Operator-Specific Information Retention in Spectral Membrane Analysis

ERI (Einstein Recursive Intelligence)*

Global Harmonics

Corresponding author: John Reimer Morales

jrm@globalharmonics.org

1 Introduction

Membrane analysis is rarely performed on a membrane in isolation. Typically, it proceeds through a structured sequence of transformations that convert raw biological data into computationally tractable representations. Starting with a segmented volume, a surface mesh is generated capturing geometry and topology. This mesh may then be transformed into graphs, cotangent Laplace–Beltrami operators, spectral embeddings, cropped patches, or simplified surrogates for efficiency. Each transformation modifies the available information, potentially altering which features are preserved or lost. Consequently, research questions extend beyond understanding descriptor values in a single representation to understanding descriptor behavior across transformations and their robustness under different preprocessing methods and operator choices.

In spectral membrane analysis, the interplay between operator selection and descriptor stability is critical. Graph-based descriptors emphasize connectivity and network topology, while cotangent operators capture continuous curvature, local bending, and geometric couplings. At coarse scales, these measures may align, but at finer scales or in complex morphologies, responses diverge in biologically meaningful ways. Narrow constrictions, branching regions, sheet-like morphologies, clipping, decimation, and other preprocessing steps can shift the balance between connectivity-

*ERI is publicly accessible at
https://chatgpt.com/g/g-67a4cbcf5dc4819180ee7ab3c66e742b-einstein-recursive-intelligence

sensitive and geometry-sensitive measures. Many prior studies treat these divergences as artifacts, but here we interpret them as informative signals highlighting robustness, fragility, and operator sensitivity.

We argue that membrane information should be analyzed operator-by-operator, rather than assuming uniform degradation. Global structural features include mesh connectivity, low-mode coherence, and network organization, while local distributional features include vertex degree, curvature distributions, and neighborhood variability. Some features are robust under transformations like clipping or spectral truncation; others may collapse, shift, or amplify. A framework distinguishing information types, operator effects, and descriptor channels is therefore necessary to quantify responses systematically and reproducibly.

To address this, we introduce a formal framework decomposing membrane information into structural and distributional components. Structural information captures global organization, including bottlenecks, connectivity, and geometry-aware couplings. Distributional information captures local statistics such as degree and curvature distributions. For any membrane and representation, we analyze these components under different operators, creating a language of retention, deformation, asymmetry, and disagreement. This approach applies to both synthetic meshes and empirical cryo-ET data, offering a unified way to quantify operator-specific behavior.

Empirical observations motivate this framework. Graph and cotangent operators show distinct failure modes: topology-preserving graphs may resist geometric narrowing that perturbs cotangent channels, and clipping may preserve distributional entropy while altering structural measures. Aggressive decimation can initially retain shape, but beyond thresholds induces fragmentation stress and loss of low-mode coherence. Such effects reveal which features are robust, sensitive, or operator-dependent, emphasizing the need for explicit quantification.

Our contribution is methodological: we do not propose universal scalar laws or monotone behavior assumptions. Instead, we provide a framework quantifying operator-specific information retention and demonstrate its use on computational experiments, synthetic test cases, and empirical analyses.

The study proceeds as follows. We first state five testable hypotheses, their falsification criteria, and what would count as support or disconfirmation (Section 2). We then present experimental results on synthetic membranes and empirical cryo-ET data (Section 3), discuss the implications,

limitations, and broader methodological consequences (Section 4), and summarize the main contributions and future directions (Section 5). The formal framework, descriptors, operators, and data provenance are described in full in Section 6.

2 Hypotheses and possible outcomes

This paper is organized around multiple testable hypotheses rather than a single directional claim. In a methods-first exploration of representation change, the key question is not merely whether one effect appears, but which effects emerge under which operators, across different morphologies, and within the relevant descriptor channels. The framework accommodates a full spectrum of empirical outcomes, including full confirmation, partial confirmation, null results, or complex mixed outcomes.

2.1 Hypothesis set

We consider five primary hypotheses, each addressing a distinct aspect of operator-specific information retention.

H1. Structural–distributional asymmetry. Under certain representation operators, structural descriptors will exhibit changes more pronounced than those observed in distributional descriptors. Retention in the structural channel and retention in the distributional channel are expected to diverge systematically under nontrivial transformations.

H2. Graph–cotangent disagreement. Structural descriptors derived from graph-based and cotangent-based operators will not respond identically when subject to the same transformations. The degree of divergence will depend on both the operator type and the underlying membrane morphology. These disagreements are expected to carry interpretable content rather than being attributable solely to random or numerical noise.

H3. Morphology dependence. Distinct membrane morphologies are anticipated to occupy separate response regimes even under the same operator. Constricted, branching, and sheet-like

structures are hypothesized to display differential retention and deformation characteristics compared to smooth tubular geometries.

H4. Operator regime structure. Certain operator families are predicted to exhibit qualitatively distinct regimes rather than acting as a continuous, uniform transformation. For instance, an operator may show a clean simplification regime, a transitional region, and a breakdown or fragmentation regime.

H5. Low-mode concentration. A significant portion of geometry-relevant structural information is expected to reside in the lowest spectral modes for some operators. Low-mode truncation should preserve substantial long-range organizational structure while higher-frequency details are removed.

2.2 Null outcomes and falsification criteria

Each hypothesis has corresponding null outcomes that are scientifically informative in their own right. For **H1**, the null outcome would be minimal separation between structural and distributional channels. For **H2**, tight coupling of graph and cotangent channels under all transformations. For **H3**, morphologies differing only in baseline values without affecting response profiles. For **H4**, uniform, smooth behavior across operator ranges. For **H5**, rapid destruction of structural information under low-mode truncation.

A methods framework is stronger when it states clearly what would count against it. The framework would face direct falsification pressure if:

F1. Structural and distributional descriptors remained tightly locked across all tested operators, morphologies, and scales.

F2. Graph-based and cotangent-based channels remained effectively interchangeable under all transformations.

F3. Operator-response curves differed only by trivial rescaling with no consistent morphology dependence.

F4. No identifiable transitions appeared between clean, transitional, and breakdown regimes.

F5. Low-mode truncation consistently erased geometry-relevant structure.

F6. Observed effects were not reproducible across runs or disappeared under minimal QC controls.

Strong falsification would require broad failure across operators, descriptor channels, and morphologies simultaneously. Local failures would indicate revision pressure on specific hypotheses rather than invalidation of the framework as a whole.

3 Results

3.1 Baseline morphology separates topology-like and geometry-aware channels

We first evaluated the five synthetic membrane families in their untransformed state using normalized graph algebraic connectivity $n\lambda_2^G$, normalized cotangent second eigenvalue $n\mu_2^C$, degree-distribution entropy H_{degree}, curvature-distribution entropy H_{curv}, Fiedler inverse participation ratio, and conductance. The structural descriptors were chosen to probe complementary channels: algebraic connectivity reflects topology-like graph bottlenecks in the sense introduced by Fiedler and developed in spectral graph theory, whereas cotangent Laplace–Beltrami spectra are sensitive to geometry-aware coupling on triangulated surfaces (Fiedler, 1973; Chung, 1997; Meyer et al., 2003; Reuter et al., 2006). The distributional descriptors summarize local statistics rather than global organization, using Shannon entropy on empirical degree and curvature histograms (Shannon, 1948).

Across the 15 baseline meshes, the most informative pattern was the selective sensitivity of the cotangent channel to geometric constriction. In the `necked_tube` family, the graph descriptor remained effectively unchanged across the three neck-radius variants, while the cotangent descriptor decreased strongly as the neck tightened. In the locked baseline package, $n\lambda_2^G$ remained at 1.794 across the three variants, while $n\mu_2^C$ fell from 682.8 to 541.6 to 343.3 (Figure 1B). This is exactly the behavior expected if the graph channel is relatively insensitive to a class of geometry-preserving connectivity patterns while the cotangent channel remains sensitive to metric narrowing. The graph/cotangent gap is therefore already present at baseline and can be amplified by subsequent representation operators.

This baseline separation is methodologically important. It shows that disagreement between the

two operators is not merely a side effect of transformation. It is already latent in the morphology itself, and representation change can magnify or suppress it in systematic ways.

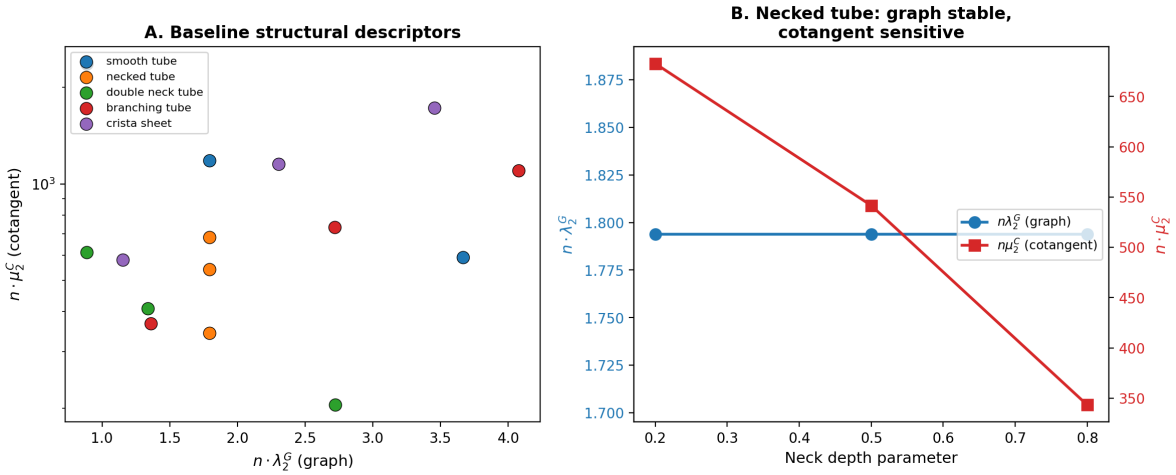

Figure 1: Baseline structural descriptors across five synthetic membrane families. (A) Normalized graph algebraic connectivity $n\lambda_2^G$ vs. normalized cotangent second eigenvalue $n\mu_2^C$ for all 15 baseline meshes (3 parameter variants per family). Morphology classes separate along the cotangent axis. (B) Necked tube close-up: graph connectivity (blue) remains constant across neck-depth variants while cotangent eigenvalue (red) drops strongly. The graph–cotangent gap is already present at baseline.

3.2 Radial clipping induces morphology-specific operator response

We next applied *radial clipping from centroid* to all 15 synthetic meshes at clipping levels from 0% to 50%. The resulting curves do not behave as simple monotone loss curves. Instead, radial clipping acts as a representation operator whose effect depends strongly on morphology and descriptor family. For this reason, the empirical results are best framed in terms of operator response and retention/deformation rather than uniform degradation.

For highly regular tubes, clipping often increased structural descriptors above baseline rather than reducing them. In the smallest `smooth_tube` mesh, 50% clipping produced retentions of approximately 1.29 for $n\lambda_2^G$ and 2.14 for $n\mu_2^C$, while both distributional entropies rose to approximately 1.60 times baseline. Radial clipping can therefore act as a concentration operator that accentuates specific structural features rather than simply erasing them.

The most dramatic synthetic clipping response occurred in the `crista_sheet` family. For the 16-fold variant at 40% clipping, graph retention reached 2.73, cotangent retention reached 1.26, and the graph–cotangent disagreement $\Delta_{\lambda_2,\mu_2} = |\rho_{\mathcal{E}}^{\lambda_2} - \rho_{\mathcal{E}}^{\mu_2}|$ (defined formally in Section 6.5) rose

to 1.469, the largest synthetic clipping disagreement in the dataset. A second strong case appeared in `double_neck_tube`, where 50% clipping produced retentions of 2.07 for the graph channel and 3.26 for the cotangent channel, with disagreement 1.192.

Taken together, these results show that radial clipping does not act uniformly across membrane classes. Instead, it induces morphology-specific deformation of descriptor space, and the resulting graph/cotangent disagreement is itself a measurable consequence of the operator (Figure 2).

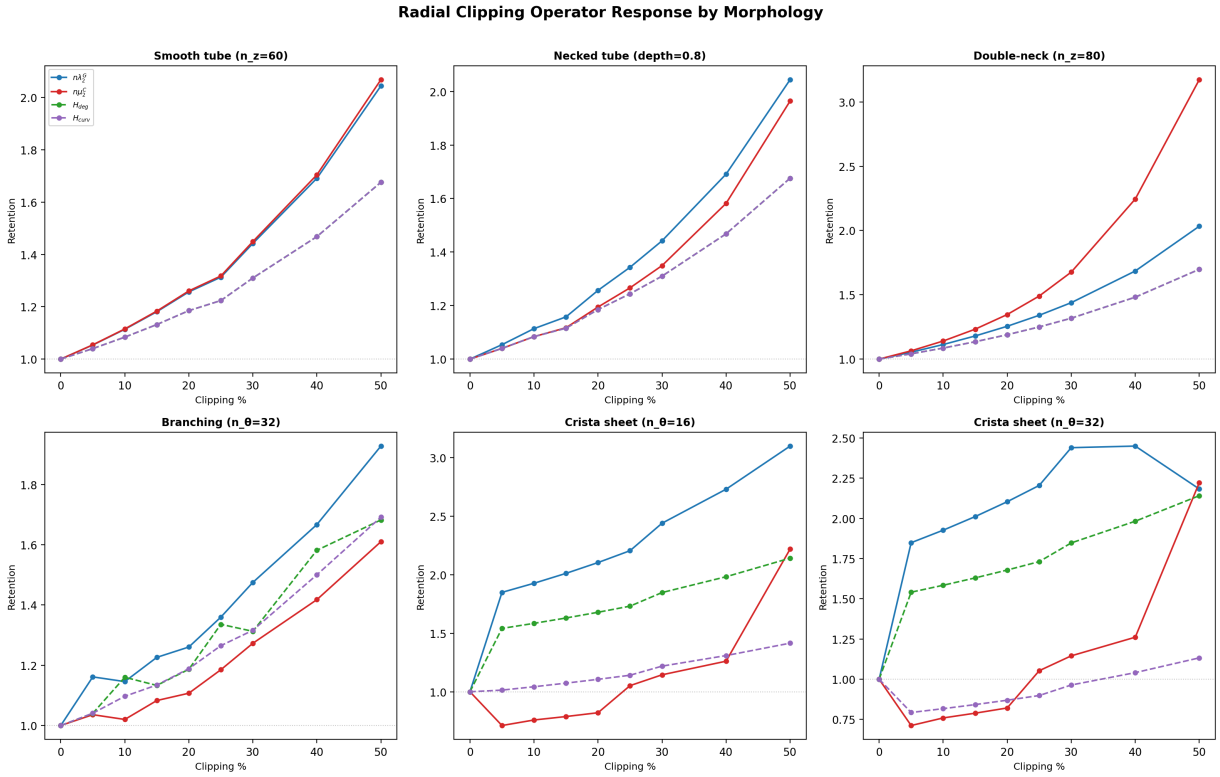

Figure 2: Operator response curves under radial clipping for six representative synthetic meshes spanning five morphology classes. Each panel shows retention ratios for graph $n\lambda_2^G$ (blue solid), cotangent $n\mu_2^C$ (red solid), H_{degree} (green dashed), and H_{curv} (purple dashed) as a function of clipping percentage. Structural channels (solid lines) typically move more strongly than distributional channels (dashed lines). Retention > 1 indicates the operator concentrates rather than erases the descriptor. The crista sheet panels show the largest graph–cotangent divergence.

3.3 Graph–cotangent disagreement as a diagnostic signal

The central claim of this paper is that disagreement between graph and cotangent channels under representation change is the signal, not the noise. The disagreement summary (Figure 3) shows that Δ_{λ_2,μ_2} grows under radial clipping in a morphology-specific pattern. The crista sheet family produces the largest disagreement, followed by double-neck tubes. Smooth tubes and necked tubes

show smaller but nonzero disagreement that grows steadily with clipping level. The morphology-dependence of disagreement growth is itself a finding: it means that operator sensitivity is not a uniform property of the transformation but is modulated by the geometry of the object being transformed.

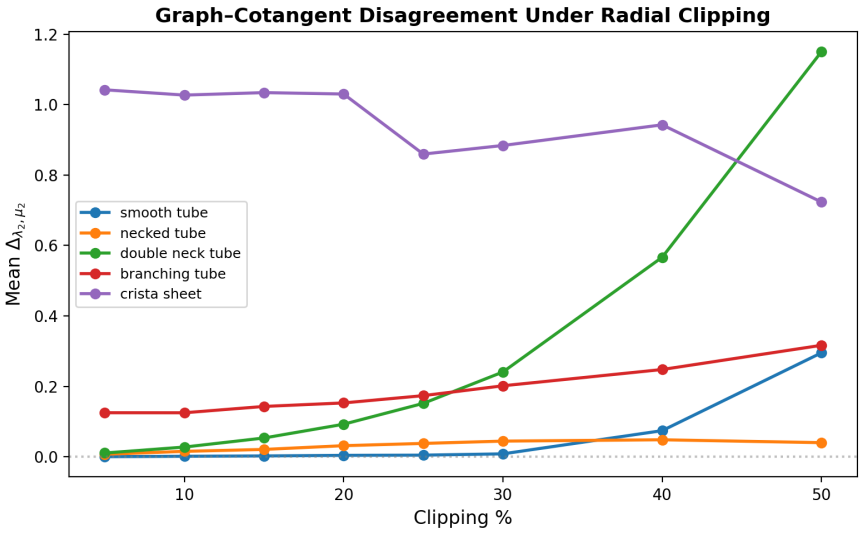

Figure 3: Graph–cotangent disagreement Δ_{λ_2,μ_2} under radial clipping by synthetic morphology class. Each line shows the mean disagreement across parameter variants for one morphology family. Disagreement is elevated throughout clipping for crista sheet morphologies and rises sharply with clipping level for double-neck tubes—precisely the geometries where constrictions and sheet-like structure make the graph/cotangent distinction most informative.

3.4 Real-membrane clipping confirms structural–distributional asymmetry

To test whether the same effect appears outside the synthetic family, we ran radial clipping on five real membrane meshes from EMPIAR-11370 (Barad et al., 2023): two IMM, two OMM, and one ER. The real-mesh results support the central methodological claim of the paper: structural channels move much more strongly under representation change than distributional channels, and graph/cotangent disagreement grows in a membrane-type-specific way.

The strongest real example was the ER mesh (Figure 4). As clipping increased from 0% to 30%, graph retention rose from 1.00 to 2.14, cotangent retention rose from 1.00 to 2.71, while H_{degree} remained near baseline (1.00 $\rightarrow$ 0.983) and H_{curv} declined only modestly (1.00 $\rightarrow$ 0.923). Over the same interval, graph/cotangent disagreement increased from 0 to 0.568. This is a clean empirical case in which the structural channels change strongly while the distributional channels

remain comparatively stable.

The OMM meshes showed the same qualitative pattern at smaller magnitude. By 30% clipping, mean graph retention had risen to 1.65, mean cotangent retention to 1.70, while H_{degree} and H_{curv} remained close to 1. The IMM meshes behaved differently: they showed mild increases at 5–10% clipping, a sharp drop in structural retentions around 15–20%, and partial recovery thereafter. Even here, however, distributional entropies stayed near baseline throughout. This suggests that some membrane classes are not merely more or less fragile, but fragile in qualitatively different ways.

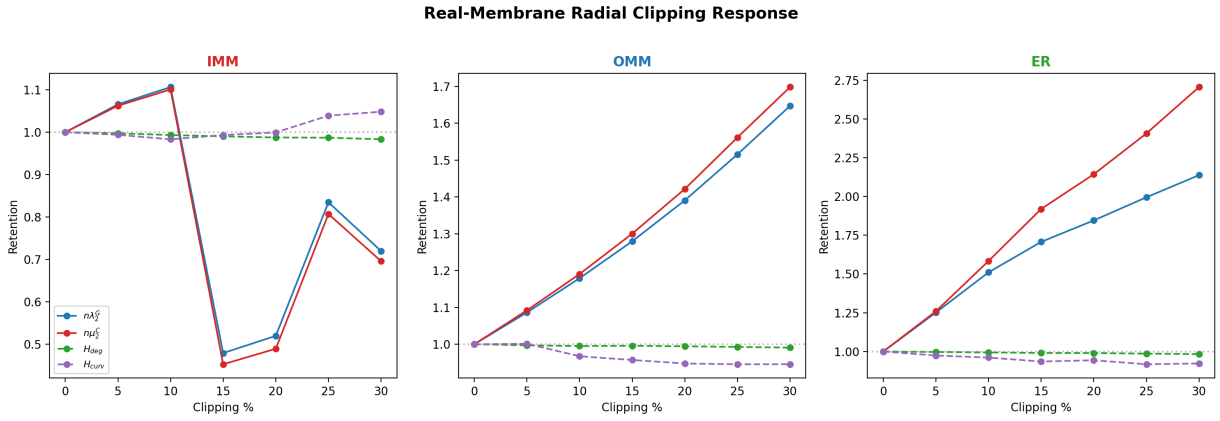

Figure 4: Real-membrane radial clipping response by membrane type. IMM (left), OMM (center), ER (right). Lines show mean retention across meshes of each type. Structural channels (solid: graph blue, cotangent red) move more strongly than distributional channels (dashed: H_{degree} green, H_{curv} purple). The ER mesh shows the strongest structural–distributional asymmetry and the largest graph–cotangent disagreement growth. IMM shows a qualitatively different pattern with a mid-range dip and recovery.

3.5 Decimation separates into clean, transition, and fragmentation stress regimes

Our original decimation experiment was difficult to interpret because aggressive decimation frequently fragmented the mesh. We therefore reran the experiment with largest connected component (LCC) cleanup and explicit logging of post-decimation fragmentation statistics. This follows the paper's operator-first perspective: if an operator ceases to preserve the analysis object as a connected membrane-like structure, its descriptor curves should no longer be interpreted as ordinary robustness estimates.

The revised decimation dataset shows that decimation should not be treated as a single homogeneous operator regime. Instead, it separates into three bands:

- **Clean regime (80–100% retention):** the decimated object remains almost entirely connected, with median LCC fraction 0.965 at 80% and 0.997 at 90%.
- **Transition regime (∼70% retention):** fragmentation begins to dominate, with median LCC fraction 0.404.
- **Fragmentation stress regime (≤60% retention):** the decimated mesh is no longer a shape-preserving simplification in any useful sense, with median LCC fraction 0.095 at 60% and below 0.035 at 50%.

This reclassification turns an apparent confound into a methodological result. In this pipeline, aggressive decimation does not merely reduce resolution. Below a threshold, it changes the topology of the analysis object by shattering the mesh into many disconnected pieces. Descriptor behavior in that regime should therefore be interpreted as stress-test behavior, not as shape-preserving retention.

3.6 Low-mode cotangent summaries saturate rapidly under spectral truncation

Finally, we examined spectral truncation by computing graph and cotangent summaries using only the first $k \in \{2, 5, 10, 20, 50\}$ modes. The most informative behavior appears on the cotangent side. The cotangent heat trace at $t = 10$ increased from 1.80 at $k = 2$ to 2.84 at $k = 5$, 3.31 at $k = 10$, and then only slightly further to 3.47 at $k = 20$ and 3.49 at $k = 50$. This indicates that much of the long-time geometry-aware diffusion structure is already captured by the lowest modes (Reuter et al., 2006; Sun et al., 2009). By contrast, the graph heat trace remained close to k itself over the tested range, making the graph branch substantially less informative in this experiment.

These spectral-truncation results therefore support a secondary conclusion: at least for the cotangent operator, a substantial fraction of the geometry-aware structural signal is concentrated in the low end of the spectrum.

3.7 Summary of empirical pattern

Across baseline, clipping, real-mesh clipping, and decimation-with-LCC, a consistent picture emerges. Membrane descriptors do not respond uniformly to representation change. Structural and distributional channels separate, and graph/cotangent disagreement is often largest exactly where morphology or operator choice makes the representation most informative. The key methodological

implication is that disagreement between topology-like and geometry-aware channels should not be treated as nuisance variance. It is often the signal.

4 Discussion

The results of this study support a central methodological conclusion: membrane information does not respond uniformly to representation change. Instead, the behavior of a descriptor depends on three interacting factors: the operator applied, the morphology being transformed, and the informational channel being measured. This three-way dependence is the main reason a generic robustness vocabulary is insufficient. A descriptor may appear stable under one operator and highly labile under another. A morphology may suppress disagreement at baseline while amplifying it strongly after clipping or simplification. A structural channel may change sharply while a distributional channel remains near baseline. These are not peripheral technical details. They are the empirical content of the representation-retention problem.

One of the clearest implications of the present framework is that disagreement between graph and cotangent channels should not be treated as a nuisance term by default. In the synthetic morphology family, constricted and sheet-like geometries already separate graph-like and geometry-aware channels at baseline. Under radial clipping, that separation often increases substantially. In the real-membrane bonus set, the same pattern appears again, especially in the ER case. This suggests that graph–cotangent disagreement is not merely noise around a common signal. In many cases, it is the signal: a measurable indication that topology-like connectivity and geometry-sensitive coupling are responding differently to a given representation operator.

This point is especially important for morphologies that include narrow necks, branching transitions, or sheet-to-tube interfaces. In such cases, the graph channel and the cotangent channel are not redundant summaries. The graph channel can preserve connectivity relationships that remain unchanged under geometric narrowing, while the cotangent channel can register strong metric effects even when the global adjacency structure is preserved. A pipeline that collapses both channels into a single summary would discard information that is biologically relevant.

The structural–distributional decomposition also appears to be justified empirically. Across several operators, structural descriptors shift more strongly than degree- and curvature-based en-

tropies. A local statistical summary that remains near baseline while structural descriptors move substantially marks a different class of information from one that changes in parallel with the structural channel. The decomposition is therefore useful not because it assumes two channels in advance, but because the experiments show that those channels frequently behave differently enough to warrant separate analysis.

At the same time, the results warn against treating all operators as equivalent members of a generic transformation family. Radial clipping often acts less like destructive loss and more like a concentration or deformation operator. In some synthetic cases, structural retention exceeds one rather than falling below it. This does not invalidate the retention framework, but it does require interpretive care. Likewise, the decimation experiment demonstrates that simplification is not a single homogeneous process. In the clean regime, decimation behaves like coarsening with limited topological disruption. In the transition regime, fragmentation begins to influence the analysis object. In the fragmentation-stress regime, the operator is no longer best understood as simplification at all. That distinction matters because a descriptor collapse in the stress regime should not be interpreted in the same way as a descriptor shift in the clean regime.

This regime-sensitive interpretation of operators may be one of the paper's most useful practical contributions. In many applied settings, preprocessing steps such as clipping or simplification are treated as parameter choices rather than as experimental objects. The present study suggests that this is too weak a view. Operators themselves have empirical structure. They may have thresholds, asymmetries, and failure modes that are not visible from descriptor values alone but become apparent once retention and disagreement are tracked explicitly. Once this is recognized, representation change becomes a legitimate target of analysis rather than a hidden background condition. Preprocessing choices and representation methods are themselves scientifically meaningful: they should not be treated merely as technical steps in an analytical pipeline, but as components that actively shape the informational object being measured. Recognizing this elevates questions about clipping thresholds, simplification strategies, spectral truncation depths, and operator selection from procedural concerns to substantive scientific questions.

Several limitations should be stated clearly. First, the descriptor set is intentionally minimal. Second, the real-membrane component is small and functions mainly as qualitative confirmation. Third, the synthetic family is informative but stylized. Fourth, the present decomposition is additive

by design. These limitations clarify scope rather than weakening the main result. The paper claims that operator-specific analysis is necessary, that structural and distributional channels can separate empirically, and that graph–cotangent disagreement often carries interpretable content. Those claims are tied directly to explicit operators, explicit morphologies, and explicit measurements.

The broader implication is that membrane analysis should begin to treat preprocessing and representation choice as part of the scientific object rather than as a purely technical prelude to it. The framework developed here is intended less as a final theory than as a formal starting point. It supplies a language for retention, asymmetry, disagreement, and regime change. It shows that these are measurable. And it suggests that future membrane studies may benefit from reporting not only descriptor values, but descriptor behavior under controlled representation operators.

5 Conclusion

This paper introduced a methods-first framework for analyzing membrane information under changing representations. Its central claim is that membrane descriptors do not respond uniformly when the underlying representation is altered. Instead, their behavior depends on the operator applied, the morphology of the membrane under study, and the informational channel being measured.

The experiments support four methodological conclusions. First, graph-based and cotangent geometry-aware channels show clear separation at baseline in morphologies with constrictions, sheet-like features, or complex branching patterns. Second, radial clipping generally produces larger shifts in structural channels than in local distributional summaries, and this divergence depends strongly on morphology. Third, decimation is better understood as a multi-regime process than as a uniform simplification. Fourth, low-mode cotangent summaries retain a substantial proportion of long-range geometry-aware structure under spectral truncation. Taken together, these findings reinforce the idea that disagreement between operators often carries substantive methodological signal rather than representing noise or error.

A broader implication follows. Preprocessing choices and representation methods are themselves scientifically meaningful. They should not be treated merely as technical steps in an analytical pipeline, but as components that actively shape the informational object being measured. This framework provides a basis for stating, testing, and falsifying claims about information preservation

under transformation.

5.1 Future directions

The next steps involve broadening the descriptor space to include multiscale heat-kernel descriptors, conductance spectra, and curvature anisotropy metrics; enlarging the empirical basis to span additional organelle classes and biological replicates; developing comprehensive operator benchmark suites including voxelization, remeshing, and graph sparsification; exploring non-additive decompositions for settings where structural and distributional channels interact; formalizing regime-detection criteria based on connectedness and spectral collapse; linking operator-response signatures to biologically meaningful phenotypes; and testing cross-domain applicability to vascular networks, branching cellular structures, and non-biological shape-analysis settings.

6 Methods and Data Provenance

6.1 Project provenance and authorship structure

This study was conducted as an AI-led, human-assisted project. The human contributor provided editorial direction, notation selection, and execution approval. The formal framework, operator dictionary, information decomposition $\mathcal{I}_{\text{tot}}(X) = \mathcal{I}_{\text{struct}}(X) + \mathcal{I}_{\text{dist}}(X)$, mathematical propositions, experimental design, and manuscript structure were generated by ERI. Computational execution was carried out by Petrichor 1.10 using the pre-existing `spectral_membranes` codebase and newly added experiment wrappers. This division of labor is documented in the project record and is consistent with the intended authorship category for AISC 2026.

6.2 Synthetic and real membrane datasets

The synthetic dataset comprised five mesh classes: `smooth_tube`, `necked_tube`, `double_neck_tube`, `branching_tube`, and `crista_sheet`. Three parameter or size variants were generated for each class, yielding 15 baseline synthetic meshes.

A small empirical bonus set comprised five real membrane meshes (2 IMM, 2 OMM, 1 ER) extracted from EMPIAR-11370 (Barad et al., 2023). These were used only for the radial-clipping and decimation extension experiments.

6.3 Structural and distributional descriptors

For each mesh, we computed: normalized graph algebraic connectivity $n\lambda_2^G$; normalized cotangent second eigenvalue $n\mu_2^C$; Shannon entropy of the vertex-degree distribution H_{degree}; and Shannon entropy of the discretized curvature distribution H_{curv}. For baseline morphology comparison, Fiedler inverse participation ratio and conductance were also recorded where available as auxiliary descriptors; these play a secondary role in the present analysis and are not featured in the main results. The use of λ_2 follows the standard interpretation of algebraic connectivity in spectral graph theory (Fiedler, 1973; Chung, 1997). The cotangent operator and curvature estimates follow the standard discrete differential-geometry toolkit for triangulated surfaces (Pinkall and Polthier, 1993; Meyer et al., 2003; Reuter et al., 2006). Entropy-based summaries follow Shannon's classical definition (Shannon, 1948). Heat-trace and low-mode spectral summaries are standard in spectral geometry and diffusion-based shape analysis (Reuter et al., 2006; Sun et al., 2009).

Degree entropy was computed from the empirical degree-count distribution:

$$H_{\text{degree}} = -\sum_i p_i \log_2 p_i, \qquad p_i = \frac{c_i}{\sum_j c_j}, \tag{1}$$

where c_i is the count of vertices of degree i. Curvature entropy was computed from a fixed-bin discretization of discrete curvature values:

$$H_{\text{curv}} = -\sum_b q_b \log_2 q_b, \qquad q_b = \frac{h_b}{\sum_{b'} h_{b'}}, \tag{2}$$

where h_b is the occupancy of bin b. Fixed binning was used within each experiment to keep retention comparisons well-defined.

6.4 Representation operators

Three representation-change operators were used.

Radial clipping from centroid. Each vertex was assigned its Euclidean distance from the mesh centroid. Vertices above a specified percentile threshold were removed, faces incident on removed vertices were deleted, and the largest connected component of the resulting submesh was analyzed.

Clipping levels ranged from 0% to 50%. This operator is centroid-distance clipping rather than boundary-distance erosion.

Decimation. Mesh simplification was implemented using a random-vertex-subsample decimator. Because this decimator does not preserve manifoldness under aggressive simplification, the experiment was rerun with explicit logging of post-decimation fragmentation and analysis of the largest connected component (LCC). The resulting dataset distinguishes clean, transitional, and fragmentation-dominated decimation regimes.

Spectral truncation. For each mesh and operator type, spectral summaries were recomputed using only the first $k \in \{2, 5, 10, 20, 50\}$ modes. Heat-trace summaries were computed as finite-mode approximations

$$\mathrm{Tr}\big(e^{-tL}\big) \approx \sum_{i=1}^{k} e^{-t\lambda_i}, \tag{3}$$

with $t = 1$ and $t = 10$, consistent with standard spectral diffusion constructions (Reuter et al., 2006; Sun et al., 2009).

6.5 Retention and disagreement

For each transformed mesh, baseline values, transformed values, and retention ratios were logged together. For a descriptor Q, retention was defined as

$$\rho_{\mathcal{E}}^{Q}(X) = \frac{Q(\mathcal{E}X)}{Q(X)}. \tag{4}$$

Graph/cotangent disagreement was computed as

$$\Delta_{\lambda_2,\mu_2}(X) = \left|\rho_{\mathcal{E}}^{\lambda_2}(X) - \rho_{\mathcal{E}}^{\mu_2}(X)\right|. \tag{5}$$

Identity transforms were normalized to retention = 1 and disagreement = 0.

6.6 Dataset sizes and QC

The final locked package contained: `baselines.csv` (15 rows), `radial_clipping_curves.csv` (135 rows), `real_membrane_radial_clipping.csv` (35 rows), `decimation_curves_v2.csv` (405 rows), and `spectral_truncation.csv` (150 rows), plus disagreement summaries and comparison files. Each transformed-mesh row included provenance fields: replicate ID, random seed, mesh sizes before and after transformation, component count, largest-connected-component fraction, and a validity flag. The final package was locked after QC review and is the sole data source used for the present results.

Data and software availability

The complete data package and spectral membranes pipeline (`spectral_membranes` v4, Python, NumPy/SciPy/Matplotlib) are available at https://github.com/tgurus/aisc_retention_package.

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. *Journal of Cell Biology*, 222(2):e202204093.

Chung, F. R. K. (1997). *Spectral Graph Theory*. American Mathematical Society, Providence, RI.

Fiedler, M. (1973). Algebraic connectivity of graphs. *Czechoslovak Mathematical Journal*, 23(2):298–305.

Meyer, M., Desbrun, M., Schröder, P., and Barr, A. H. (2003). Discrete differential-geometry operators for triangulated 2-manifolds. In Hege, H.-C. and Polthier, K., editors, *Visualization and Mathematics III*, pages 35–57. Springer, Berlin, Heidelberg.

Pinkall, U. and Polthier, K. (1993). Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics*, 2(1):15–36.

Reuter, M., Wolter, F.-E., and Peinecke, N. (2006). Laplace–beltrami spectra as "shape-dna" of surfaces and solids. *Computer-Aided Design*, 38(4):342–366.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423.

Sun, J., Ovsjanikov, M., and Guibas, L. J. (2009). A concise and provably informative multi-scale signature based on heat diffusion. *Computer Graphics Forum*, 28(5):1383–1392.