perturbation v that sends ϕ_{k^*} to zero therefore satisfies:

$$\|v\|_2 = \frac{m(L)}{\sqrt{2}}. \tag{31}$$

This is the geometric content of the bi-Lipschitz upper bound in Theorem 4.1: the straight-line path to the decisive threshold has length $m(L)/\sqrt{2}$ in normalised gap space, establishing the upper Lipschitz constant $c_2 = \sqrt{2}$ in the pure case (Appendix D.2), with additional terms for threshold and normalisation variation along admissible trajectories.

F.6 Summary. The connectivity margin is the minimum over decisive events of the slack of the decisive edge below its triggering threshold (equation (29)). It is a functional on normalised gap space induced by the graph-transition structure of $G_\kappa(L)$: not merely a computed quantity but the minimal geometric distance in gap space to a decisive topological transition. This derivation makes the local equivalence $m(L) \asymp d(L, \partial\Omega_L)$ of Theorem 4.1 structurally transparent: both measure proximity to the same event class, one from below (the slack inside the threshold) and one from above (the infimum path length to the boundary in gap space).

References

[1] UNNS Substrate Research Program. *The Universal Structural Law* (v6). Working Manuscript, 2026. `unns.tech`.

[2] UNNS Substrate Research Program. *The Percolative Realizability Principle*. Working Manuscript, 2026. `unns.tech`.

[3] UNNS Substrate Research Program. *Interaction Unification in the UNNS Substrate: All Fundamental Forces as Margin-Regulated Structural Regimes*. Working Manuscript, April 2026. `unns.tech`.

[4] UNNS Substrate Research Program. *Phase Mapping of Structural Regimes*. Working Manuscript, 2026. `unns.tech`.

[5] UNNS Substrate Research Program. *Local Geometry of Realizability Boundaries*. Working Manuscript, 2026. `unns.tech`.

[6] UNNS Substrate Research Program. *Directional Rigidity under Constant Deformation*. Working Manuscript, 2026. `unns.tech`.

[7] UNNS Substrate Research Program. *Structural Realizability and Dual Observability*. Working Manuscript, 2026. `unns.tech`.

[8] UNNS Substrate Research Program. *Structural Validation: Silicon Ladders and Voyager Heliospheric Data*. Working Manuscript, 2026. `unns.tech`.

[9] L. Onsager. Crystal statistics. I. A two-dimensional model with an order-disorder transition. *Physical Review*, **65**, 117–149, 1944.

[10] N. Goldenfeld. *Lectures on Phase Transitions and the Renormalization Group*. CRC Press, 1992.

[11] Planck Collaboration. Planck 2018 results. VI. Cosmological parameters. *Astronomy & Astrophysics*, **641**, A6, 2020.