From Isospectrality to Geometry: Sunada Transplantation, Hidden Point Structures, and the Convex Drum Problem

Xiyu Hu

August 3, 2026

Abstract

This note develops a conceptual separation between three levels of inverse spectral geometry: the spectrum of an abstract self-adjoint operator, the full heat or Green kernel after a point space has been identified, and the Euclidean geometry of a planar domain. The Sunada–Gassmann mechanism is reviewed as an algebraic construction of two non-isomorphic finite G-sets whose complex permutation modules are unitarily equivalent. In matrix language, the two families of permutation matrices are simultaneously similar by a general unitary matrix but not by a common permutation matrix. We explain how double cosets produce transplantation operators and how Schreier graphs are realized by gluing congruent tiles into planar domains.

The classical planar construction is then compared with the convex Euclidean problem. Its difficulty is not exhausted by the failure of a particular tiling to be convex. More fundamentally, the spectrum determines the Dirichlet Laplacian only as an abstract operator on a Hilbert space. It does not identify the distinguished multiplication algebra $L^\infty(\Omega)$, hence it does not identify points, supports, positivity, or locality. Once the full pointwise heat kernel is known, the intrinsic distance follows from short-time asymptotics; the missing step is the passage from trace data to a pointwise realization.

We formulate an exact operator-algebraic reformulation: for an abstract operator with a convex planar Dirichlet realization, the convex drum problem is equivalent to uniqueness, up to the commutant of the operator, of the admissible multiplication algebra representing the underlying point space. We also analyze the idea of extending distance from one point. One-point radial distance data has infinite-dimensional angular relabeling freedom, even among smooth strictly convex pointed domains. By contrast, once angular directions are identified, Euclidean flatness forces the polar metric and leaves only a boundary radial function. The note does not prove the convex drum conjecture; it isolates a precise rigidity statement whose proof would settle it.

Contents

1 The inverse spectral question and its three levels

Let $\Omega \subset \mathbb{R}^2$ be a bounded connected domain. The Dirichlet Laplacian is the nonnegative self-adjoint operator

$$A_\Omega = -\Delta_\Omega^D$$

associated with the quadratic form

$$\mathcal{E}_\Omega(u) = \int_\Omega |\nabla u|^2 \, dx, \qquad u \in H_0^1(\Omega).$$

For sufficiently regular bounded domains, its spectrum is discrete:

$$0 < \lambda_1(\Omega) \le \lambda_2(\Omega) \le \cdots \nearrow \infty,$$

where eigenvalues are repeated according to multiplicity. Kac's question asks how much of Ω is determined by this sequence [5]. General planar domains are not spectrally rigid: Gordon, Webb, and Wolpert produced noncongruent isospectral planar domains [4], and Buser–Conway–Doyle–Semmler gave explicit transplantation constructions [2]. The restriction to convex Euclidean planar domains remains the central unresolved version considered here.

It is useful to distinguish three objects.

1. The *abstract spectral operator* $(\mathcal{H}, A)$, determined up to unitary equivalence by the eigenvalues and their multiplicities.
2. The *pointwise diffusion kernel*

$$K_\Omega(t; x, y) = \sum_{j=1}^{\infty} e^{-t\lambda_j} \phi_j(x)\overline{\phi_j(y)},$$

   where the points $x, y \in \Omega$ and the eigenfunction evaluations have already been identified.
3. The *metric-measure space* (Ω, d_Ω, dx) and, in the convex Euclidean case, the embedded set $\Omega \subset \mathbb{R}^2$ up to a rigid motion.

The difficult arrow is not the last one:

$$\{\lambda_j\}_{j\ge1} \longrightarrow (\mathcal{H}, A) \xrightarrow{\text{recover points/locality}} K(t; x, y) \xrightarrow{t\downarrow 0} d(x, y) \longrightarrow \Omega.$$

The spectrum gives the first arrow. Short-time heat-kernel asymptotics give the third. The dashed arrow is the essential inverse problem.

2 What equality of spectra actually gives

Suppose Ω_1 and Ω_2 are isospectral, including multiplicities. Choose orthonormal eigenbases $\{\phi_j^{(i)}\}$ and define

$$U\phi_j^{(1)} = \phi_j^{(2)}.$$

After making arbitrary unitary choices inside repeated eigenspaces, U is a unitary map

$$U : L^2(\Omega_1) \longrightarrow L^2(\Omega_2)$$

satisfying

$$UA_{\Omega_1} = A_{\Omega_2}U.$$

Consequently,

$$Ue^{-tA_{\Omega_1}} = e^{-tA_{\Omega_2}}U, \qquad U(A_{\Omega_1} + z)^{-1} = (A_{\Omega_2} + z)^{-1}U.$$

Thus isospectrality already gives equality of the abstract diffusion operators up to a unitary change of coordinates.

This does *not* identify the pointwise kernels. An arbitrary unitary map can mix values from many spatial points. In a schematic integral form,

$$(Uf)(y) = \int_{\Omega_1} T(y, x)f(x)\,dx,$$

whereas a geometric change of variables has the much more rigid form

$$(U_F f)(y) = h(y) f(F^{-1}(y)).$$

The latter preserves supports, disjointness, and positivity after a harmless phase normalization; the former need not preserve any of them.

The heat trace is

$$Z_\Omega(t) = \mathrm{Tr}(e^{-tA_\Omega}) = \sum_j e^{-t\lambda_j} = \int_\Omega K_\Omega(t; x, x)\, dx.$$

It is a single integrated statistic of the diagonal. It is not the two-point kernel.

2.1 The spectral embedding viewpoint

For each $t > 0$, define the heat-kernel embedding

$$\Phi_t : \Omega \longrightarrow \ell^2, \qquad \Phi_t(x) = \big(e^{-t\lambda_j/2} \phi_j(x)\big)_{j \geq 1}.$$

Then

$$K_\Omega(t; x, y) = \langle \Phi_t(x), \Phi_t(y) \rangle_{\ell^2}.$$

Hence the full heat kernel is the Gram matrix of the embedded point set $\Phi_t(\Omega)$. By contrast,

$$Z_\Omega(t) = \int_\Omega \|\Phi_t(x)\|_{\ell^2}^2\, dx$$

only records an integrated second moment. The abstract spectrum fixes the axis weights $e^{-t\lambda_j/2}$ but does not identify which subset of the ambient Hilbert space should be regarded as the embedded collection of spatial points.

This provides a geometric summary of the information loss:

$$\boxed{\text{The spectrum fixes an abstract ellipsoid of modes, not the embedded point cloud.}}$$

3 The Sunada–Gassmann mechanism

3.1 Two different G-sets with the same linearization

Let G be a finite group and $H < G$ a subgroup. The left coset space

$$X_H = G/H$$

is a transitive G-set under $g \cdot aH = gaH$. Its complex permutation module is

$$\mathbb{C}[X_H] = \bigoplus_{x \in X_H} \mathbb{C} e_x,$$

with representation $\rho_H(g) e_x = e_{gx}$.

Definition 3.1. Subgroups $H_1, H_2 < G$ are *almost conjugate*, or *Gassmann equivalent*, if

$$|H_1 \cap C| = |H_2 \cap C|$$

for every conjugacy class C of G.

For $g \in G$, the character of the permutation representation is

$$\chi_H(g) = \operatorname{Tr} \rho_H(g) = \# \operatorname{Fix}_{G/H}(g).$$

A coset aH is fixed by g precisely when $a^{-1}ga \in H$, and one obtains

$$\chi_H(g) = \frac{|C_G(g)|}{|H|} \, |[g] \cap H|.$$

Therefore almost conjugacy is equivalent to

$$\chi_{H_1} = \chi_{H_2}.$$

Since finite-dimensional complex representations of finite groups are completely reducible, equality of characters is equivalent to representation isomorphism.

Proposition 3.2 (Linear equivalence without set equivalence)**.** *If H_1, H_2 are almost conjugate, then there exists a unitary map*

$$T : \mathbb{C}[G/H_1] \longrightarrow \mathbb{C}[G/H_2]$$

such that

$$T\rho_{H_1}(g) = \rho_{H_2}(g)T \qquad \text{for every } g \in G.$$

If H_1 and H_2 are not conjugate, there is no such intertwiner that is a permutation matrix.

Proof. Equality of characters gives an isomorphism of the two representations. Averaging an inner product over G and applying polar decomposition makes the intertwiner unitary. A permutation intertwiner would induce a G-equivariant bijection $G/H_1 \to G/H_2$. For transitive G-sets this is equivalent to conjugacy of the stabilizers, hence to conjugacy of H_1 and H_2. □

Thus the central algebraic distinction is

$$\boxed{\begin{array}{l}\text{one common general unitary change of basis exists;}\\ \text{one common relabeling of vertices does not exist.}\end{array}}$$

3.2 Simultaneous similarity, not similarity of one matrix

Choose generators $S = \{s_1, \ldots, s_r\}$ of G. The two coset actions are encoded by matrix tuples

$$(P_1(s_1), \ldots, P_1(s_r)), \qquad (P_2(s_1), \ldots, P_2(s_r)).$$

A G-set isomorphism requires one permutation matrix Q satisfying

$$QP_1(s_j)Q^{-1} = P_2(s_j) \qquad (1 \le j \le r).$$

A representation isomorphism requires one general invertible matrix T satisfying the same equations. Individual conjugacies

$$Q_j P_1(s_j) Q_j^{-1} = P_2(s_j)$$

with different Q_j are irrelevant: the interaction of the generators is exactly the compatibility requirement that one and the same coordinate change work for the entire tuple.

For a fixed $g \in G$, almost conjugacy implies that $P_1(g)$ and $P_2(g)$ have the same cycle type. Indeed, equality of $\operatorname{Tr}(P_i(g^k))$ for every k determines the number of cycles of every length. Thus each group element can be matched separately by some permutation. What fails is a single matching compatible with all elements at once.

3.3 Why ordinary traces cannot detect the missing set structure

For every word $w = s_{i_1} \cdots s_{i_k}$,

$$\operatorname{Tr}\big(P_i(s_{i_1}) \cdots P_i(s_{i_k})\big) = \operatorname{Tr} P_i(w).$$

Sunada equivalence makes these traces identical for the two representations. No invariant built solely from traces of ordinary matrix products can separate them as complex G-modules, because they are genuinely isomorphic as such modules.

The missing information is set-theoretic and is measured by simultaneous fixed-point data. For a subgroup $L < G$, define the *mark*

$$m_L(X) = |X^L| = \#\{x \in X : \ell x = x \text{ for all } \ell \in L\}.$$

For $L = \langle g \rangle$, this reduces to the character value $\chi_X(g)$. But for a noncyclic subgroup $L = \langle g, h \rangle$, it measures the intersection

$$|\operatorname{Fix}(g) \cap \operatorname{Fix}(h)|,$$

which need not follow from the individual cycle structures of g and h.

The full collection of marks, indexed by conjugacy classes of subgroups, determines a finite G-set up to isomorphism. Equivalently, the table of marks is the correct analogue of a complete standard form for permutation actions. The ordinary character sees only the rows corresponding to cyclic subgroups.

Example 3.3 (Equal characters but different interactions)**.** Let $G = V_4 = \{1, a, b, c\}$ and let $H_a = \langle a \rangle$, $H_b = \langle b \rangle$, $H_c = \langle c \rangle$. Consider

$$X = G/\{1\} \sqcup 2(G/G), \qquad Y = G/H_a \sqcup G/H_b \sqcup G/H_c.$$

Both have six points and the same permutation character: every nonidentity element fixes exactly two points. However,

$$|X^{\langle a,b \rangle}| = 2, \qquad |Y^{\langle a,b \rangle}| = 0.$$

Thus $\mathbb{C}[X] \cong \mathbb{C}[Y]$ as G-modules, but $X \ncong Y$ as G-sets.

In matrix coordinates, a useful basis-dependent interaction invariant is

$$\operatorname{Tr}\big(P(g) \circ P(h)\big) = |\operatorname{Fix}(g) \cap \operatorname{Fix}(h)|,$$

where $\circ$ is the Hadamard product. This is invariant under simultaneous permutation conjugacy, but not under arbitrary linear conjugacy. Its dependence on the standard coordinate basis is precisely the point.

3.4 Burnside ring versus representation ring

Finite G-sets form the Burnside ring $B(G)$, while finite-dimensional complex representations form the representation ring $R_{\mathbb{C}}(G)$. Linearization gives a natural homomorphism

$$\mathcal{L} : B(G) \longrightarrow R_{\mathbb{C}}(G), \qquad [X] \longmapsto [\mathbb{C}[X]].$$

A genuine Sunada pair satisfies

$$[G/H_1] - [G/H_2] \neq 0 \quad \text{in } B(G),$$

but

$$\mathcal{L}([G/H_1] - [G/H_2]) = 0 \quad \text{in } R_{\mathbb{C}}(G).$$

The construction exploits the kernel of linearization: distinct point configurations become identical after one forgets the distinguished idempotent basis indexed by points.

4 Double cosets and transplantation operators

All G-intertwiners between two coset modules admit a concrete description. Let

$$X_1 = G/H_1, \qquad X_2 = G/H_2.$$

Write the matrix coefficient of $T : \mathbb{C}[X_1] \to \mathbb{C}[X_2]$ as T_{bH_2,aH_1}. Equivariance gives

$$T_{gbH_2,gaH_1} = T_{bH_2,aH_1}.$$

Hence the coefficient depends only on the diagonal G-orbit of the pair (bH_2, aH_1), equivalently on the double coset

$$H_2 b^{-1} a H_1 \in H_2 \backslash G / H_1.$$

For each double coset D, define a relation matrix

$$(A_D)_{bH_2,aH_1} = \begin{cases} 1, & H_2 b^{-1} a H_1 = D, \\ 0, & \text{otherwise.} \end{cases}$$

Then

$$\mathrm{Hom}_G(\mathbb{C}[G/H_1], \mathbb{C}[G/H_2]) = \mathrm{span}\{A_D : D \in H_2 \backslash G / H_1\}.$$

A transplantation matrix has the form

$$T = \sum_D c_D A_D.$$

Finding a Sunada transplantation amounts to finding coefficients c_D for which T is invertible. Thus double cosets classify the possible relative positions of a point in one coset geometry and a point in the other.

4.1 The Fano point-line incidence map

In the classical seven-tile example, the group $G \cong PSL(3, 2)$ acts transitively on the seven points and on the seven lines of the Fano plane. Let H_P be a point stabilizer and H_L a line stabilizer. The two relative positions are

$$p \in \ell, \qquad p \notin \ell.$$

The incidence matrix

$$N_{\ell p} = \mathbf{1}_{\{p \in \ell\}}$$

is a G-intertwiner. Each point lies on three lines, and two distinct lines meet at one point, so

$$NN^* = 2I + J,$$

where J is the all-ones matrix. This is invertible, hence N gives a linear equivalence between the point and line permutation modules. But

$$Ne_p = \sum_{\ell \ni p} e_\ell$$

is a sum of three basis vectors, not one basis vector. The map is a transplantation, not a relabeling of points.

Its polar part

$$U = (NN^*)^{-1/2} N$$

is a unitary intertwiner. Any order-preserving unitary intertwiner would have to send disjoint positive coordinate atoms to disjoint positive coordinate atoms, hence would be a permutation matrix. Therefore the genuine Sunada intertwiner necessarily loses the pointwise order structure.

5 From abstract group data to planar domains

5.1 The covering and orbifold principle

Sunada's original geometric setting starts with a Riemannian manifold M on which G acts by isometries. Almost conjugate subgroups $H_1, H_2 < G$ give quotients

$$M_1 = H_1 \backslash M, \qquad M_2 = H_2 \backslash M,$$

which are isospectral under suitable hypotheses [6]. The representation-theoretic reason is that, for every G-representation V,

$$V^{H_i} \cong \mathrm{Hom}_G(\mathbb{C}[G/H_i], V),$$

and the two fixed subspaces have equal dimensions and are naturally intertwined. Taking V to be each eigenspace of the Laplacian gives equality of multiplicities.

An ordinary connected covering of a simply connected planar domain is trivial. Therefore simply connected planar counterexamples cannot arise directly from a nontrivial ordinary covering of the final domain. Gordon–Webb–Wolpert pass through orbifolds with mirror boundaries: the underlying topological space may be a disk even though its orbifold fundamental group is nontrivial [4].

5.2 Schreier graphs as gluing instructions

Choose a fundamental tile B, typically a triangle, whose sides carry labels $s_1, \ldots, s_r$. Suppose the corresponding generators act as involutions on the coset spaces. For each $x \in X_i = G/H_i$, take a congruent copy B_x. The rule

$$x \longmapsto s_j x$$

means that the side of type s_j in B_x is glued to the side of type s_j in $B_{s_j x}$ by reflection. The colored Schreier graph is the dual graph of the tiled surface:

Schreier datum	geometric datum
$x \in G/H_i$	one copy B_x
$x \xleftrightarrow{s_j} s_j x$	gluing along an s_j-edge.

An abstract gluing first produces a flat polygonal complex or orbifold. To become a genuine embedded domain in $\mathbb{R}^2$, it must satisfy additional conditions:

- the total angle around every interior vertex must be 2π;
- the developing map into $\mathbb{R}^2$ must not overlap tiles;
- the unpaired sides must form an embedded simple boundary curve.

These conditions are geometric and are not guaranteed by the group triple.

5.3 Transplanting eigenfunctions

Let f be an eigenfunction on the first tiled domain. Pull the restriction of f on each tile back to the common model tile B:

$$f \longleftrightarrow (f_x)_{x \in X_1}.$$

Given a transplantation matrix $T = (T_{y,x})$, define on the second domain

$$(\mathcal{T}f)_y(z) = \sum_{x \in X_1} T_{y,x} f_x(z), \qquad z \in B.$$

Because the Laplacian is the same on every congruent tile,

$$-\Delta(\mathcal{T}f)_y = \sum_x T_{y,x}(-\Delta f_x).$$

If $-\Delta f = \lambda f$, then every component satisfies the same equation with eigenvalue λ. The identities

$$TP_1(s_j) = P_2(s_j)T$$

ensure that edge continuation relations are preserved. For Dirichlet boundary conditions one uses odd reflection; for Neumann conditions one uses even reflection. Invertibility of T gives an isomorphism between every pair of eigenspaces, proving isospectrality. This is the explicit mechanism in [2].

6 Why Euclidean convexity is a severe extra constraint

The failure of the classical examples to be convex is not merely cosmetic. Several independent constraints collide.

6.1 No nontrivial ordinary covering structure

A convex planar domain is contractible. Consequently, any ordinary connected covering of it is trivial. A planar Sunada construction must therefore use an orbifold or reflection mechanism with fixed sets. This naturally produces mirror edges and polygonal corners rather than a free covering of a smooth convex disk.

6.2 The Euclidean angle budget

Let the model triangle have angles α, β, γ. In the Euclidean plane,

$$\alpha + \beta + \gamma = \pi.$$

Every interior vertex of the glued complex imposes an equality saying that a certain integer combination of these angles is 2π. Every boundary vertex of a convex realization imposes an inequality saying that another combination is at most π. Both companion gluings must satisfy all such constraints simultaneously.

In the hyperbolic plane the corresponding relation is

$$\alpha + \beta + \gamma < \pi,$$

which provides an angle deficit. Gordon and Webb used this additional flexibility to construct convex hyperbolic isospectral domains [3]. The same kind of combinatorial branching that creates reflex angles in the Euclidean development can be absorbed by shrinking all basic angles in the hyperbolic setting.

6.3 Global embedding and support-line constraints

Local inequalities at vertices do not by themselves imply global convexity. Each boundary edge must be a support edge: the entire domain must lie in one closed half-plane determined by its line. Moreover, the two developments must be nonoverlapping and their boundaries must be simple. These are global constraints on long reflection words, not consequences of almost conjugacy.

The group-theoretic condition controls the linearized adjacency actions. It does not control the cyclic order of boundary edges, support lines, or the placement of all vertices in convex position. This is the first reason that a successful algebraic triple need not admit a convex Euclidean realization.

6.4 Why smoothing is not harmless

Reflection constructions naturally produce polygonal boundaries because Euclidean reflections fix straight lines. Rounding a corner changes the local Laplacian and destroys the exact reflection continuation used in transplantation. Thus one cannot first construct an isospectral nonconvex polygon and then smooth or convexify it while expecting exact isospectrality to survive.

These observations explain why standard Sunada tilings meet convexity resistance, but they are not a theorem that every conceivable Sunada-type method must fail. The deeper inverse problem appears when one asks what the spectrum itself is missing.

7 Heat kernels, Green kernels, and distance

7.1 The full heat kernel recovers distance

For an appropriate Riemannian domain, the short-time heat kernel satisfies the Varadhan asymptotic

$$d_\Omega(x,y)^2 = -\lim_{t\downarrow 0} 4t \log K_\Omega(t;x,y),$$

away from the usual technical exceptional cases and with the intrinsic distance associated with the diffusion [7]. In a convex Euclidean domain, the straight segment between any two points lies in the domain, hence

$$d_\Omega(x,y) = |x-y|.$$

Therefore a point correspondence $F : \Omega_1 \to \Omega_2$ satisfying

$$K_{\Omega_1}(t;x,y) = K_{\Omega_2}(t;F(x),F(y))$$

for all t, x, y preserves Euclidean distances and must arise from a rigid motion.

The obstacle is that the spectrum gives

$$\sum_j e^{-t\lambda_j},$$

not the values

$$\sum_j e^{-t\lambda_j}\phi_j(x)\overline{\phi_j(y)}.$$

To form the latter, one must first know what the variables x and y mean inside the abstract Hilbert space.

7.2 Why the two-dimensional Green function requires caution

The zero-energy Dirichlet Green kernel is

$$G_\Omega(x,y) = \int_0^\infty K_\Omega(t;x,y)\,dt = \sum_j \frac{1}{\lambda_j}\phi_j(x)\overline{\phi_j(y)}.$$

In dimension two, the scalar Green function is conformally invariant: if $F:\Omega_1\to\Omega_2$ is conformal, then

$$G_{\Omega_2}(F(x),F(y)) = G_{\Omega_1}(x,y).$$

Thus the Green function viewed only as an abstract two-point scalar is closer to conformal structure than to Euclidean metric structure. The full heat kernel keeps the time scale and reveals metric distance through its exponential small-time behavior. Integrating over all times removes precisely this scale separation.

8 The hidden multiplication algebra

8.1 The geometric triple

Besides the Hilbert space and Laplacian, a domain carries the multiplication algebra

$$\mathcal{M}_\Omega = \{M_f : f\in L^\infty(\Omega)\}\subset B(L^2(\Omega)), \qquad M_f u = fu.$$

This is a maximal abelian von Neumann algebra. It encodes:

- measurable sets through projections $M_{\mathbf{1}_E}$;
- support and disjointness through products of projections;
- the positive cone and order structure;
- the interpretation of functions as pointwise objects.

The more faithful spectral-geometric object is therefore

$$\boxed{(L^2(\Omega),\mathcal{M}_\Omega,A_\Omega),}$$

not merely $(L^2(\Omega),A_\Omega)$.

Sunada transplantation preserves the second component A up to unitary conjugacy but generally does not preserve the multiplication algebra. A basis vector representing one tile or point is sent to a signed linear combination of several basis vectors. This is the finite-dimensional prototype of losing locality.

8.2 An exact equivalence criterion

Let an abstract positive self-adjoint operator A on a separable Hilbert space $\mathcal{H}$ have two Dirichlet realizations. Choose unitaries

$$W_i:\mathcal{H}\longrightarrow L^2(\Omega_i), \qquad W_iAW_i^{-1} = A_{\Omega_i},$$

and pull the multiplication algebras back to $\mathcal{H}$:

$$\mathfrak{M}_i = W_i^{-1}\mathcal{M}_{\Omega_i}W_i.$$

Theorem 8.1 (Multiplication-algebra criterion)**.** *Let $\Omega_1, \Omega_2 \subset \mathbb{R}^2$ be bounded connected domains regular enough for the order-rigidity theorem for Dirichlet heat semigroups. Then Ω_1 and Ω_2 are congruent if and only if there is a unitary $V \in B(\mathcal{H})$ such that*

$$VA = AV$$

and

$$V\mathfrak{M}_1 V^{-1} = \mathfrak{M}_2.$$

Proof. If the domains are congruent, pullback by a rigid motion gives the required unitary.

Conversely, set

$$T = W_2 V W_1^{-1} : L^2(\Omega_1) \to L^2(\Omega_2).$$

Then T intertwines the Dirichlet Laplacians and normalizes the multiplication algebras. A unitary normalizer of an L^∞ multiplication algebra is a weighted composition operator associated with a measure-class isomorphism. The first Dirichlet eigenvalue is simple and its eigenfunction can be chosen strictly positive. Since T maps the first eigenspace to the first eigenspace, multiplication by a constant phase converts T into an order isomorphism. The theorem of Arendt–Biegert–ter Elst then implies that the underlying Riemannian domains are isomorphic [1]. An isometry between connected Euclidean open subsets is the restriction of a Euclidean rigid motion. □

This theorem identifies the precise missing datum: an isospectral unitary equivalence is geometric exactly when it also transports the multiplication algebra.

8.3 The convex drum conjecture as uniqueness of an admissible MASA

Fix an abstract operator $(\mathcal{H}, A)$ with at least one realization as the Dirichlet Laplacian of a bounded convex planar domain. Call a maximal abelian algebra $\mathfrak{M} \subset B(\mathcal{H})$ *convex-planar admissible* if, under some identification of $(\mathcal{H}, \mathfrak{M})$ with an L^2 function space, A becomes the Dirichlet Laplacian of a bounded convex Euclidean planar domain.

Conjecture 8.2 (Convex multiplication-algebra uniqueness)**.** *Any two convex-planar admissible multiplication algebras $\mathfrak{M}_1, \mathfrak{M}_2$ for the same abstract operator A are conjugate by a unitary in the commutant*

$$\{A\}' = \{V \in B(\mathcal{H}) : VA = AV\}.$$

By the preceding theorem, this is an exact reformulation of the convex drum conjecture. It is not a weaker auxiliary statement: proving it would prove spectral rigidity of convex planar domains, and a counterexample would produce an isospectral noncongruent convex pair.

The advantage of the reformulation is conceptual. It says that the problem is not primarily to recover a metric tensor after the point space is known. It is to prove that flatness, two-dimensionality, strong locality, Dirichlet boundary behavior, and convexity select a unique pointwise multiplication structure inside an abstract spectral Hilbert space.

9 Why one-point distance data is insufficient

The idea of choosing one point and extending its distance function outward is natural, but it contains an infinite-dimensional ambiguity: a radial distance does not label directions.

9.1 A family of pointed strictly convex counterexamples

Let $R : S^1 \to (0, \infty)$ be smooth and define the star-shaped domain

$$\Omega_R = \{re^{i\theta} : 0 \leq r < R(\theta)\}.$$

The polar curve $r = R(\theta)$ is strictly convex when

$$R^2 + 2(R')^2 - RR'' > 0.$$

This condition is open in the C^2 topology and holds for $R \equiv 1$.

Choose a nonconstant R_1 sufficiently close to 1 in C^2, and choose an orientation-preserving circle diffeomorphism ψ sufficiently close to the identity but not of the form $\theta \mapsto \pm\theta + c$. Define R_2 by

$$R_2(\psi(\theta)) = R_1(\theta).$$

Both Ω_{R_1} and Ω_{R_2} are smooth strictly convex for sufficiently small perturbations. The map

$$F(re^{i\theta}) = re^{i\psi(\theta)}$$

is a pointed diffeomorphism satisfying

$$|F(x)| = |x|.$$

Thus it preserves the distance from the distinguished origin:

$$d_{\Omega_{R_2}}(0, F(x)) = d_{\Omega_{R_1}}(0, x).$$

But for two nonzero points,

$$|re^{i\theta} - se^{i\varphi}|^2 = r^2 + s^2 - 2rs\cos(\theta - \varphi),$$

whereas after applying F the angular difference becomes

$$\psi(\theta) - \psi(\varphi).$$

For a nonlinear ψ these distances are not preserved. The pointed domains are not congruent, since a pointed Euclidean isometry acts on the angular variable only by a rotation or reflection.

Hence one-point radial distance data admits at least the infinite-dimensional freedom

$$\mathrm{Diff}(S^1)/O(2).$$

The missing datum is the angular identification between different radial geodesics.

9.2 Once directions are known, flat extension is rigid

The situation changes after the tangent circle of directions at the base point has been identified. In geodesic polar coordinates around a point p, a two-dimensional Riemannian metric takes the form

$$g = dr^2 + J(r, \theta)^2 d\theta^2.$$

The transverse Jacobi field satisfies

$$J_{rr} + K(r, \theta)J = 0,$$

with initial conditions

$$J(0,\theta)=0, \qquad J_r(0,\theta)=1.$$

For a flat metric, $K=0$, and uniqueness of the ordinary differential equation gives

$$J(r,\theta)=r.$$

Therefore

$$g=dr^2+r^2d\theta^2.$$

In a convex Euclidean domain there is no interior cut locus from p: every ray continues until it reaches the boundary. Once the angular labels are known, the only remaining shape information is the stopping radius

$$R_p(\theta)=\sup\{r\geq 0 : p+r\theta\in\Omega\}.$$

Then

$$\Omega=\{p+r\theta : \theta\in S^1,\ 0\leq r<R_p(\theta)\}.$$

Thus the metric-extension step itself is rigid after the direction structure is recovered. The genuine freedom lies in recovering the angular organization of points and the boundary stopping function from spectral data.

10 Why positivity and locality cannot be read off immediately

One might hope that the positivity of the heat semigroup singles out the correct pointwise order. Abstractly, however, if $C\subset\mathcal{H}$ is a positive cone for which e^{-tA} is positive and $V\in\{A\}'$ is unitary, then

$$VC$$

is another cone with the same property. The spectrum alone cannot distinguish these transported cones.

Similarly, the one-dimensional first eigenspace gives a preferred line, and in a geometric realization one of its two rays contains the positive ground state. A single positive ray does not reconstruct the lattice operations, disjoint supports, or the projections corresponding to measurable sets.

The carré du champ illustrates the circularity. Once multiplication is known, one may write formally

$$\Gamma(f,g)=\frac{1}{2}\big(A(fg)-fAg-gAf\big),$$

which recovers the first-order metric structure of a diffusion. But the formula already uses the product fg. The operator A by itself does not specify which bilinear product on a dense subspace is the pointwise product.

The same observation appears in an eigenbasis. If, in addition to the eigenvalues, one knows all triple products

$$c_{ijk}=\int_\Omega \phi_i\phi_j\overline{\phi_k}\,dx,$$

then one knows the multiplication table

$$\phi_i\phi_j=\sum_k c_{ijk}\phi_k$$

in the appropriate weak sense. The spectrum supplies the linear operator; the triple products supply the nonlinear function algebra. This is exactly the kind of information absent from a bare eigenvalue list.

11 A precise research bottleneck

The preceding discussion reduces the problem to an intrinsic characterization question.

Problem 11.1 (Intrinsic characterization of the geometric algebra)**.** Let A be an abstract positive self-adjoint operator with compact resolvent. Characterize, using only operator-theoretic conditions involving A and a candidate maximal abelian algebra $\mathfrak{M}$, when

$$(\mathcal{H}, \mathfrak{M}, A)$$

is the Dirichlet Laplacian of a bounded convex domain in $\mathbb{R}^2$.

A successful characterization would need to encode at least:

- regularity and strong locality of the associated Dirichlet form;
- topological dimension two;
- vanishing interior curvature;
- the existence and geometry of the Dirichlet boundary;
- global convexity of the resulting Euclidean development.

The convex drum conjecture then asks whether this admissibility problem has a unique solution modulo the commutant of A.

11.1 Three tempting approaches and their obstruction

Localizing the heat trace. The trace gives $\mathrm{Tr}(e^{-tA})$. To ask for heat localized in a set E, one would use

$$\mathrm{Tr}(M_{\mathbf{1}_E} e^{-tA}).$$

But the projection $M_{\mathbf{1}_E}$ belongs to the unknown multiplication algebra. Localization is exactly what has not yet been recovered.

Recovering a billiard distance structure from the wave trace. The wave trace records singularities associated with periodic trajectories, but it is again an integrated trace. It does not label initial points and directions, and distinct contributions may interfere. Even complete length data would still require an inverse realization step from unlabelled trajectories to the boundary.

Selecting the positive cone by the ground state. The positive ground state identifies one vector, not the entire lattice cone. Commutant unitaries can preserve the operator while transporting the rest of the order structure.

These are not proofs of impossibility. They clarify why a direct passage from trace formulas to distance functions repeatedly encounters the same missing localization structure.

12 A finite-dimensional analogy for the core difficulty

The Sunada mechanism gives a useful exact analogy.

Finite G-set problem	Planar spectral problem	Lost structure
Permutation module $\mathbb{C}[X]$	Abstract Hilbert operator (L^2, A)	distinguished atoms / points
Character $\chi_X(g)$	Heat or wave trace	integrated linear statistics
General unitary intertwiner	Isospectral unitary intertwiner	may mix spatial locations
Permutation intertwiner	Weighted composition / geometric pullback	preserves locality and order
Marks $\lvert X^L \rvert$	Localized multi-point data	joint incidence structure
Pointwise algebra $\mathbb{C}^X$	Multiplication algebra $L^\infty(\Omega)$	nonlinear product and supports

In both settings, linearization forgets a preferred commutative algebra of coordinatewise multiplication. Sunada finds two inequivalent set realizations of the same linear representation. A convex planar counterexample would amount to two inequivalent convex Euclidean point realizations of the same abstract Laplacian. A proof of rigidity must show that convex Euclidean admissibility eliminates the analogue of the kernel of the Burnside-to-representation linearization map.

13 Conclusions

The discussion can be summarized in five statements.

1. Isospectrality determines the Dirichlet Laplacian as an abstract self-adjoint operator up to unitary equivalence. It does not determine its pointwise kernel representation.
2. Sunada's method constructs non-isomorphic coset G-sets whose permutation modules are unitarily equivalent. Algebraically, all generator matrices are simultaneously similar by one general unitary matrix, but not by one permutation matrix. Double cosets generate the possible transplantation operators.
3. The passage to planar domains uses reflection orbifolds and Schreier-graph tile gluings. Euclidean convexity imposes additional local angle, global embedding, support-line, and smoothness constraints that the group-theoretic data does not control.
4. Once a full heat kernel is known as a function of two points, short-time asymptotics recover distance. The spectral trace does not contain point labels. The missing object is the multiplication algebra $L^\infty(\Omega)$, which encodes points, sets, support, positivity, and locality.
5. One-point distance data has infinite-dimensional angular relabeling freedom. After directions are identified, flatness rigidly determines the polar metric and leaves only the radial boundary function. Thus the deepest problem is not metric continuation in a known coordinate structure, but recovery of that coordinate and multiplication structure from an abstract spectrum.

The resulting core conjecture is:

> For a fixed abstract Dirichlet spectrum, convexity and Euclidean flatness uniquely determine the admissible multiplication algebra, up to a unitary commuting with the Laplacian.

This statement is equivalent to the convex planar drum conjecture. It gives a concrete target for further work: characterize the geometric maximal abelian algebra intrinsically, rather than trying to extract a distance function directly from trace data before the point space has been recovered.

References

[1] W. Arendt, M. Biegert, and A. F. M. ter Elst, *Diffusion determines the manifold*, J. Reine Angew. Math. **667** (2012), 1–25.

[2] P. Buser, J. Conway, P. Doyle, and K.-D. Semmler, *Some planar isospectral domains*, Int. Math. Res. Notices **1994** (1994), no. 9, 391–400.

[3] C. S. Gordon and D. L. Webb, *Isospectral convex domains in the hyperbolic plane*, Proc. Amer. Math. Soc. **120** (1994), no. 3, 981–983.

[4] C. S. Gordon, D. L. Webb, and S. Wolpert, *Isospectral plane domains and surfaces via Riemannian orbifolds*, Invent. Math. **110** (1992), 1–22.

[5] M. Kac, *Can one hear the shape of a drum?*, Amer. Math. Monthly **73** (1966), no. 4, part II, 1–23.

[6] T. Sunada, *Riemannian coverings and isospectral manifolds*, Ann. of Math. (2) **121** (1985), no. 1, 169–186.

[7] S. R. S. Varadhan, *Diffusion processes in a small time interval*, Comm. Pure Appl. Math. **20** (1967), 659–685.