Homotopy Groups of Spheres for Machine Learning

Pavel Chekmaryov[1][0009-0008-9753-8638]

Independent Researcher, San Francisco, USA
p@occurai.com

Abstract. Homotopy groups of spheres, such as $\pi_n(S^3)$, offer a powerful topological framework for machine learning, capturing deformation-invariant properties essential for tasks with n-dimensional symmetry, like protein folding. These groups inspire novel architectures, such as Topological Neural Networks (TopNets), which leverage topological invariants to achieve superior generalization and efficiency compared to standard models like transformers. We demonstrate that TopNets, implemented in PyTorch, are not merely theoretical; they show practical performance for processing symmetric data on conventional hardware. This work bridges the gap between abstract homotopy theory and applied machine learning, presenting a concrete path for creating more robust, symmetry-aware models. We also explore the fundamental computational bottlenecks of spectral sequences, arguing that the full potential of this approach will be unlocked by novel hardware architectures, paving the way for more sophisticated geometric deep learning and advancing artificial general intelligence in topologically complex domains.

Keywords: Homotopy Groups · Topological Machine Learning · N-Dimensional Symmetry

1 Introduction

While deep learning architectures like transformers excel on many tasks, they face fundamental limitations with data possessing intrinsic n-dimensional symmetries, such as those in protein folding or robotic motion planning. Effectively processing such complex, high-dimensional structures may necessitate a paradigm shift beyond the inherent scalability limits of both transformers and classical computing architectures. We argue that a deeper, more principled approach is needed, one grounded in the mathematical field of algebraic topology.

This paper explores the use of homotopy groups of spheres, specifically $\pi_n(S^3)$, as a framework for designing a new class of models that are inherently sensitive to topological invariants. We begin with a practical implementation: Topological Neural Networks (TopNets), which use persistent homology as a scalable proxy to approximate these homotopy-inspired features [3]. These networks demonstrate the immediate value of a topological approach for building robust and efficient models on existing hardware.

However, to fully grasp the power and limitations of this paradigm, we then transition from the practical to the theoretical. We delve into the rich structure of homotopy groups themselves, such as $\pi_3(S^3) = \mathbb{Z}$ and the torsion group $\pi_4(S^3) = \mathbb{Z}_2$, to understand the profound advantages they offer for robust inference and for uncovering the intrinsic geometry of data manifolds. This theoretical exploration reveals a fundamental obstacle: the primary tool for computing these groups, the Adams spectral sequence [1], is computationally intractable on conventional architectures. This bottleneck leads us to a forward-looking conclusion: the full realization of this approach necessitates the development of novel hardware, or "topological processing units," designed for the unique demands of algebraic topology.

Our contributions are thus:

1. We present a clear narrative from the practical application of topological ideas in ML (TopNets) to the deep theoretical foundations of homotopy theory.
2. We analyze the theoretical benefits of homotopy groups for machine learning, connecting their algebraic properties to robust inference and manifold learning.
3. We identify the computational infeasibility of spectral sequences on current hardware as the key bottleneck, and propose the necessity of new hardware architectures, using the Hopf fibration [9] as a motivating benchmark.

2 TopNets: Engaging with Homotopy Groups

TopNets integrate PH with graph neural networks (GNNs) to approximate properties inspired by $\pi_n(S^3)$ [3], offering a scalable approach to topological ML.

2.1 Persistent Homology as a Tractable Proxy

While homotopy groups provide a complete description of a space's "holes," their computation is notoriously difficult. We therefore turn to persistent homology (PH), a more computationally tractable tool that offers a partial window into a space's topological structure. PH approximates topological features for a point cloud in S^3, tracking homology groups (H_0, H_1, H_2) via Rips complexes [6]. It constructs a filtration of simplicial complexes, producing a persistence diagram—a plot of feature birth and death times.

The connection between homology and homotopy is formally established by the Hurewicz theorem, which states that for a simply connected space, the first non-trivial homology and homotopy groups are isomorphic. While the conditions for this theorem do not always hold, it provides a theoretical motivation for using homology as a first-order approximation. For example, a persistent cycle detected by PH in a protein's point cloud can indicate the presence of a non-trivial element in $\pi_3(S^3)$. Thus, while PH is an imperfect proxy, it allows us to develop practical models that are sensitive to topological features, enabling more efficient inference and generalization on symmetric tasks than would be possible with purely geometric approaches.

2.2 GNN Architecture

The point cloud forms a k-NN graph, with nodes holding spherical harmonic features for SO(4)-equivariance [5]. GNNs perform message passing:

$$h_i^{(l+1)} = \phi \left(h_i^{(l)}, \sum_{j \in \mathcal{N}(i)} \psi(h_i^{(l)}, h_j^{(l)}) \right) \tag{1}$$

Training adapts weights to reflect deformation classes, aligning with homotopy-inspired patterns, such as those in robotic perception.

2.3 Proof-of-Concept and Feasibility

To test the core premise, we developed a proof-of-concept prototype in PyTorch, leveraging `torch-topological` and `torch-geometric` [7]. Initial experiments on small S^3 point clouds (100–500 points) using a 4-layer GNN (128 hidden units) confirmed the approach is computationally feasible on standard hardware, with efficient processing times ($\sim$1–5 ms/sample). While these results are preliminary and not representative of production-scale performance, they suggest the model can learn from symmetric data without the extensive data requirements of transformer attention mechanisms. This highlights the potential of the approach, while underscoring that comprehensive benchmarking remains a critical next step for future work.

3 Bridging the Gap: From Homology to Homotopy

While TopNets provide a practical starting point, bridging the gap between persistent homology and true homotopy groups is a critical research direction. This involves exploring architectures that can capture more sophisticated topological information without incurring the intractable cost of spectral sequences. Several promising avenues exist:

- **Equivariant Networks**: Enforcing symmetries like SO(4)-equivariance directly constrains a model to respect the geometry of S^3. This can implicitly capture properties related to $\pi_3(S^3)$, but often comes with a significant computational overhead [12].
- **Simplicial and Cellular Networks**: Moving beyond graphs to simplicial or cell complexes allows models to represent higher-order relations explicitly. Simplicial networks, for example, can model features like cavities and voids directly, potentially providing a more direct path to computing torsion, such as the $\mathbb{Z}_2$ structure of $\pi_4(S^3)$ [2]. These represent a "middle ground" of topological complexity.
- **Neural Fields and Parameterized Topology**: Instead of extracting features from static data, neural fields can parameterize continuous mappings on S^3 [11]. By analyzing the topology of these learned functions themselves (e.g., by studying their level sets), one could potentially infer homotopy-related properties, though this remains an open area of research.

– **Flow-Based and Diffeomorphic Models**: Techniques that model data through learnable, invertible transformations (diffeomorphisms) could, in principle, be used to classify points by determining if they can be continuously deformed into one another, which is the very definition of a homotopy.

These alternative approaches represent a rich and largely unexplored design space for creating more topologically-aware machine learning models.

4 Applications

4.1 Protein Folding

Protein configurations on S^3 leverage $\pi_3(S^3)$-inspired cycles, detected by TopNets, to model molecular symmetries with efficient inference, complementing tools like AlphaFold [10].

4.2 Deformation Detection

Training TopNets to recognize homotopy classes in S^3 mappings adapts weights to continuous deformations, enabling robust robotic perception, enhanced by simplicial networks' higher-order modeling [4].

4.3 Robotic Motion Planning

In robotic motion planning, $\pi_n(S^3)$-inspired TopNets detect stable configurations under rotational symmetries, guiding robots through complex 3D environments with real-time inference [4].

5 Homotopy Groups of the 3-Sphere

Homotopy groups $\pi_n(S^3)$ describe mappings from S^n to S^3, offering algebraic insights into topological structures [8]. Key properties include:

– $\pi_3(S^3) = \mathbb{Z}$: Captures mapping degrees, linked to the Hopf fibration [9], which models rotations in 3D space, such as molecular configurations.
– $\pi_4(S^3) = \mathbb{Z}_2$: Reflects torsion, indicating binary classification of mappings.
– Higher groups: $\pi_5(S^3) = \mathbb{Z}_2$, $\pi_6(S^3) = \mathbb{Z}_{12}$, introducing complex algebraic structures.

As SU(2), S^3 exhibits SO(4)-symmetry, ideal for tasks requiring rotational invariance. In ML, $\pi_n(S^3)$ provide deformation invariance, symmetry alignment, and topological expressivity, enabling models to capture intrinsic data structures. However, computing $\pi_n(S^3)$ is intractable due to its reliance on spectral sequences. Persistent homology approximates low-dimensional topological features, serving as a practical proxy for $\pi_n(S^3)$-inspired models.

6 Theoretical Advantages for Inference and Manifold Insights

Beyond providing a descriptive language for symmetry, homotopy groups offer profound theoretical advantages for machine learning by equipping models with a robust framework for inference and for gaining deep insights into the structure of data manifolds.

6.1 Robust Inference through Topological Invariants

Machine learning models often struggle to generalize when faced with data that has undergone continuous deformations not well-represented in the training set. Homotopy groups provide a natural solution by furnishing topological invariants—properties that remain constant under such transformations.

The core idea is that the homotopy class of a mapping is a deformation-invariant signature. For inference, this means a model sensitive to homotopy can classify data points or patterns based on their fundamental topological structure, rather than their specific geometric embedding. For instance, a model could learn to distinguish between a simple loop and a figure-eight pattern on a manifold by recognizing that one corresponds to a single generator of the fundamental group (e.g., $\pi_1(S^1) = \mathbb{Z}$) while the other represents a more complex element. This leads to models that are inherently robust to noise, rotation, and scaling, drastically improving generalization from smaller datasets. The model learns the abstract *class* of the object, not just its superficial appearance.

6.2 Uncovering Intrinsic Manifold Structure

The algebraic structure of a manifold's homotopy groups reveals its deepest intrinsic properties. By studying these groups, we can uncover hidden constraints, periodicities, and conserved quantities within the data's underlying generative process.

- **The Fundamental Group (π_1)**: A non-trivial fundamental group indicates the presence of one-dimensional "holes" or tunnels in the data manifold. In the context of a system's state space, these loops might correspond to periodic orbits or cyclical processes that are fundamental to the system's dynamics.
- **Higher Homotopy Groups ($\pi_n, n > 1$)**: These groups detect more subtle, higher-dimensional voids. For example, a non-trivial $\pi_2(X)$ implies the existence of 2-dimensional spherical voids in the manifold X. Discovering such a structure could reveal a fundamental constraint or an exclusion principle in the data that would be nearly impossible to detect with purely geometric or statistical methods.
- **Torsion**: The presence of finite (torsion) components in homotopy groups, such as $\pi_4(S^3) = \mathbb{Z}_2$, signifies that certain transformations on the manifold are finite-order. That is, repeating a sequence of deformations a specific

number of times returns the system to its original topological state. This reveals subtle, discrete symmetries within a continuous space, offering a powerful lens for understanding the quantization of states or behaviors in complex systems.

By providing a computational window into these properties, homotopy theory offers a path toward models that do not just fit data, but develop a genuine understanding of the topological landscape in which the data resides.

7 Computational Challenges of Spectral Sequences

The primary tool for computing higher homotopy groups is the Adams spectral sequence [1], which approximates $\pi_n(S^k)$ through a series of pages E_r that converge to the desired group. Each page is computed from the previous one by taking homology: $E_{r+1} = H(E_r, d_r)$. This process is computationally intensive, as the complexity of the algebraic structures on each page grows rapidly.

The nested, iterative nature of these computations—where the output of one complex homology calculation becomes the input for the next—poses a significant challenge for conventional hardware. Modern CPUs and GPUs are optimized for parallel, vectorized operations on large, flat data structures, not for the intricate, pointer-heavy data manipulations required to traverse and compute homology on the increasingly complex algebraic objects of the spectral sequence. This computational bottleneck makes it practically infeasible to compute any but the lowest-degree homotopy groups, severely limiting their direct application in machine learning.

Addressing this bottleneck requires a two-pronged approach. First, there is a pressing need for algorithmic innovation: developing new, more efficient algorithms for computing or approximating topological invariants on existing parallel hardware. Second, and in parallel, we should consider the necessity for novel hardware architectures specifically designed for algebraic topology. Such architectures would need to efficiently handle the sparse, irregular data structures and recursive computations inherent in the Adams spectral sequence. The development of "topological processing units" (TPUs) could unlock the full potential of homotopy groups for ML, enabling the direct computation and utilization of these powerful topological invariants.

A prime example of such a challenging structure is the Hopf fibration, which describes the 3-sphere S^3 as a bundle of circles over the 2-sphere S^2. This hierarchical, non-Euclidean structure is fundamental to understanding 3D rotations and the group $\pi_3(S^3)$, yet it is notoriously difficult to represent and manipulate efficiently on conventional hardware optimized for flat data grids. A future "topological processing unit" would need to handle such fiber bundle structures as first-class citizens, making the Hopf fibration a key benchmark for the development of any such novel architecture.

8 Conclusion

This paper introduces a conceptual framework for applying the deep insights of homotopy theory to machine learning. We have argued that by focusing on topological invariants, we can design models better suited for data with intrinsic symmetries. Our proof-of-concept, TopNets, demonstrates the initial feasibility of this direction, suggesting that even approximations of homotopy-inspired features can lead to practical models for symmetric tasks.

This approach opens several avenues for future work. The most immediate is the systematic exploration of the "missing middle" between simple persistent homology and the full complexity of homotopy groups, including the promising research directions of simplicial networks and diffeomorphic models. Furthermore, while we have focused on S^3, this framework could be extended to other spheres, like S^4, which are relevant in other areas of physics and mathematics.

Ultimately, we argue that the computational challenges of this field, particularly the intractability of spectral sequences on current hardware, should be seen not as a roadblock, but as a driver for innovation. The development of novel algorithms and, potentially, new hardware architectures designed for the unique demands of algebraic topology could unlock the full potential of this powerful mathematical theory, paving the way for a new generation of robust, symmetry-aware machine learning models.

Acknowledgments This work was conducted independently.

Disclosure of Interests The author has no competing interests to declare.

References

1. Adams, J.F.: Stable Homotopy and Generalised Homology. University of Chicago Press (1974)
2. Bodnar, C., Frasca, F., Otter, N., Wang, Y., Lió, P., Montúfar, G., Bronstein, M.: Weisfeiler and lehman go topological: Message passing simplicial networks. In: International Conference on Machine Learning (ICML) (2021)
3. Bodnar, C., et al.: Topological neural networks go persistent, equivariant, and continuous. Journal of Machine Learning Research **25**(1) (2024)
4. Bronstein, M.M., Bruna, J., Cohen, T., Veličkoviç:
5. Cohen, T.S., Geiger, M., Köhler, J., Welling, M.: Spherical cnns. In: International Conference on Learning Representations (ICLR) (2018)
6. Edelsbrunner, H., Harer, J.: Computational Topology: An Introduction. American Mathematical Society (2010)
7. Fey, M., Lenssen, J.E.: Fast graph representation learning with pytorch geometric. In: ICLR Workshop on Representation Learning on Graphs and Manifolds (2019)
8. Hatcher, A.: Algebraic Topology. Cambridge University Press (2002)
9. Hopf, H.: Über die abbildungen der dreidimensionalen sphäre auf die kugelfläche. Mathematische Annalen **104**(1), 637–665 (1931)
10. Jumper, J., et al.: Highly accurate protein structure prediction with alphafold. Nature **596**(7873), 583–589 (2021). https://doi.org/10.1038/s41586-021-03819-2

11. Sitzmann, V., Martel, J.N.P., Bergman, A.W., Lindell, D.B., Wetzstein, G.: Implicit neural representations with periodic activation functions. In: Advances in Neural Information Processing Systems (NeurIPS) (2020)
12. Thomas, N., Smidt, T., Kearnes, S., Yang, L., Li, L., Kohlhoff, K., Riley, P.: Tensor field networks: Rotation- and translation-equivariant neural networks. In: Advances in Neural Information Processing Systems (NeurIPS) (2018)