Compatible Priors for Model Selection of High-Dimensional Gaussian DAGs

Stefano Peluso, Guido Consonni

Abstract

Inferring the conditional independence structure of high-dimensional data is a pervasive issue in a variety of scientific settings, and graphical models represent a powerful statistical tool to address this problem. From a Bayesian perspective, a recently introduced prior for Gaussian graphical models is the DAG-Wishart. This distribution is assigned to the (modified) Cholesky decomposition parameters of a precision matrix which is Markov with respect to a given DAG. A flexible structure of its shape hyperparameters coupled with conjugacy are two desirable assets of this prior which are especially welcome for estimation and prediction. In this paper we look at the DAG-Wishart prior from the perspective of model selection, with special reference to its consistency properties in high dimensional settings. We show that Bayes factor consistency only holds when comparing two DAGs which do not belong to the same Markov equivalence class, equivalently they encode distinct conditional independencies; a similar result holds for posterior ratio consistency. We also prove that a DAG-Wishart distribution with arbitrarily chosen hyperparameters may lead to incompatible priors for model selection, because they assign different marginal likelihoods to Markov equivalent graphs. We then propose a constructive method to specify DAG-Wishart priors with suitably constrained shape hyperparameters which ensure compatibility for DAG model selection.