

Tracking topological properties of neural networks as they train

Marissa Masden

University of Oregon

mmasden@uoregon.edu

Feedforward ReLU neural networks are a class of piecewise linear function used in machine learning. The level sets of these functions correspond to the networks' performance on the practical machine learning task of separating data into two classes. However, because these functions are non-smooth, topological properties of their level sets are difficult to access theoretically. We show how it is possible to compute a cellular decomposition of the level sets of such a neural network (both in theory and in practice), and thereby track the topology of its level sets as it is trained to distinguish two subsets of \mathbb{R}^n .