From Cyclic Autocorrelation to a Layered Symplectic Thickening of Pitch-Class and Permutation Space

Jason St George

GAMUT (Geometric-Algebraic Music Theory) Research Program
jason@shapeofmusicalpossibility.org

March 2026

Abstract

We give a rigorous mathematical formalization of a hierarchical musical space built from pitch-class content, cyclic order, and local permutation structure. Let $G_n = \mathbb{Z}/n\mathbb{Z}$ and write $n = 12$ for the chromatic case. The first layer is the content layer: a pitch-class set $S \subset G_n$ is encoded by its indicator function and its cyclic autocorrelation, equivalently by the squared moduli of its Fourier coefficients. The second layer is the order layer: a rooted ordered pattern $p = (p_0, \dots, p_{k-1})$ of distinct pitch classes carries both a gap word, which isolates Slonimsky-type interval-cycle structure, and an order signal, whose Fourier transform faithfully records the cyclic order. The third layer is a symplectic thickening: for each cardinality k, the discrete pattern space embeds into a complex vector space $\mathcal{M}_{n,k} = \mathbb{C}^{n-1} \times \mathbb{C}^k$ carrying the exact symplectic form

$$\omega_{n,k}^{(\lambda)} = \frac{\mathrm{i}}{2} \sum_{j=1}^{n-1} \mathrm{d}X_j \wedge \mathrm{d}\overline{X_j} + \lambda \frac{\mathrm{i}}{2} \sum_{\ell=0}^{k-1} \mathrm{d}Y_\ell \wedge \mathrm{d}\overline{Y_\ell}, \qquad \lambda > 0.$$

On the nonvanishing-mode locus this becomes the layered action-angle form

$$\omega_{n,k}^{(\lambda)} = \sum_{j=1}^{n-1} \mathrm{d}I_j \wedge \mathrm{d}\theta_j + \lambda \sum_{\ell=0}^{k-1} \mathrm{d}J_\ell \wedge \mathrm{d}\phi_\ell.$$

We prove: (i) finite Wiener–Khinchin theorems for both content and order observables; (ii) a Fourier characterization of periodic gap words, formalizing Slonimsky's interval cycles as sparse order spectra; (iii) injective symplectic thickenings of rooted pattern spaces; (iv) Hamiltonian circle actions realizing transposition and cyclic reindexing, together with anti-symplectic inversion; and (v) a stratified symplectic structure across cardinalities. This yields a mathematically precise foundation for the geometric visualization of scale patterns, harmonic extensions, twelve-tone rows, and their local permutation fibers.

Keywords. Pitch-class set theory; cyclic autocorrelation; discrete Fourier analysis; Slonimsky cycles; symplectic geometry; Hamiltonian reduction; stratified spaces; twelve-tone systems.

1 Introduction

The chromatic pitch universe of equal temperament is naturally modeled by the finite cyclic group

$$G_n = \mathbb{Z}/n\mathbb{Z},$$

with $n = 12$ in the usual twelve-tone setting. At the most basic level, one studies subsets $S \subset G_n$ (chords, scales, aggregates, extensions) and ordered tuples $p = (p_0, \dots, p_{k-1})$ of distinct pitch

classes (rows, patterns, cyclic traversals, ordered ornaments). A longstanding difficulty is that these data support several inequivalent notions of similarity: interval content, transposition class, inversion class, cyclic order, adjacency structure, and local voice-leading. Any single naive Euclidean embedding hides some of those structures while overemphasizing others.

Two lines of thought motivate the present paper. First, Duncan emphasizes cyclic autocorrelation as a translation-invariant descriptor of interval content. In his formulation, the cyclic autocorrelation of a scale or set is the natural object that records how often each interval occurs inside the set, and thus provides a mathematically economical description of pitch-class content [5]. Second, Slonimsky's *Thesaurus of Scales and Melodic Patterns* organizes patterns by interval cycles and their interpolations, which in modern language suggests that *periodicity in a cyclic order signal* should play a primary structural role [15]. These ideas are complemented by the geometric work of Tymoczko and of Callender–Quinn–Tymoczko, who showed that musically natural quotient spaces often produce orbifold-like or stratified geometries rather than a single smooth manifold [16, 3]; by Quinn's DFT-based framework for equal-tempered harmony [9]; by Amiot's systematic treatment of Fourier analysis in music theory [10]; by Yust's applications of Fourier phase spaces to rhythm and tonality [11]; and by numerical work such as Samplaski's use of multidimensional scaling for pitch-class-set similarity [14], together with topological and graph-theoretic realizations of chord spaces by Bigo and collaborators [2]. The interplay between spectral and non-spectral transposition-invariant information, as investigated by Yust and Amiot [12], further motivates the layered approach taken here.

The aim of this paper is not merely to supply another visualization. We seek a rigorous *exegesis from seed to symplectic completion*:

(i) start with the discrete seed consisting of pitch-class content and rooted cyclic order;

(ii) separate the *content layer* from the *order layer*;

(iii) identify the correct quotient actions (transposition, inversion, cyclic reindexing);

(iv) build a continuous thickening carrying a canonical exact symplectic form;

(v) explain why the total object is stratified by cardinality rather than a single manifold;

(vi) and show how low-dimensional visualizations arise as controlled projections of this exact higher-dimensional object.

The resulting architecture is intentionally layered. The content layer is encoded by indicator functions and their Fourier moduli; the order layer is encoded by two related signals—a *gap signal*, which isolates interval cycles in the style of Slonimsky, and an *order signal*, which faithfully remembers the cyclic arrangement itself. The symplectic thickening is then obtained by placing these Fourier coordinates inside a complex vector space with its standard symplectic form.

The principal novelty is the observation that, once the correct Fourier coordinates are chosen, the full layered geometry is both simple and robust:

- autocorrelation observables become moment-map coordinates;
- the transposition action becomes Hamiltonian;
- inversion becomes anti-symplectic;
- cyclic reindexing of a pattern becomes a second Hamiltonian circle action;
- and the local permutation fiber over a content point becomes an honest symplectic vector fiber after thickening.

Thus the right continuous approximation to the discrete music-theoretic seed is not an arbitrary point cloud but a *stratified symplectic space built from spectral coordinates*. In particular, the weighted small-radius fiber renderings used in visualization arise from a genuine mathematical vector-bundle structure on each cardinality stratum.

2 Discrete seed: content, order, and quotient actions

2.1 Pitch classes and content

Fix $n \geq 2$ and write $G_n = \mathbb{Z}/n\mathbb{Z}$. We denote by

$$T_t(x) = x + t, \qquad I(x) = -x, \qquad x, t \in G_n$$

the transposition and inversion actions. Together they generate the dihedral group

$$D_n = G_n \rtimes \{\pm 1\},$$

acting on G_n by affine maps $x \mapsto \pm x + t$.

Definition 2.1 (Pitch-class content)**.** A *content set* is a nonempty subset $S \subset G_n$. Its *cardinality* is $|S| = k$. The content classes modulo transposition and modulo transposition/inversion are the orbit sets

$$\mathcal{C}_n^T = (2^{G_n} \setminus \{\varnothing\})/G_n, \qquad \mathcal{C}_n^{TI} = (2^{G_n} \setminus \{\varnothing\})/D_n.$$

When $n = 12$, the orbit count is easy to compute directly by Burnside's lemma or brute force. There are 224 T/I-classes including the empty set and 223 nonempty T/I-classes.

Remark 2.2 (Forte classification and harmonic nomenclature)**.** The standard nomenclature for these classes is due to Forte [6], who assigns each T/I-class a canonical label of the form k-n, where k is the cardinality and n is the ordinal within the cardinality group. Under this system, many of the 223 classes correspond to objects familiar to every musician: 3-11 is the major/minor triad, 4-27 is the dominant/half-diminished seventh chord, 7-35 is the diatonic scale, 4-28 is the diminished seventh chord, and 6-35 is the whole-tone collection. Others—such as 5-Z12, 6-Z44, or 8-22—lack familiar common names but occupy equally definite positions in the pitch-class universe. The prefix "Z" in Forte's notation signals that a set class belongs to a homometric pair: 5-Z12 and 5-Z36, for instance, share the same interval content yet are not T/I-equivalent (see the remark on homometry in §3 below). Of the 223 classes, approximately fifty carry widely recognized names in Western music pedagogy; the remainder chart the vast majority of pitch-class space that traditional harmonic practice left unexplored.

2.2 Ordered patterns and cyclic fibers

Definition 2.3 (Ordered pattern)**.** For $1 \leq k \leq n$, let

$$\mathcal{P}_{n,k} = \left\{ p = (p_0, \ldots, p_{k-1}) \in G_n^k \,:\, p_i \neq p_j \text{ for } i \neq j \right\}$$

be the set of ordered k-tuples of distinct pitch classes. The *content projection* is

$$c : \mathcal{P}_{n,k} \to 2^{G_n}, \qquad c(p) = \{p_0, \ldots, p_{k-1}\}.$$

Fix a content set S of size k. Then

$$\mathrm{Ord}(S) = \{p \in \mathcal{P}_{n,k} : c(p) = S\}$$

is the set of all orderings of S, hence $|\mathrm{Ord}(S)| = k!$.

There are two natural quotientings of order. The first is by cyclic reindexing:

$$\sigma \cdot (p_0, \ldots, p_{k-1}) = (p_1, \ldots, p_{k-1}, p_0),$$

which defines an action of the cyclic group C_k on $\mathrm{Ord}(S)$. The quotient $\mathrm{Ord}(S)/C_k$ records the cyclic order while forgetting the choice of starting point. The second fixes a chosen root $r \in S$ and considers only tuples with first entry r.

Definition 2.4 (Rooted cyclic fiber)**.** Let $S \subset G_n$ with $|S| = k$ and let $r \in S$. The *rooted cyclic fiber at* r is

$$\mathrm{Cyc}_r(S) = \{p \in \mathrm{Ord}(S) : p_0 = r\}.$$

Proposition 2.5 (Size of the rooted cyclic fiber)**.** *If $|S| = k$ and $r \in S$, then*

$$|\mathrm{Cyc}_r(S)| = (k-1)!.$$

Equivalently, the quotient $\mathrm{Ord}(S)/C_k$ has $(k-1)!$ elements.

Proof. Fix $p_0 = r$. The remaining $k-1$ entries may be arranged arbitrarily as a permutation of $S \setminus \{r\}$. This yields $(k-1)!$ possibilities. Since every cyclic orbit in $\mathrm{Ord}(S)$ contains exactly one representative with first entry equal to r, the same number counts $\mathrm{Ord}(S)/C_k$. □

Thus the natural *local permutation fiber* over a content point of cardinality k has size $(k-1)!$ before any further quotient by inversion or reversal.

2.3 Gap words and reconstruction

Definition 2.6 (Gap word)**.** For $p = (p_0, \ldots, p_{k-1}) \in \mathcal{P}_{n,k}$ define its *gap word*

$$\delta(p) = (\delta_0, \ldots, \delta_{k-1}) \in G_n^k, \qquad \delta_i = p_{i+1} - p_i,$$

where indices are taken modulo k.

The gap word records the successive directed intervals of the cyclic traversal. It always satisfies

$$\sum_{i=0}^{k-1} \delta_i = 0 \quad \text{in } G_n.$$

Proposition 2.7 (Reconstruction from root and gap word)**.** *Let $r \in G_n$ and let $\delta = (\delta_0, \ldots, \delta_{k-1}) \in G_n^k$. Define*

$$p_0 = r, \qquad p_m = r + \sum_{i=0}^{m-1} \delta_i \quad (1 \le m \le k-1).$$

Then $p = (p_0, \ldots, p_{k-1})$ lies in $\mathcal{P}_{n,k}$ if and only if

(a) $\sum_{i=0}^{k-1} \delta_i = 0$ in G_n, and

(b) the partial sums $0, \delta_0, \delta_0 + \delta_1, \ldots, \sum_{i=0}^{k-2} \delta_i$ are pairwise distinct in G_n.

In that case p is uniquely determined by (r, δ).

Proof. By construction,

$$p_{m+1} - p_m = \delta_m, \qquad 0 \le m \le k-2.$$

Moreover

$$p_0 - p_{k-1} = -\sum_{i=0}^{k-2} \delta_i = \delta_{k-1}$$

holds precisely when $\sum_i \delta_i = 0$. Distinctness of the entries is equivalent to distinctness of the partial sums. Uniqueness is immediate from the recursive formula. □

This elementary proposition is the first formal bridge between permutation space and interval cycles: the cyclic order is encoded by a constrained word in the group G_n, and Slonimsky-style cycle families appear as special periodic classes of those words.

3 Content layer: cyclic autocorrelation and spectral invariants

3.1 Finite Fourier transform on G_n

Let $\eta_n = e^{2\pi \mathrm{i}/n}$. For a function $f : G_n \to \mathbb{C}$, define its discrete Fourier transform by

$$\widehat{f}(j) = \sum_{x \in G_n} f(x) \eta_n^{-jx}, \qquad j \in G_n.$$

The inverse formula is

$$f(x) = \frac{1}{n} \sum_{j \in G_n} \widehat{f}(j) \eta_n^{jx}.$$

For a content set $S \subset G_n$, let χ_S be its indicator function:

$$\chi_S(x) = \begin{cases} 1, & x \in S, \\ 0, & x \notin S. \end{cases}$$

3.2 Cyclic autocorrelation

Definition 3.1 (Content autocorrelation)**.** The *cyclic autocorrelation* of S is the function

$$A_S(\tau) = \sum_{x \in G_n} \chi_S(x) \chi_S(x + \tau), \qquad \tau \in G_n.$$

Thus $A_S(\tau)$ counts the number of ordered pairs $(x, y) \in S \times S$ with $y - x = \tau$. In particular $A_S(0) = |S|$ and $A_S(-\tau) = A_S(\tau)$.

Theorem 3.2 (Finite Wiener–Khinchin for content)**.** *For every content set $S \subset G_n$,*

$$\widehat{A_S}(j) = |\widehat{\chi_S}(j)|^2 \qquad (j \in G_n).$$

Equivalently,

$$A_S(\tau) = \frac{1}{n} \sum_{j \in G_n} |\widehat{\chi_S}(j)|^2 \, \eta_n^{j\tau}.$$

Proof. We compute directly:

$$\begin{aligned}\widehat{A_S}(j) &= \sum_{\tau \in G_n} A_S(\tau)\eta_n^{-j\tau} \\ &= \sum_{\tau \in G_n} \sum_{x \in G_n} \chi_S(x)\chi_S(x+\tau)\eta_n^{-j\tau}.\end{aligned}$$

Set $y = x + \tau$, so $\tau = y - x$. Then

$$\begin{aligned}\widehat{A_S}(j) &= \sum_{x,y \in G_n} \chi_S(x)\chi_S(y)\eta_n^{-j(y-x)} \\ &= \left(\sum_{y \in G_n} \chi_S(y)\eta_n^{-jy}\right)\left(\sum_{x \in G_n} \chi_S(x)\eta_n^{jx}\right) \\ &= \widehat{\chi_S}(j)\,\overline{\widehat{\chi_S}(j)} = \left|\widehat{\chi_S}(j)\right|^2.\end{aligned}$$

The inverse formula gives the second identity. □

Corollary 3.3 (Transposition and inversion invariance)**.** *If $S' = T_t(S)$ is a transposition of S, then*

$$\widehat{\chi_{S'}}(j) = \eta_n^{-jt}\widehat{\chi_S}(j) \quad \textit{and hence} \quad A_{S'} = A_S.$$

If $S' = I(S)$ is the inversion of S, then

$$\widehat{\chi_{S'}}(j) = \widehat{\chi_S}(-j) \quad \textit{and hence} \quad A_{S'} = A_S.$$

Therefore A_S is a T/I-invariant of content.

Proof. For transposition,

$$\chi_{S'}(x) = \chi_S(x - t),$$

so

$$\widehat{\chi_{S'}}(j) = \sum_x \chi_S(x-t)\eta_n^{-jx} = \eta_n^{-jt}\sum_u \chi_S(u)\eta_n^{-ju}.$$

For inversion,

$$\chi_{S'}(x) = \chi_S(-x),$$

so

$$\widehat{\chi_{S'}}(j) = \sum_x \chi_S(-x)\eta_n^{-jx} = \sum_u \chi_S(u)\eta_n^{ju} = \widehat{\chi_S}(-j).$$

Taking moduli squared and invoking theorem 3.2 proves the claim. □

Remark 3.4 (Homometry)**.** Two content sets $S, T \subset G_n$ are homometric if $A_S = A_T$, equivalently if

$$\left|\widehat{\chi_S}(j)\right| = \left|\widehat{\chi_T}(j)\right| \qquad \text{for all } j \in G_n.$$

Thus cyclic autocorrelation is complete for interval content but not, in general, for content itself. This is exactly the finite-group form of the familiar Z-relation phenomenon in pitch-class set theory.

3.3 Moment coordinates

The previous theorem shows that autocorrelation is a function of the nonnegative numbers

$$I_j = \frac{1}{2}\left|\widehat{\chi_S}(j)\right|^2, \qquad j = 1, \dots, n-1.$$

These will reappear as action variables in the symplectic thickening. At this stage it is enough to note that the content autocorrelation is an affine linear image of the vector $(I_1, \dots, I_{n-1})$, with $|S| = \widehat{\chi_S}(0)$ serving as an additional cardinality parameter.

4 Order layer I: gap words and Slonimsky spectral sparsity

4.1 The gap signal

The gap word records directed steps in G_n. To convert it into a Fourier-compatible signal, define

$$g_p(m) = \eta_n^{\delta_m(p)} \in S^1, \qquad m \in \mathbb{Z}/k\mathbb{Z}.$$

Its Fourier transform on $\mathbb{Z}/k\mathbb{Z}$ is

$$G_p(\ell) = \sum_{m=0}^{k-1} g_p(m)\zeta_k^{-\ell m}, \qquad \zeta_k = e^{2\pi \mathrm{i}/k}, \qquad 0 \le \ell \le k-1.$$

The signal g_p does not remember the absolute root p_0, but it encodes precisely the directed interval-cycle structure of the pattern. For the Slonimsky viewpoint this is the essential order observable.

4.2 Periodicity versus Fourier support

The next theorem is simply Fourier analysis on the cyclic group $\mathbb{Z}/k\mathbb{Z}$, but its music-theoretic interpretation is decisive: *periodic interval cycles are exactly the gap signals with sparse Fourier support.*

Theorem 4.1 (Periodic gap words and sparse spectra)**.** *Let d divide k and let $g : \mathbb{Z}/k\mathbb{Z} \to \mathbb{C}$. Then the following are equivalent:*

(i) $g(m+d) = g(m)$ *for all* m*;*

(ii) $\widehat{g}(\ell) = 0$ *whenever* ℓ *is not a multiple of* k/d*.*

Proof. Let $T_d g(m) = g(m+d)$. By a standard Fourier calculation,

$$\widehat{T_d g}(\ell) = \zeta_k^{\ell d}\widehat{g}(\ell).$$

Hence $T_d g = g$ if and only if

$$(\zeta_k^{\ell d} - 1)\widehat{g}(\ell) = 0 \qquad \text{for all } \ell.$$

Now $\zeta_k^{\ell d} = 1$ if and only if $k \mid \ell d$, equivalently if and only if ℓ is a multiple of k/d. Therefore $T_d g = g$ exactly when $\widehat{g}(\ell)$ vanishes outside those multiples. □

Corollary 4.2 (Slonimsky interval cycles)**.** *If the gap signal g_p has period d, then the order spectrum G_p is supported on the harmonic lattice*

$$\{0,\ k/d,\ 2k/d,\ \ldots,\ (d-1)k/d\} \subset \mathbb{Z}/k\mathbb{Z}.$$

In particular, principal interval cycles correspond to maximally sparse gap spectra, while interpolations and superpositions may be interpreted as controlled spectral fillings of that sparse support.

Remark 4.3. This theorem gives a rigorous mathematical language for organizing Slonimsky's families. Equal-step cycles are constant or periodic gap signals; infra-, inter-, and ultrapolations appear as perturbations of those periodic signals. In modern spectral language, the Slonimsky catalogue is arranged not only by interval arithmetic but also by the support geometry of a finite Fourier transform.

4.3 Gap autocorrelation

Define the gap autocorrelation by

$$B_p(r) = \sum_{m=0}^{k-1} g_p(m)\overline{g_p(m+r)}, \qquad r \in \mathbb{Z}/k\mathbb{Z}.$$

Then the same finite Wiener–Khinchin identity gives

$$\widehat{B_p}(\ell) = |G_p(\ell)|^2 .$$

Thus interval-cycle content sits in the nonnegative coordinates $|G_p(\ell)|^2$, just as pitch-class content sits in $|\widehat{\chi_S}(j)|^2$.

Remark 4.4 (Step-word entropy as a fiber invariant)**.** The gap word $\delta(p)$ induces a step histogram $h_p(s) = \#\{i : \delta_i = s\}$ on G_n. The Shannon entropy of the normalized histogram,

$$H(p) = -\sum_{s \in G_n} \frac{h_p(s)}{k} \log_2 \frac{h_p(s)}{k},$$

measures how uniformly the pattern distributes its interval steps. Orderings with a single repeated step size (equal-division cycles) have entropy zero, while orderings with maximally irregular step distributions achieve higher entropy. Within a single permutation fiber, entropy thus provides a continuous scalar invariant that distinguishes "smooth" traversals from "jagged" ones. In the companion interactive visualization, fiber vertices are colored by this entropy: low entropy appears as blue (even step spacing), high entropy as amber (irregular intervals), giving an immediate visual cue to the internal smoothness of each cyclic ordering.

5 Order layer II: faithful order signal and exact reconstruction

The gap signal is ideal for cycle periodicity, but it forgets the root. To obtain a faithful ordered encoding we use the absolute order signal.

5.1 Order signal

Definition 5.1 (Order signal)**.** For $p = (p_0, \dots, p_{k-1}) \in \mathcal{P}_{n,k}$ define

$$s_p(m) = \eta_n^{p_m} \in S^1, \qquad m \in \mathbb{Z}/k\mathbb{Z}.$$

Its Fourier transform is

$$Y_p(\ell) = \sum_{m=0}^{k-1} s_p(m)\zeta_k^{-\ell m}, \qquad 0 \le \ell \le k-1.$$

The sequence s_p remembers the ordered pitch classes themselves, not only their successive differences.

Proposition 5.2 (Exact reconstruction from order Fourier data)**.** *For fixed n and k, the map*

$$p \longmapsto (Y_p(0), \dots, Y_p(k-1)) \in \mathbb{C}^k$$

is injective on $\mathcal{P}_{n,k}$.

Proof. By inverse Fourier transform on $\mathbb{Z}/k\mathbb{Z}$,

$$s_p(m) = \frac{1}{k} \sum_{\ell=0}^{k-1} Y_p(\ell)\zeta_k^{\ell m}.$$

Hence the vector $(Y_p(\ell))_{\ell=0}^{k-1}$ uniquely determines the sequence $(s_p(m))_{m=0}^{k-1}$. Each $s_p(m)$ is an nth root of unity, and therefore uniquely determines $p_m \in G_n$. Thus the whole ordered tuple p is recovered. □

5.2 Order autocorrelation

Define the order autocorrelation

$$C_p(r) = \sum_{m=0}^{k-1} s_p(m)\overline{s_p(m+r)}, \qquad r \in \mathbb{Z}/k\mathbb{Z}.$$

Theorem 5.3 (Finite Wiener–Khinchin for order)**.** *For every $p \in \mathcal{P}_{n,k}$,*

$$\widehat{C_p}(\ell) = |Y_p(\ell)|^2 \qquad (0 \le \ell \le k-1).$$

Proof. This is the same calculation as in theorem 3.2, applied to the function s_p on the cyclic group $\mathbb{Z}/k\mathbb{Z}$. □

Thus the order layer has its own nonnegative invariants

$$J_\ell = \frac{1}{2}\,|Y_p(\ell)|^2\,, \qquad 0 \le \ell \le k-1.$$

5.3 Symmetries of the order signal

Two order symmetries are particularly natural.

Proposition 5.4 (Cyclic reindexing)**.** *Let σ be the cyclic reindexing operator*

$$\sigma \cdot (p_0, \dots, p_{k-1}) = (p_1, \dots, p_{k-1}, p_0).$$

Then

$$Y_{\sigma^r p}(\ell) = \zeta_k^{\ell r} Y_p(\ell).$$

Proof. The order signal of $\sigma^r p$ is $s_p(m + r)$. Therefore

$$\begin{aligned} Y_{\sigma^r p}(\ell) &= \sum_{m=0}^{k-1} s_p(m+r)\zeta_k^{-\ell m} \\ &= \zeta_k^{\ell r} \sum_{u=0}^{k-1} s_p(u)\zeta_k^{-\ell u} = \zeta_k^{\ell r} Y_p(\ell). \end{aligned}$$

□

Proposition 5.5 (Pitch transposition and inversion)**.** *Let $T_t p = (p_0 + t, \dots, p_{k-1} + t)$ and $Ip = (-p_0, \dots, -p_{k-1})$. Then*

$$Y_{T_t p}(\ell) = \eta_n^t Y_p(\ell), \qquad Y_{Ip}(\ell) = \overline{Y_p(-\ell)}.$$

Proof. Transposition multiplies every term of s_p by η_n^t, so the same scalar factors out of the Fourier transform. Inversion sends $s_p(m)$ to $\overline{s_p(m)}$, hence

$$Y_{Ip}(\ell) = \sum_m \overline{s_p(m)}\zeta_k^{-\ell m} = \overline{\sum_m s_p(m)\zeta_k^{\ell m}} = \overline{Y_p(-\ell)}.$$

□

6 The layered Fourier embedding

We now assemble the content and order observables.

6.1 Content coordinates

For $p \in \mathcal{P}_{n,k}$ let $S = c(p)$ and define

$$X_p(j) = \widehat{\chi_S}(j), \qquad 1 \le j \le n-1.$$

We omit the zeroth mode because $X_p(0) = |S| = k$ is constant on the cardinality-k stratum.

6.2 The total embedding

Definition 6.1 (Layered Fourier embedding)**.** For fixed (n, k) define

$$\iota_{n,k} : \mathcal{P}_{n,k} \to \mathbb{C}^{n-1} \times \mathbb{C}^k, \qquad \iota_{n,k}(p) = \big((X_p(j))_{j=1}^{n-1}, (Y_p(\ell))_{\ell=0}^{k-1}\big).$$

Proposition 6.2 (Injectivity of the layered embedding)**.** *The map $\iota_{n,k}$ is injective.*

Proof. By theorem 5.2, the order Fourier coordinates $(Y_p(\ell))_{\ell=0}^{k-1}$ alone determine p. Therefore the larger map $\iota_{n,k}$ is injective as well. □

The significance of $\iota_{n,k}$ is conceptual rather than minimal: the X-coordinates isolate content geometry, while the Y-coordinates isolate order geometry. Their coexistence makes the layered structure visible.

7 Symplectic thickening

7.1 Ambient complex manifold

Fix $\lambda > 0$. For each k set

$$\mathcal{M}_{n,k} = \mathbb{C}^{n-1} \times \mathbb{C}^k$$

with coordinates

$$(X_1, \dots, X_{n-1}, Y_0, \dots, Y_{k-1}).$$

Definition 7.1 (Layered symplectic form)**.** Define the 2-form

$$\omega_{n,k}^{(\lambda)} = \frac{\mathrm{i}}{2} \sum_{j=1}^{n-1} \mathrm{d}X_j \wedge \mathrm{d}\overline{X_j} + \lambda \frac{\mathrm{i}}{2} \sum_{\ell=0}^{k-1} \mathrm{d}Y_\ell \wedge \mathrm{d}\overline{Y_\ell}. \tag{7.1}$$

In real coordinates $X_j = x_j + \mathrm{i}y_j$ and $Y_\ell = u_\ell + \mathrm{i}v_\ell$,

$$\omega_{n,k}^{(\lambda)} = \sum_{j=1}^{n-1} \mathrm{d}x_j \wedge \mathrm{d}y_j + \lambda \sum_{\ell=0}^{k-1} \mathrm{d}u_\ell \wedge \mathrm{d}v_\ell.$$

Theorem 7.2 (Symplectic thickening)**.** *For every n, k and $\lambda > 0$, the pair $(\mathcal{M}_{n,k}, \omega_{n,k}^{(\lambda)})$ is an exact symplectic manifold, and*

$$\iota_{n,k}(\mathcal{P}_{n,k}) \subset \mathcal{M}_{n,k}$$

is an injective discrete subset. Hence $(\mathcal{M}_{n,k}, \omega_{n,k}^{(\lambda)})$ is a symplectic thickening of the discrete pattern space $\mathcal{P}_{n,k}$.

Proof. The form (7.1) is the direct sum of standard symplectic forms on complex vector spaces, hence closed and nondegenerate. It is exact with primitive

$$\alpha_{n,k}^{(\lambda)} = \frac{\mathrm{i}}{4} \sum_{j=1}^{n-1} \left(X_j \, \mathrm{d}\overline{X_j} - \overline{X_j} \, \mathrm{d}X_j \right) + \lambda \frac{\mathrm{i}}{4} \sum_{\ell=0}^{k-1} \left(Y_\ell \, \mathrm{d}\overline{Y_\ell} - \overline{Y_\ell} \, \mathrm{d}Y_\ell \right),$$

since $\mathrm{d}\alpha_{n,k}^{(\lambda)} = \omega_{n,k}^{(\lambda)}$. Injectivity of the discrete embedding was proved in theorem 6.2. □

Proposition 7.3 (The weight parameter is inessential)**.** *For $\lambda > 0$, the linear map*

$$F_\lambda : \mathcal{M}_{n,k} \to \mathcal{M}_{n,k}, \qquad F_\lambda(X, Y) = (X, \sqrt{\lambda}\, Y)$$

is a symplectomorphism from $(\mathcal{M}_{n,k}, \omega_{n,k}^{(1)})$ to $(\mathcal{M}_{n,k}, \omega_{n,k}^{(\lambda)})$.

Proof. A direct substitution gives

$$F_\lambda^* \omega_{n,k}^{(\lambda)} = \frac{\mathrm{i}}{2} \sum_{j=1}^{n-1} \mathrm{d}X_j \wedge \mathrm{d}\overline{X_j} + \frac{\mathrm{i}}{2} \sum_{\ell=0}^{k-1} \mathrm{d}Y_\ell \wedge \mathrm{d}\overline{Y_\ell} = \omega_{n,k}^{(1)}.$$

□

Thus λ is not an intrinsic symplectic invariant; it is a scale parameter controlling the relative size of the order fiber in renderings and numerical models.

7.2 Action-angle coordinates

Let

$$U_{n,k} = \{(X, Y) \in \mathcal{M}_{n,k} : X_j \neq 0 \text{ for all } j,\ Y_\ell \neq 0 \text{ for all } \ell\}\,.$$

On $U_{n,k}$ write

$$X_j = \sqrt{2I_j}\, e^{\mathrm{i}\theta_j}, \qquad Y_\ell = \sqrt{2J_\ell}\, e^{\mathrm{i}\phi_\ell},$$

with $I_j, J_\ell > 0$ and $\theta_j, \phi_\ell \in \mathbb{R}/2\pi\mathbb{Z}$.

Theorem 7.4 (Full layered symplectic form)**.** *On $U_{n,k}$ one has*

$$\omega_{n,k}^{(\lambda)} = \sum_{j=1}^{n-1} \mathrm{d}I_j \wedge \mathrm{d}\theta_j + \lambda \sum_{\ell=0}^{k-1} \mathrm{d}J_\ell \wedge \mathrm{d}\phi_\ell. \tag{7.2}$$

Proof. This is the standard polar-coordinate identity

$$\frac{\mathrm{i}}{2}\, \mathrm{d}z \wedge \mathrm{d}\overline{z} = \mathrm{d}\left(\frac{1}{2}|z|^2\right) \wedge \mathrm{d}(\arg z)$$

applied to each coordinate separately and summed. □

Equation (7.2) is the *full layered symplectic form*: the content actions I_j and angles θ_j occupy the first layer, while the order actions J_ℓ and angles ϕ_ℓ occupy the second.

7.3 Completely integrable structure

Theorem 7.5 (Complete integrability)**.** *On each stratum $\mathcal{M}_{n,k}$, the functions*

$$I_1, \ldots, I_{n-1}, J_0, \ldots, J_{k-1}$$

pairwise Poisson commute and are functionally independent on the dense open set $U_{n,k}$. Hence $(\mathcal{M}_{n,k}, \omega_{n,k}^{(\lambda)})$ is completely integrable.

Proof. In action-angle coordinates, the Poisson bracket is

$$\{F, G\} = \sum_{j=1}^{n-1} \left(\frac{\partial F}{\partial I_j}\frac{\partial G}{\partial \theta_j} - \frac{\partial F}{\partial \theta_j}\frac{\partial G}{\partial I_j}\right) + \frac{1}{\lambda}\sum_{\ell=0}^{k-1} \left(\frac{\partial F}{\partial J_\ell}\frac{\partial G}{\partial \phi_\ell} - \frac{\partial F}{\partial \phi_\ell}\frac{\partial G}{\partial J_\ell}\right).$$

Each I_j and J_ℓ depends only on itself, so all pairwise brackets vanish. Independence on $U_{n,k}$ is obvious from the coordinate system. □

Corollary 7.6 (Autocorrelation observables Poisson commute)**.** *All content autocorrelation functions $A_S(\tau)$ and all order autocorrelation functions $C_p(r)$ are functions of the commuting actions $\{I_j\}$ and $\{J_\ell\}$ respectively. Therefore any two such observables Poisson commute on the thickened space.*

Proof. By theorems 3.2 and 5.3, both families are inverse Fourier transforms of the squared Fourier moduli, hence are functions only of the actions. □

8 Hamiltonian and anti-symplectic symmetries

8.1 Independent torus actions

The product torus

$$\mathbb{T}^{n-1} \times \mathbb{T}^k$$

acts on $\mathcal{M}_{n,k}$ by independent phase rotations in each coordinate:

$$(a_1, \dots, a_{n-1}; b_0, \dots, b_{k-1}) \cdot (X, Y) = (a_1 X_1, \dots, a_{n-1} X_{n-1}; b_0 Y_0, \dots, b_{k-1} Y_{k-1}).$$

Its moment map is

$$\mu(X, Y) = \left(\frac{1}{2}|X_1|^2, \dots, \frac{1}{2}|X_{n-1}|^2, \frac{\lambda}{2}|Y_0|^2, \dots, \frac{\lambda}{2}|Y_{k-1}|^2 \right).$$

Thus the nonnegative coordinates underlying the two autocorrelations are literally moment-map coordinates.

8.2 Transposition

Definition 8.1 (Hamiltonian transposition action)**.** Define an S^1-action τ on $\mathcal{M}_{n,k}$ by

$$\tau_\alpha(X_1, \dots, X_{n-1}; Y_0, \dots, Y_{k-1}) = (e^{-\mathrm{i}\alpha} X_1, \dots, e^{-(n-1)\mathrm{i}\alpha} X_{n-1}; e^{\mathrm{i}\alpha} Y_0, \dots, e^{\mathrm{i}\alpha} Y_{k-1}), \tag{8.1}$$

where $\alpha \in \mathbb{R}/2\pi\mathbb{Z}$.

When $\alpha = 2\pi t/n$, this restricts to the discrete pitch transposition T_t on the embedded musical seed.

Theorem 8.2 (Transposition is Hamiltonian)**.** *The action τ is Hamiltonian with moment map*

$$H_T = -\sum_{j=1}^{n-1} j\, I_j + \lambda \sum_{\ell=0}^{k-1} J_\ell. \tag{8.2}$$

Proof. The infinitesimal generator of (8.1) is

$$\mathcal{X}_T = -\sum_{j=1}^{n-1} j\, \frac{\partial}{\partial \theta_j} + \sum_{\ell=0}^{k-1} \frac{\partial}{\partial \phi_\ell}.$$

Contracting with (7.2) gives

$$\iota_{\mathcal{X}_T} \omega_{n,k}^{(\lambda)} = -\sum_{j=1}^{n-1} j\, \mathrm{d}I_j + \lambda \sum_{\ell=0}^{k-1} \mathrm{d}J_\ell = \mathrm{d}H_T.$$

□

8.3 Cyclic reindexing

Definition 8.3 (Hamiltonian cyclic reindexing action)**.** Define an S^1-action ρ on $\mathcal{M}_{n,k}$ by

$$\rho_\beta(X_1, \dots, X_{n-1}; Y_0, \dots, Y_{k-1}) = (X_1, \dots, X_{n-1}; e^{0\mathrm{i}\beta} Y_0, e^{1\mathrm{i}\beta} Y_1, \dots, e^{(k-1)\mathrm{i}\beta} Y_{k-1}).$$

When $\beta = 2\pi r/k$, this restricts to cyclic index shift by r on the order signal, as in theorem 5.4.

Theorem 8.4 (Cyclic reindexing is Hamiltonian)**.** *The action ρ is Hamiltonian with moment map*

$$H_R = \lambda \sum_{\ell=0}^{k-1} \ell \, J_\ell.$$

Proof. The infinitesimal generator is

$$\mathcal{X}_R = \sum_{\ell=0}^{k-1} \ell \, \frac{\partial}{\partial \phi_\ell}.$$

Thus

$$\iota_{\mathcal{X}_R} \omega_{n,k}^{(\lambda)} = \lambda \sum_{\ell=0}^{k-1} \ell \, \mathrm{d}J_\ell = \mathrm{d}H_R.$$

□

The commuting Hamiltonians H_T and H_R precisely reflect the two principal quotient operations: pitch transposition and choice of root on a cyclic order.

8.4 Inversion

Definition 8.5 (Inversion involution)**.** Define

$$\mathcal{I}(X_1, \dots, X_{n-1}; Y_0, \dots, Y_{k-1}) = (\overline{X_1}, \dots, \overline{X_{n-1}}; \overline{Y_0}, \overline{Y_{k-1}}, \dots, \overline{Y_1}).$$

Equivalently, in terms of order indices modulo k,

$$(\mathcal{I}Y)_\ell = \overline{Y_{-\ell}}.$$

Theorem 8.6 (Inversion is anti-symplectic)**.** *One has*

$$\mathcal{I}^2 = \mathrm{id}, \qquad \mathcal{I}^* \omega_{n,k}^{(\lambda)} = -\omega_{n,k}^{(\lambda)}.$$

Proof. Complex conjugation on each complex coordinate sends $\frac{\mathrm{i}}{2} \, \mathrm{d}z \wedge \mathrm{d}\overline{z}$ to its negative. Permutation of coordinates does not affect the form. Therefore each summand in (7.1) changes sign, and the full form does as well. □

Remark 8.7. This is the mathematically correct reason that the full T/I quotient is not globally symplectic in the naive sense: transposition is symplectic and Hamiltonian, but inversion reverses the symplectic form. The natural object is therefore a symplectic quotient by the transposition subgroup together with an anti-symplectic involution descended from inversion.

9 Reduction, fibers, and stratification

9.1 Symplectic reduction by transposition

By theorem 8.2, standard Marsden–Weinstein reduction applies on regular level sets of H_T (see [1, 13]). Thus, on each regular value c,

$$\mathcal{M}_{n,k} /\!\!/_c S^1 = H_T^{-1}(c)/S^1$$

is a symplectic orbifold of dimension

$$2(n-1+k) - 2 = 2n + 2k - 4.$$

Inside this reduced space lie the transposition classes of the embedded discrete musical patterns of cardinality k.

Likewise, reduction by the commuting reindexing action yields the unrooted cyclic order quotient, and one may reduce by the torus generated by (τ, ρ) to forget both absolute transposition and root choice.

9.2 Vector-bundle interpretation

The projection

$$\pi_{n,k} : \mathcal{M}_{n,k} \to \mathbb{C}^{n-1}, \qquad \pi_{n,k}(X, Y) = X$$

is a trivial complex vector bundle with fiber $\mathbb{C}^k$. The symplectic form splits as a product form:

$$\omega_{n,k}^{(\lambda)} = \omega_{\text{content}} \oplus \lambda\, \omega_{\text{order}}.$$

Accordingly, each cardinality-k stratum is a *content base* together with an attached *order fiber*. This is the precise mathematical version of the intuitive statement "attach each local permutation fiber to its content-class point."

Proposition 9.1 (Local permutation fibers)**.** *Fix $X \in \mathbb{C}^{n-1}$. The fiber $\pi_{n,k}^{-1}(X)$ is symplectomorphic to $\big(\mathbb{C}^k, \lambda\omega_{\text{std}}\big)$. The intersection*

$$\pi_{n,k}^{-1}(X) \cap \iota_{n,k}(\mathcal{P}_{n,k})$$

is a finite set consisting of the discrete cyclic orders compatible with the content coordinates X.

Proof. The first claim is immediate from the product decomposition. For the second, $\iota_{n,k}(\mathcal{P}_{n,k})$ is discrete by construction, and fixing X fixes the content layer. What remains is a finite subset of order coordinates corresponding to the finitely many cyclic orderings of that content. □

9.3 Stratification across cardinalities

The full musical universe is not a single manifold because different cardinalities have different order-fiber dimensions.

Definition 9.2 (Total layered space)**.** Define

$$\mathcal{M}_n = \bigsqcup_{k=1}^{n} \mathcal{M}_{n,k}.$$

Proposition 9.3 (Stratified symplectic structure)**.** *The space $\mathcal{M}_n$ is a stratified symplectic space with strata $\mathcal{M}_{n,k}$. It is not a manifold for $n \geq 2$.*

Proof. Each $\mathcal{M}_{n,k}$ is a symplectic manifold by theorem 7.2. Their dimensions are

$$\dim \mathcal{M}_{n,k} = 2(n - 1 + k),$$

which depend on k. Since a manifold has locally constant dimension, the disjoint union cannot be a single manifold when at least two strata occur with different k. □

This proposition is the rigorous form of a point often blurred in informal geometric discussions of musical space: the global object is naturally *stratified by cardinality*. The common practice of drawing everything in one three-dimensional cloud is therefore a visualization choice, not the exact geometry.

10 Visualization theorems

The exact object lives in high-dimensional symplectic coordinates, but one may obtain faithful finite-dimensional renderings by projecting the content base and shrinking the order fibers.

10.1 Four-dimensional direct-sum renderings

Choose a map

$$\Phi : \mathbb{C}^{n-1} \to \mathbb{R}^2$$

representing the content base, for example by multidimensional scaling or diffusion geometry applied to a chosen content metric [14, 4]. Choose a linear or nonlinear map

$$\Psi : \mathbb{C}^k \to \mathbb{R}^2$$

representing the order fiber. For $\varepsilon > 0$ define

$$\mathcal{V}_\varepsilon^{(4)}(X, Y) = \big(\Phi(X), \varepsilon\Psi(Y)\big) \in \mathbb{R}^4.$$

This is the canonical direct-sum rendering: two coordinates for content and two for the attached order fiber. Rotating $\mathbb{R}^4$ and projecting to $\mathbb{R}^3$ or $\mathbb{R}^2$ produces the animated "4D fiber bundle" pictures used in practice.

10.2 Three-dimensional tangent-frame renderings

Suppose instead that the content base is embedded in $\mathbb{R}^3$ by a smooth immersion

$$\Phi : \mathcal{B} \to \mathbb{R}^3$$

defined on a content chart $\mathcal{B}$. Choose an orthonormal frame (t_1, t_2) for the normal or tangent plane used to carry the order fiber. Define

$$\mathcal{V}_\varepsilon^{(3)}(X, Y) = \Phi(X) + \varepsilon\Big(\Psi_1(Y)t_1(X) + \Psi_2(Y)t_2(X)\Big).$$

This yields the familiar "fiber discs attached to base points" picture.

Proposition 10.1 (Small-radius separation)**.** *Let $B = \{b_1, \dots, b_N\} \subset \mathbb{R}^m$ be a finite set of base points and let $F_i \subset \mathbb{R}^r$ be bounded fiber images with*

$$R = \max_i \sup_{u \in F_i} \|u\| < \infty, \qquad \delta = \min_{i \neq j} \|b_i - b_j\| > 0.$$

If $\varepsilon < \delta/(2R)$, then the translated sets

$$b_i + \varepsilon F_i$$

are pairwise disjoint.

Proof. If $u \in F_i$ and $v \in F_j$, then

$$\|(b_i + \varepsilon u) - (b_j + \varepsilon v)\| \geq \|b_i - b_j\| - \varepsilon \|u - v\| \geq \delta - 2\varepsilon R > 0.$$

Hence no two translated fiber images intersect. □

This trivial estimate is nonetheless important: it justifies the "small-radius fiber bundle" principle used in visualization. For a finite collection of content classes, one can always choose a sufficiently small fiber scale so that different fibers remain visually separated.

11 Worked consequences in the twelve-tone case

For $n = 12$, the above formalism includes, without conceptual change:

- all scale and chord contents as subsets of $\mathbb{Z}_{12}$, carrying Forte labels (e.g., 3-11 for the major/minor triad, 7-35 for the diatonic scale);
- all ordered nonrepeating patterns, including segments of twelve-tone rows;
- harmonic extensions as larger content strata containing a given subset (e.g., the dominant seventh 4-27 extends the major triad 3-11);
- Slonimsky interval cycles as periodic gap signals (e.g., equal-step cycles like the diminished seventh 4-28 and whole-tone collection 6-35);
- and local permutation fibers of size $(k-1)!$ over each content class of cardinality k.

Two numerical illustrations are shown in figures 1 and 2. They are not part of the proof theory, but they are faithful shadow projections of the exact layered architecture developed above.

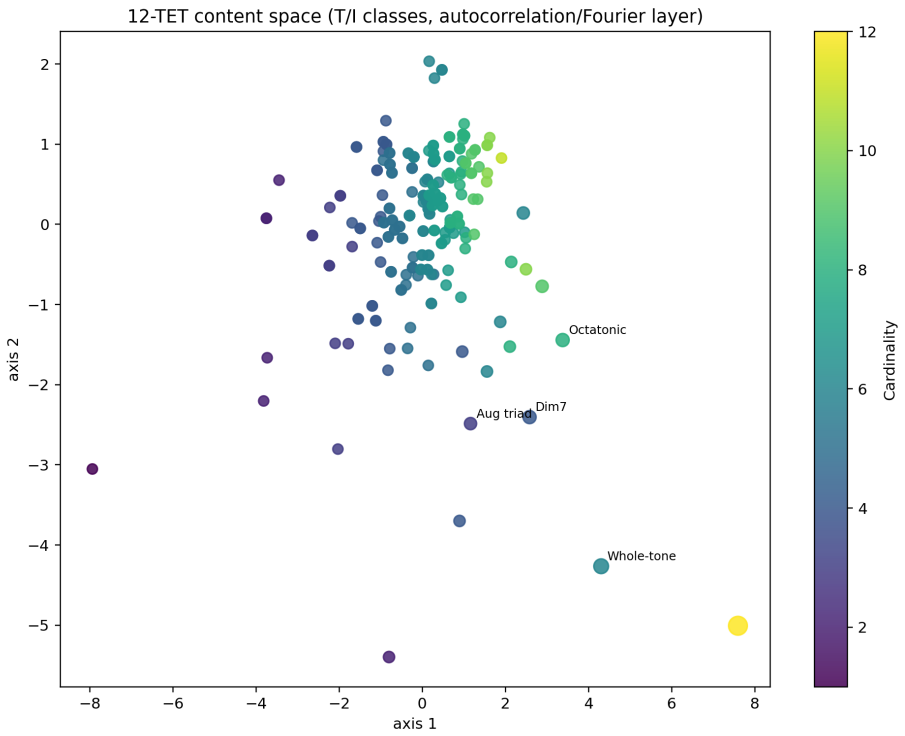

Figure 1: A two-dimensional content rendering for the nonempty T/I-classes of $\mathbb{Z}_{12}$, obtained by multidimensional scaling in the earlier computational prototype. The present paper interprets such a diagram as a low-dimensional shadow of the exact content layer.

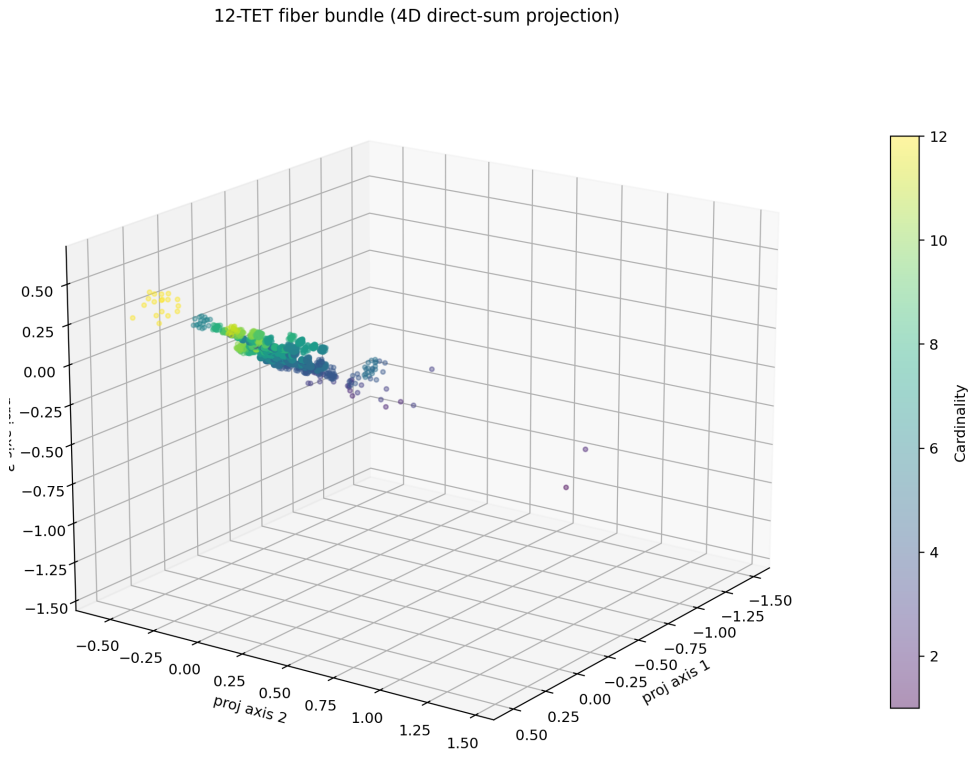

Figure 2: A projected 4D direct-sum rendering of the content base with attached order fibers from the earlier computational prototype. The exact object underlying such pictures is the stratified symplectic thickening $\mathcal{M}_{12} = \bigsqcup_{k=1}^{12} \mathcal{M}_{12,k}$.

Structure-aware renderings

The MDS scatter in figure 1 preserves pairwise content distances but does not expose the graph skeleton: it optimizes metric fidelity at the expense of combinatorial legibility. An alternative rendering treats the base layer as a quotient graph rather than an unstructured cloud, using cardinality as the radial axis and one-note addition/removal as the edge relation. Figure 3 shows this stratified shell graph for the 223 nonempty T/I-classes of $\mathbb{Z}_{12}$. Nodes are arranged by cardinality (radius), colored by cardinality, and sized by symmetry order; edges connect classes that differ by the addition or removal of a single pitch class. This layout makes the cardinality stratification of theorem 9.3 visually explicit and reveals the dense internal connectivity of the base quotient graph. Landmark classes are labeled by their common names ("Dim7" = Forte 4-28, "Whole-tone" = 6-35, "Octatonic" = 8-28, "Aug triad" = 3-12, "Aggregate" = 12-1).

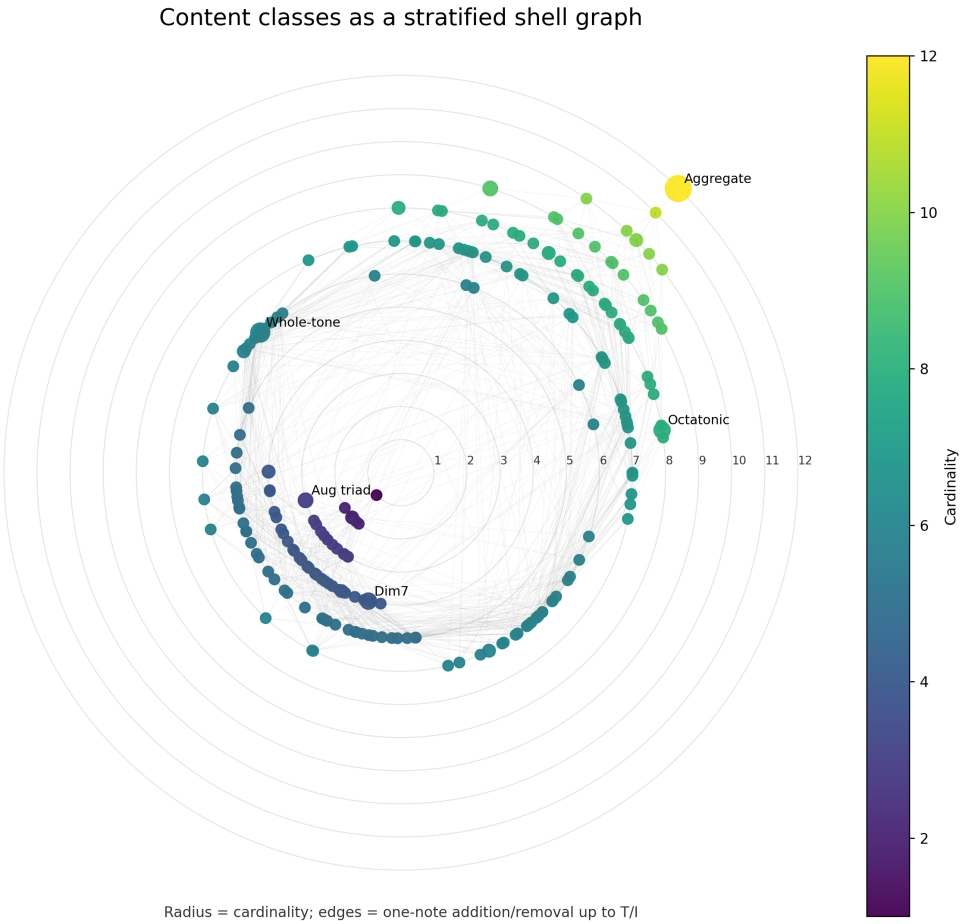

Figure 3: Content classes as a stratified shell graph for $n = 12$. Radius encodes cardinality; edges represent one-note addition/removal up to T/I-equivalence. Node size reflects symmetry order, making landmarks such as the whole-tone collection and diminished seventh chord visually prominent.

Attaching local permutation fibers to this shell graph yields a *rosette bundle map* in which each content node sprouts a small rosette of fiber points. Figure 4 shows this construction with fiber vertices colored by the cyclic step-word entropy $H(p)$ described in theorem 4.4. The entropy coloring exposes internal fiber structure at a glance: blue vertices correspond to smooth, evenly-spaced traversals, while yellow-green vertices indicate jagged, irregular orderings.

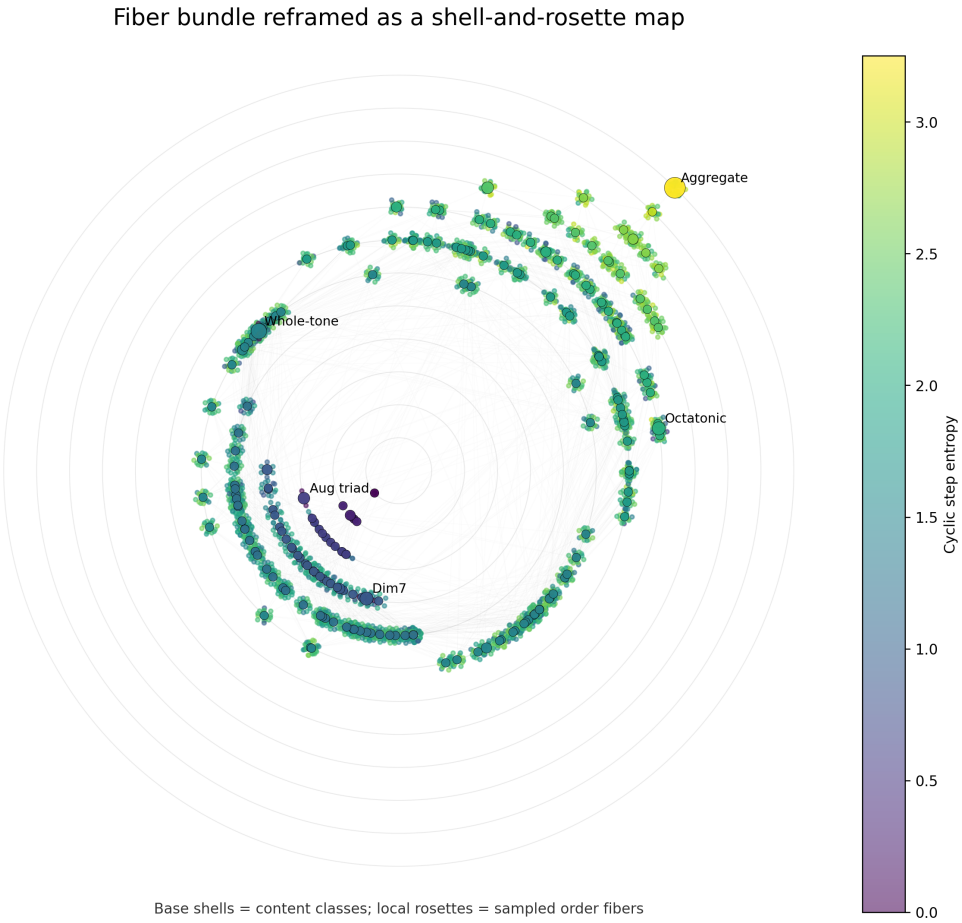

Figure 4: The fiber bundle reframed as a shell-and-rosette map. Base shells encode content classes; local rosettes represent sampled order fibers. Color encodes cyclic step entropy: blue indicates even step spacing (smooth traversals), yellow indicates irregular intervals (jagged orderings). This rendering is the structure-aware alternative to the generic 4D scatter in figure 2.

Figure 5 places these renderings side by side with their predecessors, illustrating the shift from metric-preserving point clouds to structure-aware graph maps.

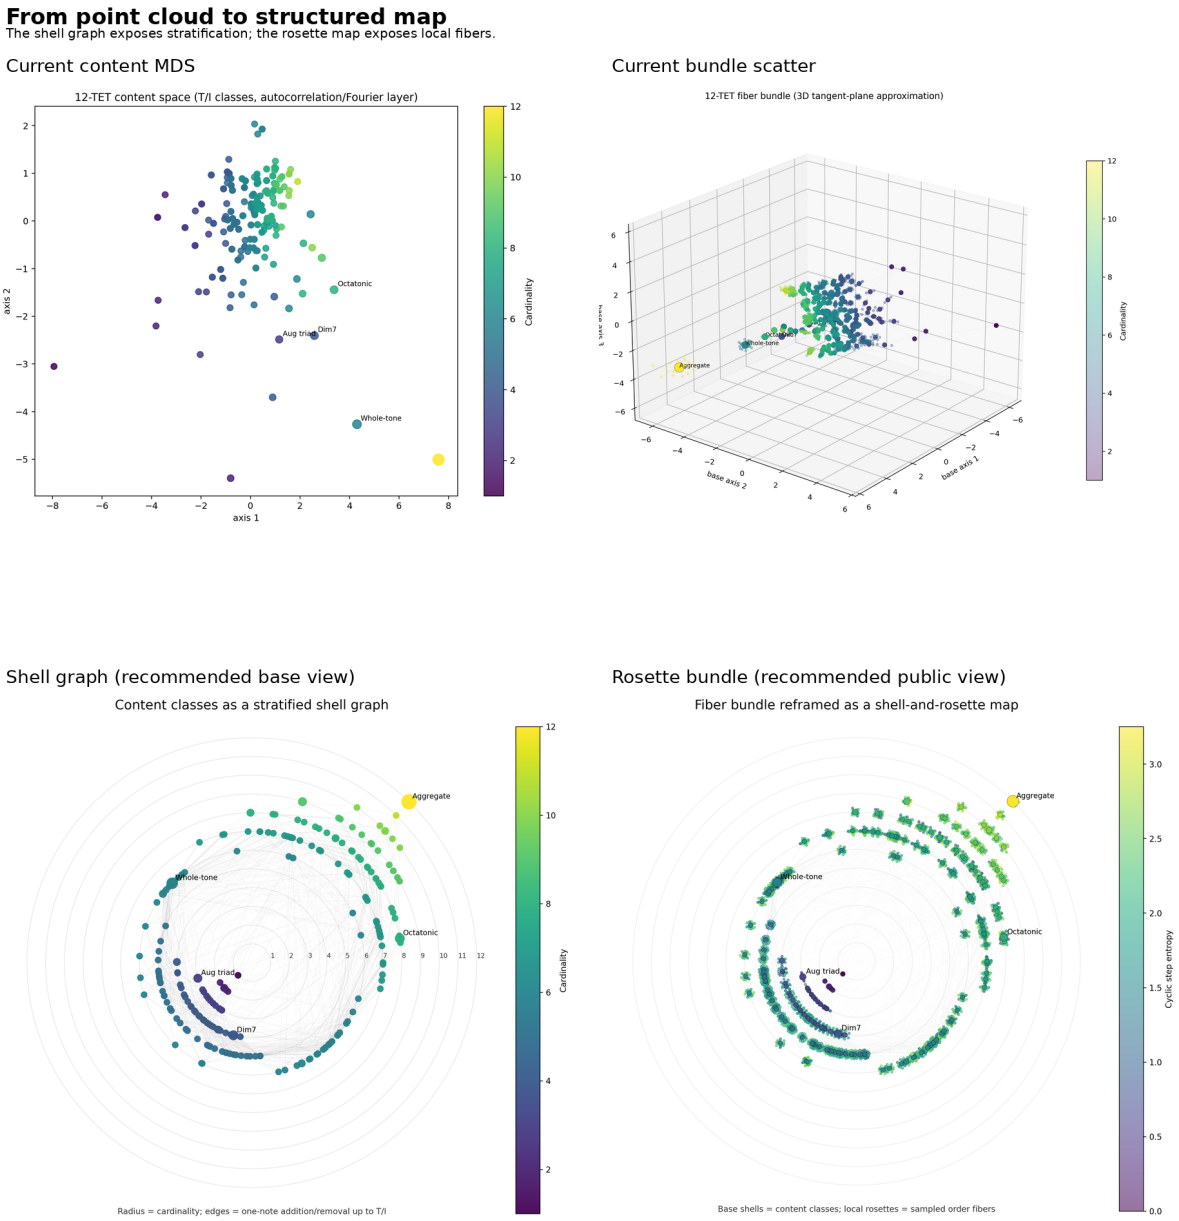

Figure 5: From point cloud to structured map. Top row: MDS content scatter and 3D bundle scatter from the earlier prototype. Bottom row: stratified shell graph (left) and rosette bundle map with entropy coloring (right). The shell graph exposes the cardinality stratification; the rosette map exposes local fiber structure.

Computational anatomy of the fiber space

The fiber construction is not merely theoretical. A companion computational pipeline generates exact or sampled fibers for every content class with three or more pitch classes (cardinalities 1 and 2 have trivial single-element fibers). Table 1 summarizes the combinatorial landscape.

Table 1: Fiber coverage by cardinality for the $n = 12$ case. For every content class of cardinality k, the exact count of rooted cyclic orderings is $(k-1)!$ (by theorem 2.5), independent of the symmetry order of the T/I-class. At cardinality 6 and above, the visualization samples at most 96 orderings per fiber.

Cardinality	T/I-classes	Orderings/fiber	Sampling	Total orderings
1	1	1	—	—
2	6	1	—	—
3	12	2	exhaustive	24
4	29	6	exhaustive	174
5	38	24	exhaustive	912
6	50	≤ 96	sampled from 120	4,800
7	38	≤ 96	sampled from 720	3,648
8	29	≤ 96	sampled from 5,040	2,784
9	12	≤ 96	sampled from 40,320	1,152
10	6	≤ 96	sampled from 362,880	576
11	1	≤ 96	sampled from 3,628,800	96
12	1	≤ 96	sampled from 39,916,800	96

The combinatorial explosion is dramatic: from 2 orderings at cardinality 3 to tens of millions at cardinality 12. Note that the rooted cyclic fiber count $(k-1)!$ does not depend on the symmetry order of the T/I-class: fixing a root eliminates translational symmetry. The diminished seventh chord (4-28, symmetry order 4) has the same $6 = (4-1)!$ rooted orderings as any other tetrachord. Symmetry enters only when one further quotients: the whole-tone collection (6-35, symmetry order 6) has the full $(6-1)! = 120$ rooted orderings, but only 20 are inequivalent modulo its sixfold translational symmetry. Even the sampled fibers preserve meaningful topology, since adjacent-swap edges connect orderings differing by a single neighbor transposition, keeping the local structure of the permutohedron visible at every scale.

The full dataset contains 223 content classes, 1,051 add/remove edges, and 216 computed fibers (the seven classes of cardinality $k \leq 2$ have trivial single-element fibers and are excluded from nontrivial fiber computation) with a total of 14,166 orderings and their associated swap-edge graphs. Each content class is identified both by its Forte number and, where one exists, its common Western name (see theorem 2.2).

Interactive companion

An interactive realization of these renderings—the *Layered Bundle Explorer*—accompanies this paper as a web-based computational artifact. It implements all four views discussed above (MDS scatter, 4D projection, shell graph, rosette map), together with semantic zoom for progressive detail revelation, a configurable hop-radius edge filter, cardinality band selection, and on-demand fiber inspection with entropy coloring and hover tooltips. Selecting any node displays its Forte number and common name alongside its pitch-class content, fiber statistics, and step-word data. A guided-tour mode animates a traversal through musically meaningful paths—from triads through diminished ladders to whole-tone and octatonic regions—offering a narrated introduction to the bundle's topology. The explorer is available at the companion project website.

12 Conclusion

The main point of this paper is structural. The musical space generated by pitch-class content and cyclic order is not best understood as a single smooth manifold. Its correct exact form is layered and stratified:

(1) a discrete seed of content sets and rooted cyclic patterns;

(2) a content layer controlled by cyclic autocorrelation and content spectra;

(3) an interval-cycle order layer controlled by periodic gap words and their sparse spectra;

(4) a faithful order layer controlled by the Fourier transform of the order signal;

(5) a symplectic thickening on the ambient spectral coordinates;

(6) Hamiltonian quotient actions for transposition and cyclic reindexing;

(7) and an anti-symplectic inversion symmetry yielding the final T/I-type identification.

This architecture explains several phenomena at once. Duncan's cyclic autocorrelation becomes the moment-coordinate description of pitch-class content. Slonimsky's interval cycles become periodic gap words, equivalently sparse Fourier supports. Local permutation fibers become honest vector fibers after thickening. The full continuous model carries a canonical exact symplectic form, so the geometry is not merely pictorial but Hamiltonian. Finally, because the cardinality-k strata have dimensions $2(n-1+k)$, the total space is inherently stratified: any human-readable 3D or 4D rendering is a controlled projection of that larger exact object.

By labeling every content class with its Forte number and, where one exists, a common name from Western tonal practice, the abstract graph acquires a legible harmonic overlay. Roughly fifty of the 223 classes carry familiar names (major triad, dominant seventh, diatonic scale, whole-tone collection, and the like); the remaining majority are named only by Forte's ordinal system. This asymmetry is itself revealing: the graph topology of single-pitch-class addition/removal connects the "named" islands through long chains of "unnamed" intermediate classes, exposing the narrow scope of traditional Western harmony within the full combinatorial landscape.

Several natural extensions remain. One may add registral information and voice-leading constraints, thereby interfacing the present content/order architecture with orbifold chord spaces in the style of Tymoczko and of Callender–Quinn–Tymoczko [16, 3]. One may also replace the Euclidean content renderings by diffusion geometry or other nonlinear spectral methods [4]. But the basic formalism is already complete: from the discrete seed of pitch classes and order, one obtains a full layered symplectic form and, with it, a mathematically coherent model of the manifold intuition behind the twelve-tone universe.

A Algorithmic extraction of the layered coordinates

For completeness we record the exact computational recipe for a pattern $p = (p_0, \dots, p_{k-1}) \in \mathcal{P}_{n,k}$.

Step 1: content coordinates. Set $S = \{p_0, \dots, p_{k-1}\}$ and compute

$$X_j = \sum_{x \in S} \eta_n^{-jx}, \qquad 1 \le j \le n-1.$$

Step 2: order coordinates. Compute the order signal

$$s_p(m) = \eta_n^{p_m}, \qquad 0 \le m \le k-1,$$

and then

$$Y_\ell = \sum_{m=0}^{k-1} s_p(m)\zeta_k^{-\ell m}, \qquad 0 \le \ell \le k-1.$$

Step 3: gap-spectrum coordinates (optional but musically informative). Compute the gap word

$$\delta_m = p_{m+1} - p_m \pmod{n},$$

the gap signal

$$g_p(m) = \eta_n^{\delta_m},$$

and its transform

$$G_\ell = \sum_{m=0}^{k-1} g_p(m)\zeta_k^{-\ell m}.$$

The support of G detects interval-cycle periodicity by theorem 4.1.

Step 4: symplectic coordinates. For nonzero modes write

$$X_j = \sqrt{2I_j}e^{\mathrm{i}\theta_j}, \qquad Y_\ell = \sqrt{2J_\ell}e^{\mathrm{i}\phi_\ell}.$$

Then the thickened point lies in $\mathcal{M}_{n,k}$ with symplectic form (7.2).

B A note on redundancy

The order signal already determines the whole ordered pattern, whereas the content signal by itself remembers only the unordered content. Accordingly, the layered embedding is redundant. This redundancy is deliberate: it separates the two musically meaningful notions of structure. The content coordinates support content autocorrelation and transposition/inversion analysis; the order coordinates support cyclic-order analysis; and the symplectic form splits into corresponding layers. The redundancy is therefore not a defect but a categorical refinement.

References

[1] V. I. Arnold, *Mathematical Methods of Classical Mechanics*, 2nd ed., Springer, New York, 1989.

[2] L. Bigo, D. Ghisi, A. Spicher, and M. Andreatta, Representation of musical structures and processes in simplicial chord spaces, *Computer Music Journal* **39** (2015), no. 3, 9–24.

[3] C. Callender, I. Quinn, and D. Tymoczko, Generalized voice-leading spaces, *Science* **320** (2008), no. 5874, 346–348.

[4] R. R. Coifman and S. Lafon, Diffusion maps, *Applied and Computational Harmonic Analysis* **21** (2006), no. 1, 5–30.

[5] A. Duncan, Combinatorial Music Theory, *Journal of the Audio Engineering Society* **39** (1991), no. 6, 427–448.

[6] A. Forte, *The Structure of Atonal Music*, Yale University Press, New Haven, 1973.

[7] R. Morris, *Composition with Pitch-Classes: A Theory of Compositional Design*, Yale University Press, New Haven, 1987.

[8] J. Rahn, *Basic Atonal Theory*, Longman, New York, 1980.

[9] I. Quinn, General equal-tempered harmony, *Journal of Music Theory* **50** (2006), no. 2.

[10] E. Amiot, *Music Through Fourier Space: Discrete Fourier Transform in Music Theory*, Springer, Cham, 2016.

[11] J. Yust, *Organized Time: Rhythm, Tonality, and Form*, Oxford University Press, Oxford, 2018.

[12] J. Yust and E. Amiot, Non-spectral transposition-invariant information in pitch-class sets and distributions, in *Mathematics and Computation in Music*, LNCS 13267, 279–291, 2022.

[13] D. McDuff and D. Salamon, *Introduction to Symplectic Topology*, 3rd ed., Oxford University Press, Oxford, 2017.

[14] A. Samplaski, Mapping the geometries of pitch-class set similarity measures via multidimensional scaling, *Music Theory Online* **11** (2005), no. 2.

[15] N. Slonimsky, *Thesaurus of Scales and Melodic Patterns*, Scribner's, New York, 1947.

[16] D. Tymoczko, The geometry of musical chords, *Science* **313** (2006), 72–74.

[17] D. Tymoczko, *A Geometry of Music: Harmony and Counterpoint in the Extended Common Practice*, Oxford University Press, Oxford, 2011.