

Discussion of “Latent Variable Graphical Model Selection Via Convex Optimization”

Steffen Lauritzen and Nicolai Meinshausen
Department of Statistics, University of Oxford, UK

February 3, 2012

We want to congratulate the authors for a thought-provoking and very interesting paper. Sparse modelling of the concentration matrix has enjoyed popularity in recent years. It has been framed as a computationally convenient convex ℓ_1 -constrained estimation problem in Yuan and Lin [2007] and can be applied readily to higher-dimensional problems. The authors argue –we think correctly– that the sparsity of the concentration matrix is for many applications more plausible after the effects of a few latent variables have been removed. The most attractive point about their method is surely that it is formulated as a convex optimisation problem. Latent variable fitting and sparse graphical modelling of the conditional distribution of the observed variables can then be obtained through a single fitting procedure.

Practical aspects. The method deserves wide adoption, but this will only be realistic if software is made available, for example as a R-package. Not many users will probably go to the trouble of implementing the method on their own, so we will strongly urge the authors to do so.

An imputation method. In the absence of readily available software, it is worth thinking whether the proposed fitting procedure can be approximated by methods involving known and well-tested computational techniques. The concentration matrix of observed and hidden variables is

$$K = \begin{pmatrix} K_O & K_{OH} \\ K_{HO} & K_H \end{pmatrix},$$

where we have deviated from the notation in the paper by omitting the asterisk. The proposed estimator $\hat{S}_n = \hat{K}_O$ of K_O was defined as

$$(\hat{K}_O, \hat{L}_n) = \operatorname{argmin}_{S,L} -\ell(S - L; \Sigma_O^n) + \lambda_n(\gamma \|S\|_1 + \operatorname{tr}(L)) \quad (1)$$

$$\text{such that } S - L \succ 0, L \succ 0, \quad (2)$$

where Σ_O^n is the empirical covariance matrix of the observed variables.

An alternative would be to replace the nuclear-norm penalization with a fixed constraint κ on the rank of the hidden variables, replacing problem (1) with

$$(\hat{K}_O, \hat{L}_n) = \operatorname{argmin}_{S,L} -\ell(S - L; \Sigma_O^n) + \lambda_n \|S\|_1 \quad (3)$$

$$\text{such that } S - L \succ 0 \text{ and } L \succ 0 \text{ and } \operatorname{rank}(L) \leq \kappa.$$

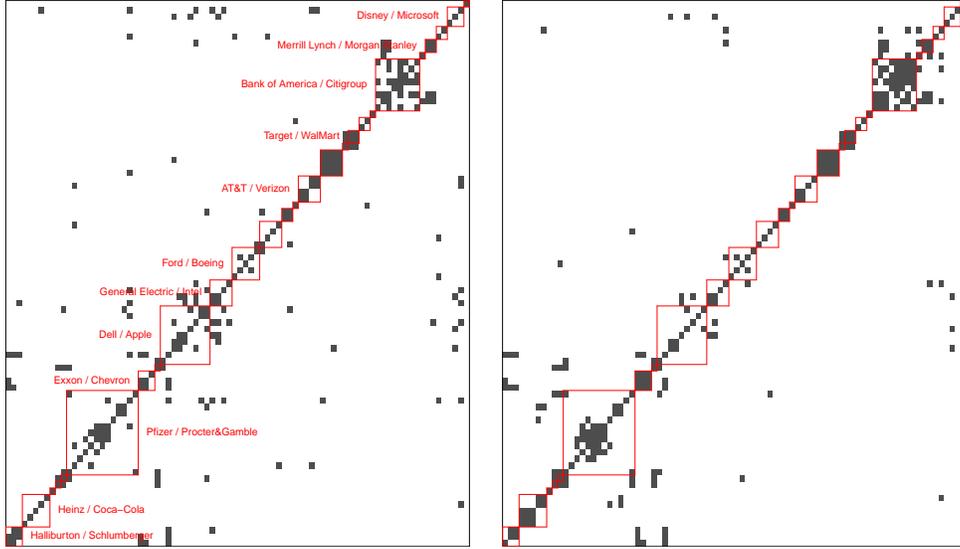


Figure 1: The non-zero entries of the concentration matrix \hat{K}_O , using the proposed procedure (1) (left) and the imputation method in (4) (right). Two representative companies are shown for some of the sectors.

This can be achieved by a missing-value formulation in combination with use of the EM algorithm, which also applies in a penalized likelihood setting [Green, 1990]. Let the hidden variables be of a fixed dimensionality κ and assume for a moment these are observed so one would find the concentration matrix \hat{K} of the joint distribution of the observed variables X_O and hidden variables X_H based on the complete data penalized likelihood as

$$\operatorname{argmin}_K -\log f_K(X_O, X_H) + \lambda \|K_O\|_1, \quad (4)$$

where f_K is the joint density of (X_O, X_H) . The EM algorithm iteratively replaces the likelihood in (4) for $t = 1, \dots, T$ by its conditional expectation and thus finds \hat{K}^{t+1} as

$$\hat{K}^{t+1} = \operatorname{argmin}_K -E_{\hat{K}^t} \left\{ \log f_K(X_O, X_H) \mid X_O \right\} + \lambda \|K_O\|_1. \quad (5)$$

The iteration is guaranteed not to increase the negative marginal penalised likelihood at every stage and will, save for unidentifiability, converge to the minimizer in (3) for most starting values. Without loss of generality one can fix the conditional concentration matrix K_H of the hidden variables to be the identity so that these are conditionally independent with variance 1, given the observed variables. Then $-K_{OH}$ is equal to the regression coefficients of the observed variables on the hidden variables. As starting value we have let $-\hat{K}_{OH}^0$ be equal to these with hidden variables determined by a principal component analysis.

The expectation in (5) can be written as the log-likelihood of a Gaussian distribution with concentration matrix K and empirical covariance matrix W^t where

$$W^t = \begin{pmatrix} \Sigma_O^n & -\Sigma_O^n \hat{K}_{OH}^t \\ -\hat{K}_{HO}^t \Sigma_O^n & \mathbf{I} + \hat{K}_{HO}^t \Sigma_O^n \hat{K}_{OH}^t \end{pmatrix}.$$

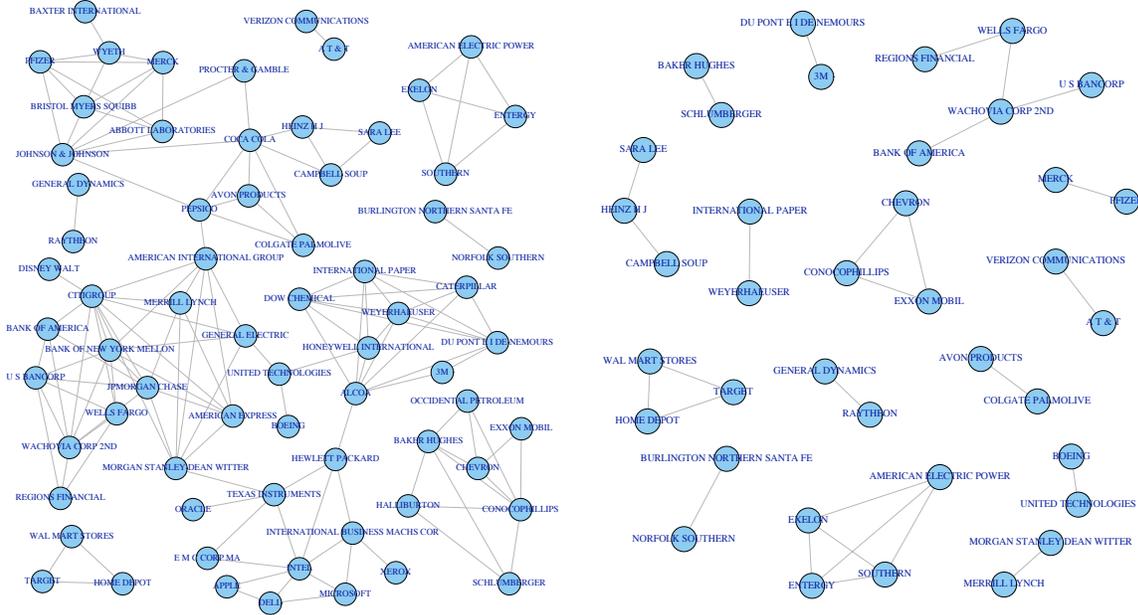


Figure 2: Left: the graph of the imputation method as in (4). Right: the graph of the stable edges. In both cases, isolated vertices have been removed from the display.

The sufficient statistics involving the missing data are thus “imputed” in W^t . Each of the updates (5) can now be computed with the *graphical lasso* [Friedman et al., 2008].

We thought it would be interesting to compare the two methods on the data example given in the paper. Figure 1 shows the solution \hat{K}_O for the stock-return example when using the proposed method (1) and the imputation method (4) with 4 iterations. The number κ of latent variables and the number of non-zero edges in \hat{K}_O is adjusted to be the same as in the original estimator.

The three pairs with the highest absolute entries in the fitted conditional concentration matrix are identical (AT&T – Verizon, Schlumberger – Baker Hughes and Merrill Lynch – Morgan Stanley) for the two methods and the 15 pairs with highest absolute entries in the off-diagonal concentration matrix have an overlap of size 12. The resulting graphs are slightly different although they share many features. Our graph has 136 edges, one more than that in the procedure described in the paper, and 77 of the edges are shared. Our graph has more isolated vertices (15 vs. 9), slightly fewer cliques (62 vs. 81) and the largest clique in our graph has six variables rather than four. The graph is displayed to the left in Figure 2 and features some clearly identified clusters of variables.

The selected graph is very unstable under bootstrap simulations. In the spirit of Meinshausen and Bühlmann [2010], we fit the graph on 2000 bootstrap samples. Only 28 edges are selected in more than half of these samples. The resulting graph is shown in Figure 2. As many as 25 of these edges appear also as edges of the estimator proposed in (1). It would have been interesting to be able to compare with the same “stability graph” of the proposed procedure but we suspect that they will match closely.

Latent directed structures. In a sense the procedure described in this paper can be seen as a modification of, or an alternative to, factor analysis, in which independent latent variables are

sought to explain all the correlations, corresponding to the graph for the observed variables being completely empty.

Methods for identifying such models can, for example, be developed using tetrad constraints [Spirtes et al, 1993, Drton et al., 2007]. Another generalisation of factor analysis is to look for sparse *directed* graphical models which has now been rather well established through e.g. the FCI algorithm [Spirtes et al, 1993, Richardson and Spirtes, 2002] with an algebraic underpinning in Sullivant [2008]. Again this could be an alternative to the procedure described in this interesting paper.

Summary. We effectively replaced the nuclear norm penalisation of L in the paper by a fixed constraint on the rank. This might be easier to do than choosing a reasonable value for the penalty on the trace of L . Using this formulation we could combine the EM algorithm with the graphical lasso, enabling us to compute the solution with readily available software. It would be interesting to see whether our procedure can be shown to recover the correct sparsity structure under similar assumptions to those in the paper. We want to congratulate the authors again for a very interesting discussion paper.

References

- M. Drton, B. Sturmfels, and S. Sullivant. Algebraic factor analysis: tetrads, pentads and beyond. *Probability Theory and Related Fields*, 138:463–493, 2007.
- J. Friedman, T. Hastie, and R. Tibshirani. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9:432, 2008.
- P. Green. On use of the EM algorithm for penalized likelihood estimation. *Journal of the Royal Statistical Society: Series B*, 52:443–452, 1990.
- N. Meinshausen and P. Bühlmann. Stability selection. *Journal of the Royal Statistical Society: Series B*, 72:417–473, 2010.
- T. Richardson and P. Spirtes. Ancestral graph Markov models. *The Annals of Statistics*, 30: 962–1030, 2007.
- P. Spirtes, C. Glymour, and R. Scheines. *Causation, Prediction and Search*. Springer-Verlag, New York, 1993.
- S. Sullivant. Algebraic geometry of Gaussian Bayesian networks. *Advances in Applied Mathematics*, 40:482–513, 2008.
- M. Yuan and Y. Lin. Model selection and estimation in the Gaussian graphical model. *Biometrika*, 94:19–35, 2007.