Alex J Gibberd & James D B Nelson: Inference for Piecewise-Constant Gaussian Graphical Models

Traditionally, inference for graphical dependency structure has assumed stationarity of the estimated distribution. Whilst popular methods for network estimation such as via covariance selection may work well in an i.i.d setting, extensions to model dynamic non-stationary systems are less well developed. We here present an extension of previous network estimation methodologies to allow consideration of dynamics. In particular we discuss a class of estimators that recovers piecewise constant Gaussian graphical models (GGM), and thus permits identification of dynamics (via changepoints) within the dependency structure.

We discuss estimating structure within a (probabilistic attributional) graphical model in the context of a statistical model selection problem. Selecting an appropriate edge structure (or network) involves to some extent a search over different graphs. Unfortunately, for hard selection methods (which penalise the number of estimated edges, such as AIC/BIC) this combinatorial search rapidly becomes intractable with data size. In the stationary setting, methods such as the graphical lasso relax the hard selection of edges penalising via an 0 norm to the 1 norm. Under the Gaussian assumption, such a relaxation allows estimation of the dependency network via a convex optimisation problem.

More widely, such penalised likelihood approaches can be discussed in terms of M-estimators and are widely adopted for their robustness properties.

By expanding a GGM to include dynamics one introduces even more degrees of freedom to an often already high-dimensional problem. To counter this, additional regularisation is required which motivates the introduction of fused and group-fused type regularisers to the network identification problem. We construct a regularised M-estimator which we refer to as the group-fused graphical lasso and demonstrate its ability to estimate piecewise constant graphical models. The potential power of such approaches is demonstrated with an application changepoint detection in gene activation dependency networks.