Reconciling modern machine learning practice and the classical bias-variance trade-off

Mikhail Belkin, Daniel Hsu, Siyuan Ma, Soumik Mandal

Breakthroughs in machine learning are rapidly changing society and science, yet fundamental understanding of learning currently lags behind. In classical machine learning, we find a model that balances under-fitting and over-fitting: complex enough to express underlying structure in data, simple enough to avoid fitting spurious patterns. In modern practice, we fit complex models like neural networks with zero training error but still obtain high accuracy on test data. Our objective is to resolve the apparent conflict between the classical understanding and modern practice. Our main result is a "double descent" risk curve that unifies the behavior in the classical and modern regimes. The mechanism underlying its emergence is posited. Our findings resolve the apparent conflict and inform how machine learning models should be designed and understood.