Topo-Consynch: Cross-Modal Topological Alignment via the Topological Nyquist Limit and Jiangqiao Topological Extraction Criterion

Lanhaijian
Independent Researcher, China

Abstract

Multimodal perception systems pair continuous visual scalar fields and discrete LiDAR point clouds, yet inherent sampling mismatch generates spurious topological structures, known as topological aliasing. Existing topological data analysis (TDA) pipelines treat topological invariants as post-hoc features and lack first-principles bounds to distinguish physically meaningful cohomological classes from sampling artifacts.

This work establishes the Topological Nyquist Limit (TNL), a fundamental sampling bound separating robust topological modes from unstable aliased components. We further formulate the Jiangqiao Topological Extraction Criterion (JTEC), an operational rule for cross-modal topological fusion. Mathematically, we construct a dual cohomological framework: the continuous visual side adopts persistent de Rham cohomology defined over sublevel-set filtrations on smooth Riemannian manifolds, while discrete LiDAR measurements induce persistent cohomology on Vietoris–Rips simplicial filtrations.

The TNL theorem proves that only $k = 0, 1$ cohomological classes maintain stable correspondence between persistent de Rham cohomology and discrete persistent cohomology under finite sampling. All $k \geq 2$ cohomological signatures arise from sampling noise and must be truncated. We introduce $\mathcal{L}_{\text{curl}} = \|\mathrm{d} \circ \mathrm{d}\,\omega\|^2$ as a cohomology consistency regularizer, directly enforcing the de Rham identity $\mathrm{d}^2 \equiv 0$ and suppressing high-order topological aliasing. The resulting Topo-Consynch framework aligns multimodal geometry at the level of cohomological equivalence classes rather than raw coordinate matching.

The theoretical paradigm unifies continuous differential topology and discrete computational topology, offering a principled alternative to heuristic topological filtering for autonomous perception, remote sensing, and embodied intelligence.

INTRODUCTION

Autonomous perception pipelines commonly fuse two heterogeneous data sources: continuous intensity fields captured by cameras and sparse discrete point samples acquired by LiDAR. Standard registration algorithms operate on pointwise geometric alignment, ignoring global topological consistency. Topological Data Analysis (TDA) provides tools to extract multi-scale structural invariants, but current applications suffer from a critical flaw: there exists no rigorous boundary to decide which topological features are reliable and which are sampling-induced pseudo-structures.

Persistent homology has been widely deployed for scene characterization. However, homology operates on chains and lacks natural differentiability for gradient-based optimization. Persistent cohomology, defined on real-valued cochains, circumvents this limitation, yet the community largely treats discrete persistent cohomology in isolation from the continuous ground-truth topology of the underlying scene manifold.

Recent advances in persistent de Rham cohomology extend classical de Rham theory to filtered smooth manifolds, constructing multi-scale cohomological evolution for continuous spaces. Nevertheless, prior work focuses on single-modality analysis and does not address cross-modal matching between continuous field filtrations and discrete point-cloud filtrations.

Our contributions are summarized as follows:

1. We establish a dual cohomological framework linking persistent de Rham cohomology (continuous visual manifold) and persistent cohomology (discrete LiDAR Rips complexes).

2. We prove the Topological Nyquist Limit (TNL): only $k = 0, 1$ cohomological classes enjoy stable approximation under finite sampling in 3D scene manifolds.

3. We propose the Jiangqiao Topological Extraction Criterion (JTEC), which systematically truncates all $k \geq 2$ cohomological components as topological aliasing.

4. We interpret $\mathcal{L}_{\text{curl}}$ as a cohomology consistency detector derived from the exterior derivative identity $\mathrm{d}^2 \equiv 0$, rather than an empirical loss term.

5. We instantiate Topo-Consynch, an end-to-end topological fusion principle built upon TNL and JTEC for multimodal perception.

PRELIMINARIES

Continuous de Rham and Persistent de Rham Cohomology

We model the physical scene as a compact C^2 Riemannian manifold $\mathcal{M} \subset \mathbb{R}^3$ with manifold reach $\tau_{\mathcal{M}}$. A visual sensor induces a smooth scalar field $f \in C^\infty(\mathcal{M})$, whose intrinsic topological structure is rigorously characterized by

de Rham cohomology:

$$H^k_{\mathrm{dR}}(\mathcal{M}) = \frac{\ker \mathrm{d}^k : \Omega^k(\mathcal{M}) \to \Omega^{k+1}(\mathcal{M})}{\operatorname{im} \mathrm{d}^{k-1} : \Omega^{k-1}(\mathcal{M}) \to \Omega^k(\mathcal{M})},$$

where $\Omega^k(\mathcal{M})$ denotes the space of smooth k-forms on $\mathcal{M}$, and d is the exterior derivative obeying the fundamental identity $\mathrm{d}^2 \equiv 0$. Classical de Rham theory provides single-scale topological invariants encoding intrinsic holes, connectivity, and global structural loops of the continuous scene manifold.

To model multi-scale topological evolution of visual observation, we adopt **persistent de Rham cohomology** over the sublevel-set filtration $\mathcal{F}_c(t) = \{x \in \mathcal{M} \mid f(x) \leq t\}$. This construction extends static de Rham cohomology to a filtered manifold embedding sequence, tracking the birth and death of continuous cohomological classes across geometric scales. The smoothed Euler characteristic $\tilde{\chi}(t)$ serves as a valid global summary statistic for the full persistent de Rham spectral evolution, enabling stable Taylor-scale decomposition of continuous topological signals.

Discrete Persistent Cohomology for LiDAR Point Clouds

Given a discrete LiDAR point sample $P \subset \mathcal{M}$ with sampling granularity h, we construct the family of Vietoris–Rips complexes $\mathcal{R}(P, \varepsilon)$ parameterized by scale radius ε. The discrete topological signature is extracted via **persistent cohomology**, defined on the cochain complex $\{C^k, \mathrm{d}^k\}$ induced by the simplicial filtration:

$$C^k = \mathbb{R}^{\#\{k\text{-simplices}\}}, \quad \mathrm{d}^k : C^k \to C^{k+1}.$$

Persistent cohomology tracks the lifespan of discrete cohomological classes, producing a persistence diagram Dgm^k that quantifies multi-scale topological structures emerging from discrete sensor sampling.

Crucially, persistent cohomology operates on real-valued cochain functions, naturally supporting differentiable optimization and gradient-based topological learning—an essential property for constructing trainable multimodal topological losses. The bottleneck stability bound

$$d_B(\mathrm{Dgm}_{\mathrm{dR}}, \mathrm{Dgm}_{\mathrm{discrete}}) \leq 2h$$

characterizes the maximal admissible discrepancy between continuous persistent de Rham classes and discrete persistent cohomology classes induced by finite sampling density.

Theoretical Interpretation of $\mathcal{L}_{\mathrm{curl}}$ Regularizer

The exterior derivative identity $\mathrm{d}^2 \equiv 0$ is an exact analytical constraint in continuous de Rham cohomology. In discrete persistent cohomology, finite sampling granularity breaks this exactness, generating spurious high-order cohomological modes that violate the de Rham closure condition.

The curl regularizer

$$\mathcal{L}_{\mathrm{curl}} = \|\mathrm{d} \circ \mathrm{d}\,\omega\|^2$$

is no longer an empirical engineering penalty—it is a **cohomology consistency detector**: it quantitatively measures the sampling-induced violation of $\mathrm{d}^2 \equiv 0$, selectively suppressing unphysical $k \geq 2$ discrete cohomological aliasing while preserving valid low-order topological consensus.

TOPOLOGICAL NYQUIST LIMIT AND JIANGQIAO TOPOLOGICAL EXTRACTION CRITERION

The core mechanism of the Topological Nyquist Limit (TNL) and Jiangqiao Topological Extraction Criterion (JTEC) is precisely formalized via dual cohomological systems:

- **Continuous side**: Persistent de Rham cohomology defines the true scale-dependent topological evolution of the scene manifold.
- **Discrete side**: Persistent cohomology provides sampled, scale-filtered topological approximations from LiDAR measurements.

- **TNL bound**: Only $k = \{0, 1\}$ cohomological classes possess lifespan $\ell_k \geq 2h$, ensuring stable correspondence between persistent de Rham and discrete persistent cohomology.
- **JTEC rule**: All $k \geq 2$ discrete cohomological components are sampling artifacts without continuous de Rham counterparts, and must be systematically truncated as topological aliasing.

This formulation upgrades the original homology-based argument to a fully consistent cohomological duality framework, eliminating theoretical asymmetry between continuous manifold geometry and discrete simplicial topology.

Scaling Law of Cohomological Lifespans

For compact Riemannian manifolds with reach $\tau_{\mathcal{M}}$, the lifespan of cohomological classes obeys

$$\ell_k \sim \mathcal{O}\big(\tau_{\mathcal{M}}^{1/(k+1)}\big).$$

As cohomological degree k increases, characteristic structural scales shrink rapidly. For $k \geq 2$, the intrinsic scale falls below the sampling threshold $2h$. Consequently, discrete sampling cannot resolve these modes without generating topological aliasing. The TNL arises directly from this geometric scaling behavior.

JTEC as a Projection Operator on Cohomology Spaces

From a functional perspective, JTEC imposes a projection operator onto the subspace spanned by $k = 0, 1$ cohomological classes. The cross-modal topological alignment loss for Topo-Consynch only evaluates discrepancy within this admissible subspace. High-order cohomological information is discarded, and $\mathcal{L}_{\text{curl}}$ penalizes any residual high-order mode activation during optimization.

TOPO-CONSYNCH PIPELINE

We outline the operational workflow:

1. Construct sublevel-set filtration $\mathcal{F}_c(t)$ from visual scalar field; compute persistent de Rham cohomology.
2. Build Vietoris–Rips filtration $\mathcal{R}(P, \varepsilon)$ from LiDAR point cloud; compute discrete persistent cohomology.
3. Apply JTEC truncation: retain only $k = 0, 1$ persistence diagrams.
4. Compute topological matching loss between the low-order cohomological signatures.
5. Optimize jointly with $\mathcal{L}_{\text{curl}}$ to suppress spurious high-order discrete cohomological modes.

DISCUSSION

Prior literature has developed persistent de Rham cohomology and Discrete Exterior Calculus (DEC) as separate mathematical tools. What distinguishes this work is the cross-modal sampling bound (TNL) and the implementable truncation criterion (JTEC) bridging continuous manifold cohomology and discrete simplicial cohomology.

The framework reveals an analogy to classical signal sampling theory: Shannon–Nyquist bounds scalar signal frequencies, while TNL bounds admissible cohomological degrees under manifold sampling. Topological aliasing is the geometric-topological counterpart of frequency aliasing in Fourier analysis.

Potential extensions include rigorous numerical validation via DEC, generalization to equivariant cohomological analysis for dynamic scenes, and integration with neural implicit scene representations.

Discussion and Future Work

This work establishes fundamental sampling bounds for cross-modal topological alignment between continuous visual manifolds and discrete point-cloud samples. The Topological Nyquist Limit (TNL) and Jiangqiao Topological Extraction Criterion (JTEC) resolve a long-standing ambiguity in topological data analysis: distinguishing physically meaningful cohomological features from sampling-induced topological aliasing. While neural implicit scene representations furnish continuous, differentiable geometric embeddings free from grid-based signal aliasing, such frameworks lack inherent constraints to filter spurious high-order cohomological artifacts. The TNL-JTEC formalism supplies this missing topological bandwidth principle, serving as a first-principles regulator applicable to any continuous geometric representation.

We outline a structured, multi-stage research roadmap for the Topo-Consynch framework, with clear hierarchical boundaries between successive iterations:

v1.x (Current Manuscript) We construct a dual cohomological framework pairing persistent de Rham cohomology for continuous scene manifolds and discrete persistent cohomology for Vietoris–Rips filtrations of point clouds. We prove the TNL theorem asserting $K_{\max} \equiv 1$ and formulate the JTEC criterion to suppress topological aliasing for $k \geq 2$. The core objects of investigation are individual cohomological classes H^k. This version lays the foundational theoretical law governing cross-modal topological matching.

v2.0 (Medium-Term Extension) Two complementary extensions will elevate the framework from abstract algebraic topology to numerically executable dynamic perception systems:

1. *DEC-based computational instantiation*: Abstract simplicial cochain complexes are replaced by Discrete Exterior Calculus (DEC). Dual meshes, discrete Hodge star operators and Hodge decompositions equip the exterior derivative d with geometric consistency satisfying a discrete Stokes theorem. Under DEC, the regularizer $\mathcal{L}_{\text{curl}} = \|\mathrm{d} \circ \mathrm{d}\,\omega\|^2$ acquires a concrete geometric interpretation as an L^2 penalty on discrete curl violations.

2. *G-equivariant cohomology for dynamic scenes*: Rigid motion group $G = \mathrm{SE}(3)$ actions are incorporated, extending static cohomology to G-equivariant cohomology $H^k_G(\mathcal{M})$. The TNL bound is generalized for equivariant topological features, requiring the effective alignment operator $\mathcal{C}_{\text{eff}}$ to satisfy G-covariance.

Neural implicit scene representations act as the primary computational carrier for v2.0 developments. The implicit field $f_\theta : \Omega \to \mathbb{R}^c$ provides a continuous, queryable geometry container where DEC operators and group-equivariant constraints can be jointly optimized.

v3.0 (Long-Term Paradigm Upgrade) The research object undergoes a qualitative dimensional elevation from individual cohomological classes to moduli spaces of cohomological configurations. We adopt the interior-exterior field paradigm: continuous de Rham cohomology of the underlying manifold forms the interior moduli space $\mathcal{H}_{\text{in}}$, while discrete-sampling-induced persistent cohomology spans the exterior moduli space $\mathcal{H}_{\text{out}}$. JTEC is redefined as a projection map

$$\Pi_{\text{JTEC}} : \mathcal{H}_{\text{in}} \times \mathcal{H}_{\text{out}} \to \mathcal{H}_{\text{low}},$$

mapping paired cohomological configurations onto the admissible low-order subspace spanned by $k = 0, 1$. At this stage, topological alignment is reinterpreted as geometric correspondence between two infinite-dimensional parameter spaces, establishing a full topological field-theoretic perspective for multimodal perception.

Future work will first implement v2.0 components atop neural implicit representations to conduct quantitative numerical verification of TNL and JTEC, laying experimental groundwork for subsequent exploration of moduli-space structures in v3.0.

CONCLUSION

This work constructs a dual cohomological framework unifying persistent de Rham cohomology (continuous visual manifolds) and persistent cohomology (discrete LiDAR samples). The Topological Nyquist Limit establishes a first-principles sampling bound, and the Jiangqiao Topological Extraction Criterion delivers a practical rule to eliminate topological aliasing by truncating all $k \geq 2$ cohomological components. The $\mathcal{L}_{\text{curl}}$ regularizer is justified from de Rham's exterior derivative identity, forming a consistent theoretical foundation for the Topo-Consynch multimodal topological alignment principle.

The paradigm provides a mathematically rigorous path to robust topological sensor fusion and connects differential geometry, computational topology, and embodied perception.

ACKNOWLEDGEMENTS

The theoretical derivation was completed in Jiangqiao Town, Jiading District, Shanghai. This article and subsequent revisions are published on Hongqiao University Technology Academic Network.

[1] C. Chen, M. K. Choi, and J. Ye, Persistent de Rham cohomology, arXiv:2102.09378 (2021).
[2] M. Desbrun, A. N. Hirani, M. Leok, and J. E. Marsden, Discrete exterior calculus, arXiv:math/0508341 (2005).
[3] F. Chazal and B. Michel, An introduction to topological data analysis: persistent homology and its applications, arXiv:1710.04024 (2017).