

Sequential matrix completion

Sergio Bacallado, University of Cambridge

The problem of completing a low-rank matrix from potentially noisy measurements of a subset of its entries has found great interest in recent years. This is motivated largely by recommendation systems, where the matrix summarizes the response of a set of customers to a set of products and the low rank assumption is due to latent features which explain customers' preferences. We consider the matrix completion problem sequentially. The goal is to define policies which sequentially select entries to measure, finding the largest entries in the matrix as early as possible. This presents an inherent exploration-exploitation tradeoff. We study a few simple policies for which it is possible to derive upper bounds on the regret, as well as approximate Thompson sampling algorithms based on a variational Bayesian analysis of matrix factor models.