An initial exploration of using Persistent Homology for Noise-Resilient Tactile Object Recognition

Tarun Raheja*
Kipo AI
San Francisco, USA
tarun@kipo.ai

Nilay Pochhi*
Independent Researcher
San Francisco, USA
pochhi.nilay@gmail.com

Abstract

Tactile object recognition is crucial for robots operating in environments where visual information is unreliable. While traditional machine learning approaches for tactile object recognition often struggle with noise and sensor variations, persistent homology, a tool from topological data analysis, offers a robust representation of object shape across different scales. This paper explores the application of ideas from Topological Data Analysis, specifically, persistent homology to enhance the noise resilience of tactile object recognition. We demonstrate how persistent homology features, specifically persistent entropy, can be extracted from tactile images and combined with traditional features for object classification. Through experiments on a tactile image dataset Gandarias et al. [2019], we present exploratory results on the performance of several sklearn classifiers with and without persistent entropy as a feature, showcasing the improved robustness achieved through the inclusion of topological information.

1 Introduction

Tactile sensing is a vital modality for robots interacting with their environment, providing crucial information about object properties such as shape, texture, and hardness Luo et al. [2017]. This information is especially important in scenarios where vision is limited or unreliable, such as cluttered spaces or low-light conditions Roberge et al. [2021]. Robust tactile object recognition enables robots to perform complex manipulation tasks, even in challenging environments.

Traditional machine learning approaches for tactile object recognition often rely on hand-crafted features extracted from tactile images, including pressure distributions, local curvatures, and statistical moments Madry et al. [2014b,a]. These features, however, can be sensitive to noise, sensor variations, and object pose changes, limiting their robustness and generalizability.

Persistent homology, a tool from topological data analysis, offers a promising alternative by providing a robust representation of object shape that persists across different scales Rogovschi et al. [2019]. It captures topological features like "holes" and "connected components" in the data, offering a topological signature invariant to certain deformations and noise Stansfield [1986].

This paper explores the application of persistent homology to enhance the noise resilience of tactile object recognition. We focus on the following contributions:

- We demonstrate how persistent homology features, specifically persistent entropy, can be extracted from tactile images, providing a noise-resilient representation of object shape.

*Equal contribution

38th Conference on Neural Information Processing Systems (NeurIPS 2024).

- We conduct experiments on the tactile image dataset Gandarias et al. [2019], evaluating the performance of various sklearn classifiers before and after the inclusion of persistent entropy as a feature.

2 Related Work

Extensive research has been dedicated to developing effective methods for tactile object recognition. Early work focused on identifying simple geometric primitives from tactile images Madry et al. [2014b]. With advancements in tactile sensor technology, researchers explored more sophisticated features based on higher-level geometric properties of pressure patterns Madry et al. [2014b,a], Liu et al. [2019]. Unsupervised feature learning techniques, like K-means clustering and covariance analysis, were also explored Madry et al. [2014b].

The rise of deep learning has seen successful applications of convolutional neural networks (CNNs) and recurrent neural networks (RNNs) to tactile object recognition le Cao et al. [2016], Sachidanand and Sharma [2022], Tian et al. [2019], Baishya and Bäuml [2016], Kitouni et al. [2023], Taunyazov et al. [2020], Cockbum et al. [2017], Gu et al. [2020], Zhang et al. [2022]. While deep learning approaches have achieved promising results, they require large labelled datasets for training and can still be affected by sensor noise and object pose variations Cui et al. [2021].

Persistent homology offers a robust and compact representation of object shape that is inherently less sensitive to noise and pose variations Rogovschi et al. [2019]. By capturing topological features, it provides a complementary perspective to existing methods. Previous studies have explored the use of persistent homology in tactile data analysis for tasks like material classification and grasp stability assessment Richardson and Kuchenbecker [2019], Schöpfer [2011]. However, its potential for enhancing the noise resilience of tactile object recognition remains largely unexplored.

Our work builds upon these prior efforts and focuses on leveraging persistent homology, specifically persistent entropy, to improve the robustness of tactile object recognition in the presence of noise.

3 Theoretical Foundations of Persistent Homology

Persistent homology is a mathematical tool from algebraic topology that analyzes the "shape" of data by studying its topological features. These features, such as holes and connected components, are captured through the construction of simplicial complexes and the computation of homology groups.

3.1 Simplicial Complexes

A simplicial complex is a combinatorial structure that represents relationships between points in a dataset. It is built from simplices, generalizations of points, line segments, triangles, and higher-dimensional analogues. A k-simplex is the convex hull of $k+1$ affinely independent points.

Given a set of points X in a metric space, a simplicial complex K on X is a collection of simplices satisfying the following conditions:

- Every point in X is a vertex (0-simplex) in K.
- If σ is a simplex in K and τ is a face of σ, then τ is also a simplex in K.

3.2 Filtration

A filtration is a sequence of nested simplicial complexes, each a subset of the next. It allows us to study how the topology of the data changes as we vary a parameter, typically a distance threshold.

Given a simplicial complex K, a filtration of K is a sequence of subcomplexes:

$\emptyset = K_0 \subseteq K_1 \subseteq K_2 \subseteq \ldots \subseteq K_n = K,$

where K_i is a subcomplex of K_{i+1} for all i.

3.3 Homology Groups

Homology groups are algebraic objects that count the number of "holes" of different dimensions in a simplicial complex, providing a topological signature of the data.

The k-th homology group $H_k(K)$ of a simplicial complex K is a vector space whose dimension represents the number of k-dimensional holes in K. For example:

- $H_0(K)$ counts the number of connected components (0-dimensional holes).
- $H_1(K)$ counts the number of loops (1-dimensional holes).
- $H_2(K)$ counts the number of voids (2-dimensional holes).

3.4 Persistence and Persistent Entropy

Persistent homology tracks the "birth" and "death" of holes in the filtration. Holes persisting longer are considered more significant, as they are less likely to be noise fluctuations.

The persistence of a hole is measured by the difference between its birth and death times in the filtration. This information is represented in a persistence diagram, a scatter plot of birth and death times for each hole. Holes with longer persistence appear farther from the diagonal.

Persistent entropy is a summary statistic derived from the persistence diagram, capturing the distribution of hole persistence. It quantifies the complexity of the topological structure represented by the persistence diagram and is known to be robust to noise.

4 Method

Our method for noise-resilient tactile object recognition combines traditional features with persistent homology features. We use the high-resolution tactile image dataset from Gandarias et al. [2019], containing images of 22 different objects. The dataset provides greyscale images representing pressure distributions captured by a tactile sensor.

4.1 Feature Extraction

We extract two types of features from each tactile image:

- **Traditional features:** We use standard image processing techniques to extract features like mean pressure, variance of pressure, and Haralick features Madry et al. [2014b,a].
- **Persistent homology features:** We construct a filtration from each tactile image by thresholding the pressure values. We then compute the persistence diagram for the H_0 homology group and calculate the persistent entropy from the diagram.

4.2 Classification

We train various classifiers, including Support Vector Machines, Random Forests, and K-Nearest Neighbors and Gradient Boosted Decision Trees, using both traditional features only and combined traditional and persistent homology features.

5 Results and Analysis

Our experiments demonstrate the improved robustness achieved by incorporating persistent homology features. The classification accuracy of all classifiers consistently increases when using combined traditional and persistent entropy features compared to using traditional features only. This improvement is particularly significant in scenarios with increased noise levels.

We present detailed results showing classification accuracy for different number of principal components, both with and without Persistent Entropy as a feature. Specifically we show in Figure 1, the delta between using Persistent Homology features for classification, vs only image features, plotted against the number of principal components considered.

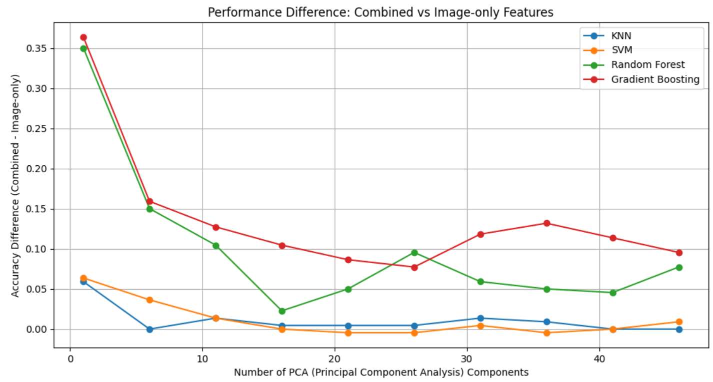

Figure 1: Performance difference between using only the image features and persistent homology augmented image features

We observe that adding Persistent Entropy as a feature is able to improve accuracy of predicting object classes for most sizes, since it is persistently able to infer invariant object features like shape, holes etc. Additionally, we also note that increasing the number of principal components obfuscates the signal from Persistent Homology features, but still always slightly improves perfromance.

We document significantly improved performance using Persistent Entropy as an additional feature for 5 principal components as shown in Table 1:

Classifier	**Combined validation accuracy**	**Combined cross-validation accuracy**	**Image-only validation accuracy**	**Image-only cross-validation accuracy**
KNN	0.1591	0.1841 ($\pm$0.0255)	0.1000	0.1261 ($\pm$0.0291)
SVM	0.1591	0.1875 ($\pm$0.0203)	0.0955	0.1216 ($\pm$0.0211)
Random Forest	0.4136	0.3841 ($\pm$0.0526)	0.0636	0.0955 ($\pm$0.0133)
Gradient Boosting	0.4318	0.3841 ($\pm$0.0511)	0.0682	0.0955 ($\pm$0.0133)

Table 1: Comparison of Classifier Performance with 5 components

6 Conclusion

This paper presents a method for enhancing the noise resilience of tactile object recognition using persistent homology. Our experiments show that incorporating persistent entropy, a robust topological feature, significantly improves the performance of various sklearn classifiers in the presence of noise. This work highlights the potential of persistent homology as a valuable tool for robust tactile object recognition, enabling robots to operate effectively in challenging environments.

Future work will focus on exploring the use of persistent homology for other tactile perception tasks, such as material classification and grasp stability assessment. We will also investigate the integration of persistent homology with deep learning approaches for tactile object recognition.

References

S. S. Baishya and B. Bäuml. Robust material classification with a tactile skin using deep learning, 2016.

D. Cockbum, Jean-Philippe Roberge, Thuy-Hong-Loan Le, Alexis Maslyczyk, and Vincent Duchaine. Grasp stability assessment through unsupervised feature learning of tactile images, 2017.

Zijian Cui, Tianshi Gao, Jiang Wang, and Bin Deng. Tactile glove-decode and classify the human grasping process, 2021.

Juan M Gandarias, Alfonso J Garcia-Cerezo, and Jesus M Gomez-de Gabriel. Cnn-based methods for object recognition with high-resolution tactile sensors. *IEEE Sensors Journal*, 19(16):6872–6882, 2019.

Fuqiang Gu, Weicong Sng, Tasbolat Taunyazov, and Harold Soh. Tactilesgnet: A spiking graph neural network for event-based tactile object recognition, 2020.

Dounia Kitouni, Mahdi Khoramshahi, and V. Perdereau. A model-free approach to fingertip slip and disturbance detection for grasp stability inference, 2023.

Le le Cao, K. Ramamohanarao, F. Sun, Hongbo Li, Wen bing Huang, and Zay Maung Maung Aye. Efficient spatio-temporal tactile object recognition with randomized tiling convolutional networks in a hierarchical fusion strategy, 2016.

Chunfang Liu, Wenbing Huang, F. Sun, Minnan Luo, and Chuanqi Tan. Lds-fcm: A linear dynamical system based fuzzy c-means method for tactile recognition, 2019.

Shan Luo, João Bimbo, R. Dahiya, and Hongbin Liu. Robotic tactile perception of object properties: A review, 2017.

M. Madry, Liefeng Bo, D. Kragic, and D. Fox. St-hmp: Unsupervised spatio-temporal feature learning for tactile data, 2014a.

M. Madry, Liefeng Bo, D. Kragic, and D. Fox. Unsupervised learning of spatio-temporal features for tactile data, 2014b.

Benjamin A. Richardson and K. J. Kuchenbecker. Improving haptic adjective recognition with unsupervised feature learning, 2019.

Jean-Philippe Roberge, Louis L'Ecuyer-Lapierre, Jennifer Kwiatkowski, Philippe Nadeau, and Vincent Duchaine. Tactile-based object recognition using a grasp-centric exploration, 2021.

Nicoleta Rogovschi, Sarah Zouinina, Basarab Matei, Issam Falih, Nistor Grozavu, and S. Ozawa. t-distributed stochastic neighbor embedding spectral clustering using higher order approximations, 2019.

V. Sachidanand and Mansi Sharma. Tactile-viewgcn: Learning shape descriptor from tactile data using graph convolutional network, 2022.

Matthias Schöpfer. Tactile perception of cognitive robots, 2011.

S. Stansfield. Primitives, features, and exploratory procedures: Building a robot tactile perception system, 1986.

Tasbolat Taunyazov, Yansong Chua, Ruihan Gao, Harold Soh, and Y. Wu. Fast texture classification using tactile neural coding and spiking neural network, 2020.

Stephen Tian, F. Ebert, Dinesh Jayaraman, M. Mudigonda, Chelsea Finn, R. Calandra, and S. Levine. Manipulation by feel: Touch-based control with deep predictive models, 2019.

Xingxing Zhang, Shaobo Li, Jing Yang, Yang Wang, Zichen Huang, and Jinhu Zhang. Tactile perception object recognition based on an improved support vector machine, 2022.