Benign Overfitting of Interpolating Linear Classifiers on non-Subgaussian mixtures

Abstract

Practical success of deep neural networks has provoked theoretically surprising phenomenon. One of the phenomenon, which has spurred intense theoretical research beyond deep neural networks, is "benign overfitting": deep neural networks seem to generalize well in over-parametrized regime even though the networks show a perfect fit to noisy training data. It is now widely known that benign overfitting also occurs in various classical statistical models. For binary linear classification, previous works have proven that benign overfitting can occur with the maximum margin classifier while assuming that data is generated from subgaussian mixtures. Our work extends previous results to non-subgaussian mixtures. Specifically, we consider binary linear classification on data generated from finite moment mixtures. Our result shows that benign overfitting can occur without assuming subgaussian mixtures even with the presence of noise. Furthermore, our work establishes benign overfitting in a new regime which was not studied in the existing literature. Finally, our work clarify the existence of a phase transition in the over-parametrized regime along with geometric intuition behind it.