Ryota Tomioka

Introduction to the analysis of learning algorithms: Does Bayesianism help?

In this lecture, we will look into machine learning from a theoretical perspective. Namely, instead of developing a new model, we ask how we can understand the behavior of a given model. The process of analyzing a model could seem hard in the beginning, not necessarily because of the math but because of the way it has been done. In this lecture, I will only assume basic knowledge of linear algebra and show how to analyze one of the most simple (but often useful) learning algorithm -- ridge regression. This algorithm is so simple that we can see everything exactly (no inequalities). We will see how the bias and variance parts behave differently with the number of samples and observe a curious "phase transition" phenomenon; i.e., more sample does not necessarily mean better performance. I will also analyze the Bayesian version of this and discuss whether and how Bayesian formulation could help from an objective viewpoint.