Projective metric geometry of tropical nuclei: gap matrices, event loci, and order chambers

Juan Luis Gastaldi[1†], Samantha Jarvis[2†], Thomas Seiller[3†], John Terilla[2,4†]

[1]ETH Zurich, Zurich, Switzerland.
[2]Queens College, City University of New York, New York, NY, USA.
[3]CNRS, Paris, France.
[4]The Graduate Center, City University of New York, New York, NY, USA.

Contributing authors: juan.luis.gastaldi@inf.ethz.ch; Samantha.jarvis@qc.cuny.edu; thomas.seiller@cnrs.fr; jterilla@gc.cuny.edu;
[†]The authors contributed equally to this work.

Abstract

The tropical row span and column span of a real matrix are, from the polyhedral point of view, different objects living in different ambient spaces. These polytopes are known to be combinatorially isomorphic as polyhedral complexes; we prove that they are isometric under a Hilbert projective metric. We show that this isometry, along with a considerable amount of additional metric and polyhedral structure, is a direct consequence of a single categorical construction: the Isbell nucleus of the matrix, viewed as a profunctor enriched over the extended reals.
The projective nucleus carries two canonical structures inherited from enrichment. The first is a Hilbert projective metric, with respect to which the Isbell conjugate maps are mutually inverse isometries—this is the Isometry Theorem. The second is a polyhedral cell decomposition cut out by the Isbell inequalities, recovering the type decomposition of tropical convexity.
These two structures are linked pointwise by the *gap matrix*. The Events Theorem identifies each positive entry of the gap matrix with the exact projective distance to the locus where the corresponding inequality becomes tight: algebraic slack in the Isbell inequalities equals geometric distance to the cell walls. Thresholding the gap matrix at successive radii produces a constructible sheaf of formal concept lattice towers, extracting discrete algebraic structure from the continuous geometry at each point.

In the square case there is generically a unique full-dimensional cell. The Centering Theorem identifies its Chebyshev center—the point maximally insulated from all cell walls—and shows that the optimal radius equals the minimum directed cycle mean of an associated digraph, connecting the projective geometry of the nucleus to the classical theory of optimal assignments.

Keywords: tropical convexity, tropical polytopes, Hilbert projective metric, polyhedral complexes, Isbell duality, enriched category theory, formal concept analysis, minimum cycle mean

2020 MSC Classification: 14T10 , 52B11 , 15A80 , 18N10 , 06A15 , 90C27

1 Introduction

A real $m \times n$ matrix M gives rise to two tropical polytopes in tropical projective spaces: a row span in $\mathbb{TP}^{n-1}$ and a column span in $\mathbb{TP}^{m-1}$. Each carries a polyhedral decomposition into cells according to combinatorial type. A foundational result of Develin and Sturmfels [1] is that these two polytopes are isomorphic as polyhedral complexes: there is a canonical bijection between their cells that preserves combinatorial type. Their proof is a direct polyhedral argument, relying on the structure of the type decomposition. The starting point of this paper is the observation that a stronger fact holds and admits a conceptual explanation: the row and column spans are not merely combinatorially isomorphic but also isometric under the Hilbert projective metric, and this isometry is an immediate consequence of the fact that both are projections of a single intrinsic object, the *Isbell nucleus* of M.

The Isbell nucleus is defined for any $\overline{\mathbb{R}}$-enriched profunctor $M : \mathcal{C}^{\mathrm{op}} \otimes \mathcal{D} \to \overline{\mathbb{R}}$, where $\overline{\mathbb{R}}$ is the monoidal poset $\overline{\mathbb{R}} = ([-\infty, \infty], \leq, +)$. The Isbell conjugates M^* and M_* form an adjunction between $\overline{\mathbb{R}}$-enriched presheaves on $\mathcal{C}$ and copresheaves on $\mathcal{D}$:

$$M^* f(d) = \min_{c \in \mathcal{C}} \bigl(M(c,d) - f(c)\bigr), \qquad M_* g(c) = \min_{d \in \mathcal{D}} \bigl(M(c,d) - g(d)\bigr).$$

The nucleus $\mathrm{Nuc}(M)$ is the fixed-point locus, the set of pairs (f, g) with $M^* f = g$ and $M_* g = f$. When $\mathcal{C}$ and $\mathcal{D}$ are finite sets and M is a real matrix, the projectivization $\mathbb{P}\mathrm{Nuc}(M)$ is a compact polyhedral space carrying two canonical structures: a *projective metric* coming from the enriched hom, and a *polyhedral cell decomposition* coming from the Isbell inequalities $f(c) + g(d) \leq M(c, d)$. The cell decomposition recovers the type decomposition of tropical convexity: a point is classified by the incidences (c, d) for which equality holds, which we call *witness pairs*. Both structures are invariant under external gauge transformations—reweighting M by row and column potentials—since these act by isometries. The nucleus is the intrinsic geometric object; the matrix M is just a coordinate presentation. In particular, $\mathbb{P}\mathrm{Nuc}(M)$ maps isometrically onto both the tropical row span and the tropical column span: there is no mystery about why the two spans are isometric, because they are projections of a single intrinsic space.

The central new tool is the *gap matrix* defined at each nucleus point (f, g) by $\delta^{(f,g)}(c, d) = M(c, d) - f(c) - g(d)$. It is a non-negative matrix with at least one zero in every row and column. Its zero entries record the witnesses and hence the combinatorial cell. Its positive entries, which measure the slack in each inequality, turn out to have a sharp geometric meaning. The three main results of the paper are the following.

Isometry Theorem (Theorem 13).

The projective Isbell maps M^* and M_* restrict to mutually inverse isometries between the presheaf and copresheaf realizations of $\mathbb{P}\mathrm{Nuc}(M)$ for the Hilbert projective metric. The tropical row and column spans are here revealed as two projections of a single intrinsic space.

Events Theorem (Theorem 40).

Each positive entry of the gap matrix equals the exact projective distance to the corresponding event locus:

$$d_{\mathbb{P}\mathrm{Nuc}}\left((f, g), \mathcal{E}_{c,d}\right) = \delta^{(f,g)}(c, d),$$

where $\mathcal{E}_{c,d}$ is the locus on which (c, d) is a witness pair. In other words, algebraic slack in the Isbell inequalities *is* geometric distance to the cell walls. This identity reflects a rigidity special to the Hilbert metric and the linear structure of the Isbell conditions: in a generic polyhedral-metric setting, constraint slack and boundary distance can differ by arbitrary distortion factors.

Centering Theorem (Theorem 53).

In the square case $|\mathcal{C}| = |\mathcal{D}| = n$, there is generically one full-dimensional cell. The Chebyshev center of this cell—the point maximizing the minimum distance to all cell walls—has optimal radius equal to the minimum directed cycle mean of an associated weighted digraph. At the center, the smallest positive gap is achieved with multiplicity at least n (generically exactly n), identifying n equidistant event loci. The minimum cycle mean is computable in $O(n^3)$ time by Karp's algorithm [2], connecting the projective geometry of the nucleus to the classical theory of optimal assignment [3, 4].

Further structures.

The Events Theorem has further structural consequences. Thresholding the gap matrix at a value $\varepsilon > 0$ records which event loci lie within projective distance ε of a given point. The resulting Boolean relation is itself a profunctor, and its Isbell nucleus is a finite lattice: a formal concept lattice in the sense of Wille. As ε increases, new witness pairs enter and the lattice grows. Since the gap matrix has finitely many distinct values, this growth factors through a finite tower of concept lattices, one for each distinct gap value. Wall-crossing between consecutive floors is governed by canonical mergers. The tower depends on the basepoint (f, g), but only through the ordering of the gap entries: within each *order chamber*—a region where the ordering is constant—the tower is invariant. Moving to a face of the chamber complex merges consecutive floors, producing canonical specialization maps. The order chambers give a refinement of the polyhedral decomposition defined by witnesses and appears to be new in tropical geometry.

This pointed thresholding procedure deserves a moment of emphasis. The naive operation of thresholding the matrix M directly ignores the geometry of the nucleus and produces no meaningful structure, as far as we can see. But thresholding the gap matrix $\delta^{(f,g)}$—which depends on one's position in the nucleus—extracts a principled family of discrete algebraic structures (formal concept lattices with joins, meets, and Galois connections) from the continuous projective geometry, organized into a constructible sheaf over the polyhedral order-chamber complex.

A running example.

To keep the discussion concrete, we repeatedly return to the following matrix. Let $C = \{c_0, c_1, c_2\}$ and $D = \{d_1, d_2, d_3, d_4\}$ and set

$$M = \begin{bmatrix} 0.7 & 1.5 & 1.7 & -1.3 \\ 1.2 & 2.6 & 0.1 & 2.2 \\ 2.0 & -1.6 & 2.0 & -2.9 \end{bmatrix}.$$

In this case $\mathbb{P}\mathrm{Nuc}(M)$ is a two-dimensional polyhedral complex. Work in the gauge slice $f(c_0) = 0$ and consider the point (f, g) with

$$f = (0, 0, 0), \qquad g = (0.7, -1.6, 0.1, -2.9).$$

Its gap matrix is

$$\delta^{(f,g)} = \begin{bmatrix} 0 & 3.1 & 1.6 & 1.6 \\ 0.5 & 4.2 & 0 & 5.1 \\ 1.3 & 0 & 1.9 & 0 \end{bmatrix},$$

whose positive values satisfy

$$0 < 0.5 < 1.3 < 1.6 = 1.6 < 1.9 < 3.1 < 4.2 < 5.1.$$

The zero pattern $\{(c_0, d_1), (c_2, d_2), (c_1, d_3), (c_2, d_4)\}$ determines the witness cell containing (f, g), while the positive entries are—by the Events Theorem—the exact projective distances from (f, g) to the surrounding cell walls. The tie at 1.6 is the value at which the order of the gap entries changes, producing a wall in the order-chamber decomposition. Figure 1 illustrates this gap-value-as-distance phenomenon.

1.1 Relation to existing work

On the tropical side, Develin and Sturmfels introduced the polyhedral theory of tropical convexity and its decomposition into combinatorial types [1]; see also [5] for

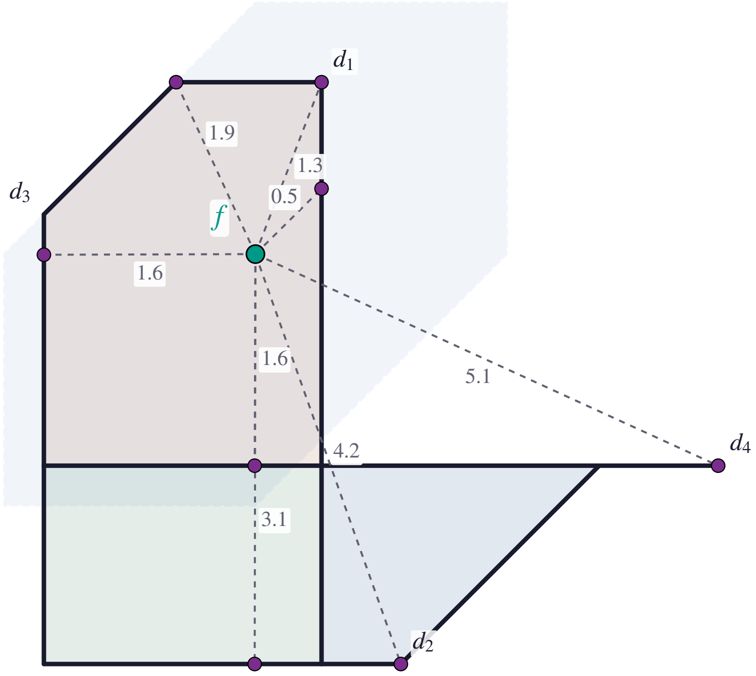

Fig. 1 The running 3×4 example in an affine chart of $\mathbb{P}\mathrm{Nuc}(M)$. The green point is the basepoint (f, g), and the shaded regions are the 2-cells of the witness decomposition. For each positive gap value, the corresponding marked event point lies on the first event locus encountered at that radius, as indicated by the dashed segment from (f, g). The repeated value 1.6 is the tie that later becomes the order wall between two adjacent chamber refinements.

background. The half-space and distance formulas that underlie our Events Theorem are close in spirit to work of Gaubert and Katz on max-plus convexity and tropical half-spaces [6, 7]; compare also [8]. Gaubert and Sergeev [9] study cyclic projectors and separation in idempotent convex geometry, using the Hilbert projective metric in the tropical setting; our 1-Lipschitz and isometry results for the Isbell maps can be seen as complements to their spectral approach.

On the categorical side, the idea of treating generalized metric spaces as categories enriched over an ordered monoid goes back to Lawvere [10]. Isbell conjugacy originates in Isbell's paper on adequate subcategories [11], and Avery and Leinster [12] give a modern treatment over an arbitrary base. In the setting of Lawvere metric spaces, Willerton developed the Isbell completion in detail [13–15], relating it to the Legendre–Fenchel transform. Via the c-transform, the same formalism also connects to optimal transport [16].

The bridge from enriched-category theory to tropical and directed-metric phenomena has been explored from several directions. Elliott and Fujii observed that Isbell-type nuclei provide a natural home for tropical polytopes [17, 18]. Bradley, Terilla, and Vlassopoulos use enrichment over $[0,1] \cong [0,\infty]$ to organize linguistic structure [19], while Gaubert and Vlassopoulos develop related directed-metric and tropical-polyhedral ideas in the setting of large language models [20]. Recent work of Bradley and Vigneaux studies magnitude and magnitude homology for categories of texts enriched by language-model probabilities [21]. Background and motivation for these applications are reviewed in [22].

The Centering Theorem connects to the combinatorial optimization literature. The minimum cycle mean of a weighted digraph is a classical invariant studied by Karp [2], and the assignment-problem duality underlying the Chebyshev LP is closely related to the Hungarian method [3, 4]; see [23, Ch. 17] for a textbook treatment. In a related but distinct direction, Akian, Gaubert, Qi, and Saadi [24] prove that the inner radius of a Hilbert ball inscribed in a tropical polyhedron equals the value of a mean payoff game; our Centering Theorem can be viewed as a pointwise refinement of this circle of ideas, identifying the Chebyshev radius of a specific cell within the nucleus with the minimum cycle mean of an explicit digraph derived from the optimal assignment.

These constructions fit into a broader program in which nuclei acquire additional compatible structure. In particular, nuclei arising from profunctors compatible with monoidal data carry further operations relevant to linear realizability; see [25–27]. The finite real case studied here exhibits the projective metric geometry, witness cells, event loci, order chambers, the associated constructible sheaf of lattice towers, and the Chebyshev centering of full-dimensional cells.

Our conventions differ slightly from some of the literature, and these differences matter for the geometry. We therefore keep Section 2 self-contained, but restrict it to the categorical material used later. The cocompletion formulas behind the Yoneda-density statements are standard; see [28].

Organization of the paper. Section 2 fixes conventions for $\overline{\mathbb{R}}$-enriched categories, profunctors, and the Isbell adjunction. Section 3 develops the projective metric geometry and proves the Isometry Theorem. Section 4 develops the witness polyhedral structure, proves the Events Theorem, introduces order chambers, and shows how pointed thresholding of the gap matrix assembles into chamberwise concept-lattice towers with canonical face-specialization maps. Section 5 specializes to the square case and proves the Centering Theorem.

2 Isbell duality over the extended reals

We fix conventions for $\overline{\mathbb{R}}$-enriched categories, the Isbell adjunction associated to a profunctor $M : \mathcal{C} \nrightarrow \mathcal{D}$, and its nucleus $\mathrm{Nuc}(M)$. The weighted-colimit formulas behind the density identities are standard; see [28].

2.1 The arithmetic of $\overline{\mathbb{R}}$

We begin with the real numbers ordered by $\leq$. Viewed as a category, its limits are infima and its colimits are suprema. Adjoining $\pm\infty$ makes $\overline{\mathbb{R}} = \mathbb{R} \cup \{\pm\infty\}$ into a poset category that is complete and cocomplete: infema and suprema of arbitrary subsets exist. Addition of real numbers defines a symmetric monoidal structure which extends to $(\overline{\mathbb{R}}, \leq, +, 0)$ by declaring $-\infty$ absorbing: $-\infty + y = -\infty$ for all $y \in \overline{\mathbb{R}}$, including $y = +\infty$. This convention is dictated by the requirement that each translation $x \mapsto x + y$ preserve colimits (suprema) and hence have a right adjoint: any other extension of addition to $\pm\infty$ would violate this. The right adjoint is *residuation* $z - y := [y, z]$, characterized by $x + y \leq z \iff x \leq z - y$ and given explicitly by

$$z - y = \sup \left\{ x \in \overline{\mathbb{R}} \mid x + y \leq z \right\} . \tag{1}$$

On finite reals this is ordinary subtraction. At the boundary one has $\infty - \infty = \infty$ and $-\infty - (-\infty) = \infty$; in particular, subtracting $-\infty$ is not the same as adding $+\infty$. Three properties are used throughout: residuation $z \mapsto z - y$ is monotone and preserves infima (being a right adjoint); the reverse map $z \mapsto x - z$ is antitone; and $-\infty$ is absorbing for addition while ∞ is absorbing for residuation.

We use $\overline{\mathbb{R}}$ rather than Lawvere's $([0, \infty], \geq, +, 0)$ because both infinite values play a geometric role. Subsets $A \subseteq C \times \mathbb{R}$ arise naturally in our applications, and they determine presheaves $C \to \overline{\mathbb{R}}$ by $c \mapsto \sup\{r \in \mathbb{R} \mid (c, r) \in A\}$: one needs $f(c) = -\infty$ when A contains no point over c and $f(c) = +\infty$ when it contains all of them. More immediately, in Lawvere's base $([0, \infty], \geq, +, 0)$ the monoidal unit 0 is the top element, so residuation is truncated: $z - y = \max(z - y, 0)$. In $\overline{\mathbb{R}}$ the unit 0 sits in the interior, residuation is ordinary subtraction, and the resulting translation action on (co)presheaves produces affine geometry—then projective geometry after quotienting by constant shifts.

2.2 Categories, functors, presheaves, and profunctors

An $\overline{\mathbb{R}}$-*category* $\mathcal{C}$ consists of a set $\mathrm{Ob}(\mathcal{C})$ and hom-values $\mathcal{C}(c, c') \in \overline{\mathbb{R}}$ satisfying $\mathcal{C}(c, c) = 0$ (identities) and $\mathcal{C}(c, c') + \mathcal{C}(c', c'') \leq \mathcal{C}(c, c'')$ (composition). The *opposite* $\mathcal{C}^{\mathrm{op}}$ reverses the hom-values: $\mathcal{C}^{\mathrm{op}}(c, c') = \mathcal{C}(c', c)$. An $\overline{\mathbb{R}}$-*functor* $F\colon \mathcal{C} \to \mathcal{D}$ is a map on objects satisfying $\mathcal{C}(c, c') \leq \mathcal{D}(Fc, Fc')$; the enriched hom between functors is $[\mathcal{D}, \overline{\mathbb{R}}](F, G) = \inf_{c \in \mathcal{C}} \mathcal{D}(Fc, Gc)$, making $[\mathcal{D}, \overline{\mathbb{R}}]$, the set of $\overline{\mathbb{R}}$-functors from $\mathcal{D}$ to $\overline{\mathbb{R}}$, into an $\overline{\mathbb{R}}$-category. If it is not ambiguous, we abbreviate $[\mathcal{D}, \overline{\mathbb{R}}](F, G)$ by $[F, G]$.

A *presheaf* on $\mathcal{C}$ is an $\overline{\mathbb{R}}$-functor $f\colon \mathcal{C}^{\mathrm{op}} \to \overline{\mathbb{R}}$; a *copresheaf* on $\mathcal{D}$ is an $\overline{\mathbb{R}}$-functor $g\colon \mathcal{D} \to \overline{\mathbb{R}}$. We regard copresheaves as objects of $[\mathcal{D}, \overline{\mathbb{R}}]^{\mathrm{op}}$ (same underlying functions, reversed pointwise order), so that the Isbell conjugates below become $\overline{\mathbb{R}}$-functors. The enriched homs are computed pointwise:

$$[\mathcal{C}, \overline{\mathbb{R}}](f, f') = \inf_{c \in \mathcal{C}} \left(f'(c) - f(c) \right), \qquad [\mathcal{D}, \overline{\mathbb{R}}]^{\mathrm{op}}(g, g') = \inf_{d \in \mathcal{D}} \left(g(d) - g'(d) \right).$$

In particular, the underlying order on presheaves is the pointwise order, while copresheaves in $[\mathcal{D},\overline{\mathbb{R}}]^{\mathrm{op}}$ carry the opposite of the pointwise order.

The enriched Yoneda lemma says that $[\mathcal{C}(-,c),f]=f(c)$ and $[g,\mathcal{D}(d,-)]_{[\mathcal{D},\overline{\mathbb{R}}]^{\mathrm{op}}}=g(d)$. In particular, $[\mathcal{C}(-,c),\mathcal{C}(-,c')]=\mathcal{C}(c,c')$ and $[\mathcal{D}(d,-),\mathcal{D}(d',-)]=\mathcal{D}(d',d)$ so the Yoneda map $y\colon \mathcal{C}\to[\mathcal{C}^{\mathrm{op}},\overline{\mathbb{R}}]$, $c\mapsto\mathcal{C}(-,c)$, and its copresheaf analogue $\mathcal{D}\to[\mathcal{D},\overline{\mathbb{R}}]^{\mathrm{op}}$, $d\mapsto\mathcal{D}(d,-)$, are fully faithful embeddings. Every presheaf and copresheaf is a pointwise supremum of representables:

$$f(x)=\sup_{c\in\mathcal{C}}\left(f(c)+\mathcal{C}(x,c)\right),\qquad g(x)=\sup_{d\in\mathcal{D}}\left(g(d)+\mathcal{D}(d,x)\right);$$

for the weighted-colimit formulation, see [28].

A *profunctor* $M\colon\mathcal{C}\nrightarrow\mathcal{D}$ is an $\overline{\mathbb{R}}$-functor $M\colon\mathcal{C}^{\mathrm{op}}\otimes\mathcal{D}\to\overline{\mathbb{R}}$, where $\mathcal{C}^{\mathrm{op}}\otimes\mathcal{D}$ has objects $\mathrm{Ob}(\mathcal{C})\times\mathrm{Ob}(\mathcal{D})$ and $(\mathcal{C}^{\mathrm{op}}\otimes\mathcal{D})((c,d),(c',d'))=\mathcal{C}(c',c)+\mathcal{D}(d,d')$. Any set S determines a discrete $\overline{\mathbb{R}}$-category with $S(s,s')=0$ if $s=s'$ and $-\infty$ otherwise; for discrete $\mathcal{C}$ and $\mathcal{D}$, every function $M\colon\mathcal{C}\times\mathcal{D}\to\overline{\mathbb{R}}$ defines a profunctor.

2.3 Isbell duality and the nucleus

Let $M\colon\mathcal{C}\nrightarrow\mathcal{D}$ be a profunctor. For each $d\in\mathcal{D}$ the column $M(-,d)$ is a presheaf on $\mathcal{C}$, and for each $c\in\mathcal{C}$ the row $M(c,-)$ is a copresheaf on $\mathcal{D}$. The Isbell conjugates extend these assignments to arbitrary (co)presheaves.

Definition 1. The *Isbell conjugates* of M are

$$M^*\colon[\mathcal{C}^{\mathrm{op}},\overline{\mathbb{R}}]\to[\mathcal{D},\overline{\mathbb{R}}]^{\mathrm{op}},$$
$$M_*\colon[\mathcal{D},\overline{\mathbb{R}}]^{\mathrm{op}}\to[\mathcal{C}^{\mathrm{op}},\overline{\mathbb{R}}],$$

defined by

$$(M^*f)(d):=\inf_{c\in\mathcal{C}}\left(M(c,d)-f(c)\right),\tag{2}$$
$$(M_*g)(c):=\inf_{d\in\mathcal{D}}\left(M(c,d)-g(d)\right).\tag{3}$$

Proposition 2. *The maps M^* and M_* are $\overline{\mathbb{R}}$-functors and satisfy $M^*\dashv M_*$, i.e.,*

$$[\mathcal{D},\overline{\mathbb{R}}]^{\mathrm{op}}(M^*f,g)=[\mathcal{C}^{\mathrm{op}},\overline{\mathbb{R}}](f,M_*g)$$

for all presheaves f and copresheaves g. Moreover, M^* and M_* are order-reversing for the pointwise order.

Proof. Functoriality of M^*: if $\delta=[f,f']=\inf_c(f'(c)-f(c))$ then $f(c)\le f'(c)-\delta$ for all c, and taking $\inf_c$ of $M(c,d)-f'(c)+\delta\le M(c,d)-f(c)$ yields $[f,f']\le[\mathcal{D},\overline{\mathbb{R}}]^{\mathrm{op}}(M^*f,M^*f')$; similarly for M_*.

For the adjunction, the key step is that residuation by $g(d)$ preserves infima:

$$\begin{aligned}[\mathcal{D},\overline{\mathbb{R}}]^{\mathrm{op}}(M^*f,g) &= \inf_d \Big(\inf_c \left(M(c,d)-f(c)\right)-g(d)\Big)\\ &= \inf_d \inf_c \left(M(c,d)-f(c)-g(d)\right)\\ &= \inf_c \Big(\inf_d \left(M(c,d)-g(d)\right)-f(c)\Big) = [\mathcal{C}^{\mathrm{op}},\overline{\mathbb{R}}](f,M_*g).\end{aligned}$$

Antitonicity: if $f \le f'$ then $M(c,d)-f'(c) \le M(c,d)-f(c)$ by antitonicity of residuation, so $(M^*f')(d) \le (M^*f)(d)$; the argument for M_* is the same. □

Since M^* and M_* are each antitone, the composites M_*M^* and M^*M_* are monotone. The adjunction $M^* \dashv M_*$ gives $f \le M_*M^*f$ for every presheaf f: indeed, the adjunction identity applied with $g = M^*f$ reads $[M^*f, M^*f] = [f, M_*M^*f]$, and the left side is ≥ 0, so $f \le M_*M^*f$. Dually $g \le M^*M_*g$. Monotonicity and expansion together imply idempotence: applying M_*M^* to $f \le M_*M^*f$ gives $M_*M^*f \le (M_*M^*)^2 f$ (the expanding direction), and applying the antitone M^* to the same inequality gives $M^*M_*M^*f \ge M^*f$, then applying the antitone M_* reverses again: $(M_*M^*)^2 f = M_*M^*M_*M^*f \le M_*M^*f$. So M_*M^* and M^*M_* are closure operators:

$$\begin{aligned} f &\le M_*M^*f, \quad (M_*M^*)^2 = M_*M^*,\\ g &\le M^*M_*g, \quad (M^*M_*)^2 = M^*M_*. \end{aligned} \tag{4}$$

In other words, each closure operator is a projection: it expands its input to the nearest fixed point, and once there, does nothing.

Definition 3. The *nucleus* of M is the $\overline{\mathbb{R}}$-category

$$\mathrm{Nuc}(M) = \{(f,g) \mid g = M^*f,\ f = M_*g\},$$

with hom-values $\mathrm{Nuc}(M)((f,g),(f',g')) = [f,f'] = [\mathcal{D},\overline{\mathbb{R}}]^{\mathrm{op}}(g,g')$, the equality holding by the adjunction.

Proposition 4. *The projection* $(f,g) \mapsto f$ *identifies* $\mathrm{Nuc}(M)$ *with* $\mathrm{Fix}(M_*M^*)$*, and* $(f,g) \mapsto g$ *identifies it with* $\mathrm{Fix}(M^*M_*)$*. Moreover,* $\mathrm{Fix}(M_*M^*) = \mathrm{im}(M_*M^*) = \mathrm{im}(M_*)$ *and* $\mathrm{Fix}(M^*M_*) = \mathrm{im}(M^*M_*) = \mathrm{im}(M^*)$*.*

Proof. The first identification is immediate: $(f,g) \in \mathrm{Nuc}(M)$ iff $f = M_*g = M_*M^*f$; dually for the second.

For any closure operator, fixed points and image coincide: every fixed point is in the image (it is its own closure), and every element of the image is fixed by idempotence. So $\mathrm{Fix}(M_*M^*) = \mathrm{im}(M_*M^*)$.

It remains to show that M_* already lands among the fixed points of M_*M^*, i.e. $M_*M^*M_* = M_*$. The expansion $g \le M^*M_*g$ and the antitonicity of M_* give $M_*M^*M_*g \le M_*g$. The reverse inequality is the expansion $M_*g \le M_*M^*M_*g$. Together: $M_*M^*M_* = M_*$, so M_* lands in $\mathrm{Fix}(M_*M^*)$. Dually $M^*M_*M^* = M^*$. □

Corollary 5. *A presheaf f is a fixed point of M_*M^* if and only if it is the largest presheaf with the same M^*-image: $M^*h = M^*f \Rightarrow h \le f$. Dually for copresheaves and M_*.*

Proof. If $f = M_*M^*f$ and $M^*h = M^*f$, then $h \le M_*M^*h = M_*M^*f = f$ by (4). Conversely, applying this to $h = M_*M^*f$ gives $M_*M^*f \le f$, and (4) gives the reverse inequality. □

3 Projective metric geometry

The nucleus constructed in Section 2 is an $\overline{\mathbb{R}}$-enriched category, and the enrichment determines a canonical projective metric. The enriched hom $[f, f'] = \inf_c(f'(c) - f(c))$ is a directed distance: $[f, f'] \ge 0$ iff $f \le f'$, and enriched composition gives the directed triangle inequality. Symmetrizing introduces a kernel: $-[f, f'] - [f', f] = 0$ exactly when $f - f'$ is constant, so the symmetrized distance descends to a genuine metric on translation classes. For real-valued presheaves, this metric has a simple form:

$$d_{\mathcal{C}}([f], [f']) = \max_c (f(c) - f'(c)) - \min_c (f(c) - f'(c)),$$

the oscillation of the difference—the tropical form of a Hilbert projective metric. The Isbell transforms, being equivariant for constant translation, respect this projectivization. For general $\overline{\mathbb{R}}$-valued presheaves, some care is needed at $\pm\infty$: first to handle fixed points of the translation action, and then because residuation at $\pm\infty$ does not behave like ordinary subtraction. We handle this in §3.2 below.

3.1 Finite index sets and the Isbell transforms

For the geometric constructions below it is convenient to work with finite, discrete $\overline{\mathbb{R}}$-categories. Thus, for the remainder of this section, $\mathcal{C}$ and $\mathcal{D}$ are finite sets (regarded as discrete $\overline{\mathbb{R}}$-categories), and a profunctor $M : \mathcal{C} \nrightarrow \mathcal{D}$ is simply a function

$$M : \mathcal{C} \times \mathcal{D} \to \overline{\mathbb{R}}.$$

In this setting, presheaves on $\mathcal{C}$ and copresheaves on $\mathcal{D}$ are just functions $\mathcal{C} \to \overline{\mathbb{R}}$ and $\mathcal{D} \to \overline{\mathbb{R}}$, and the infima in Definition 1 are minima. Accordingly,

$$\begin{aligned} (M^*f)(d) &= \min_{c \in \mathcal{C}} (M(c,d) - f(c)), \\ (M_*g)(c) &= \min_{d \in \mathcal{D}} (M(c,d) - g(d)), \end{aligned} \tag{5}$$

where $z - y$ denotes residuation in $\overline{\mathbb{R}}$ (cf. (1)).

Lemma 6. *For any finite constant $\lambda \in \mathbb{R}$ one has*

$$\begin{aligned} M^*(f + \lambda) &= M^*f - \lambda, \\ M_*(g - \lambda) &= M_*g + \lambda, \end{aligned} \tag{6}$$

where λ denotes the constant function on $\mathcal{C}$ or $\mathcal{D}$.

Proof. For the first identity, subtracting λ from each term inside the minimum gives

$$\begin{aligned}(M^*(f+\lambda))(d) &= \min_c (M(c,d) - f(c) - \lambda) \\ &= \Big(\min_c (M(c,d) - f(c))\Big) - \lambda \\ &= (M^*f)(d) - \lambda.\end{aligned}$$

The second identity is analogous. □

The equivariance (6) is the algebraic shadow of a projective symmetry: if (f,g) satisfies $g = M^*f$ and $f = M_*g$, then so does $(f+\lambda, g-\lambda)$.

3.2 Projective (co)presheaves and the Hilbert–oscillation metric

We now isolate the locus on which translation by constants acts freely.

Definition 7. Define the *finite somewhere* presheaves $[\mathcal{C}^{\mathrm{op}}, \overline{\mathbb{R}}]_{\mathrm{fs}}$ to be the full subcategory of presheaves $f\colon \mathcal{C} \to \overline{\mathbb{R}}$ for which $f(c) \in \mathbb{R}$ for at least one $c \in \mathcal{C}$. Define $[\mathcal{D}, \overline{\mathbb{R}}]^{\mathrm{op}}_{\mathrm{fs}}$ similarly for copresheaves on $\mathcal{D}$.

On these full subcategories, $(\mathbb{R}, +)$ acts freely by constant translation $f \mapsto f + \lambda$.

Definition 8. The *projective presheaf space* of $\mathcal{C}$ is the quotient

$$\mathbb{P}\mathcal{C} := [\mathcal{C}^{\mathrm{op}}, \overline{\mathbb{R}}]_{\mathrm{fs}}/\mathbb{R}.$$

Let $[f] \in \mathbb{P}\mathcal{C}$ denote the translation class of f. Similarly, the *projective copresheaf space* of $\mathcal{D}$ is

$$\mathbb{P}\mathcal{D} := [\mathcal{D}, \overline{\mathbb{R}}]^{\mathrm{op}}_{\mathrm{fs}}/\mathbb{R}.$$

For $\lambda \in \mathbb{R}$ and presheaves f, f', the enriched hom satisfies $[f+\lambda, f'] = [f, f'] - \lambda$ and $[f, f'+\lambda] = [f, f'] + \lambda$. In particular, the symmetrized quantity $-[f,f'] - [f', f]$ is $\mathbb{R}$-invariant. This is the tropical analogue of Hilbert's projective metric: it measures only the oscillation of the difference.

When f and f' are real-valued (i.e. $f, f'\colon \mathcal{C} \to \mathbb{R}$), the projective distance is simply

$$d_{\mathcal{C}}([f],[f']) = \max_c (f(c) - f'(c)) - \min_c (f(c) - f'(c)),$$

the oscillation of the difference. For general $\overline{\mathbb{R}}$-valued presheaves, some care is needed. The residuation $f(c) - f'(c)$ is not antisymmetric when both values are $+\infty$ or both are $-\infty$: one has $\infty - \infty = \infty$ and $-\infty - (-\infty) = \infty$, so subtracting $-\infty$ is not the same as adding $+\infty$. We isolate the offending indices.

Definition 9. Let $f, f' \in [\mathcal{C}^{\mathrm{op}}, \overline{\mathbb{R}}]_{\mathrm{fs}}$. Set

$$S(f,f') := \{c \in \mathcal{C} \mid f(c) = f'(c) = +\infty \text{ or } f(c) = f'(c) = -\infty\}.$$

These are the indices at which residuation fails to be antisymmetric. Define the *projective distance* between $[f], [f'] \in \mathbb{P}\mathcal{C}$ by

$$d_{\mathcal{C}}([f],[f']) := \begin{cases} \begin{aligned} &\sup_{c \notin S(f,f')} (f(c) - f'(c)) \\ &\quad - \inf_{c \notin S(f,f')} (f(c) - f'(c)) \end{aligned} & \text{if both extrema lie in } \mathbb{R}, \\ +\infty & \text{otherwise.} \end{cases} \tag{7}$$

Define $d_{\mathcal{D}}$ on $\mathbb{P}\mathcal{D}$ analogously.

Proposition 10. *The function $d_{\mathcal{C}}$ is an extended metric on $\mathbb{P}\mathcal{C}$. Moreover, whenever $d_{\mathcal{C}}([f],[f']) < \infty$ one has the identity*

$$d_{\mathcal{C}}([f],[f']) = -[f,f'] - [f',f], \tag{8}$$

where $[f,f']$ is the enriched hom in $[\mathcal{C}^{\mathrm{op}}, \overline{\mathbb{R}}]$.

Proof. Well-definedness on $\mathbb{P}\mathcal{C}$ is immediate: translating f by a finite constant λ shifts both $\sup(f - f')$ and $\inf(f - f')$ by λ, so their difference is unchanged, and $S(f,f')$ is unaffected since $\pm\infty + \lambda = \pm\infty$. Symmetry and nonnegativity are clear. If $d_{\mathcal{C}}([f],[f']) = 0$ then $f(c) - f'(c)$ is a single real constant on $\mathcal{C} \setminus S(f,f')$ and $f(c) = f'(c) \in \{\pm\infty\}$ on $S(f,f')$, so $[f] = [f']$.

Now assume $d_{\mathcal{C}}([f],[f']) < \infty$. Since $+\infty$ is absorbing for residuation, every $c \in S(f,f')$ contributes $f'(c) - f(c) = +\infty$ to the enriched hom $[f,f']$ and $f(c) - f'(c) = +\infty$ to $[f',f]$. These values cannot achieve either infimum (which is finite by hypothesis), so they drop out and

$$-[f,f'] - [f',f] = \sup_c (f(c) - f'(c)) - \inf_c (f(c) - f'(c)) = d_{\mathcal{C}}([f],[f']),$$

which is (8).

For the triangle inequality: if either summand on the right is $+\infty$ there is nothing to prove. Otherwise all three distances are finite and we use enriched composition $[f,f'] + [f',f''] \le [f,f'']$ and $[f'',f'] + [f',f] \le [f'',f]$. Negating and adding gives $d_{\mathcal{C}}([f],[f'']) \le d_{\mathcal{C}}([f],[f']) + d_{\mathcal{C}}([f'],[f''])$. □

3.3 Projective nuclei and isometries

We now impose the mild hypothesis needed to pass M^* and M_* to projective spaces.

Definition 11. We call the profunctor M *nondegenerate* if the Isbell transforms preserve the finite-somewhere condition:

$$M^*\left([\mathcal{C}^{\mathrm{op}}, \overline{\mathbb{R}}]_{\mathrm{fs}}\right) \subseteq [\mathcal{D}, \overline{\mathbb{R}}]_{\mathrm{fs}}^{\mathrm{op}},$$
$$M_*\left([\mathcal{D}, \overline{\mathbb{R}}]_{\mathrm{fs}}^{\mathrm{op}}\right) \subseteq [\mathcal{C}^{\mathrm{op}}, \overline{\mathbb{R}}]_{\mathrm{fs}}.$$

Under this hypothesis, Lemma 6 implies that M^* and M_* descend to well-defined maps

$$M^* \colon \mathbb{P}\mathcal{C} \to \mathbb{P}\mathcal{D},$$
$$M_* \colon \mathbb{P}\mathcal{D} \to \mathbb{P}\mathcal{C}.$$

Definition 12. Let $\mathrm{Nuc}(M)$ be the nucleus of M (Definition 3) and set

$$\mathrm{Nuc}(M)_{\mathrm{fs}} := \mathrm{Nuc}(M) \cap \left([\mathcal{C}^{\mathrm{op}}, \overline{\mathbb{R}}]_{\mathrm{fs}} \times [\mathcal{D}, \overline{\mathbb{R}}]_{\mathrm{fs}}^{\mathrm{op}}\right).$$

The group $\mathbb{R}$ acts on $\mathrm{Nuc}(M)_{\mathrm{fs}}$ by

$$\lambda \cdot (f, g) = (f + \lambda, g - \lambda).$$

The *projective nucleus* is the quotient

$$\mathbb{P}\mathrm{Nuc}(M) := \mathrm{Nuc}(M)_{\mathrm{fs}}/\mathbb{R}.$$

We metrize $\mathbb{P}\mathrm{Nuc}(M)$ by

$$d_{\mathbb{P}\mathrm{Nuc}}\left([(f,g)], [(f',g')]\right) := \max\left\{d_{\mathcal{C}}([f],[f']), d_{\mathcal{D}}([g],[g'])\right\}.$$

Write $\mathrm{Fix}_{Proj}(M_*M^*) \subseteq \mathbb{P}\mathcal{C}$ and $\mathrm{Fix}_{Proj}(M^*M_*) \subseteq \mathbb{P}\mathcal{D}$ for the images in projective space of the fixed-point sets of the closure operators M_*M^* and M^*M_*.

Theorem 13 (The Isometry Theorem)**.** *Let $M \colon \mathcal{C} \nrightarrow \mathcal{D}$ be a nondegenerate profunctor. Then the maps $M^* \colon \mathbb{P}\mathcal{C} \to \mathbb{P}\mathcal{D}$ and $M_* \colon \mathbb{P}\mathcal{D} \to \mathbb{P}\mathcal{C}$ are 1-Lipschitz for the metrics $d_{\mathcal{C}}$ and $d_{\mathcal{D}}$. Moreover, they restrict to mutually inverse isometries*

$$M^* \colon\ \mathrm{Fix}_{Proj}(M_*M^*) \xrightarrow{\cong} \mathrm{Fix}_{Proj}(M^*M_*),$$
$$M_* \colon\ \mathrm{Fix}_{Proj}(M^*M_*) \xrightarrow{\cong} \mathrm{Fix}_{Proj}(M_*M^*).$$

and hence identify $\mathbb{P}\mathrm{Nuc}(M)$ isometrically with either projective fixed-point set.

Proof. Functoriality of M^* in the enriched sense gives, for presheaves f, f',

$$[f, f'] \le [\mathcal{D}, \overline{\mathbb{R}}]^{\mathrm{op}}(M^*f, M^*f') = [\mathcal{D}, \overline{\mathbb{R}}](M^*f', M^*f),$$

and the same inequality with f and f' exchanged. Negating and adding yields

$$-[M^*f, M^*f'] - [M^*f', M^*f] \le -[f, f'] - [f', f].$$

If $d_{\mathcal{C}}([f],[f']) < \infty$, Proposition 10 identifies both sides with the corresponding projective metrics, giving

$$d_{\mathcal{D}}\left(M^*[f], M^*[f']\right) \le d_{\mathcal{C}}\left([f],[f']\right).$$

If $d_{\mathcal{C}}([f],[f']) = +\infty$, the inequality is tautological. Thus M^* is 1-Lipschitz, and similarly M_*.

On the projective fixed-point sets, M^* and M_* are inverse bijections (Proposition 4). Since each is 1-Lipschitz, the two inequalities

$$d_{\mathcal{D}}\left(M^*[f], M^*[f']\right) \leq d_{\mathcal{C}}\left([f],[f']\right),$$
$$d_{\mathcal{C}}\left(M_*[g], M_*[g']\right) \leq d_{\mathcal{D}}\left([g],[g']\right).$$

apply to inverse pairs and force equality. Hence both restrictions are isometries.

Finally, the identification with $\mathbb{P}\mathrm{Nuc}(M)$ is obtained by projecting $[(f,g)] \mapsto [f]$ or $[(f,g)] \mapsto [g]$. □

Consequently we obtain a diagram of metric spaces in which every arrow is an isometry:

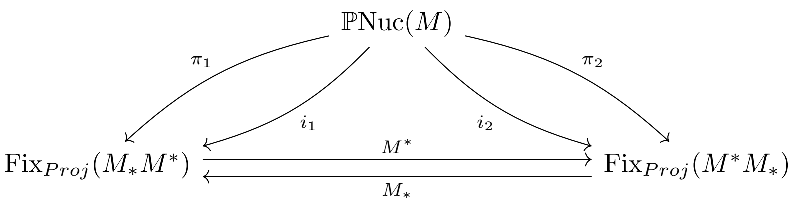

where $\pi_1([(f,g)]) = [f]$, $\pi_2([(f,g)]) = [g]$, $i_1([f]) = [(f, M^*f)]$, and $i_2([g]) = [(M_*g, g)]$.

3.4 External gauge transformations

Beyond translation by constants there is a larger symmetry: the additive group $\mathbb{R}^{\mathcal{C}} \times \mathbb{R}^{\mathcal{D}}$ acts on matrices by adding constants to rows and columns, the diagonal equivalence of tropical linear algebra. This action changes M but preserves $\mathbb{P}\mathrm{Nuc}(M)$ up to canonical isometry, so the projective metric and cell decomposition depend only on the diagonal equivalence class of M.

Definition 14. For $u \in \mathbb{R}^{\mathcal{C}}$ define $L_u \colon \mathbb{P}\mathcal{C} \to \mathbb{P}\mathcal{C}$ by $[f] \mapsto [f-u]$. For $v \in \mathbb{R}^{\mathcal{D}}$ define $R_v \colon \mathbb{P}\mathcal{D} \to \mathbb{P}\mathcal{D}$ by $[g] \mapsto [g-v]$. Given (u,v), define the *gauge transform* of M by

$$M^{(u,v)}(c,d) := M(c,d) - u(c) - v(d).$$

Lemma 15. *On projective spaces, the Isbell transforms of $M^{(u,v)}$ are conjugate to those of M:*

$$\left(M^{(u,v)}\right)^* = R_v \circ M^* \circ L_u^{-1},$$
$$\left(M^{(u,v)}\right)_* = L_u \circ M_* \circ R_v^{-1}.$$

Proof. For a representative f and any $d \in \mathcal{D}$,

$$\left(M^{(u,v)}\right)^* (f-u)(d) = \min_c \left(M(c,d) - u(c) - v(d) - (f(c) - u(c))\right)$$
$$= \min_c \left(M(c,d) - f(c)\right) - v(d)$$
$$= M^* f(d) - v(d).$$

This is precisely $(R_v \circ M^*)(f)(d)$. The statement for $\left(M^{(u,v)}\right)_*$ is analogous. □

Proposition 16. *For any $(u,v) \in \mathbb{R}^{\mathcal{C}} \times \mathbb{R}^{\mathcal{D}}$, the map*

$$\Phi_{u,v} \colon \mathbb{P}\mathrm{Nuc}(M) \to \mathbb{P}\mathrm{Nuc}\left(M^{(u,v)}\right),$$
$$[(f,g)] \mapsto [(f-u, g-v)].$$

is a well-defined isometry with inverse $[(f',g')] \mapsto [(f'+u, g'+v)]$. Consequently the projective fixed-point sets and projective nuclei of M and $M^{(u,v)}$ are canonically isometric.

Proof. The maps L_u and R_v are isometries: subtracting the same potential from both arguments leaves the difference, hence its oscillation, unchanged. Lemma 15 identifies the Isbell equations for M with those for $M^{(u,v)}$ under these isometries. Thus $(f,g) \in \mathrm{Nuc}(M)$ if and only if $(f-u, g-v) \in \mathrm{Nuc}\left(M^{(u,v)}\right)$, and the induced map on projective quotients is an isometry. □

3.5 Witness cells and the gap matrix

The projective metric of $\mathbb{P}\mathrm{Nuc}(M)$ is intrinsic and gauge-invariant (Proposition 16). When the indexing sets are finite, the Isbell inequalities also endow $\mathbb{P}\mathrm{Nuc}(M)$ with a polyhedral stratification. The bridge between metric and polyhedral structure is the *gap matrix* $\delta^{(f,g)}$: its zero entries record the witness relation and hence the combinatorial cell, while its positive entries measure slack.

We continue with $\mathcal{C}$ and $\mathcal{D}$ finite as in §3.1, though the definitions below extend verbatim to general small $\overline{\mathbb{R}}$-categories by replacing minima with infima. Fix a non-degenerate profunctor M. For a presheaf f, write $g := M^* f$ as in (5), and dually $f := M_* g$.

Definition 17. Let $f \colon \mathcal{C} \to \overline{\mathbb{R}}$ and let $g := M^* f$. An element $c \in \mathcal{C}$ is a *witness for f at $d \in \mathcal{D}$* if c realizes the minimum defining $M^* f(d)$:

$$g(d) = M(c,d) - f(c).$$

Dually, if $g \colon \mathcal{D} \to \overline{\mathbb{R}}$ and $f := M_* g$, then $d \in \mathcal{D}$ is a *witness for g at $c \in \mathcal{C}$* if d realizes the minimum defining $M_* g(c)$.

Because $M^*(f+\lambda) = M^*f - \lambda$ and $M_*(g-\lambda) = M_*g + \lambda$ for every $\lambda \in \mathbb{R}$, the witness relation depends only on the projective classes $[f] \in \mathbb{P}\mathcal{C}$ and $[g] \in \mathbb{P}\mathcal{D}$. By definition of $g = M^*f$, the inequalities

$$f(c) + g(d) \leq M(c,d), \quad c \in \mathcal{C},\ d \in \mathcal{D}$$

always hold. The gap matrix measures how far they are from equality.

Definition 18. Let $[f] \in \mathbb{P}\mathcal{C}$ and choose a representative $f\colon \mathcal{C} \to \overline{\mathbb{R}}$. Set $g := M^*f$. The *gap matrix* of $[f]$ is the function

$$\delta^f(c,d) := M(c,d) - (f(c) + g(d))\,. \tag{9}$$

Dually, for $[g] \in \mathbb{P}\mathcal{D}$ with representative g and $f := M_*g$, we set $\delta^g(c,d) := M(c,d) - (f(c) + g(d))$. If $(f,g) \in \mathrm{Nuc}(M)$, then $\delta^f = \delta^g$, and we write $\delta^{(f,g)}$.

Lemma 19. *(a) For every $\lambda \in \mathbb{R}$, one has $\delta^{f+\lambda} = \delta^f$.*

(b) For $u \in \mathbb{R}^{\mathcal{C}}$ and $v \in \mathbb{R}^{\mathcal{D}}$, let $M^{(u,v)}(c,d) := M(c,d) - u(c) - v(d)$ be the gauge transform. If $g = M^*f$, then $\left(M^{(u,v)}\right)^*(f-u) = g - v$, and the corresponding gap matrices agree:

$$\delta^f = \delta^{f-u},$$

where δ^f is computed with M and δ^{f-u} with $M^{(u,v)}$.

Proof. (a) follows from $M^*(f+\lambda) = M^*f - \lambda$ and $(f(c)+\lambda) + (g(d)-\lambda) = f(c) + g(d)$. For (b), compute

$$M^{(u,v)}(c,d) - ((f-u)(c) + (g-v)(d)) = M(c,d) - (f(c) + g(d))\,. \qquad \square$$

On the finite locus, the zeros of δ are exactly the witness pairs. The next lemma isolates the only subtlety: in $\overline{\mathbb{R}}$, a zero gap forces finiteness.

Lemma 20. *If $\delta^f(c,d) = 0$, then $f(c)$, $g(d)$, and $M(c,d)$ are all finite real numbers.*

Proof. By definition, $\delta^f(c,d) = M(c,d) - (f(c) + g(d))$ is a residuation. If $f(c) + g(d) = +\infty$ then $\delta^f(c,d) = M(c,d) - \infty \in \{-\infty, \infty\}$, never 0. If $f(c) + g(d) = -\infty$ then $\delta^f(c,d) = M(c,d) - (-\infty) = \infty$, again not 0. Thus $f(c) + g(d) \in \mathbb{R}$. Since $-\infty$ is absorbing for addition, a finite sum forces $f(c), g(d) \in \mathbb{R}$. Finally, $M(c,d) \in \mathbb{R}$ as well: if $M(c,d) = \pm\infty$ then $M(c,d) - (f(c) + g(d)) = \pm\infty$. $\square$

Proposition 21. *Let $[f] \in \mathbb{P}\mathcal{C}$ with representative f, let $g := M^*f$, and let $\delta = \delta^f$. Then:*

(a) $\delta(c,d) \geq 0$ for all $(c,d) \in \mathcal{C} \times \mathcal{D}$.

(b) If $\delta(c,d) = 0$, then c is a witness for f at d.

(c) If f and g are finite-valued, then $\delta(c,d) = 0$ if and only if c is a witness for f at d. In particular, every column contains at least one zero.

(d) If moreover $(f,g) \in \mathrm{Nuc}(M)$ and f,g are finite-valued, then $\delta(c,d) = 0$ if and only if d is a witness for g at c. In particular, every row contains at least one zero.

Proof. (a) Since $g(d) = \min_{c'} (M(c', d) - f(c'))$, we have $g(d) \le M(c, d) - f(c)$ for every c. By residuation this is equivalent to $f(c) + g(d) \le M(c, d)$, hence $\delta(c, d) \ge 0$.

(b) If $\delta(c, d) = 0$, then Lemma 20 shows the relevant entries are finite, so subtraction is ordinary: $0 = M(c, d) - f(c) - g(d)$, hence $g(d) = M(c, d) - f(c)$ and c realizes the minimum in $M^* f(d)$.

(c) If f, g are finite-valued and c is a witness for f at d, then $g(d) = M(c, d) - f(c)$ and $\delta(c, d) = 0$. The converse is (b). Since $\mathcal{C}$ is finite, every minimum defining $g(d)$ is attained, so every column contains a witness and hence a zero.

(d) Apply (c) to the dual description $f = M_* g$. □

Corollary 22. *If $(f, g) \in \mathbb{P}\mathrm{Nuc}(M)$ and f, g are finite-valued, then c is a witness for f at d if and only if d is a witness for g at c.*

Proof. This is Proposition 21(c) and (d). □

Corollary 23. *A finite-valued presheaf f is a fixed point of $M_* M^*$ if and only if every row of δ^f contains a zero. Equivalently, for every $c \in \mathcal{C}$ there exists $d \in \mathcal{D}$ such that d is a witness for $M^* f$ at c.*

Proof. If f is finite-valued and fixed by $M_* M^*$, write $g := M^* f$, so $f = M_* g$. Since $\mathcal{D}$ is finite, for each c the minimum defining $f(c)$ is attained at some d, and then $\delta^f(c, d) = 0$. Conversely, if every row contains a zero, fix c and choose d with $\delta^f(c, d) = 0$. Then $f(c) = M(c, d) - g(d)$, hence

$$(M_* g)(c) = \min_{d'} (M(c, d') - g(d')) \le M(c, d) - g(d) = f(c).$$

On the other hand $f \le M_* M^* f = M_* g$ by (4), so equality holds coordinatewise. □

For a finite-valued nucleus point, the witness pattern is exactly the zero set of the gap matrix.

Definition 24. For $(f, g) \in \mathbb{P}\mathrm{Nuc}(M)$ with f, g finite-valued, define the *witness relation*

$$Z(f, g) := \left\{ (c, d) \in \mathcal{C} \times \mathcal{D} \mid \delta^{(f,g)}(c, d) = 0 \right\}. \tag{10}$$

By Proposition 21(d), the relation $Z(f, g)$ meets every row and every column. These relations partition the finite part of $\mathbb{P}\mathrm{Nuc}(M)$: we declare $(f, g) \sim (f', g')$ if $Z(f, g) = Z(f', g')$.

Definition 25. For a relation $Z \subseteq \mathcal{C} \times \mathcal{D}$ meeting every row and column, define the *open witness cell*

$$\mathrm{Cell}^\circ(Z) := \{(f, g) \in \mathbb{P}\mathrm{Nuc}(M) \mid f, g \text{ finite-valued and } Z(f, g) = Z\}.$$

Remark 26. After choosing an affine chart for the projective quotient (for example $\min f = 0$), the closure of $\mathrm{Cell}^\circ(Z)$ is a classical polytope, cut out by the equalities $M(c, d) = f(c) + g(d)$ for $(c, d) \in Z$ together with the inequalities $M(c, d) \ge f(c) + g(d)$ for all (c, d). The open cell $\mathrm{Cell}^\circ(Z)$ is its relative interior. We develop this in Section 4.

Because $\mathcal{D}$ is discrete, the copresheaf represented by $d \in \mathcal{D}$ is the delta function $g_d(d) = 0$ and $g_d(d') = -\infty$ for $d' \neq d$. A direct computation gives $M_* g_d = M(-, d)$: the columns of the matrix M are the images of the representables.

Definition 27. For $d \in \mathcal{D}$ define the dth *anchor presheaf* $A_d : \mathcal{C} \to \overline{\mathbb{R}}$ by $A_d = M(-, d)$.

Proposition 28. *Each anchor* A_d *lies in* $\mathrm{Fix}(M_* M^*) = \mathrm{Im}(M_*)$*. Moreover, for any* $f \in \mathrm{Fix}(M_* M^*)$ *there exist weights* $\lambda_d \in \overline{\mathbb{R}}$ *such that*

$$f(c) = \min_{d \in \mathcal{D}} \left(A_d(c) - \lambda_d \right), \quad c \in \mathcal{C}. \tag{11}$$

Proof. Since $A_d = M_* g_d$, we have $A_d \in \mathrm{Im}(M_*) = \mathrm{Fix}(M_* M^*)$. Conversely, if $f \in \mathrm{Fix}(M_* M^*) = \mathrm{Im}(M_*)$ then $f = M_* g$ for some copresheaf g, hence

$$f(c) = \min_{d \in \mathcal{D}} \left(M(c, d) - g(d) \right) = \min_{d \in \mathcal{D}} \left(A_d(c) - \lambda_d \right)$$

with $\lambda_d := g(d)$. □

Equivalently, $\mathrm{Fix}(M_* M^*)$ is the closure of the anchors A_d under weighted coproducts: every fixed point f can be written as the pointwise minimum of translates $A_d - \lambda_d$, and conversely every such minimum lies in $\mathrm{Fix}(M_* M^*)$.

Proposition 29. *If* $M : \mathcal{C} \times \mathcal{D} \to \mathbb{R}$ *is a real-valued matrix on finite sets, then* $\mathbb{P}\mathrm{Nuc}(M)$ *is compact.*

Proof. Since M is real-valued, each anchor $A_d = M(-, d)$ is real-valued, and Proposition 28 then implies that every fixed point of $M_* M^*$ admits a real-valued representative. On the affine slice $\{f \mid f(c_0) = 0\}$, the Isbell inequalities $f(c) + g(d) \leq M(c, d)$ bound f and g coordinatewise, so the nucleus embeds as a closed bounded subset of $\mathbb{R}^{|\mathcal{C}|-1}$. □

4 Polyhedral geometry, event loci, and lattice towers

In Section 3.5 we associated to each finite-valued nucleus point (f, g) its *gap matrix*

$$\delta^{(f,g)}(c, d) := M(c, d) - f(c) - g(d)$$

and its zero set

$$Z(f, g) := \left\{ (c, d) \in \mathcal{C} \times \mathcal{D} \mid \delta^{(f,g)}(c, d) = 0 \right\}.$$

This section develops the polyhedral side of that data. We pass from the witness relation to explicit witness polyhedra, prove that positive gap values are exact distances to event loci, refine the witness decomposition by order chambers, and show that pointed thresholding of the gap matrix yields a chamberwise constructible sheaf of concept-lattice towers.

4.1 Witness polyhedra and admissibility

We continue with $\mathcal{C}$ and $\mathcal{D}$ finite and now assume that $M\colon \mathcal{C} \times \mathcal{D} \to \mathbb{R}$ has finite real entries. By Proposition 29, $\mathbb{P}\mathrm{Nuc}(M)$ is compact and each projective class admits a real-valued representative.

To describe the polyhedral structure concretely, we fix an affine chart. Choose $c_0 \in \mathcal{C}$ and set

$$\mathrm{Nuc}(M)_0 := \{(f, g) \in \mathrm{Nuc}(M) \mid f(c_0) = 0\}.$$

Since each class in $\mathbb{P}\mathrm{Nuc}(M)$ has a unique representative in $\mathrm{Nuc}(M)_0$, this identifies $\mathrm{Nuc}(M)_0 \cong \mathbb{P}\mathrm{Nuc}(M)$ as sets; we use this chart when explicit coordinates are needed, but the metric and cell structures are independent of the choice of c_0.

Working in this chart, the Isbell inequalities

$$f(c) + g(d) \le M(c, d) \text{ for all } c \in \mathcal{C}, d \in \mathcal{D}$$

cut out a feasibility polyhedron in $\mathbb{R}^{\mathcal{C}} \times \mathbb{R}^{\mathcal{D}}$. The Isbell fixed-point conditions for the nucleus are equivalent to the requirement that every row and every column attains equality. Given $Y \subseteq \mathcal{C} \times \mathcal{D}$, we obtain a smaller closed polyhedron by forcing equality along Y.

Definition 30. Let $Y \subseteq \mathcal{C} \times \mathcal{D}$. We say that Y *covers* $\mathcal{C}$ if for every $c \in \mathcal{C}$ there exists $d \in \mathcal{D}$ with $(c, d) \in Y$. We say that Y *covers* $\mathcal{D}$ if for every $d \in \mathcal{D}$ there exists $c \in \mathcal{C}$ with $(c, d) \in Y$.

Definition 31. For $Y \subseteq \mathcal{C} \times \mathcal{D}$, define $\mathrm{Cell}(Y)$ to be the set of pairs of real-valued functions $(f, g) \in \mathbb{R}^{\mathcal{C}} \times \mathbb{R}^{\mathcal{D}}$ satisfying the gauge condition $f(c_0) = 0$ and the constraints

$$f(c) + g(d) \le M(c, d) \text{ for all } (c, d) \in \mathcal{C} \times \mathcal{D}, \tag{12}$$

$$f(c) + g(d) = M(c, d) \text{ for all } (c, d) \in Y. \tag{13}$$

Thus $\mathrm{Cell}(Y)$ is a (possibly empty) closed polyhedron. For general Y, feasibility of (12)–(13) does not imply that (f, g) is a nucleus point. The next lemma isolates the combinatorial condition that forces the Isbell equalities in every row or column.

Lemma 32. *Let $Y \subseteq \mathcal{C} \times \mathcal{D}$ and let (f, g) satisfy (12)–(13).*

(a) If Y covers $\mathcal{D}$, then $g = M^* f$.

(b) If Y covers $\mathcal{C}$, then $f = M_* g$.

Consequently, if Y covers both $\mathcal{C}$ and $\mathcal{D}$, then every $(f, g) \in \mathrm{Cell}(Y)$ lies in $\mathrm{Nuc}(M)_0$.

Proof. (a) Fix $d \in \mathcal{D}$. From (12) we obtain $g(d) \le M(c, d) - f(c)$ for all c, hence

$$g(d) \le \min_{c \in \mathcal{C}} (M(c, d) - f(c)) = (M^* f)(d).$$

Since Y covers $\mathcal{D}$, there exists c with $(c, d) \in Y$. Then (13) gives $g(d) = M(c, d) - f(c)$, hence $g(d) \ge (M^* f)(d)$. Thus $g(d) = (M^* f)(d)$.

(b) The proof is symmetric. If Y covers both sides, (a) and (b) give $g = M^* f$ and $f = M_* g$, so $(f, g) \in \mathrm{Nuc}(M)$, and the gauge condition places it in $\mathrm{Nuc}(M)_0$. □

Covering both $\mathcal{C}$ and $\mathcal{D}$ is necessary for $\mathrm{Cell}(Y)$ to lie in the nucleus (Lemma 32), but for most covering sets the system (12)–(13) is infeasible.

Definition 33. A subset $Y \subseteq \mathcal{C} \times \mathcal{D}$ is *admissible* if it covers both $\mathcal{C}$ and $\mathcal{D}$ and $\mathrm{Cell}(Y) \neq \varnothing$.

When Y is admissible, $\mathrm{Cell}(Y)$ is a nonempty polyhedron consisting entirely of nucleus points by Lemma 32.

Lemma 34. *Let* $(f, g) \in \mathrm{Nuc}(M)_0$*, and set*

$$Z(f, g) = \{(c, d) \in \mathcal{C} \times \mathcal{D} \mid f(c) + g(d) = M(c, d)\}.$$

Then $(f, g) \in \mathrm{Cell}(Z(f, g))$*. In particular,* $Z(f, g)$ *is admissible.*

Proof. The inequalities (12) are exactly the Isbell inequalities for (f, g). The definition of $Z(f, g)$ is the equality condition (13). Since $\mathcal{C}$ and $\mathcal{D}$ are finite and (f, g) is a nucleus point, each row and each column attains equality, so $Z(f, g)$ covers both $\mathcal{C}$ and $\mathcal{D}$. □

Corollary 35. *The set* $\mathrm{Nuc}(M)_0 \cong \mathbb{P}\mathrm{Nuc}(M)$ *is a union of finitely many polytopes:*

$$\mathrm{Nuc}(M)_0 = \bigcup_{Y \text{ admissible}} \mathrm{Cell}(Y).$$

Proof. By Lemma 34, each $(f, g) \in \mathrm{Nuc}(M)_0$ lies in $\mathrm{Cell}(Z(f, g))$, and $Z(f, g)$ is admissible. Since $\mathcal{C} \times \mathcal{D}$ is finite, there are only finitely many admissible subsets Y.

Each $\mathrm{Cell}(Y)$ is a closed convex polyhedron by construction. Since M is real-valued, Proposition 29 implies that $\mathrm{Nuc}(M)_0$ is compact, hence each nonempty $\mathrm{Cell}(Y) \subseteq \mathrm{Nuc}(M)_0$ is bounded. Thus every admissible $\mathrm{Cell}(Y)$ is a polytope. □

Lemma 36. *For any* $Y, Y' \subseteq \mathcal{C} \times \mathcal{D}$ *one has*

$$\mathrm{Cell}(Y) \cap \mathrm{Cell}(Y') = \mathrm{Cell}(Y \cup Y').$$

Proof. The inequalities (12) are common to all $\mathrm{Cell}(\cdot)$, and the equalities (13) imposed by Y and by Y' together are exactly those imposed by $Y \cup Y'$. □

Corollary 37. *The admissible polytopes* $\{\mathrm{Cell}(Y)\}$ *form a finite polytopal complex inside* $\mathrm{Nuc}(M)_0 \cong \mathbb{P}\mathrm{Nuc}(M)$*: if* Y *and* Y' *are admissible, then* $\mathrm{Cell}(Y) \cap \mathrm{Cell}(Y')$ *is either empty or a common face of both.*

Proof. By Lemma 36, the intersection is $\mathrm{Cell}(Y \cup Y')$. If it is nonempty, then $Y \cup Y'$ covers both $\mathcal{C}$ and $\mathcal{D}$, hence is admissible, and $\mathrm{Cell}(Y \cup Y')$ is obtained from $\mathrm{Cell}(Y)$ and $\mathrm{Cell}(Y')$ by imposing additional linear equalities. Therefore it is a face of each. □

4.2 From witness polyhedra to witness cells

The witness polyhedra $\mathrm{Cell}(Y)$ are best regarded as closures of the open witness cells from §3.5. For $(f, g) \in \mathrm{Nuc}(M)_0$, write $\delta = \delta^{(f,g)}$ and $Z = Z(f, g) = \delta^{-1}(0)$. Then

Cell(Z) is the smallest witness polyhedron containing (f,g), and its relative interior consists of those points for which no additional inequalities become equalities:

$$\mathrm{Cell}^\circ(Z) = \left\{(f',g') \in \mathrm{Cell}(Z) \mid \delta^{(f',g')}(c,d) > 0 \text{ for all } (c,d) \notin Z\right\}.$$

In particular, for every admissible Z the witness cell $\mathrm{Cell}^\circ(Z)$ is the relative interior of the polytope $\mathrm{Cell}(Z)$.

In the classical min-plus setting this recovers the type decomposition of a tropical polytope, for instance via tropical hyperplane arrangements [1]. The gap matrix controls not only the combinatorics but also the metric geometry of this decomposition: each positive entry $\delta^{(f,g)}(c,d)$ is the exact projective distance from (f,g) to the locus where (c,d) becomes a witness pair. This is the Events Theorem, proved in §4.3 below.

4.3 The events theorem

We now show that each positive entry of the gap matrix is the exact projective distance to the locus where the corresponding inequality becomes tight. Related formulas for the distance to a half-space in Hilbert's projective metric appear in idempotent semimodule theory [8]; the result here is stronger because nucleus points are constrained by the Isbell equations, which couple all rows and columns simultaneously.

Fix a point $(f,g) \in \mathbb{P}\mathrm{Nuc}(M)$ and let $\delta = \delta^{(f,g)}$ be its gap matrix.

Definition 38. For each pair $(c,d) \in \mathcal{C} \times \mathcal{D}$, we define the corresponding *event locus*

$$\mathcal{E}_{c,d} := \left\{(f',g') \in \mathbb{P}\mathrm{Nuc}(M) \mid \delta^{(f',g')}(c,d) = 0\right\}.$$

Thus $\mathcal{E}_{c,d}$ is the locus in $\mathbb{P}\mathrm{Nuc}(M)$ where (c,d) is a witness pair.

Lemma 39. *Let $(f,g) \in \mathbb{P}\mathrm{Nuc}(M)$ and let $\lambda = \delta^{(f,g)}(c_i,d_j) > 0$. Then there exists $(f',g') \in \mathcal{E}_{c_i,d_j}$ such that*

$$d_{\mathbb{P}\mathrm{Nuc}}\left((f,g),(f',g')\right) = \lambda.$$

Proof. Work in the gauge slice $\mathrm{Nuc}(M)_0$ and choose representatives with $f(c_0) = 0$. Define $f'' : \mathcal{C} \to \mathbb{R}$ by

$$f''(c) = \begin{cases} f(c) & c \neq c_i, \\ f(c_i) + \lambda & c = c_i. \end{cases}$$

Set $g' := M^* f''$ and $f' := M_* g' = M_* M^* f''$. By the identities $M^* M_* M^* = M^*$ and $M_* M^* M_* = M_*$, we have $g' = M^* f'$ and $f' = M_* g'$, hence $(f',g') \in \mathrm{Nuc}(M)_0$.

For each $d \in \mathcal{D}$ one has

$$\begin{aligned} g'(d) &= \min\left(\min_{c \neq c_i} \left(M(c,d) - f(c)\right),\, M(c_i,d) - f(c_i) - \lambda\right) \\ &= \min\left(g(d),\, g(d) + \delta^{(f,g)}(c_i,d) - \lambda\right) \\ &= g(d) - \max\left(\lambda - \delta^{(f,g)}(c_i,d), 0\right). \end{aligned}$$

In particular,

$$g(d) - g'(d) = \max\left(\lambda - \delta^{(f,g)}(c_i, d), 0\right)$$

takes values in $[0, \lambda]$. Moreover, $g(d_j) = g'(d_j)$ because $\delta^{(f,g)}(c_i, d_j) = \lambda$. Since $(f, g) \in \mathrm{Nuc}(M)_0$, row c_i of $\delta^{(f,g)}$ contains a zero by Proposition 21(d); choose d_k with $\delta^{(f,g)}(c_i, d_k) = 0$. Then $g(d_k) - g'(d_k) = \lambda$. Therefore $d_{\mathcal{D}}([g], [g']) = \lambda$.

Since (f, g) and (f', g') are nucleus points, $f = M_* g$ and $f' = M_* g'$. The map M_* is 1-Lipschitz for the projective metrics by Theorem 13, so

$$d_{\mathcal{C}}([f], [f']) \leq d_{\mathcal{D}}([g], [g']) = \lambda.$$

By definition of $d_{\mathbb{P}\mathrm{Nuc}}$ it follows that

$$d_{\mathbb{P}\mathrm{Nuc}}((f, g), (f', g')) = \lambda.$$

Finally, at the distinguished pair (c_i, d_j) we have $g'(d_j) = g(d_j)$ and

$$f'(c_i) = \min_{d \in \mathcal{D}} (M(c_i, d) - g'(d)) \leq M(c_i, d_j) - g'(d_j) = M(c_i, d_j) - g(d_j) = f(c_i) + \lambda.$$

On the other hand, $f'' \leq f' = M_* M^* f''$ by extensivity of the closure operator $M_* M^*$, so $f'(c_i) \geq f''(c_i) = f(c_i) + \lambda$. Thus $f'(c_i) = f(c_i) + \lambda$, and

$$\delta^{(f',g')}(c_i, d_j) = M(c_i, d_j) - f'(c_i) - g'(d_j) = \delta^{(f,g)}(c_i, d_j) - \lambda = 0.$$

Hence $(f', g') \in \mathcal{E}_{c_i, d_j}$, as claimed. □

The construction in the proof has a geometric interpretation worth highlighting. The perturbation f'' moves f by λ in the c_i-coordinate direction, which generically leaves the nucleus. The closure $f' = M_* M^* f''$ projects back onto the nucleus, and the resulting point lies on the event locus $\mathcal{E}_{c_i, d_j}$ at projective distance exactly λ from (f, g). Expressing f' directly in terms of the gap matrix:

$$f'(c) = f(c) + \min_{d \in \mathcal{D}} (\delta(c, d) + \max(\lambda - \delta(c_i, d), 0)).$$

This formula, together with the analogous expression for g', determines the figures in the examples below.

Note that the closure $M_* M^*$ is nonexpansive, so it can only decrease the distance from a nucleus point: $d([f], [M_* M^* f'']) \leq d([f], [f''])$ for any presheaf f''. In general this inequality is strict. The key feature of the construction is that the perturbation $f'' = f + \lambda \mathbf{1}_{c_i}$ is calibrated so that no distance is lost: the extensivity $f'' \leq M_* M^* f''$ pins $f'(c_i)$ at $f(c_i) + \lambda$ from below, while the Isbell equation pins it from above, forcing $d([f], [f']) = \lambda$.

Lemma 39 gives the upper bound: there exists a point on $\mathcal{E}_{c,d}$ at distance exactly λ. The theorem shows this is optimal. For any subset $S \subseteq \mathbb{P}\mathrm{Nuc}(M)$, write $d_{\mathbb{P}\mathrm{Nuc}}((f, g), S) = \inf_{(f', g') \in S} d_{\mathbb{P}\mathrm{Nuc}}((f, g), (f', g'))$.

Theorem 40 (The Events Theorem)**.** *Let* $(f,g) \in \mathbb{P}\mathrm{Nuc}(M)$ *and let* $(c,d) \in \mathcal{C} \times \mathcal{D}$. *Then*

$$d_{\mathbb{P}\mathrm{Nuc}}\left((f,g), \mathcal{E}_{c,d}\right) = \delta^{(f,g)}(c,d).$$

Proof. If $\delta^{(f,g)}(c,d) = 0$, then $(f,g) \in \mathcal{E}_{c,d}$ and both sides are 0. If $M(c,d) = +\infty$, then $\delta^{(f,g)}(c,d) = +\infty$ at every nucleus point and $\mathcal{E}_{c,d} = \varnothing$, so both sides are $+\infty$. Assume $\lambda = \delta^{(f,g)}(c,d) \in (0,\infty)$. By Lemma 39 there exists $(f',g') \in \mathcal{E}_{c,d}$ with $d_{\mathbb{P}\mathrm{Nuc}}\left((f,g),(f',g')\right) = \lambda$, so $d_{\mathbb{P}\mathrm{Nuc}}\left((f,g), \mathcal{E}_{c,d}\right) \leq \lambda$.

For the reverse inequality, work in the gauge slice $\mathrm{Nuc}(M)_0$ and fix $(f_1, g_1) \in \mathcal{E}_{c,d}$. Set $a(c) = f_1(c) - f(c)$ and let

$$\alpha = \min_{c \in \mathcal{C}} a(c), \qquad \beta = \max_{c \in \mathcal{C}} a(c).$$

By definition of the projective metric on $\mathbb{P}\mathcal{C}$ one has

$$d_{\mathcal{C}}\left([f],[f_1]\right) = \beta - \alpha.$$

Since $g = M^* f$ and $g_1 = M^* f_1$, for each $d' \in \mathcal{D}$,

$$g_1(d') = \min_{c \in \mathcal{C}} \left(M(c,d') - f_1(c)\right) = \min_{c \in \mathcal{C}} \left(M(c,d') - f(c) - a(c)\right).$$

Therefore

$$g(d') - \beta \leq g_1(d') \leq g(d') - \alpha,$$

so $g_1(d') - g(d') \in [-\beta, -\alpha]$. It follows that for all $c' \in \mathcal{C}$ and $d' \in \mathcal{D}$,

$$\left(f_1(c') + g_1(d')\right) - \left(f(c') + g(d')\right) = a(c') + \left(g_1(d') - g(d')\right) \in [\alpha - \beta, \beta - \alpha],$$

hence

$$\left|\left(f_1(c') + g_1(d')\right) - \left(f(c') + g(d')\right)\right| \leq \beta - \alpha.$$

Evaluating at the distinguished pair (c,d) and using $f_1(c) + g_1(d) = M(c,d)$ gives

$$\lambda = M(c,d) - f(c) - g(d) = \left(f_1(c) + g_1(d)\right) - \left(f(c) + g(d)\right) \leq \beta - \alpha.$$

Finally, by definition of $d_{\mathbb{P}\mathrm{Nuc}}$ one has

$$d_{\mathcal{C}}\left([f],[f_1]\right) \leq d_{\mathbb{P}\mathrm{Nuc}}\left((f,g),(f_1,g_1)\right),$$

so $\lambda \leq d_{\mathbb{P}\mathrm{Nuc}}\left((f,g),(f_1,g_1)\right)$ for every $(f_1,g_1) \in \mathcal{E}_{c,d}$. Taking the infimum over $\mathcal{E}_{c,d}$ yields $d_{\mathbb{P}\mathrm{Nuc}}\left((f,g), \mathcal{E}_{c,d}\right) \geq \lambda$, hence equality. □

Corollary 41. *Let* $(f,g) \in \mathbb{P}\mathrm{Nuc}(M)$ *and write* $Z = Z(f,g)$. *Then*

$$d_{\mathbb{P}\mathrm{Nuc}}\left((f,g), \bigcup_{(c,d) \notin Z} \mathcal{E}_{c,d}\right) = \min\left\{\delta^{(f,g)}(c,d) \mid (c,d) \notin Z\right\}.$$

In particular, the smallest positive gap is the first radius at which an additional witness can appear.

Proof. Because $\mathcal{C} \times \mathcal{D}$ is finite, the distance from (f, g) to the union of the event loci is the minimum of the distances to the individual event loci. Apply Theorem 40 to each pair $(c, d) \notin Z$. $\square$

We illustrate the perturb-and-project construction on the 3×4 example from §3.5. Let $f = (0, 0, 0)$ and $g = M^* f = (0.7, -1.6, 0.1, -2.9)$, viewed in the gauge slice $\mathrm{Nuc}(M)_0$. The gap matrix is

$$\delta^{(f,g)} = \begin{bmatrix} 0 & 3.1 & 1.6 & 1.6 \\ 0.5 & 4.2 & 0 & 5.1 \\ 1.3 & 0 & 1.9 & 0 \end{bmatrix}.$$

Figure 2 isolates the event with

$$\lambda = 1.9 = \delta^{(f,g)}(c_2, d_3).$$

The shaded hexagon is the radius-λ projective ball about f in $\mathbb{P}\mathcal{C} \cong \mathbb{R}^2$. Define $f'' \colon \mathcal{C} \to \mathbb{R}$ by $f''(c) = f(c)$ for $c \neq c_2$ and $f''(c_2) = f(c_2) + \lambda$. This presheaf need not be a fixed point of $M_* M^*$, and in fact is not: the red point f'' lies on the boundary of the ball but outside the nucleus. Applying the closure produces $f' = M_* M^* f''$, shown in blue, which lies on the event locus $\mathcal{E}_{c_2,d_3}$ and remains at projective distance λ from f. In this example one finds $f' = (0.6, 0, 1.9)$, which is projectively equivalent to $(0, -0.6, 1.3)$ in the gauge $f(c_0) = 0$.

4.4 Order chambers

A witness cell is determined by the zero pattern $Z = \delta^{-1}(0)$ of the gap matrix. Theorem 40 shows that the remaining entries carry metric information: for each $(c, d) \notin Z$ the gap value $\delta(c, d)$ is the distance from (f, g) to the event locus $\mathcal{E}_{c,d}$. Keeping track only of the relative order of the positive gaps refines the witness decomposition.

Definition 42. Fix $(f, g) \in \mathbb{P}\mathrm{Nuc}(M)$ and write $\delta = \delta^{(f,g)}$. The *gap preorder* $\preceq_{f,g}$ on $\mathcal{C} \times \mathcal{D}$ is the total preorder defined by

$$(c, d) \preceq_{f,g} (c', d') \iff \delta(c, d) \leq \delta(c', d').$$

Two nucleus points in the same witness cell lie in the same *order chamber* if they induce the same gap preorder.

To describe a chamber as a polyhedron, let $\preceq$ be a total preorder on $\mathcal{C} \times \mathcal{D}$ and let $Y \subseteq \mathcal{C} \times \mathcal{D}$ be its set of minimal elements. Assume that Y covers $\mathcal{C}$ and $\mathcal{D}$, so that $\mathrm{Cell}(Y) \subseteq \mathrm{Nuc}(M)_0$. The closure of the corresponding order chamber is the subset of

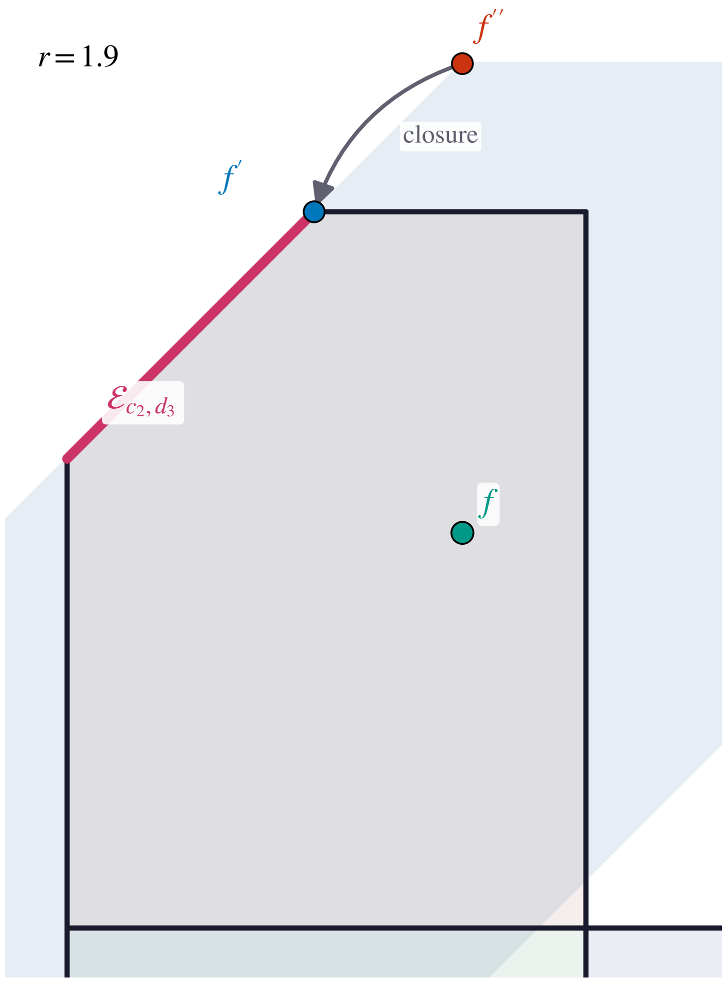

Fig. 2 The event with radius $1.9 = \delta^{(f,g)}(c_2, d_3)$. The red point f'' is obtained by moving a distance 1.9 in the c_2-direction; it lies on the radius-1.9 projective ball but outside the nucleus. Its closure $f' = M_* M^* f''$ lies on the event locus $\mathcal{E}_{c_2,d_3}$, showing that this locus is first met at radius 1.9.

Cell(Y) cut out by the weak inequalities

$$(c, d) \preceq (c', d') \Rightarrow \delta^{(f,g)}(c, d) \leq \delta^{(f,g)}(c', d').$$

The order chamber itself is the relative interior, obtained by requiring strict inequality between distinct equivalence classes. These conditions are gauge-invariant by Lemma 19.

Proposition 43. *The order chambers subdivide each witness cell into a finite polyhedral complex, cut out by a hyperplane arrangement in the gap values. A codimension-one wall corresponds to a tie $\delta(c, d) = \delta(c', d')$ between two gap values*

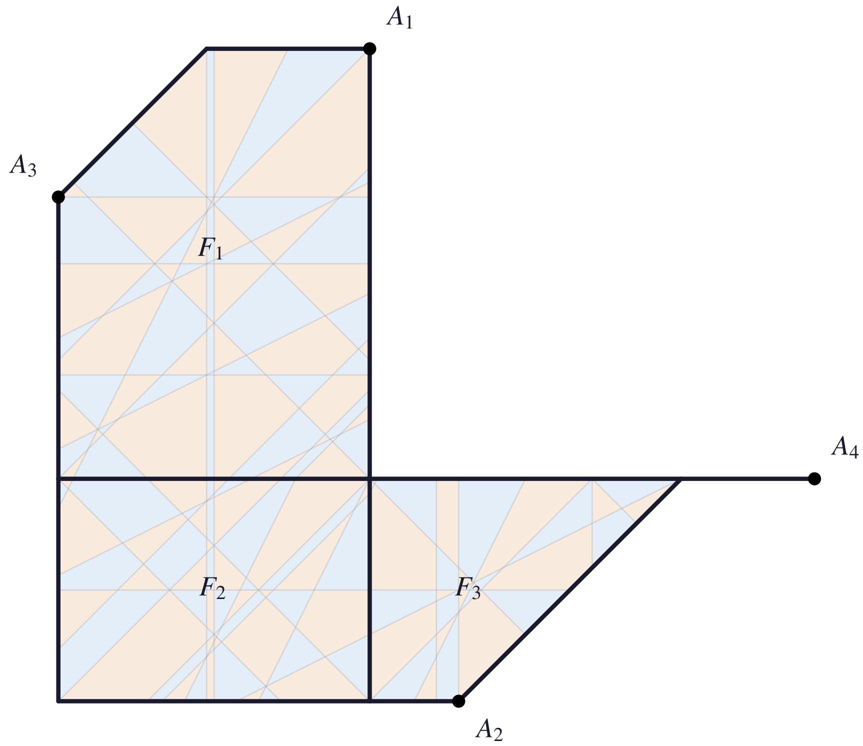

Fig. 3 The global order-chamber refinement of the witness-cell decomposition in the running example. Each witness cell is subdivided by comparing the positive gap values $\delta(c,d)$, and adjacent chambers differ by swapping the order of one neighboring pair of event radii. The two colors indicate the bipartite coloring of the chamber graph.

from distinct equivalence classes of $\preceq$; crossing the wall swaps their order and leaves the rest of the preorder unchanged. In particular, the adjacency graph of order chambers is bipartite [29].

4.5 Pointed thresholding and formal concept lattices

The nucleus of a Boolean relation is a lattice of formal concepts in the sense of Wille [30]; the nucleus of a real-valued matrix is the polyhedral metric space studied in the preceding sections. Thresholding the gap matrix connects the two: at each nucleus point (f,g) and each radius $\varepsilon > 0$, the sublevel set $\{\delta^{(f,g)}(c,d) \leq \varepsilon\}$ is a Boolean relation whose concept lattice L_ε captures the discrete Galois-closed dependencies among the incidences visible from (f,g) at resolution ε. As ε increases, the relation grows and the lattice with it, assembling a tower of concept lattices that records the order in which the entries of M are reconstructed from the witness pairs outward.

To describe the lattice L_ε explicitly, we briefly review the Boolean case. By Booleans, we mean the two-element monoidal poset $\{0,1\}$ with order $0 \leq 1$ and monoidal product $\wedge$. A Boolean-valued profunctor on discrete sets $\mathcal{C}, \mathcal{D}$ amounts to a relation $R \subseteq \mathcal{C} \times \mathcal{D}$. Pre- and copresheaves are Boolean-valued functions and may be identified with subsets. The Isbell conjugates of R are the pair of order-reversing maps between power sets

$$
\begin{aligned}
R^*\colon \mathcal{P}(\mathcal{C}) \to \mathcal{P}(\mathcal{D}), \quad & R^*(F) = \{d \in \mathcal{D} \mid (c,d) \in R \text{ for all } c \in F\}\,, \\
R_*\colon \mathcal{P}(\mathcal{D}) \to \mathcal{P}(\mathcal{C}), \quad & R_*(G) = \{c \in \mathcal{C} \mid (c,d) \in R \text{ for all } d \in G\}\,,
\end{aligned}
$$

and they form a Galois connection: $F \subseteq R_*(G) \iff G \subseteq R^*(F)$. A *formal concept* of R is a fixed point of this adjunction—a pair (F, G) with $G = R^*(F)$ and $F = R_*(G)$—exactly as in Definition 3, but enriched over Booleans rather than $\overline{\mathbb{R}}$. Following [30], F is called the *extent* and G the *intent*. The set of all formal concepts is a complete lattice, ordered by inclusion of extents (equivalently, reverse inclusion of intents):

$$
(F, G) \le (F', G') \iff F \subseteq F' \iff G' \subseteq G.
$$

Meets and joins are computed by intersecting extents and intents respectively, then re-closing:

$$
\bigwedge_i (F_i, G_i) = \left(\bigcap_i F_i, R^* \left(\bigcap_i F_i \right) \right), \quad \bigvee_i (F_i, G_i) = \left(R_* \left(\bigcap_i G_i \right), \bigcap_i G_i \right).
$$

When $R \subseteq R' \subseteq C \times D$ are nested relations, there are two canonical ways to transport a formal concept of R to one of R': re-close the extent, or re-close the intent.

Proposition 44. *Let $R \subseteq R' \subseteq C \times D$ be relations on sets of objects C and attributes D, with extensions to power sets $R^*\colon \mathcal{P}(C) \to \mathcal{P}(D)$ and $R_*\colon \mathcal{P}(D) \to \mathcal{P}(C)$, and similarly for R'. Write $L = \mathrm{Nuc}(R)$ and $L' = \mathrm{Nuc}(R')$ for their concept lattices. The maps*

$$
T^{\mathrm{ext}}_{R \to R'}, T^{\mathrm{int}}_{R \to R'} : L \to L'
$$

defined by

$$
\begin{aligned}
T^{\mathrm{ext}}_{R \to R'}(F, G) &:= (R'_*(R'^* F),\, R'^* F)\,, \\
T^{\mathrm{int}}_{R \to R'}(F, G) &:= (R'_* G,\, R'^*(R'_* G))\,.
\end{aligned}
$$

are monotone. Moreover:
(a) $T^{\mathrm{ext}}_{R \to R'}(F, G)$ is the least *concept of L' whose extent contains F.*
(b) $T^{\mathrm{int}}_{R \to R'}(F, G)$ is the greatest *concept of L' whose intent contains G.*
(c) One has the inequality $T^{\mathrm{ext}}_{R \to R'}(F, G) \le T^{\mathrm{int}}_{R \to R'}(F, G)$ in L'.

Proof. Both pairs are concepts of R' by construction, and monotonicity follows from the monotonicity of the closure operators $R'_* R'^*$ on extents and $R'^* R'_*$ on intents.

For (a), extensivity of $R'_* R'^*$ gives $F \subseteq R'_*(R'^* F)$. Now let $(F', G') \in L'$ with $F \subseteq F'$. Since R'^* is antitone, $G' = R'^* F' \subseteq R'^* F$, and applying the antitone map R'_*

gives $R'_*(R'^*F) \subseteq R'_*(G') = F'$. Thus $T^{\mathrm{ext}}_{R\to R'}(F,G)$ is the least concept of L' whose extent contains F.

For (b), extensivity of $R'^*R'_*$ gives $G \subseteq R'^*(R'_*G)$. If $(F',G') \in L'$ and $G \subseteq G'$, then antitonicity of R'_* yields $F' = R'_*(G') \subseteq R'_*(G)$, so $(F',G') \le T^{\mathrm{int}}_{R\to R'}(F,G)$.

For (c), since $(F,G) \in L$ one has $G = R^*F$. Because $R \subseteq R'$, every attribute related by R to every element of F is also related by R', so $G = R^*F \subseteq R'^*F$. Applying R'_* gives $R'_*(R'^*F) \subseteq R'_*(G)$, which is the extent inequality $T^{\mathrm{ext}}_{R\to R'}(F,G) \le T^{\mathrm{int}}_{R\to R'}(F,G)$. □

4.6 Chamberwise lattice towers

Return to the $\overline{\mathbb{R}}$-profunctor M on finite discrete sets $\mathcal{C}$ and $\mathcal{D}$, and fix a nucleus point $(f,g) \in \mathrm{Nuc}(M)$. For $\varepsilon \ge 0$, the sublevel set of the gap matrix defines a Boolean relation.

Definition 45. For $\varepsilon \ge 0$, define a Boolean profunctor $R^{(f,g)}_\varepsilon : \mathcal{C}^{\mathrm{op}} \otimes \mathcal{D} \to \{0,1\}$ by

$$R^{(f,g)}_\varepsilon(c,d) = \begin{cases} 1 & \text{if } \delta^{(f,g)}(c,d) \le \varepsilon, \\ 0 & \text{if } \delta^{(f,g)}(c,d) > \varepsilon. \end{cases}$$

Let $L_\varepsilon(f,g) := \mathrm{Nuc}(R^{(f,g)}_\varepsilon)$ denote its concept lattice.

An element of $L_\varepsilon(f,g)$ is a pair (F,G) with $F = (R_\varepsilon)_*G$ and $G = (R_\varepsilon)^*F$. The Events Theorem gives these threshold relations a direct metric interpretation.

Proposition 46. *For every $(f,g) \in \mathrm{Nuc}(M)$, every $\varepsilon \ge 0$, and every $(c,d) \in \mathcal{C} \times \mathcal{D}$,*

$$(c,d) \in R^{(f,g)}_\varepsilon \iff \delta^{(f,g)}(c,d) \le \varepsilon \iff d_{\mathbb{P}\mathrm{Nuc}}\left((f,g), \mathcal{E}_{c,d}\right) \le \varepsilon.$$

Thus the threshold relation records exactly which event loci lie within projective distance ε of (f,g).

Proof. The first equivalence is the definition of $R^{(f,g)}_\varepsilon$, and the second is Theorem 40. □

For brevity, write $R_\varepsilon := R^{(f,g)}_\varepsilon$ and $L_\varepsilon := L_\varepsilon(f,g)$. If $\varepsilon \le \varepsilon'$ then $R_\varepsilon \subseteq R_{\varepsilon'}$, and Proposition 44 provides two monotone transport maps

$$T^{\mathrm{ext}}_{\varepsilon,\varepsilon'} \;\le\; T^{\mathrm{int}}_{\varepsilon,\varepsilon'} : L_\varepsilon \to L_{\varepsilon'},$$

the first preserving extents minimally, the second preserving intents maximally.

For a fixed numerical value of ε, the relation $R^{(f,g)}_\varepsilon$ need not be locally constant as (f,g) varies in an order chamber, since the gap values $\delta^{(f,g)}(c,d)$ move continuously. What *is* locally constant on an order chamber is the *order* in which incidences enter as ε increases. It is therefore natural to reindex the construction by the equivalence classes of the chamber preorder.

Fix an order chamber Q with total preorder $\preceq_Q$ on $\mathcal{C} \times \mathcal{D}$, and let

$$E_0 \prec_Q E_1 \prec_Q \cdots \prec_Q E_m$$

be the equivalence classes, ordered from smallest to largest. For $0 \leq k \leq m$ define

$$R_k^Q := \bigcup_{i=0}^{k} E_i \subseteq \mathcal{C} \times \mathcal{D},$$

giving a finite chain of relations $R_0^Q \subseteq R_1^Q \subseteq \cdots \subseteq R_m^Q$. Setting $L_k^Q := \mathrm{Nuc}(R_k^Q)$, each inclusion $R_k^Q \subseteq R_{k'}^Q$ carries a pair of monotone transport maps $T_{k,k'}^{\mathrm{ext}} \leq T_{k,k'}^{\mathrm{int}} \colon L_k^Q \to L_{k'}^Q$ from Proposition 44, giving a finite tower of concept lattices

$$L_0^Q \to L_1^Q \to \cdots \to L_m^Q$$

with canonical lower and upper structure maps at each step.

Proposition 47. *Let Q be an order chamber with preorder classes $E_0 \prec_Q \cdots \prec_Q E_m$. For every $(f, g) \in Q$ and every $\varepsilon \geq 0$, there is a unique index k such that $R_\varepsilon^{(f,g)} = R_k^Q$. Consequently the real-parameter family $\varepsilon \mapsto L_\varepsilon(f, g)$ factors through the finite tower $L_0^Q \to \cdots \to L_m^Q$, and this tower depends only on the chamber Q, not on the chosen point.*

Proof. Points of the same order chamber determine the same total preorder on the gap values. The sublevel condition $\delta^{(f,g)}(c, d) \leq \varepsilon$ therefore selects an initial segment of the preorder classes, and every initial segment arises for a unique threshold value. □

4.7 The constructible sheaf of lattices

The chamberwise towers assemble into a global structure over the order-chamber complex.

Proposition 48. *Let Q' be a face of the closure of an order chamber Q. Then $\preceq_{Q'}$ is obtained from $\preceq_Q$ by merging consecutive equivalence classes, so each class for $\preceq_{Q'}$ is a union of consecutive classes for $\preceq_Q$. Consequently, the chain $\left\{R_\ell^{Q'}\right\}$ is a subsequence of $\left\{R_k^Q\right\}$, and the lattice tower for Q' is a coarsening of the tower for Q: it retains the same concept lattices at the merged thresholds, with structure maps given by the direct transport for the corresponding (larger) inclusion.*

Proof. Passing from Q to a face imposes equalities among neighboring gap values while preserving their order relative to the remaining classes. Hence only consecutive preorder blocks merge, and the relation chain and lattice tower coarsen accordingly. □

$$
\begin{array}{ccccccccccc}
L_0^Q & \longrightarrow & L_1^Q & \longrightarrow & L_2^Q & \longrightarrow & L_3^Q & \longrightarrow & L_4^Q & \longrightarrow & L_5^Q \\
\| & & & & \| & & & & \| & & \| \\
L_0^{Q'} & & \xrightarrow{\quad T_{0,2} \quad} & & L_1^{Q'} & & \xrightarrow{\quad T_{2,4} \quad} & & L_2^{Q'} & \xrightarrow{T_{4,5}} & L_3^{Q'}
\end{array}
$$

Fig. 4 Specialization to a face. If $Q' \leq \overline{Q}$ merges consecutive preorder blocks, the tower over Q' retains the lattices at the merged thresholds, with structure maps given by the direct transport for the coarsened inclusions.

Propositions 47 and 48 say that the assignment

$$Q \longmapsto \left(L_0^Q \to L_1^Q \to \cdots \to L_m^Q\right)$$

is a *constructible sheaf of lattice towers* over the order-chamber complex: each open chamber receives a finite tower of concept lattices, and specialization to a face coarsens the tower by merging consecutive floors. The two transport maps $T^{\text{ext}} \leq T^{\text{int}}$ furnish canonical lower and upper structure maps at each step of each tower; both are preserved by the coarsening.

We illustrate on the running example. Let $\mathcal{C} = \{c_0, c_1, c_2\}$, $\mathcal{D} = \{d_1, d_2, d_3, d_4\}$, and let

$$M = \begin{bmatrix} 0.7 & 1.5 & 1.7 & -1.3 \\ 1.2 & 2.6 & 0.1 & 2.2 \\ 2.0 & -1.6 & 2.0 & -2.9 \end{bmatrix}.$$

Work in the gauge slice $f(c_0) = 0$, so $f = (0, x, y) \in \mathbb{R}^3$. Consider three nucleus points

$$f_1 = (0, 0, 0), \qquad f_2 = (0, -0.1, 0), \qquad f_3 = (0, 0.1, 0),$$

with $g_i := M^* f_i$ and $\delta_i(c, d) := M(c, d) - f_i(c) - g_i(d)$. One computes

$$\begin{aligned} g_1 &= (0.7, -1.6, 0.1, -2.9), \\ g_2 &= (0.7, -1.6, 0.2, -2.9), \\ g_3 &= (0.7, -1.6, 0.0, -2.9) \end{aligned}$$

and gap matrices

$$\delta_1 = \begin{bmatrix} 0.0 & 3.1 & 1.6 & 1.6 \\ 0.5 & 4.2 & 0.0 & 5.1 \\ 1.3 & 0.0 & 1.9 & 0.0 \end{bmatrix}, \qquad \delta_2 = \begin{bmatrix} 0.0 & 3.1 & 1.5 & 1.6 \\ 0.6 & 4.3 & 0.0 & 5.2 \\ 1.3 & 0.0 & 1.8 & 0.0 \end{bmatrix}, \qquad \delta_3 = \begin{bmatrix} 0.0 & 3.1 & 1.7 & 1.6 \\ 0.4 & 4.1 & 0.0 & 5.0 \\ 1.3 & 0.0 & 2.0 & 0.0 \end{bmatrix}.$$

In particular,

$$\delta_2(c_0, d_3) < \delta_2(c_0, d_4), \qquad \delta_1(c_0, d_3) = \delta_1(c_0, d_4), \qquad \delta_3(c_0, d_4) < \delta_3(c_0, d_3),$$

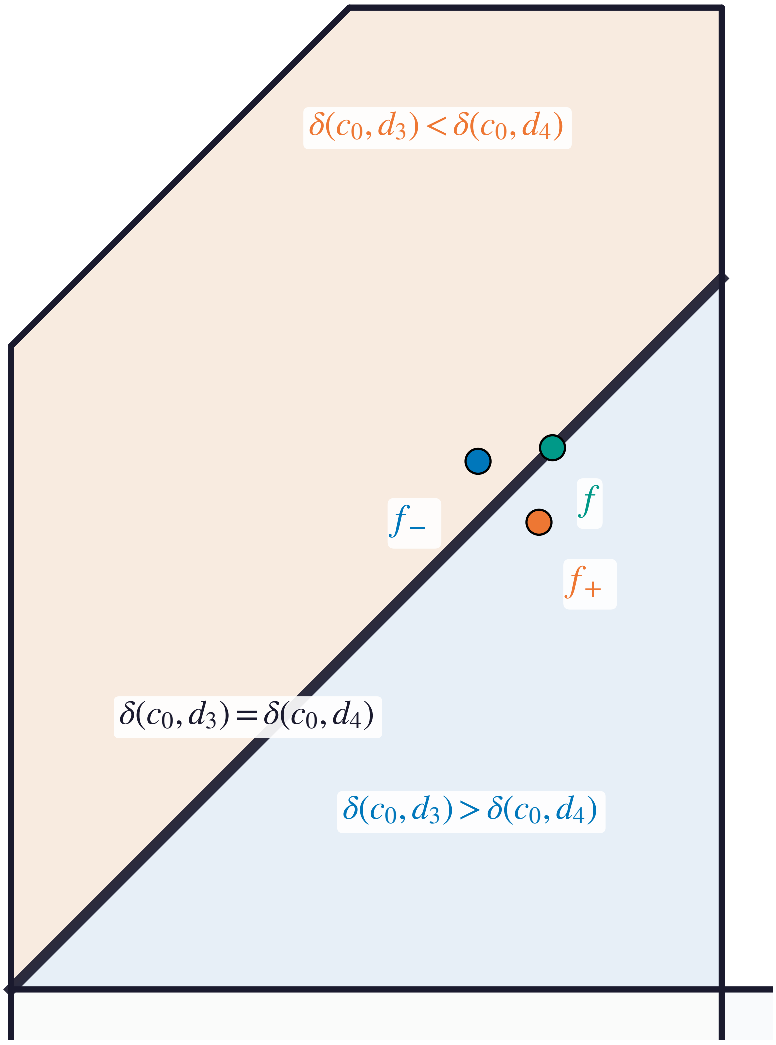

Fig. 5 Local wall crossing at the 1.6 tie. Along the diagonal wall one has $\delta(c_0, d_3) = \delta(c_0, d_4)$. The points f_2 and f_3 lie in the adjacent chambers where these two event radii are ordered oppositely, while f_1 lies on the wall.

so f_2 and f_3 lie in adjacent order chambers Q_2 and Q_3, separated by the wall $Q_1 = \overline{Q_2} \cap \overline{Q_3}$ through f_1. Figure 5 shows these three points.

Let $R_\varepsilon^{(f_i,g_i)} := \{(c,d) \mid \delta_i(c,d) \le \varepsilon\}$ be the threshold relations and write $L(R) := \mathrm{Nuc}(R)$ for the concept lattice of a relation $R \subseteq \mathcal{C} \times \mathcal{D}$. Define the relations

$$
\begin{aligned}
R_0 &:= \{(c_0,d_1),(c_1,d_3),(c_2,d_2),(c_2,d_4)\},\\
R_1 &:= R_0 \cup \{(c_1,d_1)\}, \qquad R_2 := R_1 \cup \{(c_2,d_1)\},\\
R_{3a} &:= R_2 \cup \{(c_0,d_3)\}, \qquad R_{3b} := R_2 \cup \{(c_0,d_4)\},\\
R_4 &:= R_2 \cup \{(c_0,d_3),(c_0,d_4)\},\\
R_5 &:= R_4 \cup \{(c_2,d_3)\}, \qquad R_6 := R_5 \cup \{(c_0,d_2)\},\\
R_7 &:= R_6 \cup \{(c_1,d_2)\}, \qquad R_8 := R_7 \cup \{(c_1,d_4)\}.
\end{aligned}
$$

The chamberwise towers are:

$$
\begin{aligned}
&\text{at } f_2 \in Q_2: && R_0 \subset R_1 \subset R_2 \subset R_{3a} \subset R_4 \subset R_5 \subset R_6 \subset R_7 \subset R_8\\
&\text{at } f_1 \in Q_1: && R_0 \subset R_1 \subset R_2 \subset R_4 \subset R_5 \subset R_6 \subset R_7 \subset R_8\\
&\text{at } f_3 \in Q_3: && R_0 \subset R_1 \subset R_2 \subset R_{3b} \subset R_4 \subset R_5 \subset R_6 \subset R_7 \subset R_8.
\end{aligned}
$$

The only combinatorial difference between the two chambers is the order in which (c_0,d_3) and (c_0,d_4) enter; on the wall they enter simultaneously.

The corresponding concept lattices are:

$$
\begin{aligned}
L(R_2) = \{&(\varnothing \mid \{d_1,d_2,d_3,d_4\}),\ (\{c_1\} \mid \{d_1,d_3\}),\\
&(\{c_2\} \mid \{d_1,d_2,d_4\}),\ (\{c_0,c_1,c_2\} \mid \{d_1\})\},
\end{aligned}
$$

$$
\begin{aligned}
L(R_{3a}) = \{&(\varnothing \mid \{d_1,d_2,d_3,d_4\}),\ (\{c_2\} \mid \{d_1,d_2,d_4\}),\\
&(\{c_0,c_1\} \mid \{d_1,d_3\}),\ (\{c_0,c_1,c_2\} \mid \{d_1\})\},\\
L(R_{3b}) = \{&(\varnothing \mid \{d_1,d_2,d_3,d_4\}),\ (\{c_1\} \mid \{d_1,d_3\}),\ (\{c_2\} \mid \{d_1,d_2,d_4\}),\\
&(\{c_0,c_2\} \mid \{d_1,d_4\}),\ (\{c_0,c_1,c_2\} \mid \{d_1\})\},\\
L(R_4) = \{&(\varnothing \mid \{d_1,d_2,d_3,d_4\}),\ (\{c_0\} \mid \{d_1,d_3,d_4\}),\ (\{c_2\} \mid \{d_1,d_2,d_4\}),\\
&(\{c_0,c_1\} \mid \{d_1,d_3\}),\ (\{c_0,c_2\} \mid \{d_1,d_4\}),\ (\{c_0,c_1,c_2\} \mid \{d_1\})\}.
\end{aligned}
$$

For $k \ge 4$ the lattices are the same in all three towers: $L(R_5)$ is a 3-element chain, $L(R_6)$ and $L(R_7)$ are 2-element chains, and $L(R_8)$ is the one-point lattice.

In this example, the direct transport map $T_{2,4}\colon L(R_2) \to L(R_4)$ agrees with both composites $T_{3a,4} \circ T_{2,3a}$ and $T_{3b,4} \circ T_{2,3b}$, for either choice of T^{ext} or T^{int}. Figure 6 displays this as a commutative diamond.

5 The square case: optimal assignments and Chebyshev centering

When $\mathcal{C} = \mathcal{D}$ and $|\mathcal{C}| = n$, the projective nucleus of a real $n \times n$ matrix M acquires canonical structure that has no analogue in the rectangular case: a distinguished scalar

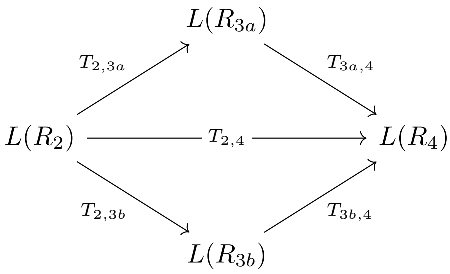

Fig. 6 Wall specialization at the 1.6 tie. The adjacent chambers Q_2 and Q_3 correspond to the two strict refinements of the wall preorder: in Q_2 the incidence (c_0, d_3) enters before (c_0, d_4) (intermediate relation R_{3a}), while in Q_3 the order is reversed (R_{3b}). On the wall, the tie merges these floors, yielding the direct inclusion $R_2 \subset R_4$. In this example the diamond commutes for both T^{ext} and T^{int}.

val(M) (the tropical determinant), at most one full-dimensional cell (and generically exactly one), and within that cell a canonical center point whose inscribed radius is a computable invariant of M. This section develops these invariants: §5.1–§5.2 identify the full-dimensional cell via the tropical determinant, and §5.4–§5.5 solve the Chebyshev centering problem, expressing the optimal radius as a minimum directed cycle mean.

5.1 Full-dimensional cells are permutation cells

The projective space $\mathbb{P}\mathcal{C}$ has dimension $n-1$, so every witness polyhedron Cell(Y) has dimension at most $n-1$. We call Cell(Y) *full-dimensional* if it achieves this bound. Any subset $Y \subseteq \mathcal{C} \times \mathcal{C}$ that covers both factors has $|Y| \geq n$, with equality if and only if Y is the graph of a permutation $\sigma \in S_n$:

$$\Gamma_\sigma := \{(c, \sigma(c)) \mid c \in \mathcal{C}\}\,.$$

Proposition 49. *If* Cell(Y) *is full-dimensional, then* $Y = \Gamma_\sigma$ *for a permutation* $\sigma \in S_n$.

Proof. Suppose Y contains two pairs (c, d) and (c', d) with the same second coordinate d. The witness equalities (13) give $f(c) + g(d) = M(c, d)$ and $f(c') + g(d) = M(c', d)$, hence $f(c') - f(c) = M(c', d) - M(c, d)$ throughout Cell(Y). In the gauge $f(c_0) = 0$, this is a nontrivial affine relation among the $n-1$ free coordinates of f, so Cell(Y) has dimension at most $n-2$. Therefore each $d \in \mathcal{C}$ appears in at most one pair of Y. Since Y covers $\mathcal{D} = \mathcal{C}$, each d appears exactly once, giving $|Y| \leq n$. But Y also covers $\mathcal{C}$, so $|Y| \geq n$, and Y is the graph of a bijection. $\square$

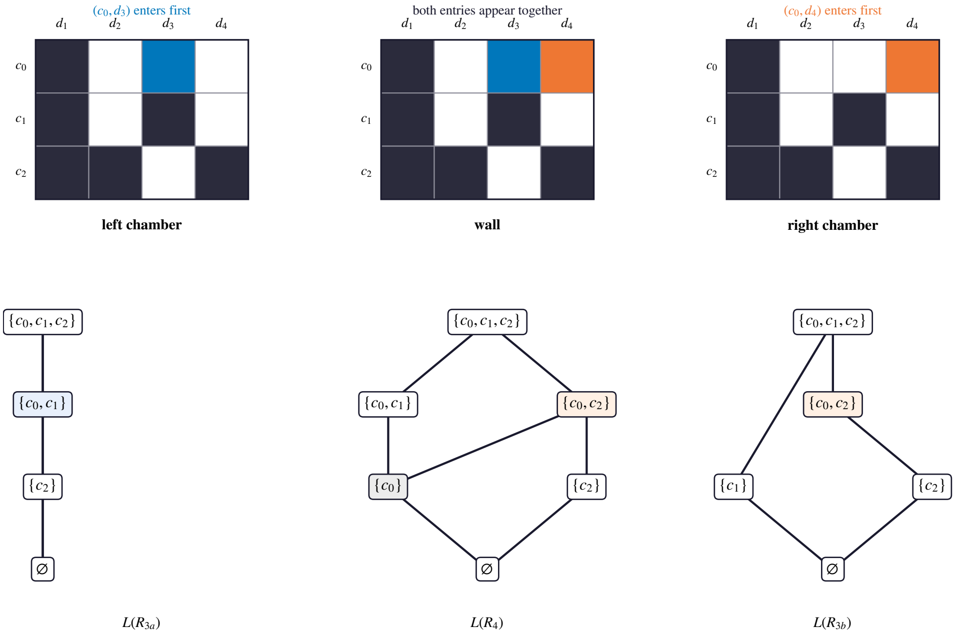

Fig. 7 Thresholded relations and Boolean nuclei across the 1.6 tie. In the left chamber the new incidence (c_0, d_3) enters first, giving R_{3a}; on the wall both incidences appear together, giving R_4; in the right chamber (c_0, d_4) enters first, giving R_{3b}. The lattices below show the corresponding change in the Boolean nuclei.

5.2 The tropical determinant and the full-dimensional cell

For $\sigma \in S_n$, the *M-cost* of σ is $\sum_{c\in\mathcal{C}} M(c, \sigma(c))$. The *tropical determinant* of M is the minimum cost:

$$\mathrm{val}(M) := \min_{\sigma\in S_n} \sum_{c\in\mathcal{C}} M(c, \sigma(c)).$$

The matrix M is *tropically nonsingular* if the minimum is achieved by a unique permutation. A permutation achieving $\mathrm{val}(M)$ is *optimal*; finding one is the linear assignment problem [3, 4]; see also [5] for the tropical perspective.

By Proposition 49, every full-dimensional cell has the form $\mathrm{Cell}(\Gamma_\sigma)$. The admissible permutations turn out to be exactly the optimal ones:

Proposition 50. *The permutation graph Γ_σ is admissible if and only if σ achieves $\mathrm{val}(M)$.*

This is a restatement, in Isbell-duality language, of the classical LP-duality characterization of optimal assignments [23, Ch. 17]. We include the proof here because it gives an explicit construction of a nucleus point in $\mathrm{Cell}(\Gamma_\sigma)$. This is useful because the full-dimensional cell is difficult to reach by naive geometric means: projecting a generic point of $\mathbb{P}\mathcal{C}$ onto $\mathbb{P}\mathrm{Nuc}(M)$ typically lands on a low-dimensional boundary cell rather than the full-dimensional interior. The shortest-path construction in the proof reappears, in strengthened form, in the Centering Theorem.

Proposition 51. *The permutation graph Γ_σ is admissible if and only if σ achieves* $\mathrm{val}(M)$.

Proof. Suppose Γ_σ is admissible and choose $(f, g) \in \mathrm{Cell}(\Gamma_\sigma)$. Summing the witness equalities $f(c) + g(\sigma(c)) = M(c, \sigma(c))$ over c gives

$$\sum_c f(c) + \sum_d g(d) = \sum_c M(c, \sigma(c)), \tag{14}$$

since σ is a bijection. For any other permutation τ, summing the Isbell inequalities $f(c) + g(\tau(c)) \leq M(c, \tau(c))$ over c gives $\sum_c f(c) + \sum_d g(d) \leq \sum_c M(c, \tau(c))$. Comparing with (14) shows σ achieves $\mathrm{val}(M)$.

Now assume σ achieves $\mathrm{val}(M)$. We construct $(f, g) \in \mathrm{Cell}(\Gamma_\sigma)$ via shortest-path potentials.

Step 1: the auxiliary digraph. Form the complete directed graph on $\mathcal{C}$ with edge weights

$$w(c \to c') := M(c', \sigma(c)) - M(c, \sigma(c)).$$

The weight $w(c \to c')$ measures the change in cost when column $\sigma(c)$ is reassigned from row c to row c'.

Step 2: no negative cycles. For a directed cycle $x_0 \to x_1 \to \cdots \to x_k = x_0$, the edge weights telescope:

$$\sum_{i=0}^{k-1} w(x_i \to x_{i+1}) = \sum_{i=0}^{k-1} \big[M(x_{i+1}, \sigma(x_i)) - M(x_i, \sigma(x_i))\big].$$

If this were negative, replacing σ on the cycle by $\tau(x_{i+1}) = \sigma(x_i)$ would yield a permutation of strictly lower cost, contradicting optimality.

Step 3: shortest-path potentials. Fix a base vertex c_0. Since there are no negative cycles, shortest-path distances from c_0 are well-defined. Set $f(c)$ to be the shortest-path distance from c_0 to c, so $f(c_0) = 0$ and $f(c') \leq f(c) + w(c \to c')$ for all c, c'. Define g by

$$g(\sigma(c)) := M(c, \sigma(c)) - f(c).$$

The witness equalities $f(c) + g(\sigma(c)) = M(c, \sigma(c))$ hold by construction. For the off-diagonal Isbell inequality at $(c', \sigma(c))$ with $c' \neq c$:

$$f(c') + g(\sigma(c)) = f(c') + M(c, \sigma(c)) - f(c) \leq M(c', \sigma(c)),$$

since $M(c', \sigma(c)) - M(c, \sigma(c)) = w(c \to c') \geq f(c') - f(c)$. As σ is a bijection, every $d \in \mathcal{C}$ is $\sigma(c)$ for a unique c, so $f(c') + g(d) \leq M(c', d)$ for all pairs. Thus $(f, g) \in \mathrm{Cell}(\Gamma_\sigma)$. □

Proposition 52. *If σ and τ are distinct permutations both achieving* $\mathrm{val}(M)$*, then* $\mathrm{Cell}(\Gamma_\sigma) = \mathrm{Cell}(\Gamma_\tau) = \mathrm{Cell}(\Gamma_\sigma \cup \Gamma_\tau)$*, and this cell has dimension at most $n - 3$.*

Proof. Let $(f, g) \in \mathrm{Cell}(\Gamma_\sigma)$. Summing the witness equalities for σ gives $\sum_c f(c) + \sum_d g(d) = \mathrm{val}(M)$ as in (14). Summing the Isbell inequalities along τ gives $\sum_c f(c) + \sum_d g(d) \le \mathrm{val}(M)$, using that τ is a bijection achieving the same value. The left-hand sides are identical, so every inequality $f(c) + g(\tau(c)) \le M(c, \tau(c))$ is an equality. Thus $(f, g) \in \mathrm{Cell}(\Gamma_\tau)$, giving $\mathrm{Cell}(\Gamma_\sigma) \subseteq \mathrm{Cell}(\Gamma_\tau)$; by symmetry the two cells coincide. Since $\sigma \neq \tau$ differ on at least two elements of $\mathcal{C}$, we have $|\Gamma_\sigma \cup \Gamma_\tau| \ge n + 2$, so the cell satisfies at least $n + 2$ witness equalities and has dimension at most $n - 3$. ☐

Combining the three propositions: $\mathbb{P}\mathrm{Nuc}(M)$ has a full-dimensional cell if and only if M is tropically nonsingular, and in that case there is exactly one.

5.3 Chebyshev centering

Assume now that M is tropically nonsingular with unique optimal permutation σ. On the full-dimensional cell $\mathrm{Cell}(\Gamma_\sigma)$, the witness equalities $f(c) + g(\sigma(c)) = M(c, \sigma(c))$ pin g to f:

$$g(\sigma(c)) = M(c, \sigma(c)) - f(c), \qquad c \in \mathcal{C}. \tag{15}$$

The gap matrix on this cell is

$$\delta^{(f,g)}(c, d) = M(c, d) - f(c) - g(d),$$

with $\delta^{(f,g)}(c, \sigma(c)) = 0$ for all c and $\delta^{(f,g)}(c, d) > 0$ for $d \neq \sigma(c)$ at every interior point of $\mathrm{Cell}(\Gamma_\sigma)$. The Events Theorem identifies each positive gap entry with the projective distance to the corresponding event locus. List the $n(n-1)$ positive gap values in nondecreasing order:

$$e_1(f, g) \le e_2(f, g) \le \cdots \le e_{n(n-1)}(f, g).$$

Then e_1 is the distance from (f, g) to the nearest cell wall. A *Chebyshev center* of $\mathrm{Cell}(\Gamma_\sigma)$ is any point maximizing e_1, and the *Chebyshev radius* is

$$r^* := \max_{(f,g) \in \mathrm{Cell}(\Gamma_\sigma)} e_1(f, g).$$

Since $e_1 = 0$ on $\partial\,\mathrm{Cell}(\Gamma_\sigma)$ and $e_1 > 0$ in the interior, every Chebyshev center lies in the open cell $\mathrm{Cell}^\circ(\Gamma_\sigma)$.

5.4 The Chebyshev LP

Substituting (15) into the gap formula, the gap at an off-diagonal pair $(c, \sigma(c'))$ with $c \neq c'$ is

$$\delta^{(f,g)}(c, \sigma(c')) = \alpha(c, c') + f(c') - f(c), \tag{16}$$

where

$$\alpha(c, c') := M(c, \sigma(c')) - M(c', \sigma(c')) \tag{17}$$

depends only on M and σ. Note that $\alpha(c, c') = w(c' \to c)$, the edge weight of the auxiliary digraph from the proof of Proposition 51 with the direction reversed: the shortest-path potentials constructed there satisfy $\alpha(c, c') + f(c') - f(c) \geq 0$, which is exactly the condition $\delta \geq 0$ on the off-diagonal gaps.

Since σ is a bijection, the off-diagonal pairs (c, d) with $d \neq \sigma(c)$ are exactly the pairs $(c, \sigma(c'))$ with $c \neq c'$. In the gauge slice $f(c_0) = 0$, the Chebyshev problem becomes the linear program

$$\begin{aligned} \text{maximize} \quad & r \\ \text{subject to} \quad & \alpha(c, c') + f(c') - f(c) \geq r, \quad c \neq c', \\ & f(c_0) = 0. \end{aligned} \tag{18}$$

This LP has n variables—$f(c_1), \ldots, f(c_{n-1})$ and r—and $n(n-1)$ inequality constraints.

5.5 Statement and proofs

Theorem 53 (The Centering Theorem)**.** *Let M be a real $n \times n$ matrix with a unique optimal permutation σ, and let $\alpha(c, c') = M(c, \sigma(c')) - M(c', \sigma(c'))$.*

(a) *The Chebyshev radius r^* equals the minimum directed cycle mean of the digraph $(\mathcal{C}, \alpha)$:*

$$r^* = \min_{\gamma} \frac{1}{|\gamma|} \sum_{(c,c') \in \gamma} \alpha(c, c'),$$

where γ ranges over all directed cycles in the complete graph on $\mathcal{C}$.

(b) *At every vertex of the (convex) set of Chebyshev centers, the multiplicity of e_1 is at least n.*

(c) *For M outside a proper algebraic subset of $\mathbb{R}^{n \times n}$, the Chebyshev center is unique and the multiplicity of e_1 is exactly n.*

Proof of (a). The constraints in (18) read $f(c) - f(c') \leq \alpha(c, c') - r$ for all $c \neq c'$. This is a system of *difference constraints*: upper bounds on pairwise differences of the variables $f(c)$. By a standard result in combinatorial optimization [23, Corollary 8.3b], such a system is feasible if and only if the weighted digraph with edge weights $\alpha(c, c') - r$ has no negative-weight directed cycle. A cycle γ of length $|\gamma|$ has nonnegative total weight if and only if $\sum_{(c,c') \in \gamma} (\alpha(c, c') - r) \geq 0$, equivalently $r \leq (1/|\gamma|) \sum_{(c,c') \in \gamma} \alpha(c, c')$, the cycle mean of γ. This holds for every directed cycle if and only if $r \leq \min_\gamma \operatorname{mean}(\gamma)$. Hence $r^* = \min_\gamma \operatorname{mean}(\gamma)$. □

Proof of (b). The LP (18) lives in $\mathbb{R}^n$ with $n(n-1)$ inequality constraints. The set of optimal solutions is a face of the feasible polyhedron. At any vertex of this face, at least n linearly independent inequality constraints are active. Each active constraint $\alpha(c, c') + f^*(c') - f^*(c) = r^*$ corresponds, via (16), to an off-diagonal pair $(c, \sigma(c'))$ with $\delta^{(f^*, g^*)}(c, \sigma(c')) = r^* = e_1$. Hence the multiplicity of e_1 is at least n. □

Proof of (c). LP nondegeneracy—the condition that no vertex of the feasible polyhedron has more than n active inequality constraints—fails on a proper algebraic subset of the space of α-coefficients, hence of $\mathbb{R}^{n\times n}$. For nondegenerate instances, the optimal vertex is unique and exactly n constraints are active. □

5.6 A square example

The running 3×4 example is not square. To illustrate the Centering Theorem, consider the 3×3 matrix

$$M = \begin{bmatrix} 1 & 1 & 6 \\ 6 & 3 & 1 \\ 1 & 6 & 5 \end{bmatrix}.$$

The unique optimal permutation is the 3-cycle $\sigma = (1\,2\,0)$, with $\mathrm{val}(M) = M(0,1) + M(1,2) + M(2,0) = 1 + 1 + 1 = 3$, so M is tropically nonsingular and $\mathrm{Cell}(\Gamma_\sigma)$ is the unique full-dimensional cell. On this cell the three witness equalities $f(c) + g(\sigma(c)) = M(c, \sigma(c))$ pin g to f; in the gauge $f(c_0) = 0$ with $f_1 = f(c_1)$, $f_2 = f(c_2)$, the six off-diagonal gap functions become

$$\delta(0,0) = f_2, \qquad \delta(1,1) = 2 - f_1, \qquad \delta(2,1) = 5 - f_2,$$
$$\delta(0,2) = 5 + f_1, \qquad \delta(1,0) = 5 - f_1 + f_2, \qquad \delta(2,2) = 4 + f_1 - f_2.$$

Requiring each to be non-negative gives the cell. Four of the six inequalities are active:

$$f_2 \geq 0, \qquad f_1 \leq 2, \qquad f_2 \leq 5, \qquad f_2 - f_1 \leq 4;$$

the remaining two ($f_1 \geq -5$ and $f_1 - f_2 \leq 5$) are redundant. The cell is therefore a quadrilateral with vertices $(-4, 0)$, $(2, 0)$, $(2, 5)$, $(1, 5)$.

The α-matrix $\alpha(c, c') = M(c, \sigma(c')) - M(c', \sigma(c'))$ is

$$\alpha = \begin{bmatrix} \cdot & 5 & 0 \\ 2 & \cdot & 5 \\ 5 & 4 & \cdot \end{bmatrix}.$$

The five directed cycle means are:

$$\mu(0 \to 1 \to 0) = \tfrac{5+2}{2} = 3.5, \qquad \mu(0 \to 2 \to 0) = \tfrac{0+5}{2} = 2.5, \qquad \mu(1 \to 2 \to 1) = \tfrac{5+4}{2} = 4.5,$$
$$\mu(0 \to 1 \to 2 \to 0) = \tfrac{5+5+5}{3} = 5, \qquad \mu(0 \to 2 \to 1 \to 0) = \tfrac{0+4+2}{3} = 2.$$

The minimum is $r^* = 2$, achieved uniquely by the 3-cycle $0 \to 2 \to 1 \to 0$. Solving the difference-constraint system $f(c) - f(c') \leq \alpha(c, c') - r^*$ via shortest-path potentials gives the Chebyshev center $f^* = (0, 0, 2)$, $g^* = (-1, 1, 1)$. The gap matrix at the center

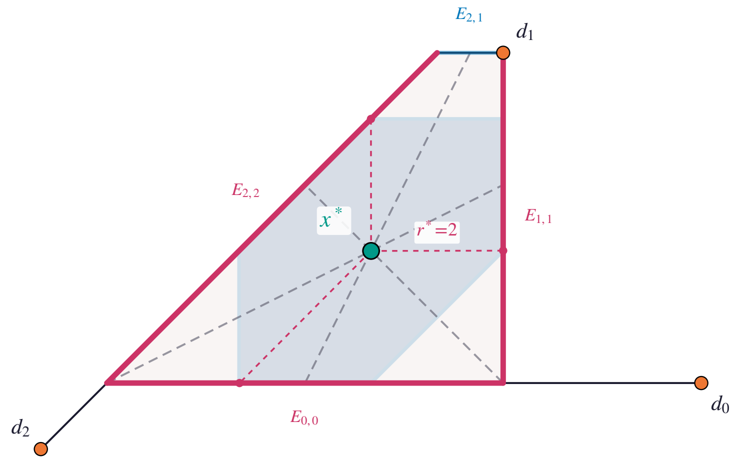

Fig. 8 The Chebyshev center of $\mathrm{Cell}(\Gamma_\sigma)$ for the 3×3 matrix M of §5.6, shown in the gauge $f(c_0) = 0$. The full-dimensional cell (shaded quadrilateral) has four walls, one for each active Isbell inequality. Three walls (bold, $E_{0,0}$, $E_{1,1}$, $E_{2,2}$) are tight at the center x^*, meaning $\delta = r^* = 2$; the fourth ($E_{2,1}$, lighter) has $\delta = 3 > r^*$. The inscribed hexagon is the spread-metric ball of radius r^*. Dashed lines from x^* to the three tangent points confirm equidistance. Three order-chamber walls through x^* partition the cell into regions of constant gap ordering. Column anchors d_0, d_1, d_2 are the images of the columns of M in $\mathbb{TP}^2$.

is

$$\delta^{(f^*,g^*)} = \begin{bmatrix} 2 & 0 & 5 \\ 7 & 2 & 0 \\ 0 & 3 & 2 \end{bmatrix},$$

whose zero entries sit at the σ-positions $(0,1)$, $(1,2)$, $(2,0)$. As predicted by part (b) of the Centering Theorem, exactly $n = 3$ off-diagonal gaps achieve the value $r^* = 2$: these are $\delta(0,0)$, $\delta(1,1)$, and $\delta(2,2)$, corresponding to the event loci $E_{0,0}$, $E_{1,1}$, and $E_{2,2}$. The center is equidistant from these three walls in the Hilbert projective metric, and the inscribed hexagonal ball of radius r^* is tangent to each (Figure 8).

6 Conclusion

A real matrix M is a coordinate presentation of an intrinsic geometric object: its Isbell nucleus. The projective nucleus $\mathbb{P}\mathrm{Nuc}(M)$ carries a Hilbert projective metric and a witness-cell decomposition, both arising directly from the enriched-categorical structure, and both invariant under the external gauge transformations that express the freedom in choosing a coordinate presentation. The gap matrix $\delta^{(f,g)}$ ties these two geometries together pointwise.

6.1 What the gap matrix organizes

For a nucleus point (f, g), the gap matrix $\delta^{(f,g)}(c, d) = M(c, d) - f(c) - g(d)$ is a nonnegative matrix that simultaneously plays three roles. Its zero pattern determines the witness cell containing (f, g). Its positive entries are the exact projective distances to the event loci where new witnesses appear (the Events Theorem). And thresholding it at successive radii extracts a tower of formal concept lattices, capturing the discrete combinatorial structure visible from (f, g) at each resolution.

The gap matrix can therefore be regarded as a local coordinate system for the polyhedral geometry, expressed in the projective metric, that at the same time organizes the lattice-theoretic shadow of the surrounding cell structure.

6.2 Discrete algebraic structure from continuous geometry

The lattice-tower construction in Section 4 extracts discrete algebraic objects—formal concept lattices with joins, meets, and Galois connections—from the continuous projective geometry of $\mathbb{P}\mathrm{Nuc}(M)$. The key point is that the correct thresholding is not performed on the matrix M directly, but on the gap matrix $\delta^{(f,g)}$ at a chosen basepoint. This pointed thresholding is geometrically meaningful: it records which event loci lie within a given projective radius. Proposition 46 makes this precise, and Proposition 47 shows that on each order chamber the real-parameter family factors through a finite tower of concept lattices with canonical floor mergers. Moving to a face of the chamber complex merges consecutive floors, producing canonical specialization maps (Proposition 48).

The order-chamber complex thus carries a constructible sheaf of lattice towers, determined by the projective geometry of $\mathbb{P}\mathrm{Nuc}(M)$, in which each lattice encodes the Galois-closed dependencies among row–column incidences visible at a given resolution.

6.3 Chebyshev centering and minimum cycle means

The Centering Theorem (Theorem 53) adds a further layer to the interaction between metric and polyhedral structure. While the Events Theorem measures how far a given point is from each cell wall, the Centering Theorem identifies the point that maximizes the minimum such distance. The optimal radius turns out to equal the minimum directed cycle mean of an edge-weighted digraph derived from the matrix, connecting the projective geometry of the nucleus to the classical theory of optimal assignment [3, 4] and to Karp's characterization of minimum cycle means [2]. At the Chebyshev center, the smallest gap e_1 is achieved with multiplicity at least n—generically exactly n—so the center is equidistant from n event loci, analogous to the incenter of a simplex. The pair sums $S(c, c') = M(c, \sigma(c')) + M(c', \sigma(c)) - M(c, \sigma(c)) - M(c', \sigma(c'))$, being intrinsic to the matrix, obstruct higher multiplicities for e_2 and explain why only e_1 generically exhibits this phenomenon.

6.4 Further directions

Several directions remain open. Beyond the finite discrete case treated here, it would be interesting to understand what replaces the witness-cell and order-chamber stratifications for general small $\overline{\mathbb{R}}$-categories, and whether the Events Theorem has an analogue in that setting.

Within the discrete setting, the chamberwise lattice towers suggest computable invariants of the nucleus that go beyond the cell structure alone. As the basepoint moves, concepts are born and die at different event radii; tracking these births and deaths produces persistence-type summaries of the lattice changes. In applications where the matrix M arises from data, the lattice tower at a point provides a structured decomposition of the data into formal concepts at varying resolutions, with the Galois connections at each level encoding principled algebraic operations (joins and meets) on the resulting clusters. Developing these invariants in examples is a natural next step.

A further direction concerns nuclei endowed with additional compatible structure. When the profunctor is compatible with monoidal data, the nucleus inherits richer operations relevant to linear realizability; see [25–27]. The interaction between these monoidal structures and the projective metric geometry of the present paper remains to be explored.

References

[1] Develin, M., Sturmfels, B.: Tropical convexity. Documenta Mathematica **9**, 1–27 (2004) https://doi.org/10.4171/DM/154

[2] Karp, R.M.: A characterization of the minimum cycle mean in a digraph. Discrete Mathematics **23**(3), 309–311 (1978) https://doi.org/10.1016/0012-365X(78)90011-0

[3] Kuhn, H.W.: The hungarian method for the assignment problem. Naval Research Logistics Quarterly **2**(1–2), 83–97 (1955) https://doi.org/10.1002/nav.3800020109

[4] Burkard, R., Dell'Amico, M., Martello, S.: Assignment Problems. Society for Industrial and Applied Mathematics, Philadelphia (2009)

[5] Maclagan, D., Sturmfels, B.: Introduction to Tropical Geometry. Graduate Studies in Mathematics, vol. 161. American Mathematical Society, Providence, RI (2015). https://doi.org/10.1090/gsm/161

[6] Gaubert, S., Katz, R.: Max-plus convex geometry. In: Relations and Kleene Algebra in Computer Science (RelMiCS 2006). Lecture Notes in Computer Science, vol. 4136, pp. 192–206. Springer, ??? (2006). https://doi.org/10.1007/11828563_13

[7] Gaubert, S., Katz, R.: Minimal half-spaces and external representation of tropical

polyhedra. Journal of Algebraic Combinatorics **33**(3), 325–348 (2011) https://doi.org/10.1007/S10801-010-0246-4

[8] Nitica†, V., Singer, I.: Max-plus convex sets and max-plus semispaces. i. Optimization **56**(1-2), 171–205 (2007) https://doi.org/10.1080/02331930600819852 https://doi.org/10.1080/02331930600819852

[9] Gaubert, S., Sergeev, S.: Cyclic projectors and separation theorems in idempotent convex geometry. Journal of Mathematical Sciences **155**(6), 815–829 (2008) https://doi.org/10.1007/s10958-008-9243-8 . Translated from Fundamentalnaya i Prikladnaya Matematika

[10] Lawvere, F.W.: Metric spaces, generalized logic, and closed categories. Rendiconti del Seminario Matematico e Fisico di Milano **43**, 135–166 (1973) https://doi.org/10.1007/BF02924844 . Reprinted in Reprints in Theory and Applications of Categories, No. 1 (2002), pp. 1–37

[11] Isbell, J.R.: Adequate subcategories. Illinois Journal of Mathematics **4**(4), 541–552 (1960) https://doi.org/10.1215/ijm/1255456274

[12] Avery, T., Leinster, T.: Isbell conjugacy and the reflexive completion. Theory and Applications of Categories **36**(12), 306–347 (2021) https://doi.org/10.70930/tac/r1jknjot

[13] Willerton, S.: Tight spans, Isbell completions and semi-tropical modules. Theory and Applications of Categories **28**(22), 696–732 (2013) https://doi.org/10.70930/tac/rkp3zgxc

[14] Willerton, S.: Galois Correspondences and Enriched Adjunctions. Blog post (2014). https://golem.ph.utexas.edu/category/2014/02/galois_correspondences_and_enr.html Accessed 2025-12-08

[15] Willerton, S.: The Legendre-Fenchel transform from a category theoretic perspective. arXiv e-prints (2015) https://doi.org/10.48550/arXiv.1501.03791 arXiv:1501.03791 [math.CT]

[16] Ambrosio, L., Gigli, N.: A user's guide to optimal transport. In: Modelling and Optimisation of Flows on Networks. Lecture Notes in Mathematics, vol. 2062, pp. 1–155. Springer, Berlin, Heidelberg (2013). https://doi.org/10.1007/978-3-642-32160-3_1

[17] Elliott, J.A.: On the fuzzy concept complex. PhD thesis, University of Sheffield, Sheffield, UK (2017). https://etheses.whiterose.ac.uk/18342/

[18] Fujii, S.: Enriched categories and tropical mathematics. arXiv e-prints (2019) https://doi.org/10.48550/arXiv.1909.07620 arXiv:1909.07620 [math.CT]

[19] Bradley, T.-D., Terilla, J., Vlassopoulos, Y.: An enriched category theory of language: From syntax to semantics. La Matematica **1**(2), 551–580 (2022) https://doi.org/10.1007/s44007-022-00021-2

[20] Gaubert, S., Vlassopoulos, Y.: Directed metric structures arising in large language models. arXiv e-prints (2024) https://doi.org/10.48550/arXiv.2405.12264 arXiv:2405.12264 [cs.LG]

[21] Bradley, T.-D., Vigneaux, J.P.: The magnitude of categories of texts enriched by language models. Theory and Applications of Categories **44**(37), 1256–1281 (2025) https://doi.org/10.48550/arXiv.2501.06662 arXiv:2501.06662 [math.CT]

[22] Bradley, T., Gastaldi, J.L., Terilla, J.: The structure of meaning in language: Parallel narratives in linear algebra and category theory. Notices of the American Mathematical Society **71**(2), 174–185 (2024) https://doi.org/10.1090/noti2868

[23] Schrijver, A.: Combinatorial Optimization: Polyhedra and Efficiency. Algorithms and Combinatorics, vol. 24. Springer, Berlin (2003)

[24] Akian, M., Gaubert, S., Qi, Y., Saadi, O.: Tropical linear regression and mean payoff games: or, how to measure the distance to equilibria. SIAM Journal on Discrete Mathematics **37**(2), 632–674 (2023) https://doi.org/10.1137/21M1428297

[25] Jarvis, S.K.: A novel closed monoidal structure on the nucleus of a profunctor. PhD thesis, The Graduate Center, City University of New York (June 2025). Doctoral dissertation (Ph.D.), Mathematics; advisor: John Terilla. https://academicworks.cuny.edu/gc_etds/6231

[26] Seiller, T.: Mathematical Informatics. PhD thesis, Sorbonne Paris Nord University (2024). Habilitation thesis. https://theses.hal.science/tel-04616661

[27] Gastaldi, J.L., Jarvis, S., Seiller, T., Terilla, J.: Linear realizability and structures in ℝ-enriched adjunctions. Preprint, available from the authors (2025)

[28] Kelly, G.M.: Basic Concepts of Enriched Category Theory. London Mathematical Society Lecture Note Series, vol. 64. Cambridge University Press, ??? (1982). Reprinted in: Reprints in Theory and Applications of Categories, No. 10, 2005

[29] Aguiar, M., Mahajan, S.: Topics in Hyperplane Arrangements. Mathematical Surveys and Monographs, vol. 226. American Mathematical Society, Providence, RI (2017)

[30] Wille, R.: Restructuring lattice theory: An approach based on hierarchies of concepts. In: Rival, I. (ed.) Ordered Sets: Proceedings of the NATO Advanced Study Institute Held at Banff, Canada, August 28 to September 12, 1981. NATO Advanced Study Institutes Series, vol. 83, pp. 445–470. Springer, Dordrecht (1982).

https://doi.org/10.1007/978-94-009-7798-3_15