

Spectral connectivity descriptors separate inner from outer mitochondrial membranes in cryo-electron tomography

John Reimer Morales

Independent Researcher

Correspondence: jrm@globalharmonics.org

1 Introduction

Cellular cryo-electron tomography (cryo-ET) enables three-dimensional reconstruction of organellar architecture at subnanometer resolution in near-native conditions. The development of automated segmentation workflows and surface mesh reconstruction algorithms has made it possible to model and quantify membrane ultrastructure across large numbers of organelles ([Barad et al., 2023](#)). The Surface Morphometrics pipeline generates triangulated surface meshes from voxel segmentations using the Screened Poisson algorithm ([Kazhdan and Hoppe, 2013](#)) and computes local morphological features including curvedness, inter- and intra-membrane distances, relative orientation, and—most recently—membrane thickness ([Medina et al., 2026](#)).

These local descriptors have proven powerful for detecting condition-dependent changes in membrane ultrastructure. However, local morphometrics are inherently point-wise: they characterize membrane properties at each mesh vertex independently and do not capture the global connectivity architecture of the membrane as a whole. A crista junction, for instance, is not merely a region of high curvature—it is a topological and geometric bottleneck that constrains diffusion and transport between the crista lumen and the inter-membrane space. Curvature describes the shape at the junction; connectivity describes its role in the membrane’s global organization.

Spectral graph theory provides a natural mathematical framework for quantifying such global connectivity. Given a triangulated surface mesh, one can construct a graph Laplacian operator

whose eigenvalues and eigenvectors encode the mesh’s connectivity structure at all scales simultaneously (Chung, 1997; von Luxburg, 2007). The second-smallest eigenvalue (λ_2 , the algebraic connectivity; Fiedler 1973) measures the global stiffness of the connectivity graph—how easily the mesh can be partitioned into weakly connected subdomains. The associated eigenvector (the Fiedler vector) identifies the “softest cut” on the surface: the location where the membrane is most weakly connected. Laplace–Beltrami eigenvalues have been used as intrinsic shape descriptors (“Shape-DNA”; Reuter et al. 2006) for surface classification. On inner mitochondrial membranes, the Fiedler vector is expected to localize near candidate crista-junction or constriction regions.

A complementary perspective is provided by the cotangent Laplace–Beltrami operator, which incorporates the geometric embedding of the mesh—not just its adjacency structure—through cotangent weights derived from triangle angles and areas (Meyer et al., 2003; Pinkall and Polthier, 1993). Where the graph Laplacian asks “is there a connection?”, the cotangent operator asks “how wide is the connection?” Because mesh-adjacency graph spectra are primarily sensitive to connectivity structure, whereas cotangent spectra also depend on geometric embedding, agreement between the two operators indicates shared bottleneck structure while divergence isolates geometry-sensitive constriction. This distinction is biologically meaningful: crista junctions are topologically present (the connection exists) but can be geometrically narrow (the passage is constricted).

Spectral approaches to biological shape analysis have precedent. Laplacian eigenvectors have been used for 3D cell shape classification (Du et al., 2013) and for segmentation of volumetric data in medical imaging. However, their systematic application to cryo-ET membrane surfaces has not been explored. The key question is empirical: do spectral connectivity descriptors capture biologically meaningful variation in membrane architecture, and does a dual-operator approach provide information that a single operator does not?

Here we present a spectral morphometrics pipeline that extends Surface Morphometrics with graph Laplacian and cotangent Laplace–Beltrami descriptors. We apply this pipeline to inner mitochondrial membrane (IMM), outer mitochondrial membrane (OMM), and endoplasmic reticulum (ER) surface meshes extracted from publicly deposited cryo-ET data (Barad et al., 2023). We report four principal findings: (1) in matched within-mitochondrion comparisons, the cotangent operator perfectly separates inner from outer membranes while the graph operator captures a related but not identical signal; (2) across the IMM cohort, both operators recover a strongly

consistent bottleneck ranking while diverging in biologically interpretable special cases; (3) across membrane classes, both operators recover the biologically expected connectivity hierarchy; and (4) vertex-count normalization removes mesh-density confounds while preserving biological signal.

2 Results

2.1 The cotangent operator separates inner from outer mitochondrial membranes within the same organelle

The strongest test of whether spectral descriptors capture biologically meaningful membrane architecture is a within-organelle comparison. Each of the 15 mitochondria in the dataset yielded both an IMM and an OMM surface mesh, providing 15 matched pairs that control for cell identity, tomogram acquisition, and reconstruction method. All spectral analyses were performed on the largest connected component (LCC) of each mesh, because algebraic connectivity is undefined on disconnected graphs (see Section 2.3). Although preprocessing and boundary fraction still differ somewhat between the two membrane layers, the paired design eliminates most confounds, isolating the membrane-identity variable.

The cotangent operator separated the two membrane layers in all 15 pairs: $\mu_2(\text{OMM}) > \mu_2(\text{IMM})$ in every case (Wilcoxon signed-rank $p = 6.1 \times 10^{-5}$; median OMM/IMM ratio = $3.24\times$, range $1.36\text{--}20.70\times$). The normalized graph operator also separated the membrane layers, but less consistently: $n \cdot \lambda_2(\text{OMM}) > n \cdot \lambda_2(\text{IMM})$ in 11 of 15 pairs (Wilcoxon $p = 0.018$; median ratio = $1.68\times$). The four pairs where the graph operator did not separate (TE12, TE6, TE7_mito_1, TE9) showed a consistent pattern: by the graph metric, the inner membrane was not more bottlenecked than the outer, whereas by the cotangent metric it was, with cotangent ratios of $1.6\text{--}3.0\times$. This is consistent with mitochondria whose crista-associated constrictions are topologically present but geometrically narrow—passages that exist in the connectivity graph but whose physical width is restricted. The cotangent operator is sensitive to the narrowness; the graph operator is not.

The OMM/IMM ratio was not significantly correlated between operators across pairs ($\rho = 0.41$, $p = 0.13$), indicating that the graph and cotangent branches measure partially independent aspects of the IMM–OMM structural difference. If both operators always agreed proportionally, the dual-operator approach would be redundant. The fact that they agree on direction but disagree on

magnitude provides the scientific justification for deploying both.

TE4_mito_1 was the most extreme pair: graph ratio $10.19\times$ and cotangent ratio $20.70\times$. This mitochondrion has an inner membrane that is both topologically weakly connected relative to its OMM and geometrically constricted. The TE4 paired comparison, together with its within-IMM divergence described in the following section, carries two messages simultaneously: (1) the IMM is dramatically more bottlenecked than its OMM, and (2) topological and geometry-aware bottlenecking are not identical notions even within the same membrane.

Restricting the paired analysis to the seven pairs where the IMM LCC retained $\geq 70\%$ of original vertices (quality tier A; defined in Methods, Section 4.2) did not weaken the result: cotangent μ_2 separated the membrane layers in all 7 pairs (Wilcoxon $p = 0.016$), and graph $n \cdot \lambda_2$ separated 5 of 7. The main finding is thus robust to the fragmentation-related quality stratification.

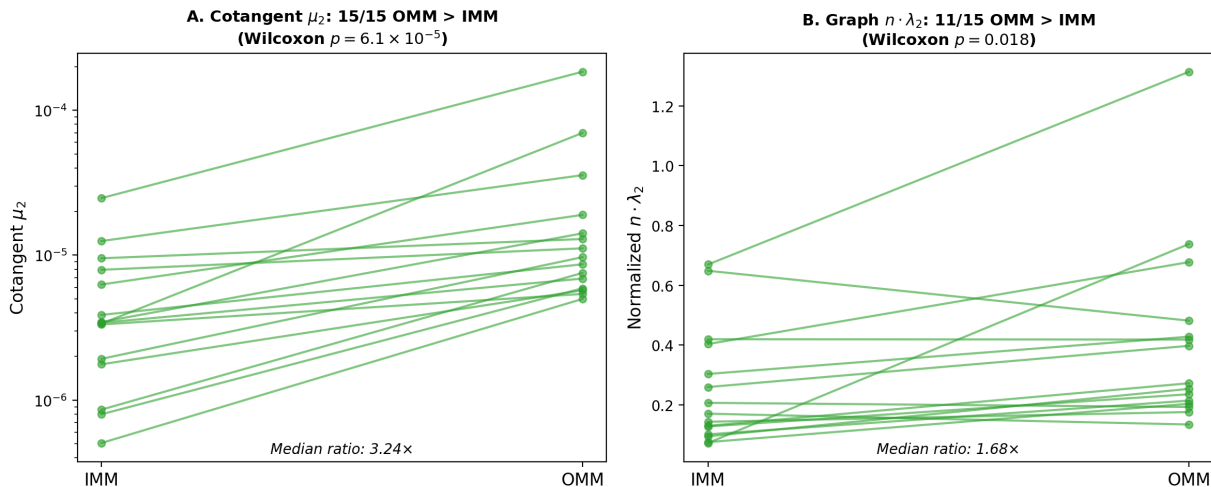


Figure 1: Paired within-mitochondrion IMM vs. OMM comparison. For each of 15 matched mitochondria, cotangent μ_2 (A) and normalized graph $n \cdot \lambda_2$ (B) are compared between inner and outer membranes. Green lines: OMM > IMM (expected direction); red lines: IMM \geq OMM (reversal). The cotangent operator separates the two membrane layers in all 15 pairs (Wilcoxon $p = 6.1 \times 10^{-5}$); the graph operator separates 11 of 15 ($p = 0.018$). The four reversal cases (labeled in panel B) are mitochondria where the graph metric does not register a stronger IMM bottleneck but the cotangent metric does.

2.2 Both operators recover a consistent bottleneck ranking across IMM membranes, with biologically interpretable divergence

Across the 15 IMM meshes, graph λ_2 and cotangent μ_2 were strongly correlated (Pearson $r = 0.977$, $p = 4.0 \times 10^{-10}$; Spearman $\rho = 0.889$, $p = 9.2 \times 10^{-6}$). Both operators recovered a strongly

consistent bottleneck ranking across the IMM cohort: membranes ranked as most bottlenecked by one operator were consistently ranked as most bottlenecked by the other. This consistency establishes that the spectral descriptors are tracking a real, shared structural property of these membranes.

However, the operators were not interchangeable. The most striking divergence was TE4_mito_1_IMM, which the graph Laplacian ranked as the most bottlenecked membrane in the cohort (lowest $n \cdot \lambda_2 = 0.072$) while the cotangent operator ranked it 6th ($\mu_2 = 3.37 \times 10^{-6}$). TE4 also exhibited the lowest spectral gap ratio ($\lambda_2/\lambda_3 = 0.065$), indicating two nearly equivalent competing soft cuts. By the graph metric, TE4 behaves like a strong topological bottleneck, whereas by the cotangent metric its constriction is more moderate, indicating a divergence between topological and geometry-sensitive bottleneck severity—precisely the distinction that a dual-operator approach is designed to detect.

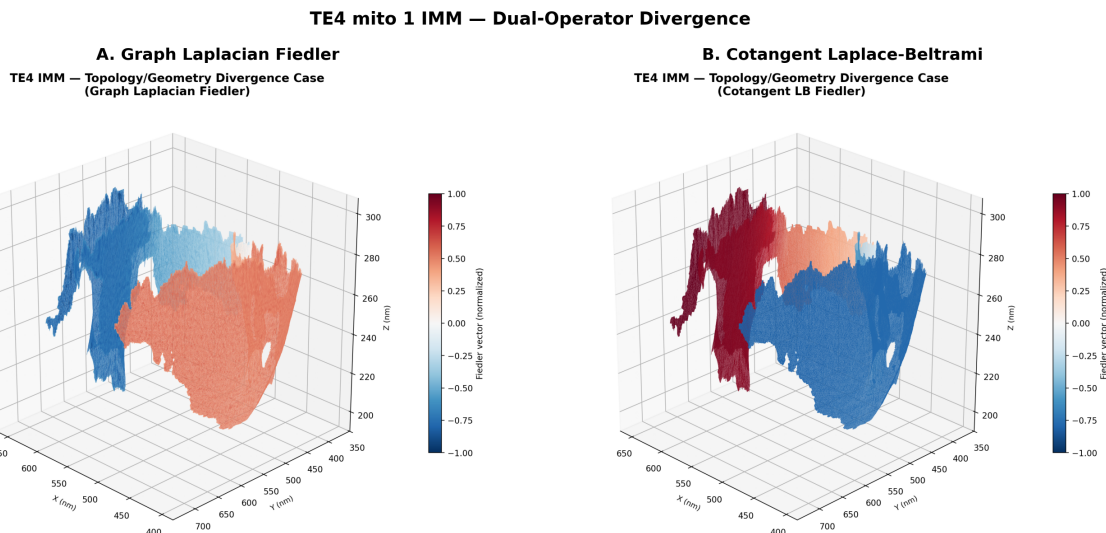


Figure 2: Dual-operator divergence on TE4_mito_1_IMM. (A) Graph Laplacian Fiedler vector. (B) Cotangent Laplace–Beltrami second mode. Both panels show the same mesh (45,823 LCC vertices, 9.5% boundary fraction) but the partitions differ: by the graph metric, TE4 has the strongest topological bottleneck in the IMM cohort; by the cotangent metric, its geometric constriction is moderate. TE4 has the lowest spectral gap ratio ($\lambda_2/\lambda_3 = 0.065$), indicating two nearly equivalent competing soft cuts. Diverging colormap (RdBu_r), normalized to $[-1, 1]$.

2.3 Real cryo-ET membrane meshes require explicit connected-component handling for spectral interpretability

We extracted 15 IMM, 15 OMM, and 7 ER surface meshes from 11 cryo-electron tomograms of mouse embryonic fibroblasts deposited in EMPIAR-11370 (Barad et al., 2023). A critical preprocessing requirement emerged: the deposited IMM meshes contained a mean of 17.5 disconnected components per mesh (range: 7–52). Spectral quantities—particularly λ_2 —are not interpretable on disconnected graphs, where the algebraic connectivity is trivially zero. We therefore adopted a largest connected component (LCC) extraction step as standard preprocessing.

LCC vertex fractions ranged from 28.3% to 97.7% across the 15 IMMs (Table 1). The LCC step should be understood as a choice of analysis object for spectral interpretability, not as an assertion that the biological membrane is fully captured by a single continuous deposited mesh. Fragmentation in the deposited meshes is consistent with reconstruction and segmentation incompleteness, and should not be interpreted as biological discontinuity. Alternatives—such as analyzing all components separately or applying graph union operators—were considered but not adopted here because they would complicate the interpretation of λ_2 as a global connectivity measure. The paired IMM–OMM conclusions remain valid under LCC analysis because both membrane layers of each mitochondrion undergo the same preprocessing: any bias introduced by LCC extraction applies symmetrically.

For each LCC, we computed the symmetric normalized graph Laplacian and the cotangent Laplace–Beltrami operator, solving for the smallest eigenvalues using a shift-invert sparse eigensolver. Both operators produced well-conditioned solutions on all 37 meshes spanning vertex counts from 3,907 to 78,784. Automated quality control flagged near-degenerate triangle angles (minimum angle $\approx 0^\circ$) on all meshes—a known Screened Poisson artifact (Kazhdan and Hoppe, 2013)—but degenerate-face removal produced no change in cotangent eigenvalues ($n = 6$ meshes tested, maximum $\Delta\mu_2 = 0.00\%$), confirming robustness to this mesh-quality limitation.

Figure 3 shows the graph Laplacian Fiedler vector projected onto TE11_mito_1_IMM (23,241 LCC vertices, 6.8% boundary fraction). The mode partitions the membrane into two subcompartments (red and blue in a diverging colormap) with the zero-crossing localizing to a narrow region consistent with a candidate crista-junction-like transition. This localization was not imposed by

the analysis—it emerges from the spectral decomposition of the mesh connectivity graph. Similar spatially coherent, low-frequency partitions were observed on all IMM meshes where the LCC retained sufficient structural complexity.

The cotangent Laplace–Beltrami operator (Meyer et al., 2003; Pinkall and Polthier, 1993) produced qualitatively similar but not identical mode patterns. Figure 2 shows both operators applied to TE4_mito_1_IMM, where the graph and cotangent Fiedler modes partition the membrane differently—a divergence analyzed quantitatively in the preceding section.

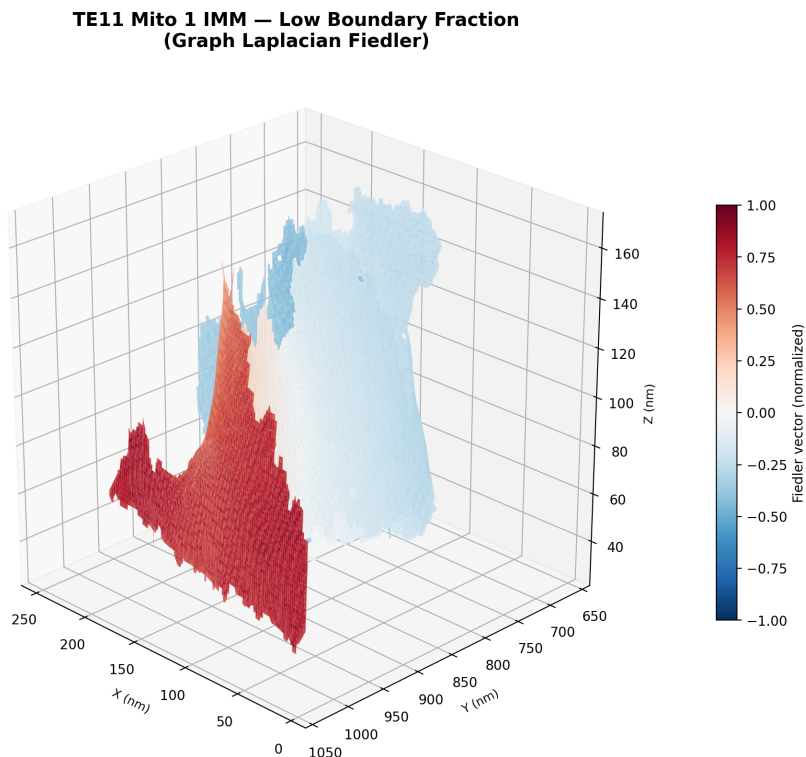


Figure 3: Graph Laplacian Fiedler vector overlay on a real cryo-ET inner mitochondrial membrane. The largest connected component of TE11_mito_1_IMM (23,241 vertices, 6.8% boundary fraction) is colored by the Fiedler vector, with red and blue indicating the two sides of the softest partition. The zero-crossing localizes to a region consistent with a candidate crista-junction-like transition between membrane subcompartments. Diverging colormap (RdBu.r), normalized to $[-1, 1]$.

2.4 Cross-membrane-type comparison

Beyond the paired IMM–OMM comparison, the spectral descriptors discriminated all three membrane classes in the expected biological order. Median normalized algebraic connectivity ($n \cdot \lambda_2$) was 0.170 for IMM, 0.272 for OMM, and 0.448 for ER (Kruskal–Wallis $p = 0.040$; exploratory

Table 1: Per-mesh spectral descriptors for 15 IMM membranes from EMPIAR-11370. LCC%: fraction of original vertices retained in the largest connected component. BF%: boundary fraction. Quality tier A: LCC $\geq 70\%$ of original vertices ($n = 7$); tier B: $< 70\%$ ($n = 8$).

Mesh ID	LCC V	LCC%	BF%	$n \cdot \lambda_2$	μ_2	$A \cdot \mu_2$	Tier
TE10_mito_1	69,786	79	29	0.260	7.99e-07	0.73	A
TE11_mito_1	23,241	36	7	0.670	2.48e-05	1.91	B
TE11_mito_2	63,898	98	9	0.128	1.93e-06	0.38	A
TE12_mito_1	34,382	41	19	0.207	3.33e-06	0.59	B
TE2_mito_1	72,339	81	31	0.075	5.06e-07	0.28	A
TE3_mito_1	63,872	80	11	0.143	1.77e-06	0.45	A
TE4_mito_1	45,823	59	9	0.072	3.37e-06	0.24	B
TE5_mito_1	8,687	45	14	0.101	9.52e-06	0.32	B
TE5_mito_2	12,141	28	13	0.130	7.93e-06	0.39	B
TE6_mito_1	38,103	47	12	0.170	3.42e-06	0.50	B
TE7_mito_1	56,052	82	10	0.420	6.28e-06	1.28	A
TE7_mito_2	26,429	47	11	0.404	1.25e-05	1.20	B
TE7_mito_3	23,052	48	14	0.095	3.46e-06	0.29	B
TE8_mito_1	78,784	93	23	0.303	8.60e-07	0.84	A
TE9_mito_1	75,446	93	15	0.649	3.87e-06	1.93	A

pairwise Mann–Whitney, uncorrected: IMM $<$ OMM, $p = 0.046$; IMM $<$ ER, $p = 0.039$; OMM vs. ER not significant, $p = 0.37$). The cotangent operator provided sharper separation: median μ_2 was 3.42×10^{-6} for IMM ($A \cdot \mu_2 = 0.50$), 9.69×10^{-6} for OMM ($A \cdot \mu_2 = 0.92$), and 4.27×10^{-5} for ER ($A \cdot \mu_2 = 1.35$) (Kruskal–Wallis $p = 9.4 \times 10^{-4}$; exploratory pairwise, uncorrected: IMM $<$ OMM, $p = 0.002$; IMM $<$ ER, $p = 0.002$; OMM vs. ER not significant, $p = 0.27$). This ordering reflects the known membrane architecture: inner mitochondrial membranes are the most compartmentalized, outer mitochondrial membranes form a smoother envelope, and endoplasmic reticulum membranes have the most open organization. Cross-type medians are reported as raw μ_2 with the scale-independent product $A \cdot \mu_2$ (defined in Methods) shown in parentheses; the IMM $<$ OMM $<$ ER ordering is preserved under area normalization, confirming that it is not an artifact of systematic surface-area differences.

Because boundary fraction differs systematically across membrane types (IMM $>$ OMM; Wilcoxon $p = 1.2 \times 10^{-4}$), we tested whether the cross-type spectral ordering survived after controlling for this covariate. In a linear model (spectral metric \sim membrane type + boundary fraction), the membrane-type effect on log-transformed cotangent μ_2 remained significant (F-test $p = 0.003$), while the effect on $n \cdot \lambda_2$ did not ($p = 0.13$). The graph operator’s cross-type separation is thus partially confounded by systematic boundary-fraction differences, whereas the cotangent opera-

tor retains significant cross-type separation after adjusting for boundary fraction. This further supports the cotangent branch as the more biologically informative cross-type discriminant. The paired IMM–OMM comparison, which controls for boundary fraction by design, is unaffected by this covariate.

Heat trace and spectral entropy—multiscale descriptors derived from the full eigenvalue spectrum—were highly stable within membrane classes ($CV < 0.5\%$ for heat trace, $< 0.002\%$ for spectral entropy within IMM) and significantly discriminative across classes (Kruskal–Wallis $p = 0.005$ for heat trace, $p = 0.006$ for spectral entropy). Within a membrane class, these descriptors serve as reproducibility benchmarks; across classes, they reflect differences in multiscale membrane organization.

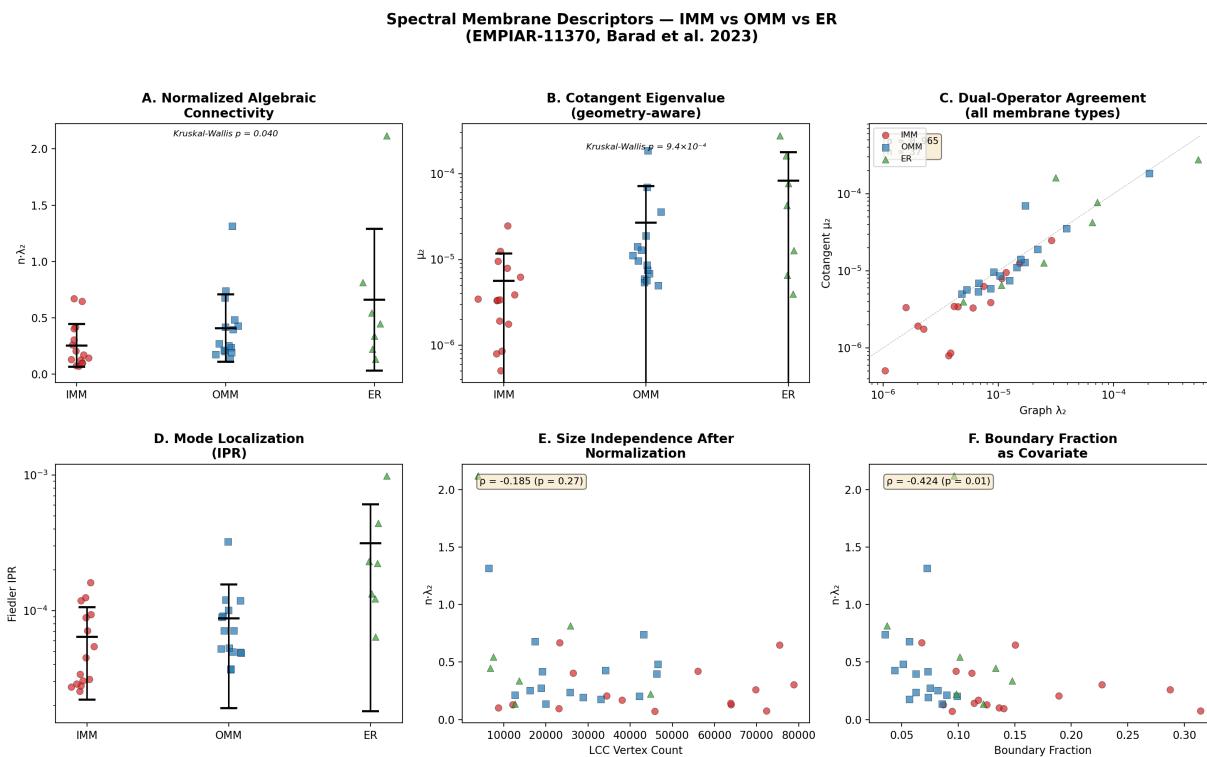


Figure 4: Cross-membrane-type spectral descriptors. (A) Normalized algebraic connectivity $n \cdot \lambda_2$ by membrane type (Kruskal–Wallis $p = 0.040$). (B) Cotangent μ_2 by membrane type ($p = 9.4 \times 10^{-4}$). (C) Dual-operator agreement (log-log scatter). (D) Fiedler inverse participation ratio. (E) $n \cdot \lambda_2$ vs. vertex count, showing size-independence after normalization. (F) $n \cdot \lambda_2$ vs. boundary fraction (covariate check; see text for within-type and ANCOVA analysis). 15 IMM (red circles), 15 OMM (blue squares), 7 ER (green triangles).

2.5 Normalization, boundary sensitivity, and robustness

Vertex-count normalization was essential for cross-mesh comparison. Raw λ_2 was significantly confounded with mesh density (Spearman $\rho = -0.59$, $p = 0.020$ within IMM). The $n \cdot \lambda_2$ normalization removed this dependence ($\rho = 0.19$, $p = 0.51$).

Boundary fraction—the proportion of mesh edges on the open boundary—was a significant covariate across membrane types (pooled $\rho = -0.42$, $p = 0.009$) but not a significant predictor of $n \cdot \lambda_2$ within the IMM cohort ($\rho = -0.12$, $p = 0.68$). Within IMM, cotangent μ_2 showed a marginally significant boundary sensitivity ($\rho = -0.55$, $p = 0.035$). The graph and cotangent branches thus have different sensitivity profiles to open-boundary truncation: the normalized graph metric is largely stable within IMM, whereas the geometry-aware cotangent branch remains partially boundary-sensitive. We recommend reporting boundary fraction alongside all spectral descriptors and interpreting the cotangent branch with this covariate in view.

Sensitivity analysis on three representative IMM meshes (TE4, TE11_mito_1, TE2) confirmed robustness to mesh perturbation. At 10% face dropout, graph λ_2 shifted by less than 2% and cotangent μ_2 by less than 6%. Vertex perturbation at $\sigma = 0.1$ (relative to median edge length) produced less than 0.1% change in λ_2 and less than 4% in μ_2 . The graph Laplacian is essentially perturbation-invariant, reflecting its purely topological nature; the cotangent operator is more geometry-sensitive but remains stable within the range of mesh-quality variation encountered in cryo-ET data.

3 Discussion

We have shown that spectral connectivity descriptors derived from graph Laplacian and cotangent Laplace–Beltrami operators provide a biologically informative layer of analysis for cryo-ET membrane surfaces. The central empirical finding is the paired IMM–OMM comparison: the geometry-aware cotangent operator separated inner from outer mitochondrial membranes within the same organelle in all 15 pairs, capturing a membrane-layer distinction consistent with known crista-associated bottlenecks in inner mitochondrial membranes more sensitively than the topology-only graph operator.

This result addresses a gap in the existing Surface Morphometrics toolkit. Curvature, thickness,

and inter-membrane distance are local measurements—they characterize the membrane at each point independently. Spectral descriptors, by contrast, are global: λ_2 and μ_2 encode the membrane’s worst bottleneck anywhere on the surface, and the Fiedler vector localizes that bottleneck without requiring prior knowledge of where to look. The two levels of description—local and global—are complementary rather than competing. Whether the spectral layer adds quantitative predictive power beyond local morphometrics—for example, in regression models predicting functional or condition-dependent outcomes—remains to be tested in future work.

The dual-operator design proved essential rather than ornamental. At the population level, the graph and cotangent operators agreed strongly on the bottleneck hierarchy across 15 IMM membranes ($r = 0.977$). But the four reversal cases in the paired IMM–OMM comparison—where the graph operator failed to separate while the cotangent operator succeeded—demonstrate that the two operators capture partially independent structural information. The graph Laplacian measures connectivity; the cotangent operator also measures how wide connections are. For membranes whose constrictions are topologically present but geometrically narrow, only the geometry-aware operator detects the constriction. A single-operator approach would miss this distinction.

A methodological finding deserves emphasis: the deposited cryo-ET membrane meshes in EMPIAR-11370 were substantially fragmented, with IMMs averaging 17.5 disconnected components. This fragmentation is consistent with the Screened Poisson surface reconstruction applied to incomplete segmentations. Because spectral descriptors require connected graphs, explicit connected-component handling is necessary—not as an optional cleanup step but as a prerequisite for spectral interpretability. We adopted LCC extraction as the default, reported fragmentation statistics for every mesh, and note that future work should examine whether fragmentation patterns carry independent biological information.

Several limitations should be noted. First, our dataset comprises 11 tomograms from a single cell type (mouse embryonic fibroblasts) under control conditions. The paired design is robust, but condition-dependent changes (e.g., responses to stress, disease states, or pharmacological perturbation) remain to be tested. Second, boundary fraction is a covariate for the cotangent operator: truncated meshes may underestimate geometric connectivity. We report boundary fraction for all analyses and recommend it as a standard covariate in spectral membrane studies. Third, we have not yet demonstrated that the spectral descriptors add quantitative predictive value beyond exist-

ing local morphometrics in a formal ablation or regression framework; the present work establishes that they capture a qualitatively different—global rather than local—aspect of membrane organization. Fourth, because the present study uses public deposition meshes generated by a fixed reconstruction workflow, part of the future value of the method will be to test how the descriptors behave under alternate segmentation and meshing pipelines.

In summary, spectral connectivity descriptors are computationally lightweight, straightforwardly compatible with existing Surface Morphometrics outputs, and provide a global connectivity characterization that complements local morphometric measurements. The paired IMM–OMM result demonstrates genuine biological resolving power: these descriptors recover a membrane-layer distinction consistent with known crista-associated bottlenecks in inner mitochondrial membranes. We propose spectral descriptors as a practical extension of the cryo-ET membrane analysis toolkit.

4 Materials and Methods

4.1 Data source and mesh extraction

Surface meshes were extracted from cryo-ET reconstructions deposited in EMPIAR-11370 (Barad et al., 2023). This dataset comprises tomograms of mouse embryonic fibroblasts (MEF) treated with vehicle control, from which we analyzed 11 tomograms (designated TE2–TE12) containing segmented mitochondrial membranes and endoplasmic reticulum. Meshes were generated by the Surface Morphometrics pipeline using Screened Poisson surface reconstruction (Kazhdan and Hoppe, 2013) with curvature estimation via pycurv (Salfer et al., 2020), and exported in PLY format. A total of 37 meshes were analyzed: 15 IMM, 15 OMM, and 7 ER segments. Selection criteria: all tomograms containing at least one segmented mitochondrion with both IMM and OMM surfaces were included.

4.2 Preprocessing: connected-component extraction

The mesh adjacency graph was constructed for each extracted surface. Connected components were identified by breadth-first search. The largest connected component (LCC) was extracted as the primary analysis object for all spectral computations. Vertex indices were remapped, and the

LCC was verified to be connected. No smoothing or topology-altering operations were applied. Fragmentation statistics (component count, LCC vertex fraction) were recorded for every mesh. IMM meshes were stratified post hoc into quality tier A (LCC retaining $\geq 70\%$ of original vertices, $n = 7$) and quality tier B ($< 70\%$, $n = 8$) for sensitivity to fragmentation.

4.3 Graph Laplacian spectral analysis

The symmetric normalized graph Laplacian $L^{\text{sym}} = I - D^{-1/2}WD^{-1/2}$ was constructed from the unweighted adjacency matrix W of the LCC. Eigenvalues and eigenvectors were computed using SciPy’s ARPACK shift-invert sparse eigensolver ($\sigma = 10^{-8}$, $k = 10$). The algebraic connectivity λ_2 (second-smallest eigenvalue), spectral gap ratio λ_2/λ_3 , Fiedler vector, inverse participation ratio (IPR), and Fiedler-derived conductance were computed. Vertex-count normalization $n \cdot \lambda_2$ was applied for all cross-mesh comparisons. For graphs embedded in surfaces of fixed geometric complexity, algebraic connectivity is expected to scale approximately as $1/n$; the product $n \cdot \lambda_2$ removes this scaling at leading order, yielding a quantity that reflects bottleneck geometry rather than mesh density.

4.4 Cotangent Laplace–Beltrami spectral analysis

The cotangent stiffness matrix C and lumped mass matrix M were assembled from triangle geometry following the standard finite-element discretization (Meyer et al., 2003). The generalized eigenvalue problem $C\varphi = \mu M\varphi$ was solved using SciPy’s ARPACK shift-invert sparse eigensolver ($\sigma = 10^{-8}$, $k = 6$). The second eigenvalue μ_2 and associated eigenvector were extracted. For scale-independent comparison, the area-normalized quantity $A \cdot \mu_2$ was computed, where A is the total LCC surface area; this product is dimensionless and invariant under uniform rescaling of the mesh, since μ_2 scales as inverse area while A scales as area. No mesh coarsening was required: the shift-invert solver handled all LCCs (up to 78,784 vertices) within seconds per mesh on a standard workstation.

4.5 Heat-kernel features

A geometric diffusion-time grid τ was derived independently for each mesh from its median edge length h : $\tau \in [0.5 h^2, 50 h^2]$, 20 log-spaced values. This mesh-specific grid ensures that summary

scalars at the midpoint of the τ -grid are compared at matched relative diffusion scales across meshes. Heat trace $K(\tau) = \sum_i \exp(-\lambda_i \tau)$, spectral entropy $H(\tau) = -\sum_i p_i \log p_i$ (where $p_i = \exp(-\lambda_i \tau)/K(\tau)$), and spectral dimension $d_s(\tau) = -2 d \log K / d \log \tau$ were computed at each scale.

4.6 Quality control

An automated QC module evaluated 10 checks per mesh: vertex count, boundary fraction, degenerate faces, minimum triangle angle, edge-length uniformity, connectivity, λ_2 near zero, spectral gap, and Fiedler localization. Flags were severity-graded (PASS, NOTE, WARN, FAIL). All meshes were flagged for near-degenerate triangle angles; this was confirmed to not affect cotangent eigenvalues (see Results). Meshes with boundary fraction $> 20\%$ were flagged as heavily truncated.

4.7 Sensitivity analysis

Three representative IMM meshes (TE4, TE11_mito_1, TE2) were subjected to perturbation analysis. Face dropout (1%, 5%, 10% random face removal with LCC re-extraction) and vertex perturbation (Gaussian noise at $\sigma = 0.01, 0.05, 0.10$ relative to median edge length, applied along vertex normals) were tested with 2–3 replicates per condition. All three meshes were classified as stable ($< 15\%$ maximum change at the mildest perturbation level).

4.8 Statistical analysis

Paired within-mitochondrion comparisons used the Wilcoxon signed-rank test (two-sided). Unpaired cross-membrane-type comparisons used the Kruskal–Wallis test with pairwise Mann–Whitney U follow-up (reported in text). Cross-type covariate analysis used ordinary least-squares linear models of the form *spectral metric* \sim *membrane type* + *boundary fraction*, with membrane type encoded as dummy variables (IMM as reference) and log-transformed cotangent μ_2 as the response variable; the membrane-type effect was assessed by an F-test comparing the full model to a boundary-fraction-only reduced model. Correlations were assessed by both Pearson and Spearman coefficients. All analyses were performed in Python 3.11 using SciPy 1.11 and NumPy.

4.9 Software availability

The spectral membranes pipeline is implemented as a lightweight Python package (`spectral_membranes` v4) with dependencies limited to NumPy, SciPy, and Matplotlib. The code consumes `.ply` or `.obj` mesh files produced by the Surface Morphometrics pipeline and outputs per-mesh spectral feature tables in CSV format. Source code is available at [https://github.com/\[to-be-deposited\]](https://github.com/[to-be-deposited]) and will receive a permanent DOI via Zenodo upon publication.

References

- Barad, B.A., Medina, M., Fuentes, D., Wiseman, R.L., and Grotjahn, D.A. (2023). Quantifying organellar ultrastructure in cryo-electron tomography using a surface morphometrics pipeline. *J. Cell Biol.* **222**, e202204093.
- Chung, F.R.K. (1997). *Spectral Graph Theory*. CBMS Regional Conference Series in Mathematics, Vol. 92. American Mathematical Society.
- Du, C.-J., Hawkins, P.T., Stephens, L.R., and Bretschneider, T. (2013). 3D time series analysis of cell shape using Laplacian approaches. *BMC Bioinformatics* **14**, 296.
- Fiedler, M. (1973). Algebraic connectivity of graphs. *Czech. Math. J.* **23**, 298–305.
- Kazhdan, M., and Hoppe, H. (2013). Screened Poisson surface reconstruction. *ACM Trans. Graph.* **32**, 29.
- Medina, M., Chang, Y.-T., Rahmani, H., Frank, M., Khan, Z., Fuentes, D., Heberle, F.A., Waxham, M.N., Barad, B.A., and Grotjahn, D.A. (2026). Surface Morphometrics reveals local membrane thickness variation in organellar subcompartments. *J. Cell Biol.* **225**, e202505059.
- Meyer, M., Desbrun, M., Schröder, P., and Barr, A.H. (2003). Discrete differential-geometry operators for triangulated 2-manifolds. In *Visualization and Mathematics III*, pp. 35–57. Springer.
- Pinkall, U., and Polthier, K. (1993). Computing discrete minimal surfaces and their conjugates. *Exp. Math.* **2**, 15–36.

Reuter, M., Wolter, F.-E., and Peinecke, N. (2006). Laplace–Beltrami spectra as “Shape-DNA” of surfaces and solids. *Comput. Aided Des.* **38**, 342–366.

Salfer, M., Steigärter, J.F., et al. (2020). Reliable estimation of membrane curvature for cryo-electron tomography. *PLoS Comput. Biol.* **16**, e1007962.

von Luxburg, U. (2007). A tutorial on spectral clustering. *Stat. Comput.* **17**, 395–416.