
LELAND MCINNES, Tutte Institute for Mathematics and Computing
Persistent Homology in High Dimensions

Persistent homology has proven itself to be a powerful tool for topological data analysis. It allows for shape analysis of "point-cloud" data in varying dimensions. In practice many applications of persistent homology have been on relatively low dimensional data. With the rise of deep learning, vast new troves of data have been unlocked – either through "embedding vectors" associated to unstructured datasets, or through the patterns of activations of the neural network itself. Such data sets typically have hundreds or thousands of dimensions. How well does persistent homology perform in such cases? What methods can we use to improve the results of persistent homology for such data? In this talk we'll explore these questions with simple example cases, and look at two different methods to make persistent homology more effective in high dimensions.