Topological Parallax: a Geometric Specification for Deep Perception Models

Paul Bendich
Geometric Data Analytics, and Duke University, United States
bendich@math.duke.edu

For safety and robustness of AI systems, we introduce topological parallax as a theoretical and computational tool that compares a trained model to a reference dataset to determine whether they have similar multi-scale geometric structure.

Our proofs and examples show that this geometric similarity between dataset and model is essential to trustworthy interpolation and perturbation, and we conjecture that this new concept will add value to the current debate regarding the unclear relationship between "overfitting" and "generalization" in applications of deep-learning.

In typical DNN applications, an explicit geometric description of the model is impossible, but parallax can estimate topological features (components, cycles, voids, etc.) in the model by examining the effect on the Rips complex of geodesic distortions using the reference dataset. Thus, parallax indicates whether the model shares similar multiscale geometric features with the dataset.

Parallax presents theoretically via topological data analysis [TDA] as a bi-filtered persistence module, and the key properties of this module are stable under perturbation of the reference dataset.

Joint work with Abraham D. Smith, Gabrielle Angeloro, Michael Catanzaro and Nirav Patel.