

Benign overfitting

Peter Bartlett

University of California, Berkeley, Department of Statistics, USA

Deep learning methodology has revealed some major surprises from the perspective of statistical complexity: even without any explicit effort to control model complexity, these methods find prediction rules that give a near-perfect fit to noisy training data and yet exhibit excellent prediction performance in practice. In this talk, we survey some recent work on this phenomenon of ‘benign overfitting.’ In the setting of linear prediction, we give a characterization of linear regression problems for which the minimum norm interpolating prediction rule has near-optimal prediction accuracy. The characterization shows that overparameterization is essential: the number of directions in parameter space that are unimportant for prediction must significantly exceed the sample size. We discuss implications for deep networks and for robustness to adversarial examples, and we describe extensions to ridge regression and barriers to analyzing benign overfitting via model-dependent generalization bounds.