

Graphical Model Selection In High Dimensions: Practical Methods And Fundamental Limits

Martin Wainwright

University of California, Berkeley

Given samples from a Markov random field, how to determine the unknown graph? This graph selection problem is important for many applications of graphical models, and has received a great deal of attention over the past years. In this talk, we present a simple polynomial-complexity method, based on pseudolikelihood and ℓ_1 -regularization, for estimating graph structure. We show that it can recover the correct graph structure high probability as long as $n = \Omega(d^2 \log p)$, where n is the sample size, p is the number of vertices, and d is the maximum degree. Using information-theoretic methods, we show that no method can recover the graph correctly when the sample size is much smaller than this critical amount. We then discuss extensions of these methods to problems with noisy, missing and/or dependent data.

Based on joint works with John Lafferty, Po-Ling Loh, and Pradeep Ravikumar.