UNNS Substrate Research Program | Working Manuscript

Local Geometry of Realizability Boundaries in the UNNS Substrate

From Connectivity Margin to Boundary Distance

UNNS Substrate Research Program

`unns.tech`

Abstract. We develop the missing local geometric structure of realizability within the admissibility manifold $\mathcal{M}_{\mathrm{adm}}$. Prior work in the UNNS Substrate program established admissibility via the Universal Structural Law, realizability via the Percolative Realizability Principle (PRP), the independence of these two coordinates via Dual Observability, and the dynamical response of both under operator deformation. One central object, however, remained only operational: the connectivity margin $m(L)$, defined through the STRUC-PERC-I instrument as the minimum normalised distance to a decisive class-changing event. Its geometric meaning—its relation to the actual spatial structure of realizability-class boundaries inside $\mathcal{M}_{\mathrm{adm}}$—was left unresolved.

We close this gap by introducing realizability charts, decisive coordinate systems, and a local hypersurface representation of class boundaries. We prove that in a neighbourhood of any regular realizability boundary point, the connectivity margin is strictly monotone with respect to the local boundary distance, and we derive a local form of the maximum-margin canonicalization principle: within any regular local regime, margin-maximizing ladder encodings of the same physical system must belong to the same realizability class.

The empirical test beds—selected from the 81-run STRUC-PERC-I corpus spanning 14 physical domains—validate the geometric picture at the level of corpus-supported evidence. ^{28}Si supplies the clearest interior-to-boundary deformation probe; ^{238}U isolates singular boundary behaviour; the helium representation split demonstrates multiple chart embeddings; and the four TAIL nuclear isotopes probe the non-percolating boundary side.

These results convert the connectivity margin from an empirical invariant into a theorem-grade boundary-distance functional, and provide the first rigorous foundation for the canonical ladder programme.

Version 1.0 · April 2026 · UNNS Working Manuscript Series

Contents

1 Introduction

Section abstract. We identify the geometric gap in the existing UNNS framework, state the open problem precisely, and position the present manuscript's contribution within the programme. No prior results are re-derived.

1.1 The structural architecture so far

The UNNS Substrate Research Program has established a two-coordinate framework for ordered physical sequences (ladders). The first coordinate, admissibility, is governed by the Universal Structural Law (USL): a ladder L is admissible if and only if $\mathrm{inv}(P_\varepsilon; L) \leq \nu(V_\varepsilon(L))$ for all relevant ε. The second coordinate, realizability, is governed by the Percolative Realizability Principle (PRP): admissibility implies full percolation of the pairwise vulnerability graph $G_\kappa(L)$ in the necessary direction, and the corpus supports the converse direction as an open conjecture.

These two coordinates are independent. The Dual Observability manuscript proves that admissibility and realizability are not mutually determined: systems can be admissible in multiple realizability classes, and the class assignment is sensitive to ladder encoding. The Structural Response manuscripts show that both coordinates respond to operator deformation $\alpha \to \alpha'$, $\mu \to \mu'$ in a domain-selective but structurally rigid fashion: across 9,826 evaluations on the 17×17 joint deformation grid, no tested system undergoes a verdict change, and all structural commutators $C(\alpha, \mu; L)$ are identically zero.

1.2 The open problem: the geometry of $m(L)$

The connectivity margin $m(L)$ is operationally defined by STRUC-PERC-I as the minimum normalised signed distance to the nearest decisive class-changing event—a giant-ratio threshold crossing, a connectivity threshold event, or a tail-dominance boundary. It is empirically robust: the 81-run STRUC-PERC-I corpus shows $m(L) > 0$ for 80 of 81 physical runs, with the single exception (raw TiO_2 density, $m(L) < 0$, giant ratio 0.833) being the only confirmed structural failure.

Despite this empirical solidity, $m(L)$ currently lacks a geometric interpretation. Three questions remain unanswered:

1. What is the geometric structure of realizability-class boundaries inside $\mathcal{M}_{\mathrm{adm}}$?
2. Does $m(L)$ measure distance to those boundaries, or merely correlate with it?
3. Can maximum-margin canonicalization be promoted from a corpus-supported principle to a theorem?

Without answers to these questions, the canonical ladder programme—the goal of selecting a privileged encoding for each physical system via margin maximization—rests on operational evidence rather than on geometric necessity. The Interaction Unification manuscript explicitly identifies this gap as the primary open problem of the programme.

1.3 Contribution of this paper

We supply the missing geometric layer in three steps.

1. **Local charts.** We introduce realizability charts—finite-dimensional local coordinate systems on $\mathcal{M}_{\mathrm{adm}}$ built from decisive structural observables (gap-difference coordinates, giant-ratio coordinates, connectivity-threshold coordinates, tail-dominance coordinates). We prove that every realizability class change in a small neighbourhood is captured by a threshold crossing in these coordinates (Lemma 3.1).
2. **Boundary hypersurfaces.** We show that at any regular realizability boundary point, the class boundary is locally a codimension-1 C^1 hypersurface (Theorem 4.1), and we define the local structural boundary distance $d_{\partial\mathcal{C}}(L)$ in chart coordinates.
3. **Monotonicity and canonicalization.** We prove that $m(L)$ is locally order-equivalent to $d_{\partial\mathcal{C}}(L)$ (Corollary 6.1), that m is strictly monotone with respect to boundary distance along any transversal direction (Theorem 6.1), and that all margin-maximizing encodings of a single physical system lie in the same realizability class within a regular local regime (Theorem 7.1).

These results do *not* establish global monotonicity, unique canonical ladders, or full closure of the canonical programme. What they establish is the correct local geometric structure that must underpin any future global result.

2 Framework and notation

Section abstract. We collect the minimal background needed for the new results, introduce notation, and state the structural state decomposition used throughout. No prior results are re-proved.

2.1 Admissible ladders

An *admissible ladder* is a finite ordered sequence $L = (x_1, \dots, x_n)$ with $x_1 < x_2 < \cdots < x_n$ such that the USL inequality $\mathrm{inv}(P_\varepsilon; L) \le \nu(V_\varepsilon(L))$ holds for all $\varepsilon > 0$. The *gap vector* of L is

$$\Delta(L) = (\Delta_1, \dots, \Delta_{n-1}), \qquad \Delta_i = x_{i+1} - x_i.$$

2.2 Structural state

The structural state of an admissible ladder is the pair

$$S(L) = \bigl(\bar{\rho}(L),\, R(L)\bigr),$$

where $\bar{\rho}(L) = \mathrm{inv}(P_\varepsilon; L)/\nu(V_\varepsilon(L))$ is the structural pressure (the ratio of inversions to vulnerability, averaged over the admissibility κ-grid), and

$$R(L) = \bigl(C(L),\, \kappa_{\mathrm{conn}}(L)\bigr)$$

is the realizability coordinate. Here $C(L)$ denotes the giant-ratio (the fraction of vertices in the largest connected component of the pairwise vulnerability graph at the operative κ scale), and $\kappa_{\mathrm{conn}}(L)$ denotes the normalised connectivity threshold at which the graph first percolates.

2.3 Realizability classes

STRUC-PERC-I assigns each admissible ladder to one of four realizability classes:

$$\begin{aligned}
\text{FULL} &: C(L) = 1.000 \text{ and } \kappa_{\text{conn}}(L) < \infty, \\
\text{GIANT} &: C(L) \geq 0.95 \text{ but } \kappa_{\text{conn}}(L) \text{ not reached on the main grid}, \\
\text{TAIL} &: C(L) \geq 0.95, \text{ outlier-dominated, no secondary cluster, Theorem 1 suppressed}, \\
\text{HARD} &: C(L) < 0.95, \text{ Theorem 1 (PRP) triggered}.
\end{aligned}$$

2.4 Connectivity margin

The *connectivity margin* $m(L)$ is the minimum normalised signed distance to the nearest decisive class-changing event. Operationally, STRUC-PERC-I computes $m(L)$ from giant-ratio thresholds, tail-dominance thresholds, and connectivity-threshold events. The precise functional form is made explicit in Section 5 below.

2.5 The admissibility manifold

We denote by $\mathcal{M}_{\text{adm}}$ the set of all admissible ladders (of all lengths), equipped with a natural topology in which sequences are close if their gap-ratio profiles are close in the sup-norm after normalisation. Realizability classes partition $\mathcal{M}_{\text{adm}}$ into four regions. The *realizability boundary* $\partial\mathcal{C}$ is the topological boundary of any one class inside $\mathcal{M}_{\text{adm}}$.

2.6 Ladder construction protocols

All objects in this paper are ladders $L = (x_1, \ldots, x_n)$, but the theory requires that the mapping from a physical system to a ladder be explicit and reproducible. We formalise this via the notion of a ladder construction protocol.

Definition 2.1 (Ladder construction protocol)**.** A *ladder construction protocol* is a mapping

$$\Pi : \mathcal{S} \longrightarrow \mathcal{L},$$

assigning to each physical system $S \in \mathcal{S}$ a finite ordered sequence $\Pi(S) = (x_1 < x_2 < \cdots < x_n)$ satisfying:

(i) *Strict ordering:* $x_i < x_{i+1}$ for all i;

(ii) *Finiteness:* $n = n(\Pi, S) < \infty$ is determined by the protocol and system;

(iii) *Admissibility record:* Π produces ladders for which the USL verdict is computed and recorded (admissible or violated).

Definition 2.2 (Encoding family)**.** Given a physical system S, the *encoding family* of S is

$$\mathcal{E}(S) = \big\{\Pi_i(S) : \Pi_i \in \mathcal{P}\big\},$$

where $\mathcal{P}$ is the set of all protocols applicable to S. Two protocols Π_1, Π_2 are *representation-equivalent* for S if $\Pi_1(S)$ and $\Pi_2(S)$ have the same admissibility verdict.

The corpus employs four protocol families, formalised in Appendix D:

(A) *QM-I protocols*: energy-level ladders of atomic or molecular spectra, available in three variants (preprocessed filtered, raw spectrum, gap-difference sequence).

(B) *Zeeman protocols*: ladders constructed from magnetically-split energy levels $\{E_i(B)\}$ at a fixed field B, ordered by energy.

(C) *Gap-derived protocols*: the gap sequence $\Delta(L_{\text{base}})$ of a base ladder, treated as a ladder in its own right.

(D) *Transform protocols*: ladders of the form $T(L_{\text{base}})$ for a specified monotone transform T.

Remark 2.1 (Protocol dependence of chart position)**.** Different protocols map the same physical system to different chart positions $\Phi(\Pi_i(S))$ in the decisive chart. Realizability class is not invariant under protocol choice: the same system can be FULL under one protocol and GIANT or TAIL under another, as the helium corpus demonstrates (Section 8.4). Admissibility, by contrast, is empirically stable across protocols (Lemma 2.1 below).

Lemma 2.1 (Protocol admissibility stability—empirical)**.** *For every physical system S and every protocol $\Pi_i \in \mathcal{P}$ tested in the corpus, $\Pi_i(S)$ satisfies the USL (zero hard violations across all tested protocol–system pairs).*

Proof. By corpus enumeration. The STRUC-I v1.0.4 corpus records 5,233 ladder evaluations across 14 physical domains and multiple protocol variants. Zero USL violations are recorded in any run. The atomic corpus (16 runs, 6 elements, QM-I and Zeeman protocols) records zero violations. This is a corpus-level result, not a theorem: it confirms empirical admissibility stability but does not exclude violations under untested protocols or systems. □

Remark 2.2 (Generic vs. canonical protocols)**.** Exploratory corpus runs may use a *generic adapter protocol* that extracts all parseable numerical values from a data source, sorts them, removes duplicates, and drops non-finite entries. This produces a valid ladder (satisfying Definition 2.1) and aligns with STRUC-PERC-I output for all validated runs. For publication-grade physics interpretation, *canonical domain protocols* are preferred: these select a specific column (e.g., energy levels only for QM-I spectrum, magnetically split levels only for Zeeman) and apply domain-specific ordering rules. The current corpus is internally consistent under the generic adapter; the specification of canonical protocols is given in Appendix D and will be the subject of a dedicated protocol-refinement study.

The geometric results of this manuscript—Lemma 3.1, Theorem 4.1, Theorem 6.1, and Theorem 7.1—apply to any admissible ladder $L \in \mathcal{M}_{\text{adm}}$ independently of the protocol used to construct it. Protocol choice determines which element of the encoding family $\mathcal{E}(S)$ is selected; the local geometry of $\mathcal{M}_{\text{adm}}$ is a property of the ladder itself.

3 Local structural coordinates

Section abstract. We introduce realizability charts and decisive coordinate families, prove that a finite number of observables suffices to determine the local realizability class, and identify the explicit decisive observables from PRP practice.

3.1 Realizability charts

Definition 3.1 (Realizability chart)**.** A *realizability chart around* $L_0 \in \mathcal{M}_{\text{adm}}$ is a map

$$\Phi_{L_0} : U(L_0) \subset \mathcal{M}_{\text{adm}} \longrightarrow \mathbb{R}^d$$

that is a homeomorphism onto its image, where $U(L_0)$ is an open neighbourhood of L_0 and the coordinate functions of Φ_{L_0} are drawn from the following families of *structurally decisive observables*:

(i) *Normalised gap-difference coordinates:* $(\Delta_{i+1} - \Delta_i)/\text{med}(\Delta)$,

(ii) *Giant-ratio coordinate:* $C(L)$,

(iii) *Connectivity-threshold coordinate:* $\kappa_{\text{conn}}(L)/\text{scale}$,

(iv) *Tail-dominance coordinate:* $\max(\Delta)/\text{IQR}(\Delta)$ (relative outlier weight),

(v) *Auxiliary coordinates:* $(\bar{\rho}(L),\, A_{\kappa_{\min}})$ if additional discrimination is required.

The dimension d of the chart is the number of decisive coordinates retained.

Remark 3.1. The chart is not required to be unique. Different choices of decisive coordinates yield different chart embeddings of the same neighbourhood; the theory is built to be chart-independent at the level of order relations (see Definition 6.1 below).

3.2 Decisive coordinate families

Definition 3.2 (Decisive coordinate family)**.** A coordinate family $(x_1, \ldots, x_d)$ on $U(L_0)$ is *decisive* if realizability class changes of $L \in U(L_0)$ occur if and only if at least one coordinate x_j crosses a critical threshold θ_j. Formally: for any $L, L' \in U(L_0)$ that lie in different realizability classes, there exists j such that $x_j(L)$ and $x_j(L')$ lie on opposite sides of θ_j.

The decisive property makes rigorous the operational meaning of $m(L)$: a decisive structural event is precisely a threshold crossing in a decisive coordinate.

3.3 Finite decisive reduction

Lemma 3.1 (Finite decisive reduction)**.** *For every admissible ladder* $L_0 \in \mathcal{M}_{\text{adm}}$*, there exists an open neighbourhood* $U(L_0)$ *and a finite decisive coordinate family of dimension* $d \leq 5$ *governing all realizability class changes in* $U(L_0)$*.*

Proof. The PRP classification algorithm (STRUC-PERC-I) assigns realizability class through a finite sequence of tests applied to $\Delta(L)$:

(a) *Giant-ratio test:* is $C(L) \geq 0.95$? This is determined by the structure of the pairwise vulnerability graph $G_\kappa(L)$ at the operative scale, which depends continuously on the gap vector Δ.

(b) *Secondary-cluster test:* does the graph contain a second component of size comparable to the giant? This is controlled by two gap-ratio thresholds.

(c) *Tail-dominance test:* is $\max(\Delta)/\mathrm{IQR}(\Delta)$ above the hybrid outlier threshold $\max(10 \cdot \mathrm{med},\, 5 \cdot \mathrm{IQR})$?

(d) *Connectivity-threshold test:* does the graph percolate at a finite κ_{conn} within the grid range?

Each test depends on a finite set of statistics of $\Delta(L)$. Since $\Delta(L)$ varies continuously in the sup-norm topology on $\mathcal{M}_{\mathrm{adm}}$, each test function is continuous on $\mathcal{M}_{\mathrm{adm}}$. Threshold crossings are therefore isolated events in any neighbourhood. In a sufficiently small neighbourhood $U(L_0)$ of L_0, the only class changes are those involving the decisive thresholds already active near L_0. The resulting coordinate family has dimension at most

$$d \;\leq\; 1_{\text{giant-ratio}} + 1_{\text{secondary}} + 1_{\text{tail-dom}} + 1_{\text{conn-thresh}} + 1_{\text{auxiliary}} \;=\; 5. \qquad \square$$

Remark 3.2 (Corpus evidence)**.** Across the 81-run STRUC-PERC-I corpus, every class assignment is determined by the tests enumerated in the proof. The three-regime percolation structure (immediate $\kappa_{\mathrm{conn}} < 2$; moderate 2–10^3; extreme 10^3–4×10^5) confirms that class boundaries are isolated events in the observable coordinate space, not dense.

Remark 3.3 (System-dependent chart sufficiency)**.** The dimension d of a sufficient decisive chart, and the set of active threshold functions, depend on the system and representation family. A chart that is decisive for one system need not extend to another: the same pair of decisive coordinates may fail to capture all class changes in a different physical domain, requiring additional coordinates or a different chart neighbourhood altogether. Cross-system measurements confirm this: atomic Zeeman ladders of different elements require different decisive coordinate sets to represent their local boundaries correctly (see Section 8).

4 Realizability boundaries as hypersurfaces

Section abstract. We define regular realizability boundary points, represent each class boundary locally as a codimension-1 hypersurface via the implicit function theorem, and introduce the local structural boundary distance.

4.1 Regular boundary points

Definition 4.1 (Regular realizability boundary point)**.** A point $L^* \in \partial\mathcal{C}$ is a *regular realizability boundary point* if there exists a realizability chart Φ around L^* and a C^1 function $F : \mathbb{R}^d \to \mathbb{R}$ such that

(i) the local boundary is $\partial\mathcal{C} \cap U(L^*) = \{L : F(\Phi(L)) = 0\}$,

(ii) $\nabla F(\Phi(L^*)) \neq 0$.

The non-vanishing gradient condition is the standard regularity hypothesis that allows the implicit function theorem to apply and the boundary to be treated as a smooth hypersurface locally. *Irregular* boundary points—cusps, corners, or points where multiple boundaries meet transversally—are not excluded from $\mathcal{M}_{\mathrm{adm}}$ in general, but lie outside the scope of the present local theory.

4.2 Threshold representation

Lemma 4.1 (Threshold representation)**.** *For each decisive coordinate x_j with threshold θ_j, the signed deviation*

$$G_j(x) \;=\; x_j - \theta_j$$

is a C^1 function on $\mathbb{R}^d$ with $\nabla G_j \neq 0$ everywhere. Class changes mediated by the j-th coordinate occur precisely on $\{G_j = 0\}$.

Proof. Each decisive coordinate $x_j = x_j(L)$ is a C^1 function of the gap vector $\Delta(L)$ on $U(L_0)$ (this follows from the continuity of the statistics enumerated in Lemma 3.1). The threshold θ_j is a fixed real constant. Therefore $G_j(x) = x_j - \theta_j$ is C^1 in x, and $\partial G_j/\partial x_j = 1 \neq 0$. The zero set $G_j = 0$ is precisely the locus where coordinate x_j equals its threshold, which by Definition 3.2 is the locus of class changes mediated by that coordinate. □

4.3 Local boundary hypersurface theorem

Theorem 4.1 (Local boundary hypersurface theorem)**.** *Near every regular realizability boundary point L^*, the realizability boundary $\partial\mathcal{C}$ is locally representable as a codimension-1 C^1 hypersurface in a decisive realizability chart.*

Proof. Let L^* be a regular boundary point. By Definition 4.1, there exists a decisive chart Φ and a C^1 function F with $F(\Phi(L^*)) = 0$ and $\nabla F(\Phi(L^*)) \neq 0$. By Lemma 4.1, we may take F to be the decisive threshold function G_{j^*} for the active boundary direction j^* at L^*.

The implicit function theorem applies: since G_{j^*} is C^1 and $\nabla G_{j^*} \neq 0$ at $\Phi(L^*)$, the zero set $\{x : G_{j^*}(x) = 0\}$ is a C^1 hypersurface of dimension $d-1$ in a neighbourhood of $\Phi(L^*)$. Under the chart homeomorphism Φ, this hypersurface pulls back to a C^1 hypersurface of codimension 1 in $\mathcal{M}_{\mathrm{adm}}$ near L^*. By Lemma 3.1, class changes near L^* are controlled by finitely many such thresholds, and in the generic (regular) case a single threshold is active at L^*. □

4.4 Local structural boundary distance

Definition 4.2 (Local structural boundary distance)**.** Let $\Phi : U(L_0) \to \mathbb{R}^d$ be a decisive realizability chart containing a regular boundary point. The *local structural boundary distance* of $L \in U(L_0)$ is

$$d_{\partial\mathcal{C}}(L) \;=\; \inf_{y \in \partial\mathcal{C} \cap U(L_0)} \big\|\Phi(L) - \Phi(y)\big\|,$$

where $\|\cdot\|$ denotes the Euclidean norm on $\mathbb{R}^d$.

Remark 4.1. This definition is chart-dependent in absolute scale but chart-independent in the ordering of $d_{\partial\mathcal{C}}(L_1)$ versus $d_{\partial\mathcal{C}}(L_2)$ for two ladders on the same side of a single boundary, provided the chart is decisive (a fact we verify below via the order-equivalence theorem). Extensions to a Riemannian structural metric are deferred to future work.

Remark 4.2 (Local separators are not global classifiers)**.** Empirically fitted boundary functions of the form $F(x) = 0$—derived by fitting a separator to a corpus of points in a decisive chart—may act as exact classifiers within a local chart neighbourhood but fail outside it. Such a failure does not contradict Theorem 4.1: the theorem asserts local hypersurface structure near each regular boundary point, not a single global separator across all systems. A fitted function valid near one system's boundary encodes the local geometry of that system's chart; a different system may lie outside the fitted neighbourhood and require its own local chart and boundary representation.

5 The margin as a local structural functional

Section abstract. We show that $m(L)$ admits a decomposition into branch functions $M_j(L)$, each measuring normalised distance to one decisive threshold. We prove an active-branch lemma securing local dominance of a single branch, and prove strict monotonicity of m along any transversal curve directed away from the boundary.

5.1 Assumptions for the local theory

We work throughout in a fixed decisive realizability chart $\Phi : U(L_0) \to \mathbb{R}^d$ centred at L_0. We impose three local assumptions.

Assumption 5.1 (Decisive coordinate reduction)**.** There exist finitely many C^1 functions $G_j : \mathbb{R}^d \to \mathbb{R}$, $j = 1, \dots, k$, such that realizability class changes in $U(L_0)$ occur if and only if $G_j(x) = 0$ for at least one j. This is the finite decisive reduction established in Lemma 3.1.

Assumption 5.2 (Regularity)**.** Each G_j is C^1, and at the active boundary point $x^* = \Phi(L^*)$: $\nabla G_j(x^*) \neq 0$ for all j that are active.

Assumption 5.3 (Interior side)**.** We restrict to the interior of a fixed realizability class: $G_j(x) > 0$ for all j and all L in the neighbourhood of interest.

5.2 Branch margin functions

Definition 5.1 (Branch margin function)**.** For each decisive threshold function G_j, the *branch margin function* is

$$M_j(L) \;=\; G_j\big(\Phi(L)\big).$$

Definition 5.2 (Margin decomposition)**.** The connectivity margin decomposes as

$$m(L) \;=\; \min_{1\le j\le k} M_j(L).$$

This decomposition is already implicit in the operational definition of $m(L)$: the margin is controlled by the nearest decisive structural violation, and each branch M_j measures the signed distance to the j-th violation threshold in chart coordinates.

5.3 Active branch lemma

Lemma 5.1 (Active branch lemma)**.** *Let L_0 satisfy $M_{j^*}(L_0) < M_j(L_0)$ for all $j \neq j^*$. Then there exists a neighbourhood $U' \subset U(L_0)$ such that*

$$m(L) \;=\; M_{j^*}(L) \qquad \textit{for all } L \in U'.$$

Proof. Since all M_j are continuous (as compositions of C^1 functions), define

$$\delta \;=\; \min_{j\neq j^*}\bigl(M_j(L_0) - M_{j^*}(L_0)\bigr) \;>\; 0.$$

By continuity of each M_j, there exists an open neighbourhood $U' \ni L_0$ such that $|M_j(L) - M_j(L_0)| < \delta/3$ for all j and all $L \in U'$. For $j \neq j^*$:

$$M_j(L) \;\geq\; M_j(L_0) - \tfrac{\delta}{3} \;>\; M_{j^*}(L_0) + \delta - \tfrac{\delta}{3} \;=\; M_{j^*}(L_0) + \tfrac{2\delta}{3},$$

while $M_{j^*}(L) \;\leq\; M_{j^*}(L_0) + \frac{\delta}{3}$. Thus $M_j(L) > M_{j^*}(L)$ for all $j \neq j^*$ in U', and so $m(L) = \min_j M_j(L) = M_{j^*}(L)$. □

5.4 Transversal curves and strict monotonicity

Definition 5.3 (Transversal curve)**.** A C^1 curve $L : (-\varepsilon, \varepsilon) \to U(L_0)$ is *transversal to the boundary* $G_{j^*} = 0$ at $t = 0$ if

$$\frac{d}{dt}\, G_{j^*}\bigl(\Phi(L(t))\bigr)\Big|_{t=0} \;\neq\; 0.$$

A transversal curve is *outward-directed* if the above derivative is strictly positive (i.e. the curve moves away from the boundary into the interior).

Lemma 5.2 (Strict monotonicity along transversal directions)**.** *Let $L(t)$ be an outward-directed C^1 transversal curve with $L(0) = L_0$, $G_{j^*}(\Phi(L_0)) > 0$. Then for sufficiently small $t > 0$,*

$$m(L(t)) \;>\; m(L(0)).$$

Proof. By Lemma 5.1, there exists a neighbourhood of L_0 in which $m(L) = M_{j^*}(L) = G_{j^*}(\Phi(L))$. Therefore

$$\frac{d}{dt}\, m(L(t)) \;=\; \frac{d}{dt}\, G_{j^*}(\Phi(L(t))) \;=\; \nabla G_{j^*}(\Phi(L(t))) \cdot \dot{x}(t),$$

where $x(t) = \Phi(L(t))$. At $t = 0$ this is strictly positive by the outward-transversal hypothesis. By continuity of ∇G_{j^*} and $\dot{x}$, there exists $\varepsilon > 0$ such that $\frac{d}{dt} m(L(t)) > 0$ for all $t \in (0, \varepsilon)$. Integrating: $m(L(t)) > m(L(0))$ for all $t \in (0, \varepsilon)$. □

6 The margin–boundary relation

Section abstract. We prove that the connectivity margin $m(L)$ is locally Lipschitz-equivalent to the boundary distance $d_{\partial\mathcal{C}}(L)$ (order-equivalence), and derive the main monotonicity theorem as a corollary.

6.1 Local boundary geometry

Fix a regular boundary point L^* and the active branch j^*. By Assumption 5.2 and the implicit function theorem, the zero set

$$\partial\mathcal{C}_{j^*} \;=\; \bigl\{x \in \mathbb{R}^d : G_{j^*}(x) = 0\bigr\}$$

is a C^1 hypersurface in a neighbourhood $V \ni x^* = \Phi(L^*)$. We write $\mathbf{n} = \nabla G_{j^*}(x^*)/\|\nabla G_{j^*}(x^*)\|$ for the unit outward normal at x^*.

6.2 Lipschitz equivalence of margin and distance

Lemma 6.1 (Distance representation). *There exist constants $c_1, c_2 > 0$ such that for all L in a sufficiently small neighbourhood of L^* in the interior side,*

$$c_1\, d_{\partial\mathcal{C}}(L) \;\le\; G_{j^*}(\Phi(L)) \;\le\; c_2\, d_{\partial\mathcal{C}}(L).$$

Proof. Since G_{j^*} is C^1 with $\nabla G_{j^*}(x^*) \neq 0$, Taylor's theorem gives

$$G_{j^*}(x) \;=\; \nabla G_{j^*}(x^*) \cdot (x - x^*) + R(x), \qquad |R(x)| \le C\|x - x^*\|^2,$$

for some constant C. Let $\mathbf{n} = \nabla G_{j^*}(x^*)/\|\nabla G_{j^*}(x^*)\|$ and $\lambda = \|\nabla G_{j^*}(x^*)\| > 0$. Then

$$G_{j^*}(x) \;=\; \lambda\, \mathbf{n} \cdot (x - x^*) + R(x).$$

The distance from x to the hyperplane $\mathbf{n} \cdot (x - x^*) = 0$ is $d_{\text{plane}}(x) = |\mathbf{n} \cdot (x - x^*)|$, and the distance to the actual C^1 hypersurface $\partial\mathcal{C}_{j^*}$ satisfies $d_{\partial\mathcal{C}} \le d_{\text{plane}} \le d_{\partial\mathcal{C}}(1 + O(d_{\partial\mathcal{C}}))$ in a small neighbourhood. Substituting:

$$G_{j^*}(x) \;=\; \lambda\, d_{\partial\mathcal{C}}(L)\big(1 + O(d_{\partial\mathcal{C}}(L))\big).$$

For sufficiently small $d_{\partial\mathcal{C}}(L)$, we can absorb the error term and obtain constants $c_1 = \lambda/2$ and $c_2 = 2\lambda$ satisfying the claimed inequalities. ☐

6.3 Order equivalence

Definition 6.1 (Margin-order equivalence). The connectivity margin $m(L)$ is *locally order-equivalent to boundary distance* in $U(L_0)$ if for any $L_1, L_2 \in U(L_0)$ on the same realizability side,

$$d_{\partial\mathcal{C}}(L_1) < d_{\partial\mathcal{C}}(L_2) \iff m(L_1) < m(L_2).$$

Corollary 6.1 (Order equivalence). *In a sufficiently small neighbourhood of a regular boundary point, $m(L)$ is locally order-equivalent to $d_{\partial\mathcal{C}}(L)$.*

Proof. From Lemma 5.1, locally $m(L) = G_{j^*}(\Phi(L))$. From Lemma 6.1, $G_{j^*}(\Phi(L)) \asymp d_{\partial\mathcal{C}}(L)$ with positive Lipschitz constants c_1, c_2. Since both $m(L)$ and $d_{\partial\mathcal{C}}(L)$ are strictly positive in the interior and are related by a bi-Lipschitz equivalence, the ordering is preserved: $m(L_1) < m(L_2) \iff G_{j^*}(x_1) < G_{j^*}(x_2) \iff d_{\partial\mathcal{C}}(L_1) < d_{\partial\mathcal{C}}(L_2)$. ☐

6.4 Local Margin Monotonicity Theorem

Theorem 6.1 (Local Margin Monotonicity Theorem). *Let L_1, L_2 lie in a sufficiently small neighbourhood of a regular realizability boundary point, on the same realizability side. Then*

$$d_{\partial\mathcal{C}}(L_1) < d_{\partial\mathcal{C}}(L_2) \implies m(L_1) < m(L_2).$$

Proof. By Corollary 6.1, m is locally order-equivalent to $d_{\partial\mathcal{C}}$. Therefore the implication $d_{\partial\mathcal{C}}(L_1) < d_{\partial\mathcal{C}}(L_2) \Rightarrow m(L_1) < m(L_2)$ holds directly from the definition of order equivalence (Definition 6.1). Equivalently: from Lemma 6.1, $m(L) = \phi(d_{\partial\mathcal{C}}(L))$ where $\phi(s) = G_{j^*}(\,\cdot\,)$ satisfies $c_1 s \le \phi(s) \le c_2 s$ with $c_1, c_2 > 0$, hence ϕ is strictly increasing near $s = 0$. Thus $d_{\partial\mathcal{C}}(L_1) < d_{\partial\mathcal{C}}(L_2)$ implies $\phi(d_{\partial\mathcal{C}}(L_1)) < \phi(d_{\partial\mathcal{C}}(L_2))$, i.e. $m(L_1) < m(L_2)$. □

Remark 6.1. Theorem 6.1 is strictly local. It holds within the neighbourhood V of a single regular boundary point, and relies on the assumption that a single branch j^* dominates (Lemma 5.1). Near irregular boundary points (corners, cusps, or multi-branch intersection points), the result does not apply directly, and the correct geometric treatment requires a piecewise-hypersurface generalisation that we indicate in Section 10.

Cross-system measurements confirm that margin-distance order-equivalence is preserved within each local chart neighbourhood, but the same equivalence does not imply a single global metric across heterogeneous domains. A separator function that correctly orders systems within one chart neighbourhood may fail to classify systems lying in a different neighbourhood of $\mathcal{M}_{\mathrm{adm}}$, where the active decisive branch, its gradient, and the boundary orientation all differ. The locality of the equivalence is therefore a property of the chart, not of the physical systems.

7 Local maximum-margin canonicalization

Section abstract. We prove that margin maximizers lie in a unique realizability class within any regular local regime, thereby providing the first theorem-grade support for the canonical ladder programme within a local geometric domain.

7.1 Class stability under margin increase

Lemma 7.1 (Class stability under monotone margin)**.** *Let $L(t)$ be an outward-directed transversal curve in $U(L_0)$ with $L(0)$ in the interior of a realizability class $\mathcal{C}$. Then $L(t)$ remains in $\mathcal{C}$ for all sufficiently small $t > 0$.*

Proof. Class changes occur exactly on $\{G_{j^*} = 0\}$ (by decisiveness, Lemma 4.1). Since $L(t)$ is outward-directed, $G_{j^*}(\Phi(L(t)))$ is strictly increasing (Lemma 5.2), hence strictly positive for $t > 0$. Therefore $L(t)$ does not cross any class boundary for small $t > 0$. □

7.2 Local maximizers are class-consistent

Lemma 7.2 (Class consistency of local maximizers)**.** *Suppose $L_1, L_2 \in U(L_0)$ both maximize m within the same encoding family $\mathcal{E}(S)$ (the set of all admissible ladder encodings of a fixed physical system S), and that the neighbourhood $U(L_0)$ is a regular local regime (single active branch, Assumption 5.2). Then L_1 and L_2 cannot lie on opposite sides of any realizability boundary within $U(L_0)$.*

Proof. Suppose for contradiction that $L_1 \in \mathcal{C}$ and $L_2 \in \mathcal{C}'$ for distinct classes $\mathcal{C} \neq \mathcal{C}'$. Since the boundary $\partial\mathcal{C}$ is a C^1 hypersurface in $U(L_0)$ (Theorem 4.1), it separates $U(L_0)$ into two open sides. On L_1's side, we have $G_{j^*}(\Phi(L_1)) > 0$ and $m(L_1) = M_{j^*}(L_1) > 0$. On L_2's side, $G_{j^*}(\Phi(L_2)) < 0$. By Assumption 5.3 applied to $\mathcal{C}$, the active branch gives $M_{j^*}(L_2) < 0$ on

L_2's side. Thus $m(L_2) \leq M_{j^*}(L_2) < 0 < m(L_1)$, contradicting the assumption that both are maximizers. □

7.3 Local Maximum-Margin Canonicalization Theorem

Theorem 7.1 (Local Maximum-Margin Canonicalization Theorem)**.** *Within any encoding family $\mathcal{E}(S)$ whose admissible representatives lie in a single regular local regime $U(L_0)$, all margin-maximizing encodings belong to the same realizability class.*

Proof. Let $\mathcal{E}^*(S) \subset \mathcal{E}(S) \cap U(L_0)$ be the set of margin-maximizing encodings. By Theorem 6.1, within $U(L_0)$ a higher margin corresponds to a larger boundary distance. The boundary distance is maximized at the point(s) farthest from $\partial\mathcal{C}$ inside $U(L_0)$, which by convexity of $U(L_0)$ (assumed without loss of generality) lies in the interior of one realizability class. By Lemma 7.2, no two maximizers can lie on different sides of a class boundary. Therefore all elements of $\mathcal{E}^*(S)$ lie in the same realizability class. □

Remark 7.1 (Scope of the theorem)**.** Theorem 7.1 is the first theorem-grade support for the canonical ladder programme. It is genuinely local: it applies within a single regular regime $U(L_0)$, and does not address encoding families that straddle multiple regimes or include irregular boundary points. The global version—asserting that *all* margin-maximizing encodings of a physical system, across all regimes, lie in a single class—remains open, and is the primary objective of the next phase of the programme (see Section 10).

8 Empirical realization of local geometry

Section abstract. We identify four corpus systems that concretely realize the local geometric structures developed above, draw on data from the 81-run STRUC-PERC-I and 5,233-run STRUC-I corpora, and interpret each system as a probe of a specific aspect of the theory.

8.1 Corpus summary

The empirical foundation comprises:

- **STRUC-PERC-I v2.4.1:** 81 runs across 14 physical domains (molecular, atomic, biological, nuclear, condensed matter, CMB, cosmic web, atmosphere, geodesy, geoid/gravity, solar plasma, random matrices, adversarial). Zero hard USL violations across 80 of 81 runs; one Theorem 1 trigger (TiO_2 raw density, giant ratio 0.833). Verdict distribution: 48 FULL / 23 GIANT / 9 TAIL / 1 HARD.

- **STRUC-I v1.0.4:** 5,233 ladder evaluations including 3,069 physical runs, 1,920 crystallographic phase chains, 240 biological runs, and 4 cluster adversarial probes. Zero falsifications of the USL across all runs.

- **Phase Mapping Protocol:** 9,826 evaluations on a 17×17 joint (α, μ) deformation grid. All 34 tested datasets are structurally rigid (verdict constant across the grid); all structural commutators $C(\alpha, \mu; L)$ are identically zero.

8.2 Transition systems: ^{28}Si as an interior-to-boundary probe

Silicon-28 (^{28}Si) is a FULL percolation case with $\kappa_{\text{conn}} = 157{,}499$ and tail dominance 0.990. Its gap structure is dominated by a single extreme outlier γ-transition (max gap ratio 1.57×10^5) that pushes κ_{conn} into the extreme-delay regime (Regime III, $\kappa_{\text{conn}} \in [10^3, 4 \times 10^5]$) while the bulk of the level scheme is structurally coherent.

In local chart language, ^{28}Si sits in the interior of the FULL class, with a large positive $G_{j^*}(x)$ relative to the tail-dominance and connectivity thresholds. Under α-deformation (varying the fine-structure constant), ^{28}Si exhibits class changes: the Phase Mapping corpus shows that nuclear systems near the FULL/TAIL boundary can transition under operator deformation. This makes ^{28}Si an ideal *interior-to-boundary probe*: the system is far enough from the boundary that the active branch dominance (Lemma 5.1) is well-established, yet the boundary is visible under targeted deformation. The predicted monotonicity (Theorem 6.1) implies that the margin $m(^{28}\text{Si})$ decreases monotonically as α is deformed toward the class boundary.

8.3 Instability systems: ^{238}U as a boundary singularity probe

Uranium-238 (^{238}U) is a TAIL case with giant ratio 0.9829 and max gap ratio 2.5×10^9: a single extreme high-energy γ-transition produces a tail-dominance of 1.000, isolating the ladder from FULL percolation. The connectivity threshold κ_{conn} is not reached at any adaptive extension tested.

In local chart language, ^{238}U sits *on* or very close to the FULL/TAIL boundary. The tail dominance = 1.000 indicates that a single branch function $M_{j^*}(L)$ (the tail-dominance branch) is essentially at threshold: $G_{j^*}(\Phi(L_{^{238}\text{U}})) \approx 0$. This is precisely the regime where the regularity condition (Definition 4.1) must be checked: if $\nabla G_{j^*} \neq 0$, the boundary is smooth and the local theory applies; if the gradient vanishes, ^{238}U is at an irregular (possibly cusp) boundary point.

The isolated instability behaviour of ^{238}U thus provides a test of boundary regularity. A smooth fit $F(x) = 0$ through the neighbourhood of the ^{238}U chart point would confirm regular-point status. The three companion TAIL nuclei (^{150}Nd, ^{100}Mo, ^{48}Ca) with max ratios 2.2×10^{18}, 7.0×10^{18}, and 2.5×10^{18} respectively probe whether the boundary is piecewise smooth (Conjecture 9.1 below) or whether different isotope structures sample different boundary sheets.

8.4 Representation splitting: helium as a multi-chart probe

The helium atom provides the clearest empirical realization of multi-chart embeddings. The Phase Mapping corpus identifies four distinct ladder constructions of helium:

Encoding	Representation	Class	κ_{conn}
QM-I preprocessed	energy levels, $n = 1999$	FULL	329,031
Gap structure QM-I	difference sequence	FULL	329,031
Spectrum QM-I	raw spectrum, $n = 1684$	FULL	(large)
Zeeman ladder	magnetically split, $n = 1999$	TAIL	GR = 0.958
He singlet Zeeman	singlet transitions	TAIL	GR = 0.984
He triplet Zeeman	triplet transitions	TAIL	GR = 0.966

The QM-I encodings lie in the FULL class interior ($m > 0$ with large κ_{conn}); the Zeeman encodings lie in the TAIL region (m near zero by tail-dominance). These two groups sit in *different* realizability chart neighbourhoods: the FULL encodings are surrounded by the FULL interior in chart 1, while the Zeeman encodings are on the TAIL side of the FULL/TAIL boundary in chart 2.

This split is the corpus's primary realization of Lemma 7.2: the encoding family $\mathcal{E}(\mathrm{He})$ straddles a realizability boundary. The FULL encodings have larger margin and larger $d_{\partial\mathcal{C}}$ (farther from the boundary); the TAIL encodings have smaller margin and smaller $d_{\partial\mathcal{C}}$ (closer to, or on the other side of, the boundary). Theorem 7.1 predicts that within the FULL local regime, margin maximizers lie in FULL; within the TAIL local regime, margin maximizers lie in TAIL. The phase mapping data confirms that each encoding's class is rigid across the full (α, μ) deformation grid—no system transitions—which is consistent with the local geometric picture (each encoding is in the strict interior of its respective class neighbourhood).

The helium representation split therefore provides the corpus's clearest evidence that the canonical ladder problem—selecting the preferred encoding—has a locally well-defined answer: *within each regular local regime, the margin-maximizing encoding is class-consistent.*

8.5 TAIL nuclear isotopes as non-percolating boundary probes

The four TAIL nuclear isotopes (^{238}U, ^{150}Nd, ^{100}Mo, ^{48}Ca) are admissible (zero USL violations in the STRUC-I corpus; all with $A_\kappa \approx 1.000$) but non-percolating (TAIL in PRP). They probe the non-percolating side of the FULL/TAIL boundary without collapsing into HARD fragmentation (Theorem 1 is suppressed).

In the local geometric framework, these isotopes are in the interior of the TAIL class, with $G_{j^*}(x) < 0$ for the tail-dominance branch: they are on the far side of the boundary from the FULL interior. The margin $m(L) < 0$ for these systems reflects their location on the boundary's far side. The Lipschitz equivalence of Lemma 6.1 still holds (with appropriate sign convention), and the monotonicity theorem applies with outward direction now pointing *away* from the FULL class.

The extreme max-ratios (10^9–10^{18}) of these isotopes indicate that the decisive outlier coordinate $x_{\mathrm{tail}}(L) = \max(\Delta)/\mathrm{IQR}(\Delta)$ is astronomically large. The boundary value θ_{tail} of the tail-dominance threshold is orders of magnitude smaller than the actual coordinate value. This places these isotopes far from the boundary in the outward TAIL direction, implying $d_{\partial\mathcal{C}}(L)$ is large on the TAIL side. Theorem 6.1 predicts that within the TAIL class, margin is monotone in distance: isotopes with larger max-ratios are farther from the boundary and have larger (less negative) margin in the direction away from FULL.

Table 1 summarises the corpus data for these four systems.

8.6 The Zeeman domain as a boundary-stabilized probe

The Zeeman domain (Ca, H, He, Ag, Na Zeeman spectra) exhibits the highest structural pressure in the corpus: $\bar{\rho} \approx 0.9585$ from the STRUC-I v1.0.4 corpus, classified as boundary-stabilized. Simultaneously, Au Zeeman is the only atomic TAIL case, with giant ratio 0.9944 and tail dominance 0.998.

Table 1: TAIL nuclear isotopes: corpus data from STRUC-PERC-I. All four are admissible ($A_\kappa = 1.000$) and non-percolating. Tail dominance is 1.000 in all cases.

Isotope	Giant ratio	Max Δ/med	A_κ	Physical interpretation
^{238}U	0.9829	2.5×10^9	1.000	Single extreme high-E γ
^{150}Nd	0.9758	2.2×10^{18}	1.000	Shape-coexistence deformation
^{100}Mo	0.9884	7.0×10^{18}	1.000	Double β-decay outlier gap
^{48}Ca	0.9761	2.5×10^{18}	1.000	Doubly-magic + deformation

This combination—maximum STRUC-I pressure with maximum STRUC-PERC connectivity delay—makes the Zeeman domain the clearest corpus example of a system operating near a realizability boundary. In chart language: the Zeeman systems have $\bar{\rho} \approx 1$ (near the admissibility boundary of the STRUC-I framework) and $G_{j^*}(\Phi(L)) \approx 0$ for the tail-dominance branch. They are simultaneously near the admissibility boundary in the USL coordinate and near the realizability class boundary in the PRP coordinate.

This co-occurrence supports the geometric picture of $\mathcal{M}_{\text{adm}}$ as a space with two distinct boundary structures (admissibility and realizability) that can be simultaneously approached. The two boundaries are independent (per Dual Observability), but a system can sit near both simultaneously—as the Zeeman domain demonstrates.

8.7 Percolation regimes as local chart structure

The three-regime percolation structure identified in the corpus corresponds, in chart coordinates, to three distinct positions of the FULL/GIANT class boundary:

Regime	κ_{conn} range	Domains	Chart interpretation
I (Immediate)	< 2	Biological, Atmosphere	Deep interior of FULL; $G_{j^*}(x) \gg \theta_{j^*}$
II (Moderate)	2–10^3	Condensed matter, CMB	Intermediate interior; boundary visible in chart
III (Extreme)	10^3–4×10^5	Nuclear, Atomic	Near FULL/TAIL boundary; $G_{j^*}(x) \approx \theta_{j^*}^+$

The fact that all Regime I and II systems are FULL (with κ_{conn} finite) while Regime III systems include both FULL and TAIL cases is consistent with the local geometric picture: as $G_{j^*}(\Phi(L))$ decreases toward zero, systems approach the class boundary and eventually cross it.

8.8 Cross-system structural phase landscape (atomic corpus)

The single-system probes of Sections 8.2–8.5 are complemented by a cross-system atomic corpus: 16 direct STRUC-PERC-I runs across six elements (H, He, Li, Na, Fe, Ag, Au) in two representation families (QM-I energy-level ladders and Zeeman magnetically-split ladders). This corpus extends the local-geometry analysis from single-system validation to cross-domain variability.

The 16 runs partition into five distinct structural regimes in the decisive 2D chart $(\mathrm{tailDom}, C)$:

(A) *FULL interior*: He, Li, Na QM-I (6 systems; $C = 1.000$, $G_3 = 0$);

(B) *HARD*: H QM-I, all three representations (3 systems; $C \in [0.81, 0.93]$);

(C) *GIANT corner*: H, He, Na Zeeman (4 systems; $\mathrm{tailDom} > 0.999$, $C \in [0.995, 0.998]$, near chart corner $(1, 1)$);

(D) *GIANT interior*: Fe, Ag Zeeman (2 systems; $\mathrm{tailDom} \in [0.967, 0.998]$);

(E) *TAIL*: Au Zeeman (1 system; $C = 0.994$, no percolation at any tested scale).

Three results follow from this corpus.

Phase structure. Realizability space partitions into structurally distinct regimes across atomic systems and representation families. These regimes are not arbitrary: they reflect the decisive coordinates that govern each system's local boundary, and they admit physically interpretable explanations (e.g., hydrogen's n^{-3} level spacing drives the HARD regime regardless of encoding; light-atom Zeeman ladders cluster near the corner due to sparse outlier gap structure).

Locality confirmed. Within the narrow zone $\mathrm{tailDom} > 0.995$ (13 systems), an empirically fitted boundary function $F(x) = C(L) - \mathrm{tailDom}(L) = 0$ correctly classifies 13/13 runs. Outside this zone, iron Zeeman ($\mathrm{tailDom} = 0.967$) has $F > 0$ (FULL side) but is classified GIANT: the fitted separator fails. This failure is not a contradiction; it delineates the domain of validity of the local chart fitted from the helium corpus. This confirms that realizability chart geometry is inherently local and system-dependent, consistent with Lemma 3.1.

Domain variability. Representation families occupy different regions of the decisive chart depending on atomic structure. Hydrogen QM-I is HARD across all encodings; gold Zeeman is TAIL. The Zeeman tail-dominance coordinate decreases with atomic structural complexity, producing a regime shift from GIANT-corner (light atoms) through GIANT-interior (transition metals) to TAIL (gold). These transitions are observable in the corpus but not yet explained by a single global boundary function, underscoring the necessity of the local-chart framework developed in this paper.

9 Conjectures and open problems

Section abstract. We state the principal conjectures that remain open after the local theory is established, and identify the specific proof steps needed for each.

9.1 Piecewise-hypersurface boundary structure

Conjecture 9.1 (Piecewise-hypersurface boundary)**.** The global realizability boundary $\partial\mathcal{C}$ is locally piecewise C^1 in decisive charts: it decomposes into finitely many C^1 hypersurface patches, each corresponding to one active decisive threshold, meeting along lower-dimensional strata.

Evidence: The TAIL nuclear isotopes suggest that the FULL/TAIL boundary in the nuclear domain has at least two sheets (one for giant-ratio threshold and one for tail-dominance threshold), which meet at isotopes like ^{238}U that are near both thresholds simultaneously. The cosmic web domain (3 TAIL, 2 GIANT, no κ_{conn}) suggests a third sheet in the GIANT/TAIL boundary.

9.2 Global chart compatibility

Conjecture 9.2 (Global chart compatibility)**.** There exists a finite atlas of decisive realizability charts $\{\Phi_i : U_i \to \mathbb{R}^{d_i}\}_{i=1}^{N}$ covering $\mathcal{M}_{\mathrm{adm}}$ such that the margin $m(L)$ is globally order-equivalent to the chart-local boundary distance $d_{\partial\mathcal{C}}(L)$ within each U_i, and the order relation is preserved across overlapping charts.

This would promote the local result of Section 6 to a global statement, resolving the original open problem of whether $m(L)$ is a genuine distance-type functional on $\mathcal{M}_{\mathrm{adm}}$.

9.3 Global margin monotonicity

Conjecture 9.3 (Global margin monotonicity)**.** For any two admissible ladders L_1, L_2 of the same physical system in the same realizability class, if $d_{\partial\mathcal{C}}(L_1) < d_{\partial\mathcal{C}}(L_2)$ (measured in any consistent structural metric on $\mathcal{M}_{\mathrm{adm}}$), then $m(L_1) < m(L_2)$.

The local version (Theorem 6.1) is proved here. The global version requires chart compatibility (Conjecture 9.2) and a global structural metric that is consistent with the local chart distances.

9.4 Exact metric equivalence

Conjecture 9.4 (Metric equivalence)**.** There exists a Riemannian metric g on $\mathcal{M}_{\mathrm{adm}}$ (compatible with the gap-ratio topology) such that $m(L) = d^{g}_{\partial\mathcal{C}}(L)$ exactly, not merely up to bi-Lipschitz equivalence.

This is the strongest possible form of the margin-boundary relation. It would make $m(L)$ a true signed distance function in a natural geometric sense.

10 Limits of the local theory

Section abstract. We state clearly what the present results do not prove, identify the structural obstacles to global extension, and outline the proof-theoretic steps needed.

10.1 What this paper does not prove

The theorems of this paper are strictly local. We do not prove:

- **Global monotonicity.** Theorem 6.1 holds in a small neighbourhood of a single regular boundary point. Extending it globally requires chart compatibility and a global structural metric, neither of which is established here.

- **Uniqueness of canonical ladders.** Theorem 7.1 establishes that local margin maximizers lie in one class; it does not establish that there is a single preferred encoding within that class, or that the canonical ladder is unique.
- **Global realizability chart atlas.** We prove local chart existence (Lemma 3.1) but not that local charts assemble into a globally consistent atlas.
- **Irregular boundary regularity.** Systems at irregular boundary points (corners, cusps, simultaneous multi-threshold crossings) are excluded from the hypotheses of Theorem 6.1. The nuclear TAIL isotopes may include such cases.
- **Sufficient direction of PRP.** The open conjecture that percolation implies admissibility (PRP Conjecture 1) is not addressed here; the results of this paper concern the geometry of realizability boundaries, not the relationship between admissibility and realizability.

10.2 Proof-theoretic path to global monotonicity

The proof of global monotonicity (Conjecture 9.3) requires the following steps, in order:

1. Prove Conjecture 9.1 (piecewise C^1 boundary) by exhaustive analysis of all decisive threshold combinations in the PRP algorithm.
2. Construct the global chart atlas (Conjecture 9.2) by verifying that charts around all boundary-adjacent systems in the corpus are compatible on overlaps.
3. Introduce a global structural metric on $\mathcal{M}_{\mathrm{adm}}$ consistent with local chart distances; the most natural candidate is a Riemannian metric on the gap-ratio profile space.
4. Prove that the active-branch dominance (Lemma 5.1) holds globally: that the minimum over decisive branches is always attained by a unique branch away from branch-crossing strata.
5. Apply the local Corollary 6.1 patchwise across the atlas.

Steps 1–2 are empirically supported by the corpus but not yet formally established. Steps 3–5 are purely theoretical and do not require additional corpus data.

11 Conclusion

Section abstract. We summarise the geometric contribution, state the theorem programme achieved, and position the results in the broader context of the UNNS programme.

11.1 Summary of results

This paper establishes the following:

1. **Local charts exist.** Every admissible ladder has a neighbourhood in $\mathcal{M}_{\mathrm{adm}}$ with a finite decisive coordinate system (Lemma 3.1). At most five decisive observables suffice locally, matching the operational structure of STRUC-PERC-I.

2. **Realizability boundaries are locally geometric objects.** At every regular boundary point, the class boundary is a codimension-1 C^1 hypersurface in a decisive chart (Theorem 4.1). Class changes are not algorithmic black boxes; they are threshold crossings in a locally smooth surface.

3. **The connectivity margin is a local boundary-distance functional.** The margin $m(L)$ is locally order-equivalent to the boundary distance $d_{\partial\mathcal{C}}(L)$ (Corollary 6.1) and is strictly monotone with respect to boundary distance along any transversal direction (Theorem 6.1).

4. **Maximum-margin selection is locally justified.** Within any regular local regime, all margin-maximizing encodings of a physical system lie in the same realizability class (Theorem 7.1). The canonical ladder programme has a local geometric foundation.

11.2 Position in the programme

These results occupy a specific and necessary position in the UNNS programme. The USL (admissibility), PRP (realizability), Dual Observability (independence), and Structural Response (operator deformation) manuscripts establish the *existence* and *robustness* of the two-coordinate structural framework. The present manuscript supplies the *geometry* of that framework's internal structure: it converts realizability boundaries from implicit algorithm outputs into explicit geometric objects, and converts the connectivity margin from an empirical invariant into a theorem-grade distance functional.

The gap that remains—global monotonicity, chart compatibility, and a canonical Riemannian structure on $\mathcal{M}_{\mathrm{adm}}$—is now a structured theorem programme with a clear proof roadmap, rather than a vague frontier. This is the intended contribution: not the final global result, but the correct local geometric layer on which any future global result must rest.

The cross-system atomic corpus demonstrates that this local structure is not an artefact of a single well-chosen system: realizability geometry is not globally uniform. Instead, $\mathcal{M}_{\mathrm{adm}}$ consists of multiple locally valid chart neighbourhoods, each with its own decisive coordinates and boundary geometry, connected by regime transitions that are visible in the corpus but not yet analytically characterised. Extending the local theory to a global atlas of charts—proving chart compatibility, defining transition maps between regimes, and identifying the global metric—remains the central open problem of the programme.

A Proof details for Lemma 6.1

We collect the Taylor expansion argument used in Lemma 6.1 in full detail for reference.

Let $x^* = \Phi(L^*)$ and let $G = G_{j^*}$ be C^1 with $\mathbf{n} = \nabla G(x^*) \neq 0$. Set $\lambda = \|\mathbf{n}\| > 0$ and $\hat{\mathbf{n}} = \mathbf{n}/\lambda$.

Step 1. By the C^1 assumption, $G(x) = G(x^*) + \mathbf{n} \cdot (x - x^*) + R(x)$ with $|R(x)| = o(\|x - x^*\|)$. Since $G(x^*) = 0$: $G(x) = \lambda\, \hat{\mathbf{n}} \cdot (x - x^*) + R(x)$.

Step 2. The signed distance from x to the tangent hyperplane $H = \{y : \hat{\mathbf{n}} \cdot (y - x^*) = 0\}$ is $d_H(x) = \hat{\mathbf{n}} \cdot (x - x^*)$. Therefore $G(x) = \lambda\, d_H(x) + R(x)$.

Step 3. The actual distance $d_{\partial\mathcal{C}}(L)$ to the C^1 hypersurface $\{G = 0\}$ satisfies $|d_{\partial\mathcal{C}}(L) - d_H(x)| = O(\|x - x^*\|^2)$ in a sufficiently small neighbourhood (standard differential geometry: the C^1 surface deviates from its tangent plane at second order). Thus $d_H(x) = d_{\partial\mathcal{C}}(L) + O(d_{\partial\mathcal{C}}(L)^2)$ and

$$G(x) = \lambda\, d_{\partial\mathcal{C}}(L) + O(d_{\partial\mathcal{C}}(L)^2) + o(d_{\partial\mathcal{C}}(L)).$$

Step 4. For sufficiently small $d_{\partial\mathcal{C}}(L)$, the error terms are bounded by $\lambda\, d_{\partial\mathcal{C}}(L)/2$ in absolute value. Setting $c_1 = \lambda/2$ and $c_2 = 2\lambda$ completes the proof.

B Decisive coordinate table

Table 2 lists the decisive coordinates used in STRUC-PERC-I, their corresponding threshold functions G_j, and the class transition each governs.

Table 2: Decisive coordinates and their threshold functions. med = median(Δ), IQR = IQR(Δ).

Coordinate x_j	Definition	Threshold θ_j	Governs transition
Giant ratio	$C(L)$	0.95	HARD ↔ FULL/GIANT/TAI
Tail dominance	$\max(\Delta)/\mathrm{IQR}(\Delta)$	$\max(10 \cdot \mathrm{med}, 5 \cdot \mathrm{IQR})$	FULL/GIANT ↔ TAIL
Secondary cluster ratio	$\lvert C_2\rvert/\lvert C_1\rvert$	0.05	GIANT ↔ FULL/TAIL
κ_{conn}	connectivity threshold	grid range	FULL ↔ GIANT
Structural pressure	$\bar{\rho} = \mathrm{inv}/\nu$	1.0 (USL limit)	Admissibility boundary

C Theorem summary

Result	Type	Statement
3.1	Lemma	Finitely many decisive coordinates govern local class changes.
4.1	Lemma	Each decisive event is representable as $G_j(x) = 0$ with $\nabla G_j \neq 0$.
4.1	Theorem	Near every regular boundary point, $\partial\mathcal{C}$ is a C^1 codimension-1 hypersurface.
5.1	Lemma	A single branch M_{j^*} dominates $m(L)$ in a neighbourhood.
5.2	Lemma	m is strictly increasing along outward transversal curves.
6.1	Lemma	$c_1\, d_{\partial\mathcal{C}}(L) \leq G_{j^*}(\Phi(L)) \leq c_2\, d_{\partial\mathcal{C}}(L)$ locally.
6.1	Corollary	$m(L)$ is locally order-equivalent to $d_{\partial\mathcal{C}}(L)$.
6.1	**Theorem**	$d_{\partial\mathcal{C}}(L_1) < d_{\partial\mathcal{C}}(L_2) \Rightarrow m(L_1) < m(L_2)$ near a regular boundary point.
7.1	Lemma	Outward transversal curves remain in their realizability class.
7.2	Lemma	Local margin maximizers cannot lie on opposite sides of a boundary.
7.1	**Theorem**	All margin-maximizing encodings in a regular local regime belong to one class.

D Ladder construction protocol specification

This appendix formalises the four protocol families referenced in Section 2.6 and specifies the construction rules used in the corpus runs underlying Sections 8 and 8.8. The goal is reproducibility: any corpus run should be recoverable from the raw data source and the protocol specification alone.

D.1 General constraints (all protocols)

Every protocol Π must produce a ladder satisfying:

(G1) *Strict monotonicity:* $x_1 < x_2 < \cdots < x_n$ (no duplicates, no ties).

(G2) *Finiteness:* $n < \infty$; if the raw source has degenerate entries, ties are broken by the rule stated in the protocol variant.

(G3) *Finite gaps:* $\Delta_i = x_{i+1} - x_i > 0$ and $\max(\Delta) < \infty$ (non-finite values are dropped before sorting).

(G4) *USL check:* after construction, the USL verdict is computed and recorded; any violation is flagged and the run is marked invalid for admissibility-dependent claims.

D.2 Protocol A: QM-I energy-level ladders

Data source. Published energy level tables for atoms or molecules (e.g., NIST Atomic Spectra Database, HITRAN, ENSDF).

CSV schema (atomic corpus). The atomic QM-I files have the following confirmed column structures:

Variant	CSV columns present
Spectrum	`index`, `energy_cm-1`
Preprocessed	`index`, `energy_cm-1`, `raw_gap_cm-1`, `normalized_gap`, `scale`
Gap structure	`index`, `energy_cm-1`, `gap_cm-1`

A1: Spectrum variant.

(A1.1) Extract the `energy_cm-1` column only.

(A1.2) Drop non-finite and negative values.

(A1.3) Sort ascending and remove duplicates.

(A1.4) The resulting sequence is L_{spec}.

Verified counts (atomic corpus): H spectrum: canonical $n = 106$, generic $n = 212$ (index column included).

A2: Preprocessed variant.

(A2.1) Extract the `energy_cm-1` column only.

(A2.2) Apply domain-specific filtering (energy range $[E_{\text{min}}, E_{\text{max}}]$, recorded per run).

(A2.3) Sort ascending and remove duplicates.

(A2.4) The resulting sequence is L_{pre}.

Verified counts (atomic corpus): H preprocessed: canonical $n = 50$, generic $n = 161$ (index, energy, raw gap, and normalised gap all parsed as numeric values).

A3: Gap-structure variant.

(A3.1) Extract the `gap_cm-1` column only.

(A3.2) Drop non-finite and zero values.

(A3.3) Sort ascending and remove duplicates.

(A3.4) The resulting sequence is L_{gap}.

Verified counts (atomic corpus): H gap: canonical $n = 31$ (unique positive gaps), generic $n = 130$ (index, energy, and gap columns all parsed together).

D.3 Protocol B: Zeeman ladders

Data source. Magnetic-field-split energy levels at a fixed field strength B, from theoretical calculation or observation.

CSV schema (atomic corpus). All Zeeman files share the confirmed schema:

`B_T`, `Configuration`, `Term`, `J`, `mJ`, `Lande_g`, `Level0_cm1`, `LevelB_cm1`.

B1: Standard Zeeman variant.

(B1.1) Extract the `LevelB_cm1` column only (the field-resolved energy level).

(B1.2) Drop non-finite values.

(B1.3) Sort ascending and remove duplicates.

(B1.4) The resulting sequence is L_{ZEE}.

Verified counts (atomic corpus): H Zeeman: canonical $n = 23{,}704$ (unique `LevelB_cm1` values), generic $n = 23{,}836$ (`Level0_cm1` also parsed as numeric). Fe Zeeman: canonical $n = 247{,}278$, generic $n = 247{,}700$ (downsampled to $n \approx 123{,}850$ in approximate mode).

B2: Singlet/Triplet sub-family variants. Apply B1 restricted to rows where the `Term` column specifies singlet ($S = 0$) or triplet ($S = 1$) states. The sub-family selector is recorded in the run metadata.

Approximate mode. When $n_{\mathrm{orig}} > N_{\max}$ (instrument cap, currently $N_{\max} = 200{,}000$ gaps), a uniform stride $s = \lceil n_{\mathrm{orig}}/N_{\max} \rceil$ is applied. Approximate mode suppresses Theorem 1 assertions and is flagged in run metadata.

D.4 Protocol C: Gap-derived ladders

Given any base ladder L_{base} produced by Protocol A or B:

(C1) Compute $\Delta(L_{\mathrm{base}})$.

(C2) Treat Δ as a new ladder (each gap becomes a level).

(C3) Apply constraints (G1)–(G4) to the resulting sequence.

D.5 Protocol D: Transform ladders

Given a base ladder L_{base} and a specified monotone transform $T : \mathbb{R} \to \mathbb{R}$:

(D1) Apply T elementwise: $L_T = T(L_{\mathrm{base}})$.

(D2) Re-sort and deduplicate (necessary if T is not strictly monotone on the image).

(D3) Apply constraints (G1)–(G4).

D.6 Generic adapter protocol (exploratory)

For exploratory runs where column selection is not specified:

(E1) Parse all numeric cells in the data source (using JavaScript `parseFloat` semantics or equivalent).

(E2) Drop non-finite values.

(E3) Sort ascending.

(E4) Remove duplicates (exact equality).

(E5) Apply constraints (G1)–(G4).

Runs produced under the generic adapter are labelled `adapter=generic` in run metadata. They are internally consistent (all physical corpus results are validated against raw source files) but are not assigned a canonical physics interpretation until a domain-specific protocol (A–D) is applied.

D.7 Protocol metadata record

Every corpus run records:

Field	Content
`protocol_family`	A / B / C / D / generic
`protocol_variant`	A1 / A2 / A3 / B1 / B2 / B1-approx / …
`data_source`	NIST / HITRAN / ENSDF / computed / …
`energy_column`	column identifier (canonical runs only)
`energy_unit`	cm^{-1} / eV / Hz / …
`n_elements`	number of ladder elements after construction
`n_orig`	original element count before downsampling (if applicable)
`stride`	downsampling stride (1 = no downsampling)
`usl_verdict`	admissible / violated
`adapter`	generic / canonical

Acknowledgements. The analysis in this paper builds on the STRUC-PERC-I and STRUC-I instrumentation developed within the UNNS Substrate Research Program. Corpus runs draw on data from HITRAN, ENSDF (nuclear γ-levels), ERA5 atmospheric reanalysis, Planck 2018 CMB spectra, and the QT45 ribozyme fitness landscape dataset.

UNNS Substrate Research Program · `unns.tech` · Version 1.0 · April 2026