

A path following algorithm for Sparse Pseudo-Likelihood Inverse Covariance Estimation (SPLICE)

Guilherme V. Rocha, Peng Zhao, Bin Yu

July 24, 2008

Abstract

Given n observations of a p -dimensional random vector, the covariance matrix and its inverse (precision matrix) are needed in a wide range of applications. Sample covariance (e.g. its eigenstructure) can misbehave when p is comparable to the sample size n . Regularization is often used to mitigate the problem.

In this paper, we proposed an ℓ_1 penalized pseudo-likelihood estimate for the inverse covariance matrix. This estimate is sparse due to the ℓ_1 penalty, and we term this method SPLICE. Its regularization path can be computed via an algorithm based on the homotopy/LARS-Lasso algorithm. Simulation studies are carried out for various inverse covariance structures for $p = 15$ and $n = 20, 1000$. We compare SPLICE with the ℓ_1 penalized likelihood estimate and a ℓ_1 penalized Cholesky decomposition based method. SPLICE gives the best overall performance in terms of three metrics on the precision matrix and ROC curve for model selection. Moreover, our simulation results demonstrate that the SPLICE estimates are positive-definite for most of the regularization path even though the restriction is not enforced.

Acknowledgments

The authors gratefully acknowledge the support of NSF grant DMS-0605165, ARO grant W911NF-05-1-0104, NSFC (60628102), and a grant from MSRA. B. Yu also thanks the Miller Research Professorship in Spring 2004 from the Miller Institute at University of California at Berkeley and a 2006 Guggenheim Fellowship. G. Rocha also acknowledges helpful comments by Ram Rajagopal, Garvesh Raskutti, Pradeep Ravikumar and Vincent Vu.

1 Introduction

Covariance matrices are perhaps the simplest statistical measure of association between a set of variables and widely used. Still, the estimation of covariance matrices is extremely data hungry, as the number of fitted parameters grows rapidly with the number of observed variables p . Global properties of the estimated covariance matrix, such as its eigenstructure, are often used (e.g. Principal Component Analysis, Jolliffe, 2002). Such global parameters may fail to be consistently estimated when the number of variables p is non-negligible in the comparison to the sample size n . As one example, it is a well-known fact that the eigenvalues and eigenvectors of an estimated covariance matrix are inconsistent when the ratio $\frac{p}{n}$ does not vanish asymptotically (Marchenko and Pastur, 1967; Paul et al., 2008). Data sets with a large number of observed variables p and small number of observations n are now a common occurrence in statistics. Modeling such data sets creates a need for regularization procedures capable of imposing sensible structure on the estimated covariance matrix while being computationally efficient.

Many alternative approaches exist for improving the properties of covariance matrix estimates. *Shrinkage methods* for covariance estimation were first considered in Stein (1975, 1986) as a way to correct the overdispersion of the eigenvalues of estimates of large covariance matrices. Ledoit and Wolf (2004) present a shrinkage estimator that is the asymptotically optimal convex linear combination of the sample covariance matrix and the identity matrix with respect to the Froebenius norm. Daniels and Kass (1999, 2001) propose alternative strategies using shrinkage toward diagonal and more general matrices. *Factorial models* have also been used as a strategy to regularize estimates of covariance matrices (Fan et al., 2006). *Tapering* the covariance matrix is frequently used in time series and spatial models and have been used recently to improve the performance of covariance matrix estimates used by classifiers based on linear discriminant analysis (Bickel and Levina, 2004) and in Kalman filter ensembles (Furrer and Bengtsson, 2007). Regularization of the covariance matrix can also be achieved by *regularizing its eigenvectors* (Johnstone and Lu, 2004; Zou et al., 2006).

Covariance selection methods for estimating covariance matrices consist of imposing sparsity on the precision matrix (i.e., the inverse of the covariance matrix). The Sparse Pseudo-Likelihood

Inverse Covariance Estimates (SPLICE) proposed in this paper fall into this category. This family of methods was introduced by Dempster (1972). An advantage of imposing structure on the precision matrix stems from its close connections to linear regression. For instance, Wu and Pourahmadi (2003) use, for a fixed order of the random vector, a parametrization of the precision matrix \mathbf{C} in terms of a decomposition $\mathbf{C} = \mathbf{U}'\mathbf{D}\mathbf{U}$ with \mathbf{U} upper-triangular with unit diagonal and \mathbf{D} a diagonal matrix. The parameters \mathbf{U} and \mathbf{D} are then estimated through p linear regressions and Akaike’s Information Criterion (AIC, Akaike, 1973) is used to promote sparsity in \mathbf{U} . A similar covariance selection method is presented in Bilmes (2000). More recently, Bickel and Levina (2008) have obtained conditions ensuring consistency in the operator norm (spectral norm) for precision matrix estimates based on banded Cholesky factors. Two disadvantages of imposing the sparsity in the factor \mathbf{U} are: sparsity in \mathbf{U} does not necessarily translates into sparsity of \mathbf{C} and; the sparsity structure in \mathbf{U} is sensitive to the order of the random variables within the random vector. The SPLICE estimates proposed in this paper constitute an attempt at tackling these issues.

The AIC selection criterion used in Wu and Pourahmadi (2003) requires, in its exact form, that the estimates be computed for all subsets of the parameters in \mathbf{U} . A more computationally tractable alternative for performing parameter selection consists in penalizing parameter estimates by their ℓ_1 -norm (Breiman, 1995; Tibshirani, 1996; Chen et al., 2001), popularly known as the LASSO in the context of least squares linear regression. The computational advantage of the ℓ_1 -penalization over penalization by the dimension of the parameter being fitted (ℓ_0 -norm) – such as in AIC – stems from its convexity (Boyd and Vandenberghe, 2004). Homotopy algorithms for tracing the entire LASSO regularization path have recently become available (Osborne et al., 2000; Efron et al., 2004). Given the high-dimensionality of modern days data sets, it is no surprise that ℓ_1 -penalization has found its way into the covariance selection literature.

Huang et al. (2006) propose a covariance selection estimate corresponding to an ℓ_1 -penalty version of the Cholesky estimate of Wu and Pourahmadi (2003). The off-diagonal terms of \mathbf{U} are penalized by their ℓ_1 -norm and cross-validation is used to select a suitable regularization parameter. While this method is very computationally tractable (an algorithm based on the homotopy algorithm for linear regressions is detailed below in Appendix B.1), it still suffer from the deficien-

cies of Cholesky-based methods. Alternatively, Banerjee et al. (2005), Banerjee et al. (2007), Yuan and Lin (2007), and Friedman et al. (2008) consider an estimate defined by the Maximum Likelihood of the precision matrix for the Gaussian case penalized by the ℓ_1 -norm of its off-diagonal terms. While these methods impose sparsity directly in the precision matrix, no path-following algorithms are currently available for them. Rothman et al. (2007) analyze the properties of estimates defined in terms of ℓ_1 -penalization of the exact Gaussian neg-loglikelihood and introduce a permutation invariant method based on the Cholesky decomposition to avoid the computational cost of semi-definite programming.

The SPLICE estimates presented here impose sparsity constraints directly on the precision matrix. Moreover the entire regularization path of SPLICE estimates can be computed by homotopy algorithms. It is based on previous work by Meinshausen and Bühlmann (2006) for neighborhood selection in Gaussian graphical models. While Meinshausen and Bühlmann (2006) use p separate linear regressions to estimate the neighborhood of one node at a time, we propose merging all p linear regressions into a single least squares problem where the observations associated to each regression are weighted according to their conditional variances. The loss function thus formed can be interpreted as a pseudo neg-loglikelihood (Besag, 1974) in the Gaussian case. To this pseudo-negloglikelihood minimization, we add symmetry constraints and a weighted version of the ℓ_1 -penalty on off-diagonal terms to promote sparsity. The SPLICE estimate can be interpreted as an approximate solution following from replacing the exact neg-loglikelihood in Banerjee et al. (2007) by a quadratic surrogate (the pseudo neg-loglikelihood).

The main advantage of SPLICE estimates is algorithmic: by use of a proper parametrization, the problem involved in tracing the SPLICE regularization path can be recast as a linear regression problem and thus amenable to be solved by a homotopy algorithm as in Osborne et al. (2000) and Efron et al. (2004). To avoid computationally expensive cross-validation, we use information criteria to select a proper amount of regularization. We compare the use of Akaike's Information criterion (AIC, Akaike, 1973), a small-sample corrected version of the AIC (AIC_C, Hurvich et al., 1998) and the Bayesian Information Criterion (BIC, Schwartz, 1978) for selecting the proper amount of regularization.

We use simulations to compare SPLICE estimates to the ℓ_1 -penalized maximum likelihood estimates (Banerjee et al., 2005, 2007; Yuan and Lin, 2007; Friedman et al., 2008) and to the ℓ_1 -penalized Cholesky approach in Huang et al. (2006). We have simulated both small and large sample data sets. Our simulations include model structures commonly used in the literature (ring and star topologies, AR processes) as well as a few randomly generated model structures. SPLICE had the best performance in terms of the quadratic loss and the spectral norm of the precision matrix deviation ($\|\mathbf{C} - \hat{\mathbf{C}}\|_2$). It also performed well in terms of the entropy loss. SPLICE had a remarkably good performance in terms of selecting the off-diagonal terms of the precision matrix: in the comparison with Cholesky, SPLICE incurred a smaller number of false positives to select a given number of true positives; in the comparison with the penalized exact maximum likelihood estimates, the path following algorithm allows for a more careful exploration of the space of alternative models.

The remainder of this paper is organized as follows. Section 2 presents our pseudo-likelihood surrogate function and some of its properties. Section 3 presents the homotopy algorithm used to trace the SPLICE regularization path. Section 4 presents simulation results comparing the SPLICE estimates with some alternative regularized methods. Finally, Section 5 concludes with a short discussion.

2 An approximate loss function for inverse covariance estimation

In this section, we establish a parametrization of the precision matrix Σ^{-1} of a random vector \mathbb{X} in terms of the coefficients in the linear regressions among its components. We emphasize that the parametrization we use differs from the one previously used by Wu and Pourahmadi (2003). Our alternative parametrization is used to extend the approach used by Meinshausen and Bühlmann (2006) for the purpose of estimation of sparse precision matrices. The resulting loss function can be interpreted as a pseudo-likelihood function in the Gaussian case. For non-Gaussian data, the minimizer of the empirical risk function based on the loss function we propose still yields consistent estimates. The loss function we propose also has close connections to linear regression and lends itself well for a homotopy algorithm in the spirit of Osborne et al. (2000) and Efron et al. (2004). A comparison of this approximate loss function to its exact counterpart in the Gaussian case suggests

that the approximation is better the sparser the precision matrix.

In what follows, \mathbf{X} is a $\mathbb{R}^{n \times p}$ matrix containing in each of its n rows observations of the zero-mean random vector \mathbb{X} with covariance matrix Σ . Denote by \mathbb{X}_j the j -th entry of \mathbb{X} and by \mathbb{X}_{J^*} the $(p-1)$ dimensional vector resulting from deleting \mathbb{X}_j from \mathbb{X} . For a given j , we can permute the order of the variables in \mathbb{X} and partition Σ to get:

$$\text{cov} \left(\begin{bmatrix} \mathbb{X}_j \\ \mathbb{X}_{J^*} \end{bmatrix} \right) = \begin{bmatrix} \sigma_{j,j} & \Sigma_{j,J^*} \\ \Sigma_{J^*,j} & \Sigma_{J^*,J^*} \end{bmatrix}$$

where J^* corresponds to the indices in \mathbb{X}_{J^*} , so $\sigma_{j,j}$ is a scalar, Σ_{j,J^*} is a $p-1$ dimensional row vector and Σ_{J^*,J^*} is a $(p-1) \times (p-1)$ square matrix. Inverting this partitioned matrix (see, for instance Hocking, 1996) yield:

$$\begin{bmatrix} \sigma_{j,j} & \Sigma_{j,J^*} \\ \Sigma_{J^*,j} & \Sigma_{J^*,J^*} \end{bmatrix}^{-1} = \begin{bmatrix} \frac{1}{d_j^2} & -\frac{1}{d_j^2} \beta_j \\ \mathbf{M}_1 & \mathbf{M}_2^{-1} \end{bmatrix}, \quad (1)$$

where:

$$\begin{aligned} \beta_j &= \left(\beta_{j,1} \quad \dots, \beta_{j,j-1}, \beta_{j,j+1}, \dots, \beta_{j,p} \right) = -\Sigma_{j,J^*} \Sigma_{J^*,J^*}^{-1} \in \mathbb{R}^{(p-1)}, \\ d_j &= \sqrt{\left(\sigma_{j,j} - \Sigma_{j,J^*} \Sigma_{J^*,J^*}^{-1} \Sigma_{J^*,j} \right)} \in \mathbb{R}_+, \\ \mathbf{M}_2 &= \Sigma_{J^*,J^*} - \Sigma_{J^*,j} \sigma_{j,j}^{-1} \Sigma_{j,J^*}, \\ \mathbf{M}_1 &= -\mathbf{M}_2^{-1} \cdot \left(\Sigma_{J^*,J^*}^{-1} \Sigma_{J^*,j} \right) = -\frac{1}{d_j^2} \beta_j', \quad (\text{the second equality due to symmetry}). \end{aligned}$$

We will focus on the d_j and β_j parameters in what follows.

The parameters β_j and d_j^2 correspond respectively to the coefficients and the expected value of the squared residuals of the best linear model of \mathbb{X}_j based on \mathbb{X}_{J^*} , irrespectively of the distribution of \mathbb{X} . In what follows, we will let β_{jk} denote the coefficient corresponding to \mathbb{X}_k in the linear model of \mathbb{X}_j based on \mathbb{X}_{J^*} . We define:

- **D**: a $p \times p$ diagonal matrix with d_j along its diagonal and,

- \mathbf{B} : a $p \times p$ matrix with zeros along its diagonal and off-diagonal terms given by β_{jk} .

Using (1) for $j = 1, \dots, p$ yields:

$$\Sigma^{-1} = \mathbf{D}^{-2}(\mathbf{I}_p - \mathbf{B}) \quad (2)$$

Since Σ^{-1} is symmetric, (2) implies that the following constraints hold:

$$d_k^2 \beta_{jk} = d_j^2 \beta_{kj}, \quad \text{for } j, k = 1, \dots, p. \quad (3)$$

Equation (2) shows that the sparsity pattern of Σ^{-1} can be inferred from sparsity in the regression coefficients contained in \mathbf{B} . Meinshausen and Bühlmann (2006) exploit this fact to estimate the neighborhood of a Gaussian graphical model. They use the LASSO (Tibshirani, 1996) to obtain sparse estimates of β_j :

$$\hat{\beta}_j(\lambda_j) = \arg \min_{b_j \in \mathbb{R}^{p-1}} \|\mathbf{X}_j - \mathbf{X}_{J^*} b_j\|^2 + \lambda_j \|b_j\|_1, \quad \text{for } j = 1, \dots, p \quad (4)$$

The neighborhood of the node \mathbb{X}_j was then estimated based on the entries of the $\hat{\beta}_j$ that were set to zero. Minor inconsistencies could occur as the regressions are run separately. As an example, one could have $\hat{\beta}_{jk}(\lambda_j) = 0$ and $\hat{\beta}_{kj}(\lambda_k) \neq 0$, which Meinshausen and Bühlmann (2006) solve by defining AND and OR rules for defining the estimated neighborhood.

To extend the framework of Meinshausen and Bühlmann (2006) to the estimation of precision matrices the parameters d_j^2 must also be estimated and the symmetry constraints in (3) must be enforced. We use a pseudo-likelihood approach (Besag, 1974) to form a surrogate loss function involving all terms of \mathbf{B} and \mathbf{D} . For Gaussian \mathbb{X} , the negative log-likelihood function of \mathbb{X}_j given \mathbb{X}_{J^*} is:

$$\log \left[p(\mathbf{X}_j | \mathbf{X}_{J^*}, d_j^2, \beta_j) \right] = -\frac{n}{2} \log(2\pi) - \frac{n}{2} \log(d_j^2) - \frac{1}{2} \left(\frac{\|\mathbf{X}_j - \mathbf{X}_{J^*} \beta_j\|^2}{d_j^2} \right). \quad (5)$$

The parameters d_j^2 and β_j can be consistently estimated by minimizing (5). A pseudo-neg-

loglikelihood function can be formed as:

$$\begin{aligned}\mathcal{L}(\mathbf{X}; \mathbf{D}, \mathbf{B}) &= \log \left(\prod_{j=1}^p p(\mathbf{X}_j | \mathbf{X}_{J^*}, d_j^2, \beta_j) \right) \\ &= -\frac{np}{2} \log(2\pi) - \frac{n}{2} \log \det(\mathbf{D}^2) - \frac{1}{2} \text{tr} [\mathbf{X}(\mathbf{I}_p - \mathbf{B}') \mathbf{D}^{-2} (\mathbf{I}_p - \mathbf{B}) \mathbf{X}'].\end{aligned}\tag{6}$$

An advantage of the surrogate $\mathcal{L}(\mathbf{X}; \mathbf{D}, \mathbf{B})$ is that, for fixed \mathbf{D} , it is a quadratic form in \mathbf{B} . To promote sparsity on the precision matrix, we propose using a weighted ℓ_1 -penalty on \mathbf{B} :

$$\begin{aligned}(\hat{\mathbf{D}}(\lambda), \hat{\mathbf{B}}(\lambda)) &= \arg \min_{(\mathbf{B}, \mathbf{D})} \{n \log \det(\mathbf{D}^2) + \text{tr} [\mathbf{X}(\mathbf{I}_p - \mathbf{B}') \mathbf{D}^{-2} (\mathbf{I}_p - \mathbf{B}) \mathbf{X}']\} + \lambda \|\mathbf{B}\|_{w,1} \\ \text{s.t.} &\begin{cases} b_{jj} &= 0 \\ d_{kk}^2 b_{jk} &= d_{jj}^2 b_{kj} \\ d_{kj}^2 &= 0, \quad \text{for } k \neq j \\ d_{jj}^2 &\geq 0 \end{cases}\end{aligned}\tag{7}$$

where $\|\mathbf{B}\|_{w,1} = \sum_{j,k=1}^p w_{jk} |b_{jk}|$. From (2), the precision matrix estimate $\hat{\mathbf{C}}(\lambda)$ is then given by:

$$\hat{\mathbf{C}}(\lambda) = \hat{\mathbf{D}}^{-2}(\lambda) [\mathbf{I}_p - \hat{\mathbf{B}}(\lambda)].\tag{8}$$

The weights w_{jk} in (7) can be adjusted to accommodate differences in scale between the b_{jk} parameters. In the remainder of this paper, we fix $w_{jk} = 1$ for all j, k such that $j \neq k$ (notice that the weights w_{jj} are irrelevant since $b_{jj} = 0$ for all j).

The main advantage of the pseudo-likelihood estimates as defined in (7) is algorithmic. Fixing \mathbf{D} , the minimizer with respect to the \mathbf{B} parameter is the solution of a ℓ_1 -penalized least squares problem. Hence, for fixed \mathbf{D} it is possible to adapt the homotopy/LARS-LASSO algorithm (Osborne et al., 2000; Efron et al., 2004) to obtain estimates for all values of λ . For each fixed \mathbf{B} , the minimizer with respect to \mathbf{D} has a closed form solution. The algorithm presented in Section 3 is based on performing alternate optimization steps with respect to \mathbf{B} and \mathbf{D} to take advantage of these properties.

One drawback of the precision matrix estimate $\hat{\mathbf{C}}(\lambda)$ in (7) is that it cannot be ensured to be positive semi-definite. However, from the optimality conditions to (7) it can be proven that

for a large enough λ , all terms in the $\hat{\mathbf{B}}(\lambda)$ matrix are set to zero. For such highly regularized estimates, $\hat{\mathbf{C}}(\lambda)$ is diagonal. Therefore, continuity of the regularization path implies existence of $\lambda^* < \inf\{\lambda : \hat{\mathbf{B}}(\lambda) = 0\}$ for which $\hat{\mathbf{C}}(\lambda^*)$ is diagonal dominant and thus positive semi-definite.

We return to the issue of positive semi-definiteness later in the paper. We will prove in Section 2.4 that, when the ℓ_1 -norm is replaced by the ℓ_2 -norm, positive semi-definiteness is ensured for any value of the regularization parameter. That suggests that a penalty similar to the elastic net (Zou and Hastie, 2005) can make the estimates along the ℓ_1 -norm regularization path positive semi-definite for a larger stretch. In addition, in Section 4 we present evidence that the ℓ_1 -norm penalized estimates are positive semi-definite for most of the regularization path.

2.1 Alternative normalization

Before we move on, we present a normalization of the \mathbf{X} data matrix that leads to a more natural representation of the precision matrix $\hat{\mathbf{C}}$ in terms of $\hat{\mathbf{B}}$ and $\hat{\mathbf{D}}$ while resulting in more convenient symmetry constraints.

The symmetry constraints in (3) can be rewritten in a more insightful form:

$$\frac{d_k}{d_j}\beta_{jk} = \frac{d_j}{d_k}\beta_{kj}, \text{ for all } j \neq k. \quad (9)$$

This alternative representation, suggests that the symmetry constraints can be more easily applied to a renormalized version of the data. We define:

$$\begin{aligned} \tilde{\mathbf{X}} &= \mathbf{X}\mathbf{D}^{-1}, & \text{and} \\ \tilde{\mathbf{B}} &= \mathbf{D}^{-1}\mathbf{B}\mathbf{D}. \end{aligned} \quad (10)$$

Under this renormalization, $\tilde{\mathbf{B}}$ is symmetric, a fact that will be explored in the algorithmic section below to enforce the symmetry constraint within the homotopy algorithm used to trace regularization paths for SPLICE estimates.

Another advantage of this renormalization is that the precision matrix estimate can be written

as:

$$\hat{\mathbf{C}}(\lambda) = \hat{\mathbf{D}}^{-1} \left[\mathbf{I}_p - \hat{\tilde{\mathbf{B}}}(\lambda) \right] \hat{\mathbf{D}}^{-1}, \quad (11)$$

making the analysis of the positive semi-definiteness of $\hat{\mathbf{C}}(\lambda)$ easier: $\hat{\mathbf{C}}(\lambda)$ is positive semi-definite if and only if $\mathbf{I}_p - \hat{\tilde{\mathbf{B}}}(\lambda)$ is positive semi-definite. This condition is satisfied as long as the eigenvalues of $\hat{\tilde{\mathbf{B}}}(\lambda)$ are smaller than 1.

2.2 Comparison of exact and pseudo-likelihoods in the Gaussian case

Banerjee et al. (2005), Yuan and Lin (2007), Banerjee et al. (2007) and Friedman et al. (2008) have considered estimates defined as the minimizers of the exact likelihood penalized by the ℓ_1 -norm of the off-diagonal terms of the precision matrix \mathbf{C} :

$$\begin{aligned} \hat{\mathbf{C}}_{exact}(\lambda) &= \arg \min_{\mathbf{C}} n \log \det(\mathbf{C}) + \text{tr}[\mathbf{X}\mathbf{C}\mathbf{X}'] + \lambda \|\mathbf{C}\|_1 \\ &\text{s.t.} \quad \mathbf{C} \text{ is symmetric positive semi-definite.} \end{aligned} \quad (12)$$

A comparison of the exact and pseudo-likelihood functions in terms of the $(\mathbf{D}, \tilde{\mathbf{B}})$ parametrization suggests when the approximation will be appropriate. In terms of $(\mathbf{D}, \tilde{\mathbf{B}})$, the pseudo-likelihood function in (6) is:

$$\mathcal{L}(\mathbf{X}; \mathbf{D}, \mathbf{B}) = -\frac{np}{2} \log(2\pi) - \frac{n}{2} \log \det(\mathbf{D}^2) - \frac{1}{2} \text{tr} \left[\tilde{\mathbf{X}} (\mathbf{I}_p - \tilde{\mathbf{B}})^2 \tilde{\mathbf{X}}' \right]. \quad (13)$$

In the same parametrization, the exact likelihood function is:

$$\begin{aligned} \mathcal{L}_{exact}(\mathbf{X}; \mathbf{D}, \mathbf{B}) &= -\frac{np}{2} \log(2\pi) - \frac{n}{2} \log \det \left[\left(\mathbf{I}_p - \tilde{\mathbf{B}} \right) \right] \\ &\quad - \frac{n}{2} \log \det(\mathbf{D}^2) - \frac{1}{2} \text{tr} \left[\tilde{\mathbf{X}} \left(\mathbf{I}_p - \tilde{\mathbf{B}} \right) \tilde{\mathbf{X}}' \right] \end{aligned} \quad (14)$$

A comparison between (13) and (14) reveals two differences in the exact and pseudo neg-loglikelihood functions. First, the exact expression involves one additional $\log \det \left[\left(\mathbf{I}_p - \tilde{\mathbf{B}} \right) \right]$ term not appearing in the surrogate expression. Secondly, in the squared deviation term, the surrogate expression has

the weighting matrix $\left[\left(\mathbf{I}_p - \tilde{\mathbf{B}} \right) \right]$ squared in the comparison to the exact likelihood. For $\tilde{\mathbf{B}} = \mathbf{0}$, the log det vanishes and the weighting term equals identity and thus idempotent. Since these two functions are continuous in $\tilde{\mathbf{B}}$, we can expect this approximation to work better the smaller the off-diagonal terms in $\tilde{\mathbf{B}}$. In particular, in the completely sparse case, the two functions coincide and the approximation is exact.

2.3 Properties of the pseudo-likelihood estimate

In the classical setting where p is fixed and n grows to infinity, the unregularized pseudo-likelihood estimate $\hat{\mathbf{C}}(0)$ is clearly consistent. The unconstrained estimates for β_j and d_j^2 are all consistent. In addition, the symmetry constraints we impose are observed in the population version and thus introduce no asymptotic bias in the (symmetry) constrained estimates. That in conjunction with the results from Knight and Fu (2000) can be used to prove that, as long as $\lambda = o_p(n)$ the ℓ_1 -penalized estimates in (7) are consistent.

Still in the classical setting, Yuan and Lin (2007) point out that the unpenalized estimate $\hat{\mathbf{C}}(0)$ is not efficient as it does not coincide with the maximum likelihood estimate. However, the comparison of the exact and pseudo-likelihoods presented in Section 2.2 suggests that the loss in efficiency should be smaller the sparser the true precision matrix. In addition, the penalized pseudo-likelihood estimate lends itself better for path following algorithms and can be used to select an appropriate λ while simultaneously providing a good starting point for algorithms computing the exact solution.

One interesting question we will not pursue in this paper concerns the quality of the pseudo-likelihood approximation in the non-classical setting where p_n is allowed to grow with n . In that case, the classical efficiency argument favoring the exact maximum likelihood over the pseudo-likelihood approximation no longer holds. To the best of our knowledge, it is an open question whether the penalized exact ML has advantages over a pseudo-likelihood approximation in the non-parametric case (p_n growing with n).

2.4 Penalization by ℓ_2 -norm and Positive Semi-definiteness

As mentioned above, the algorithm we propose in Section 3 suffers from the drawback of not enforcing a positive semi-definite constraint. Imposing such constraint is costly in computational terms and would slow down our path following algorithm. As we will see in the experimental Section 4 below, the unconstrained estimate is positive semi-definite for the greater part of the regularization path even in nearly singular designs.

Before we review such experimental evidence, however, we study an alternative penalization that does result in positive semi-definite estimates for all levels of regularization. Consider the penalized pseudo-likelihood estimate defined by:

$$\begin{aligned} \hat{\mathbf{B}}_2(\lambda_2) &= \arg \min_{\tilde{\mathbf{B}}} \left\{ \text{tr} \left[(\mathbf{I}_p - \tilde{\mathbf{B}}) \tilde{\mathbf{X}}' \tilde{\mathbf{X}} (\mathbf{I}_p - \tilde{\mathbf{B}}') \right] + \lambda_2 \cdot \text{tr} \left[\tilde{\mathbf{B}}' \tilde{\mathbf{B}} \right] \right\} \\ \text{s.t.} & \quad \begin{cases} b_{jj} = 0 \\ \tilde{b}_{jk} = \tilde{b}_{kj} \end{cases} \end{aligned} \quad (15)$$

Our next result establishes that the estimate defined by is positive semi-definite along its entire path:

Theorem 1. *Let $\hat{\mathbf{C}}(\lambda_2) = \hat{\mathbf{D}}_2(\lambda_2)^{-1} \left(\mathbf{I}_p - \hat{\mathbf{B}}_2(\lambda_2) \right) \hat{\mathbf{D}}_2(\lambda_2)^{-1}$ be the precision matrix estimate resulting from (15). For any $\hat{\mathbf{D}}_2(\lambda_2)$ and $\lambda_2 > 0$, the precision matrix estimate $\hat{\mathbf{C}}_2(\lambda_2)$ is positive semi-definite.*

This result suggests that the ℓ_2 -penalty may be useful in inducing positive semi-definiteness. So an estimate incorporating both ℓ_2 - and ℓ_1 -penalties,

$$\begin{aligned} \hat{\mathbf{B}}_{EN}(\lambda_2, \lambda_1) &= \arg \min_{\tilde{\mathbf{B}}} \left\{ \text{tr} \left[(\mathbf{I}_p - \tilde{\mathbf{B}}) \tilde{\mathbf{X}}' \tilde{\mathbf{X}} (\mathbf{I}_p - \tilde{\mathbf{B}}') \right] + \lambda_2 \cdot \text{tr} \left[\tilde{\mathbf{B}}' \tilde{\mathbf{B}} \right] + \lambda_1 \cdot \left| \tilde{\mathbf{B}}' \tilde{\mathbf{B}} \right|_{w,1} \right\} \\ \text{s.t.} & \quad \begin{cases} b_{jj} = 0 \\ \tilde{b}_{jk} = \tilde{b}_{kj} \end{cases}, \end{aligned} \quad (16)$$

can have improved performance, over the ℓ_1 -penalty alone, in terms of positive semi-definiteness. The subscript *EN* is a reference to the Elastic Net penalty of Zou and Hastie (2005). In our experimental section below (Section 4), we have come across non-positive semi-definite precision

matrix estimate in none other than a few replications. We hence left a detailed study of the Elastic Net precision matrix estimate defined in (16) for future research.

3 Algorithms

We now detail an iterative algorithm for computing the regularization path for SPLICE estimates. The computational advantage of that approximate loss function used in (7) follows from two facts:

- For a fixed \mathbf{D} , the optimization problem reduces to that of tracing a constrained LASSO path;
- For a fixed \mathbf{B} , the optimizer $\hat{\mathbf{D}}$ has a closed form solution;

Based on these two facts, we propose an iterative algorithm for obtaining estimates $\hat{\mathbf{D}}(\hat{\lambda})$ and $\hat{\mathbf{B}}(\hat{\lambda})$. Within each step, the path for $\hat{\mathbf{D}}$ and $\hat{\mathbf{B}}$ as a function of the regularization parameter λ is traced and the choice of the regularization parameter λ is updated. We now describe the steps of our SPLICE algorithm.

1. Let $k = 1$ and $\hat{\mathbf{D}}_0 = \text{diag} \left(\frac{\|\mathbf{X}_j\|_2^2}{n} \right)_{j=1}^p$.

2. Repeat until *convergence*:

- (a) Set $\tilde{\mathbf{X}} = \mathbf{X}\hat{\mathbf{D}}_{k-1}^{-1}$ and use a path-tracing algorithm for the regularized problem:

$$\begin{aligned} \hat{\mathbf{B}}(\lambda) &= \arg \min_{\mathbf{B}} \left\{ \text{tr} \left[\tilde{\mathbf{X}}(\mathbf{I} - \tilde{\mathbf{B}}')(\mathbf{I} - \tilde{\mathbf{B}})\tilde{\mathbf{X}}' \right] + \lambda \cdot \|\tilde{\mathbf{B}}\|_{w,1} \right\} \\ \text{s.t.} \quad &\tilde{b}_{jk} = \tilde{b}_{kj} \\ &\tilde{b}_{jj} = 0 \end{aligned}$$

- (b) Compute the corresponding path for \mathbf{D} :

$$\hat{\mathbf{D}}_k(\lambda) = \text{diag} \left(\frac{\|(\mathbf{I}_p - \hat{\mathbf{B}}(\lambda))\mathbf{X}_j\|_2^2}{n} \right)_{j=1}^p$$

- (c) Select a value for the regularization parameter $\hat{\lambda}_k$;

(d) Set $\hat{\mathbf{B}}_k = \hat{\mathbf{B}}(\hat{\lambda}_k)$ and $\hat{\mathbf{D}}_k = \hat{\mathbf{D}}(\hat{\lambda}_k)$;

3. Return the estimate $\hat{\mathbf{C}} = \hat{\mathbf{D}}_k^{-1}(\mathbf{I} - \hat{\mathbf{B}}_k)\hat{\mathbf{D}}_k^{-1}$;

A subtle but important detail is the weighting vector \tilde{w} used when applying the penalty in step 2.(a). Since the path is traced in terms of the normalized $\tilde{\mathbf{B}}$ parameter instead of \mathbf{B} , a correction must be made in these weights. This can be determined by noticing that:

$$\|\mathbf{B}\|_{w,1} = \sum_{j,k=1}^p w_{jk} |b_{jk}| = \sum_{j,k=1}^p w_{jk} \frac{d_j}{d_k} \left| \frac{d_k}{d_j} b_{jk} \right| = \sum_{j,k=1}^p \tilde{w}_{jk} |\tilde{b}_{jk}| = \|\tilde{\mathbf{B}}\|_{\tilde{w},1},$$

as long as $\tilde{w}_{jk} = w_{jk} \frac{d_j}{d_k}$. As mentioned before, we fix $w_{jk} = 1$ throughout this paper, so we set $\tilde{w}_{jk} = w_{jk} \frac{d_j}{d_k}$ in our experiments. Of course, d_j are unknown so the current estimate is plugged-in every time step 2.(a) is executed.

In the remainder of this section, we show how to adapt the homotopy/LARS-LASSO algorithm to enforce the symmetry constraints $d_k^2 b_{jk} = d_j^2 b_{kj}$ along the regularization path, how to select $\hat{\mathbf{B}}$ and $\hat{\mathbf{D}}$ from the path, and discuss some convergence issues related to the algorithm above.

3.1 Enforcing the symmetry constraints along the path

The expression defining the regularization path in step 2.(a) of the SPLICE algorithm above can be rewritten as:

$$\begin{aligned} \hat{\mathbf{B}}(\lambda) &= \arg \min_{\mathbf{B}} \text{vec} \left(\tilde{\mathbf{X}}(\mathbf{I}_p - \mathbf{B}') \right)' \text{vec} \left(\tilde{\mathbf{X}}(\mathbf{I}_p - \mathbf{B}') \right) + \lambda \|\mathbf{B}\|_{w,1} \\ \text{s.t.} \quad &b_{jj} = 0 \\ &b_{jk} = b_{kj}, \end{aligned} \tag{17}$$

which is a quadratic form in \mathbf{B} penalized by a weighted version of its ℓ_1 -norm. To enforce the equality restriction in (17), we massage the data into an appropriate form so the homotopy/LARS-LASSO algorithm can be used in its original form.

The optimization problem in (17) corresponds to a constrained version of a penalized regression

of modified response (\mathbf{Y}) and predictors (\mathbf{Z}) given by:

$$\mathbf{Y} = \begin{bmatrix} \tilde{\mathbf{X}}_1 \\ \tilde{\mathbf{X}}_2 \\ \vdots \\ \tilde{\mathbf{X}}_p \end{bmatrix} \quad \text{and}$$

$$\mathbf{Z} = \begin{bmatrix} \tilde{\mathbf{X}}_{1*} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \tilde{\mathbf{X}}_{2*} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \tilde{\mathbf{X}}_{3*} & \cdots & \mathbf{0} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots & \tilde{\mathbf{X}}_{p*} \end{bmatrix}$$

Since the “model” for \mathbf{Y} given \mathbf{Z} is additive, we can force $b_{jk} = b_{kj}$ by creating a modified design matrix $\tilde{\mathbf{Z}}$ where the columns corresponding to b_{jk} and b_{kj} are summed into a single column. More precisely, the column corresponding to \tilde{b}_{jk} with $j < k$ in the $\tilde{\mathbf{Z}}$ design matrix has all elements set to zero except for the rows corresponding to \mathbf{X}_k and \mathbf{X}_j in the \mathbf{Y} vector. These rows must be set to $\tilde{\mathbf{X}}_j$ and $\tilde{\mathbf{X}}_k$ respectively:

$$\tilde{\mathbf{Z}} = \begin{bmatrix} \tilde{\mathbf{X}}_2 & \tilde{\mathbf{X}}_3 & \cdots & \tilde{\mathbf{X}}_{p-1} & \tilde{\mathbf{X}}_p & \mathbf{0} & \cdots & \mathbf{0} & \cdots & \mathbf{0} \\ \tilde{\mathbf{X}}_1 & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{0} & \tilde{\mathbf{X}}_3 & \cdots & \tilde{\mathbf{X}}_p & \cdots & \mathbf{0} \\ \mathbf{0} & \tilde{\mathbf{X}}_1 & \cdots & \mathbf{0} & \mathbf{0} & \tilde{\mathbf{X}}_2 & \cdots & \mathbf{0} & \cdots & \mathbf{0} \\ \vdots & & & & & & & & & \\ \mathbf{0} & \mathbf{0} & \cdots & \tilde{\mathbf{X}}_1 & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{0} & \cdots & \tilde{\mathbf{X}}_p \\ \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} & \tilde{\mathbf{X}}_1 & \cdots & \mathbf{0} & \mathbf{0} & \cdots & \tilde{\mathbf{X}}_{p-1} \end{bmatrix}$$

The path for the constrained $\tilde{\mathbf{B}}(\lambda)$ can then be traced by simply applying a weighted version of the LASSO algorithm to \mathbf{Y} and $\tilde{\mathbf{Z}}$. We emphasize that, even though the design matrix $\tilde{\mathbf{Z}}$ has $O(np^3)$ elements, it is extremely sparse. It can be proved that it contains only $O(np^2)$ non-zero terms. It is thus crucial that the implementation of the homotopy/LARS-LASSO algorithm used to trace this path make use of the sparsity of the design matrix.

3.2 Computational complexity of one step of the algorithm

The algorithm proposed above to trace the regularization path of $\hat{\mathbf{B}}(\lambda)$ uses the homotopy/LARS-LASSO algorithm with a modified regression with $\frac{p(p-1)}{2}$ regressors and np observations. The computational complexity of the k -th step of the homotopy algorithm in step 2.(a) is of order $O(a_k^2 + a_k np)$, where a_k denotes the number of non-zero terms on the upper triangular, off-diagonal of $\tilde{\mathbf{B}}(\lambda_k)$ (a detailed analysis is presented in Efron et al., 2004). Clearly, the computational cost increases rapidly as more off-diagonal terms are added to the precision matrix.

When p grows with n keeping a constant ratio p/n , we find it plausible that most selection criteria will pick estimates at the most regularized regions of the path: a data sample can only support so many degrees of freedom. Thus, incorporating early stopping criteria into the path tracing at step 2.(a) of the SPLICE algorithm can greatly reduce the computational cost of obtaining the path even further without degrading the statistical performance of the resulting estimates.

If one insists in having the entire precision matrix path, the SPLICE algorithm is still polynomial in p and n . In the case where no variables are dropped and the variables are added one at a time, the complexity of the first K steps of the path is given by $O(K^3 + K^2 np)$. Under the same assumptions and setting $K = 0.5 \cdot p(p-1)$, the SPLICE algorithm has cost of order $O(p^6 + np^5)$ to compute the entire path of solutions to the pseudo-likelihood problem. As a comparison, an algorithm presented by Banerjee et al. (2005) has a cost of order $O(\frac{p^{4.5}}{\eta})$ to compute an approximate solution for the penalized exact likelihood at a single point of the path, where η represents the desired level of approximation.

3.3 Selection criteria for λ :

Different criteria can be used to pick a pair \mathbf{B}, \mathbf{D} from the regularization path (cf. 2(c) in the SPLICE algorithm). In the experiment section below, we show the results of using Akaike's Information Criterion (AIC, Akaike, 1973), the Bayesian information Criterion (BIC, Schwartz, 1978) and a corrected version of the AIC criterion (AIC_C, Hurvich et al., 1998) using the unbiased estimate of the degrees of freedom of the LASSO along the regularization path. More precisely, we

set:

$$\hat{\lambda} = \arg \min_{\lambda} \mathcal{L}_{exact}(\mathbf{X}, \hat{\mathbf{D}}, \hat{\mathbf{B}}(\lambda)) + K(n, \hat{\text{df}}(\hat{\mathbf{C}}(\lambda)))$$

where:

$$K(n, \hat{\text{df}}(\hat{\mathbf{C}}(\lambda))) = \begin{cases} 2\hat{\text{df}}(\hat{\mathbf{C}}(\lambda)), & \text{for AIC,} \\ \frac{1 + \frac{\hat{\text{df}}(\hat{\mathbf{C}}(\lambda))}{n}}{1 - \frac{\hat{\text{df}}(\hat{\mathbf{C}}(\lambda)) + 2}{n}}, & \text{for AIC}_c, \\ \log(n)\hat{\text{df}}(\hat{\mathbf{C}}(\lambda)), & \text{for BIC,} \end{cases}$$

As our estimate of the degrees of freedom along the path, we follow Zou et al. (2007) and use p plus the number of non-zero terms in the upper triangular section of $\hat{\mathbf{B}}(\lambda)$, that is,

$$\hat{\text{df}}(\lambda) = p + \left| \left\{ (i, j) : i < j \text{ and } \hat{\mathbf{B}}_{ij}(\lambda) \neq 0 \right\} \right|. \quad (18)$$

The p term in expression (18) stems from the degrees of freedom used for estimating the means.

3.4 Stopping criterion

We now determine the convergence criterion we used to finish the loop started in step (2) of the SPLICE algorithm. To do that, we look at the variation in the terms of the diagonal matrix D .

We stop the algorithm once:

$$\max_{1 \leq j \leq p} \left\{ \log \left(\frac{[\hat{\mathbf{D}}_{k+1}]_j}{[\hat{\mathbf{D}}_k]_j} \right) \right\} < 10^{-2},$$

or when the number of iterations exceeds a fixed number Q , whichever occurs first.

We have also observed that letting λ_{k+1} differ from λ_k often resulted in oscillating estimates caused solely by small variations in λ from one step to the next. To avoid this issue, we fixed the regularization parameter λ after a number $M < Q$ of “warm-up” rounds.

We have set $M = 6$ and $Q = 100$ for the simulations in Section 4. For most cases, the selected value of λ and the $\hat{\mathbf{D}}$ matrix had become stable and the algorithm had converged before either the

maximum number M of “warm-up” steps was reached.

4 Numerical Results

In this section, we compare the results obtained by SPLICE in terms of estimation accuracy and model selection performance to other covariance selection methods, namely the Cholesky covariance selection procedure proposed by Huang et al. (2006) (Cholesky) and the ℓ_1 -penalized maximum likelihood estimate studied previously by Yuan and Lin (2007); Banerjee et al. (2005) and Friedman et al. (2008) (Exact PML). We compare the effectiveness of three different selection criteria – AIC, AIC_C and BIC (see Section 3.3)– in picking an estimate from the SPLICE, Cholesky and Penalized Exact Log-Likelihood regularization paths. We also compare the model selection performance over the entire regularization path using ROC curves (for details, see 4.1.2). Finally, we study the positive semi-definiteness of the SPLICE estimates along the regularization path in a near-singular design.

4.1 Comparison of SPLICE to alternative methods

Figure 1 shows the graphical models corresponding to the simulated data sets we will be using to compare SPLICE, Cholesky and Exact PML in terms of estimation accuracy and covariance selection. All designs involve a 15-dimensional zero-mean Gaussian random vector ($p = 15$) with precision matrix implied by the graphical models shown in Figure 1. A relatively small sample size is used to emulate the effect of high-dimensionality. For all comparisons, the estimates are computed based on either 20 or 1,000 independent samples from each distribution (small sample case: $n = 20$, large sample case: $n = 1,000$). The results presented here are based on $r = 200$ replications for each case.

Before we show the results, a few words on our choice of simulation cases:

- The star model (see Figure 1 panel a) provides an interesting example where the ordering of the variables can have a great influence in the results of an order-dependent method such as Cholesky. In the *Direct Star* case, the hub variable is the first entry in the 15-dimensional

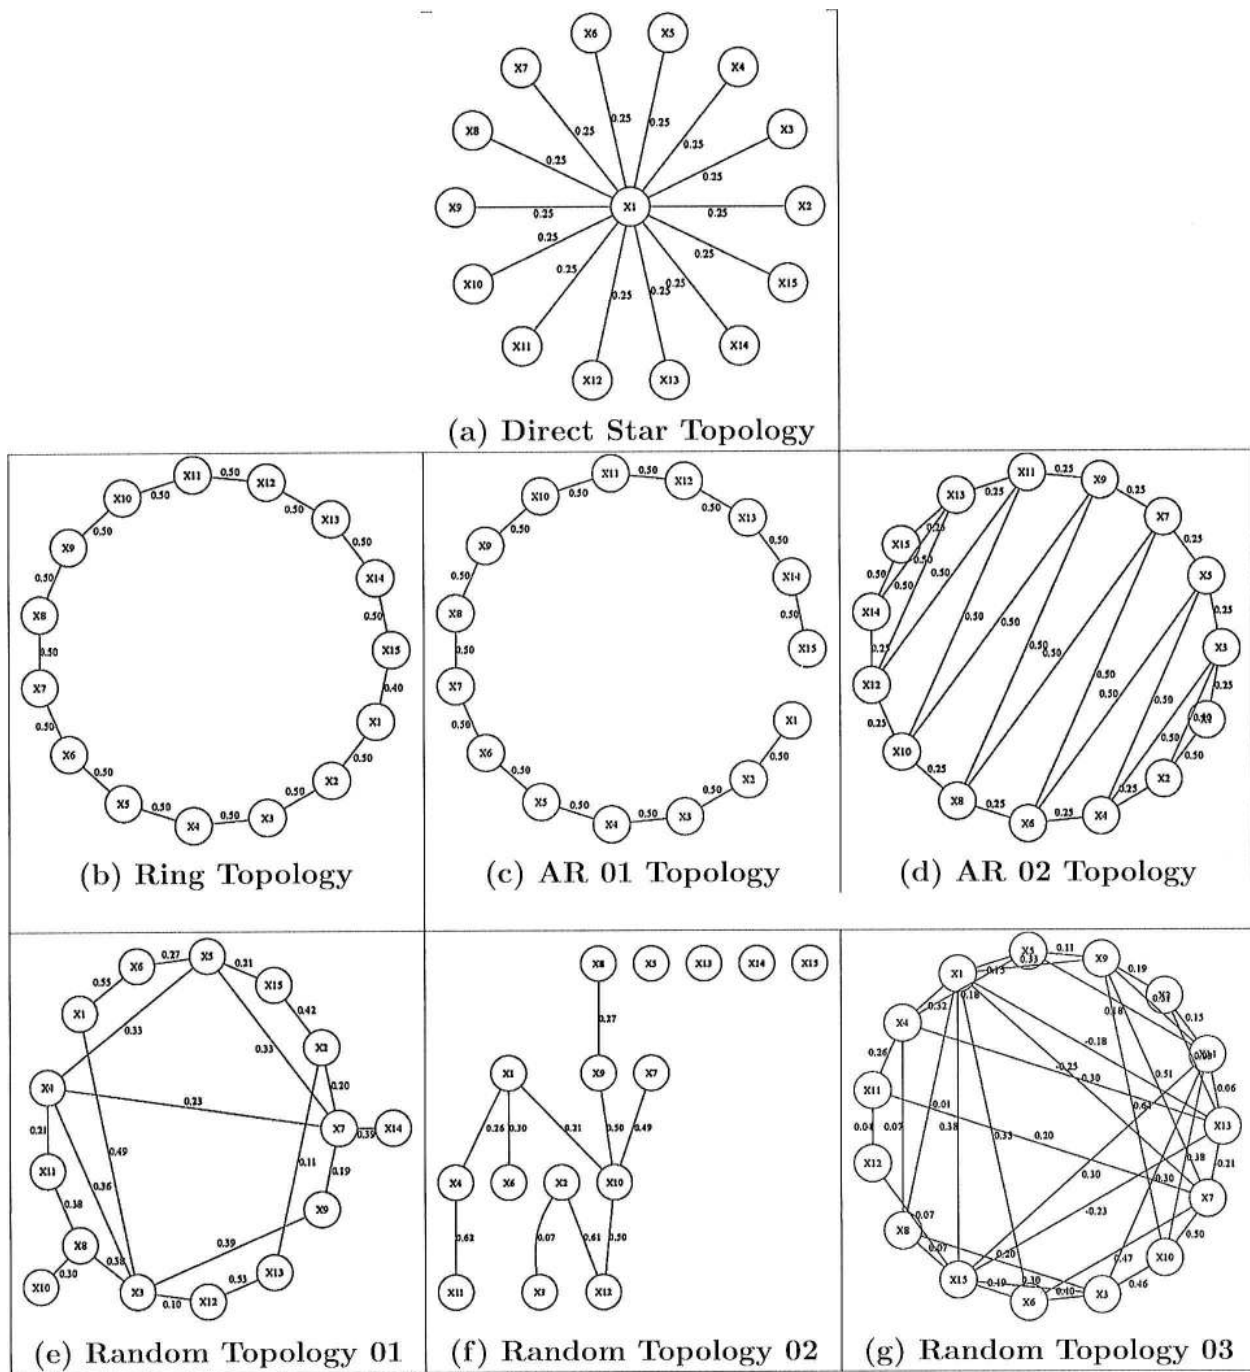


Figure 1: **Simulated cases:** In this table, the topology of the precision matrix for all simulated cases is shown. In each case, the precision matrix is normalized to have unit diagonal. The edges show the value of c_{ij} whenever it differs from zero.

random vector (as shown in the figure). Conditioning on the first variable (X_1) makes all other variables independent (X_2, \dots, X_{15}). Meanwhile, in the *inverted star* topology, the hub variable is put at the last position of the 15-dimensional random vector. As a result, no conditional independence is present until the last variable is added to the conditioning set.

- In the “AR-like” family of models (see Figure 1 panels b, c and d) each 15-dimensional random vector corresponds (panels c and d) or is very similar (panel b) to 15 observations from an auto-regressive process. This family of models tends to give some advantage to the Cholesky procedure as the ordering of the variables within the vector contain some information about the dependency structure among the variables. The cases in this family are loosely based on some of the simulation designs used in Yuan and Lin (2007);
- The “random” designs (see Figure 1 panels e, f and g) were obtained by randomly choosing precision matrices as described in Appendix C. We used these designs to make sure our results are valid in somewhat less structured environments.

4.1.1 Estimation accuracy of SPLICE, Cholesky and Exact PML

We evaluate the accuracy of the precision matrix estimates according to following four metrics.

- The quadratic loss of the estimated precision matrix, defined as

$$\Delta_2(\mathbf{C}, \hat{\mathbf{C}}) = \text{tr} \left(\mathbf{C} \hat{\mathbf{C}}^{-1} - \mathbf{I}_p \right)^2.$$

- The entropy loss at the estimated precision matrix, defined as

$$\Delta_e(\mathbf{C}, \hat{\mathbf{C}}) = \text{tr} \left(\mathbf{C} \hat{\mathbf{C}}^{-1} \right) - \log \left(\mathbf{C} \hat{\mathbf{C}}^{-1} \right) - n.$$

- The spectral norm of the deviation of the estimated precision matrix $\left(\hat{\mathbf{C}} - \mathbf{C} \right)$ where the spectral norm of a square matrix \mathbf{A} is defined as $\|\mathbf{A}\|_2 = \sup_x \frac{\|x' \mathbf{A} x\|}{\|x\|^2}$.
- The spectral norm of the deviation of the estimated covariance matrix $\left(\hat{\Sigma} - \Sigma \right)$.

For each of Cholesky, SPLICE and Exact PML, we compute estimates taken from their paths using the selection criteria mentioned in 3.3: AIC, AIC_C and BIC. The Cholesky and SPLICE estimates are chosen from the breakpoints of their respective regularization paths. The path tracing algorithm we have used for the Cholesky estimate is sketched in Appendix B.1. The path following algorithm for SPLICE is the one described in Section 3. The Exact ML estimate is chosen by minimizing the selection criterion over an equally spaced 500-point λ -grid between 0 and the maximum absolute value of the off-diagonal terms of the sample covariance matrix. We used the implementation of the ℓ_1 -penalized Exact log-likelihood for Matlab made available at Prof. Alexandre D’Aspremont’s web site (<http://www.princeton.edu/~aspremon/>).

Boxplots of the different accuracy measures for each of the methods and selection criteria are shown in Figures 2, 3 and 4 for the small sample case ($n = 20$), and in Figures 5, 6 and 7 for the large sample case ($n = 1,000$). For larger sample sizes, the Cholesky estimates do suffer some deterioration in terms of the entropy and quadratic losses when an inappropriate ordering of the variables is used. As we will later see, an inappropriate ordering can also degrade the Cholesky performance in terms of selecting the right covariance terms.

A comparison of the different methods reveals that the best method to use depends on whether the metric used is on the inverse covariance (precision matrix) or covariance matrix.

- With respect to all the three metrics on the inverse covariance (quadratic, entropy and spectral norm losses), the best results are achieved by SPLICE. In the case of the quadratic loss, this result can be partially attributed to the similarity between the quadratic loss and the pseudo-likelihood function used in the SPLICE estimate. In terms of the spectral norm on the precision matrix ($\|\mathbf{C} - \hat{\mathbf{C}}\|_2$), SPLICE performed particularly well in larger sample sizes ($n = 1,000$). For the quadratic and entropy loss functions, AIC was the criterion picking the best SPLICE estimates. In terms of the spectral norm loss, SPLICE performs better when coupled with AIC_C for small sample sizes ($n = 20$) and when coupled with BIC in larger sample sizes ($n = 1,000$).
- In terms of the spectral norm of the covariance estimate deviation ($\|\Sigma - \hat{\Sigma}\|_2$), the best performance was achieved by Exact PML. The performance of Exact PML was somewhat

Performance Metric	Recommended procedure
$\Delta_2(\mathbf{C}, \hat{\mathbf{C}})$	SPLICE + AIC
$\Delta_e(\mathbf{C}, \hat{\mathbf{C}})$	SPLICE + AIC
$\ \hat{\mathbf{C}} - \mathbf{C}\ _2$	SPLICE + AIC _C (for smaller sample size, $n = 20$) SPLICE + BIC (for larger sample size, $n = 1,000$)
$\ \hat{\Sigma} - \Sigma\ _2$	Exact penalized ML + AIC

Table 1: **Suitable estimation methods for different covariance structures and performance metrics:** The results shown here are a summary of the results shown in Figures 2 through 7. For each metric, we show the best combination of estimation method and selection criterion based on our simulations.

insensitive to the selection criterion used in many cases: this may be caused by the uniform grid on λ missing regions where the penalized models rapidly change as λ varies. Based on the cases where the selection criterion affected the performance of Exact PML, BIC seems to yield the best results in terms of $\|\Sigma - \hat{\Sigma}\|_2$.

For ease of reference, these results are collected in Table 1.

4.1.2 Model Selection performance of SPLICE

To evaluate the model selection performance of the different covariance selection methods, we compare their Relative Operating Characteristic (ROC) curves defined as a curve containing in its horizontal axis the minimal number of false positives that is incurred on average for a given number of true number of positives, shown on the vertical axis, to be identified. The ROC curve for a method shows its model selection performance over all possible choices of the tuning parameter λ . We have chosen to compare ROC curves instead of the results for particular selection criterion as different applications may penalize false positives and negatives differently.

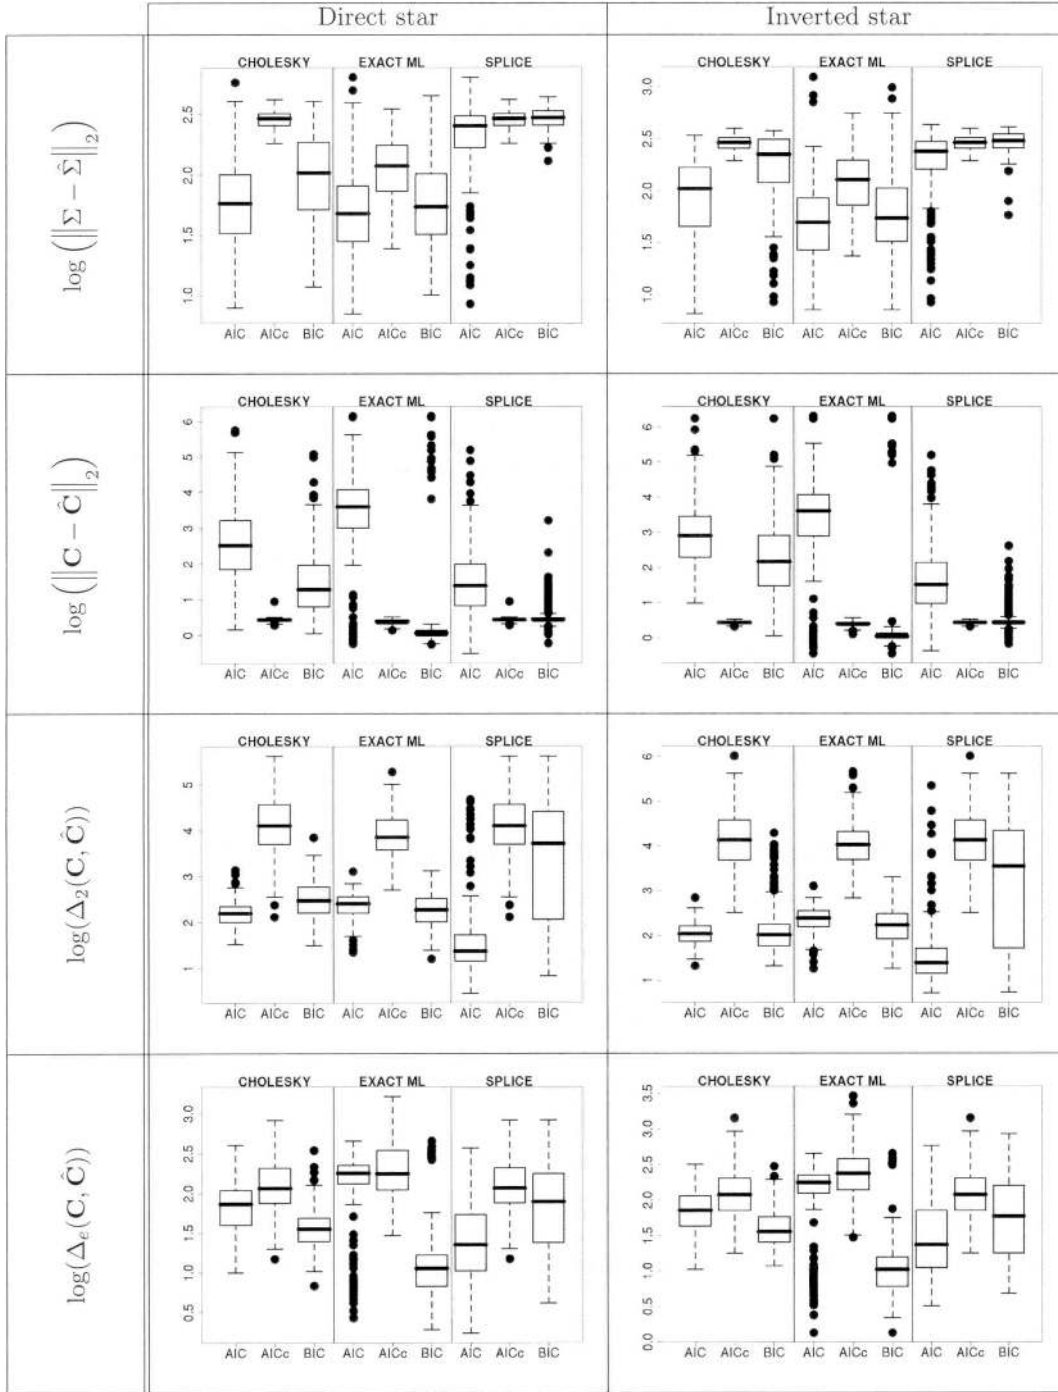


Figure 2: Accuracy metrics for precision matrix estimates in the Star cases for $p = 15$ and $n = 20$

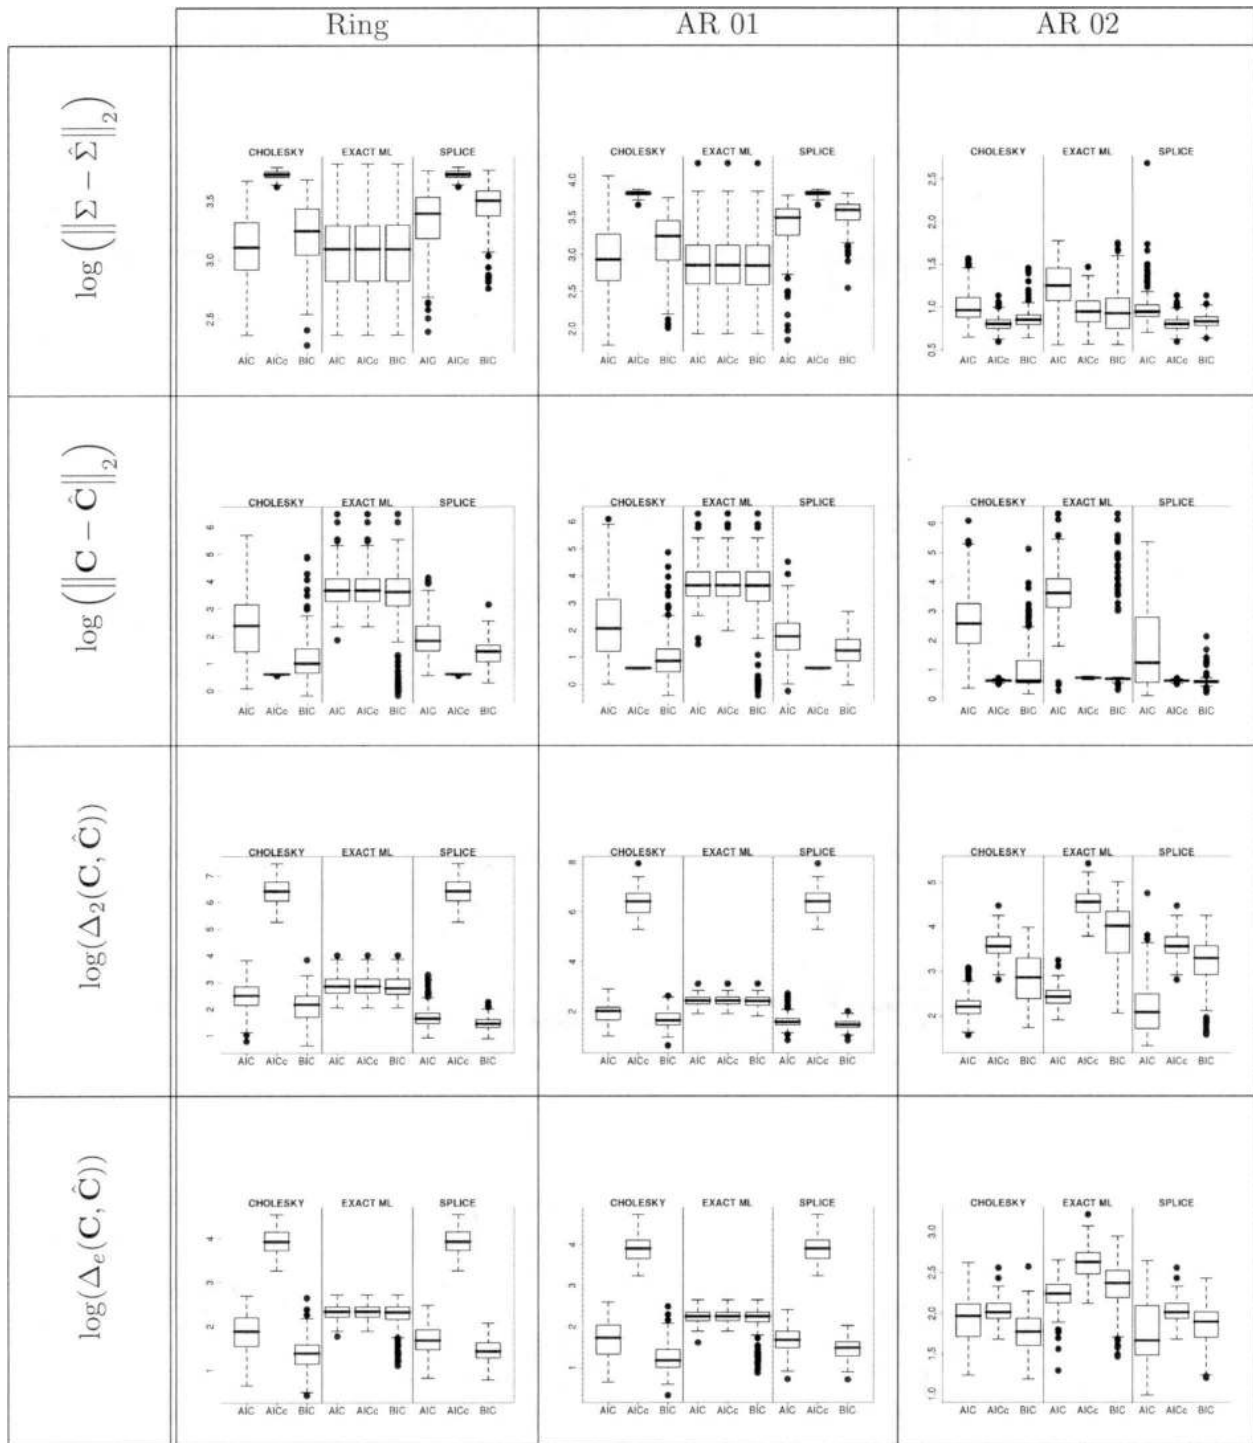


Figure 3: Accuracy metrics for precision matrix estimates in the AR-like cases for $p = 15$ and $n = 20$

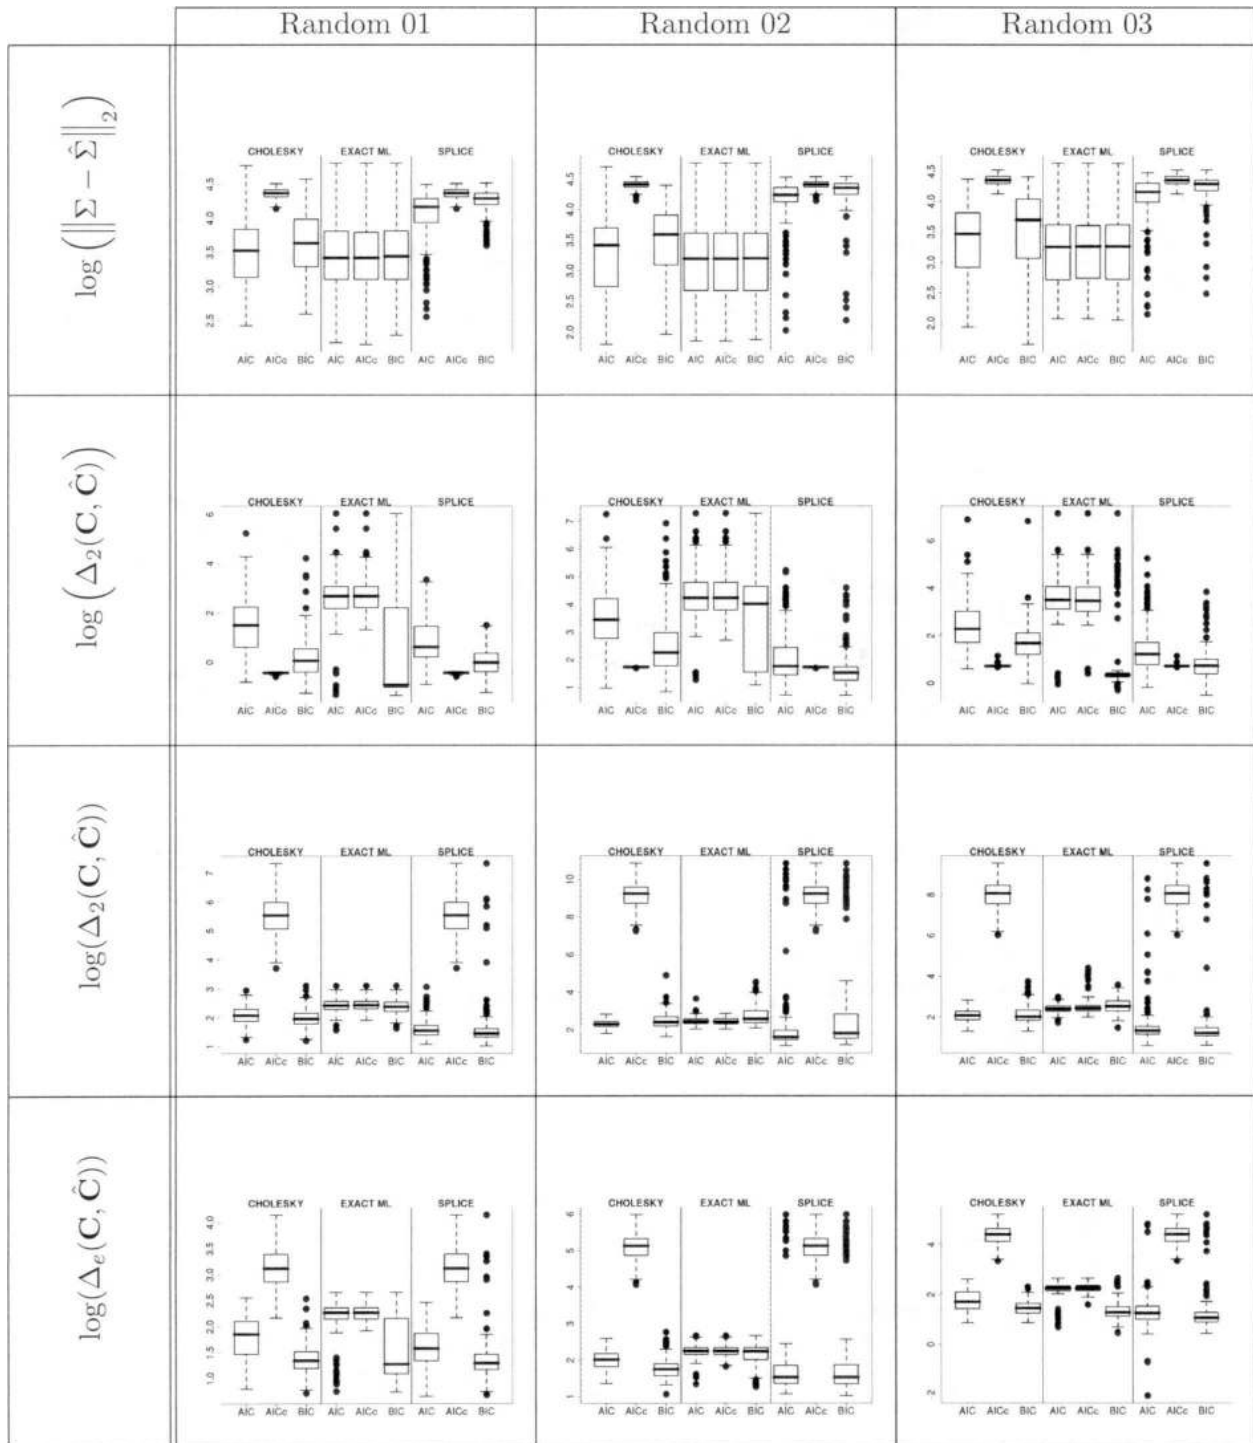


Figure 4: Accuracy metrics for precision matrix estimates in the randomly generated designs for $p = 15$ and $n = 20$

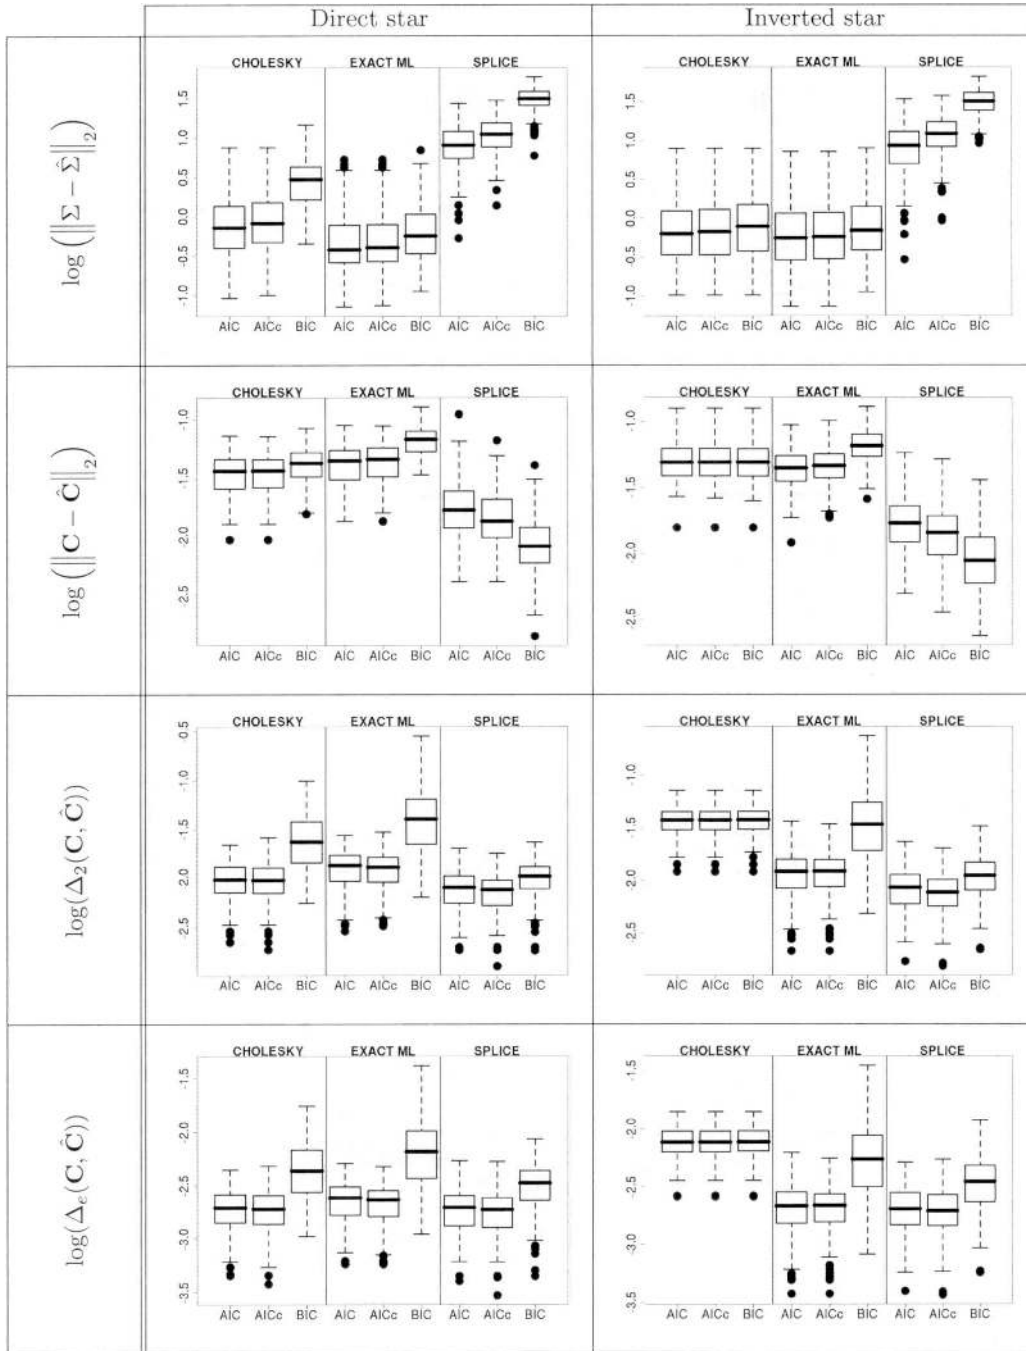


Figure 5: Accuracy metrics for precision matrix estimates in the Star cases for $p = 15$ and $n = 1,000$:

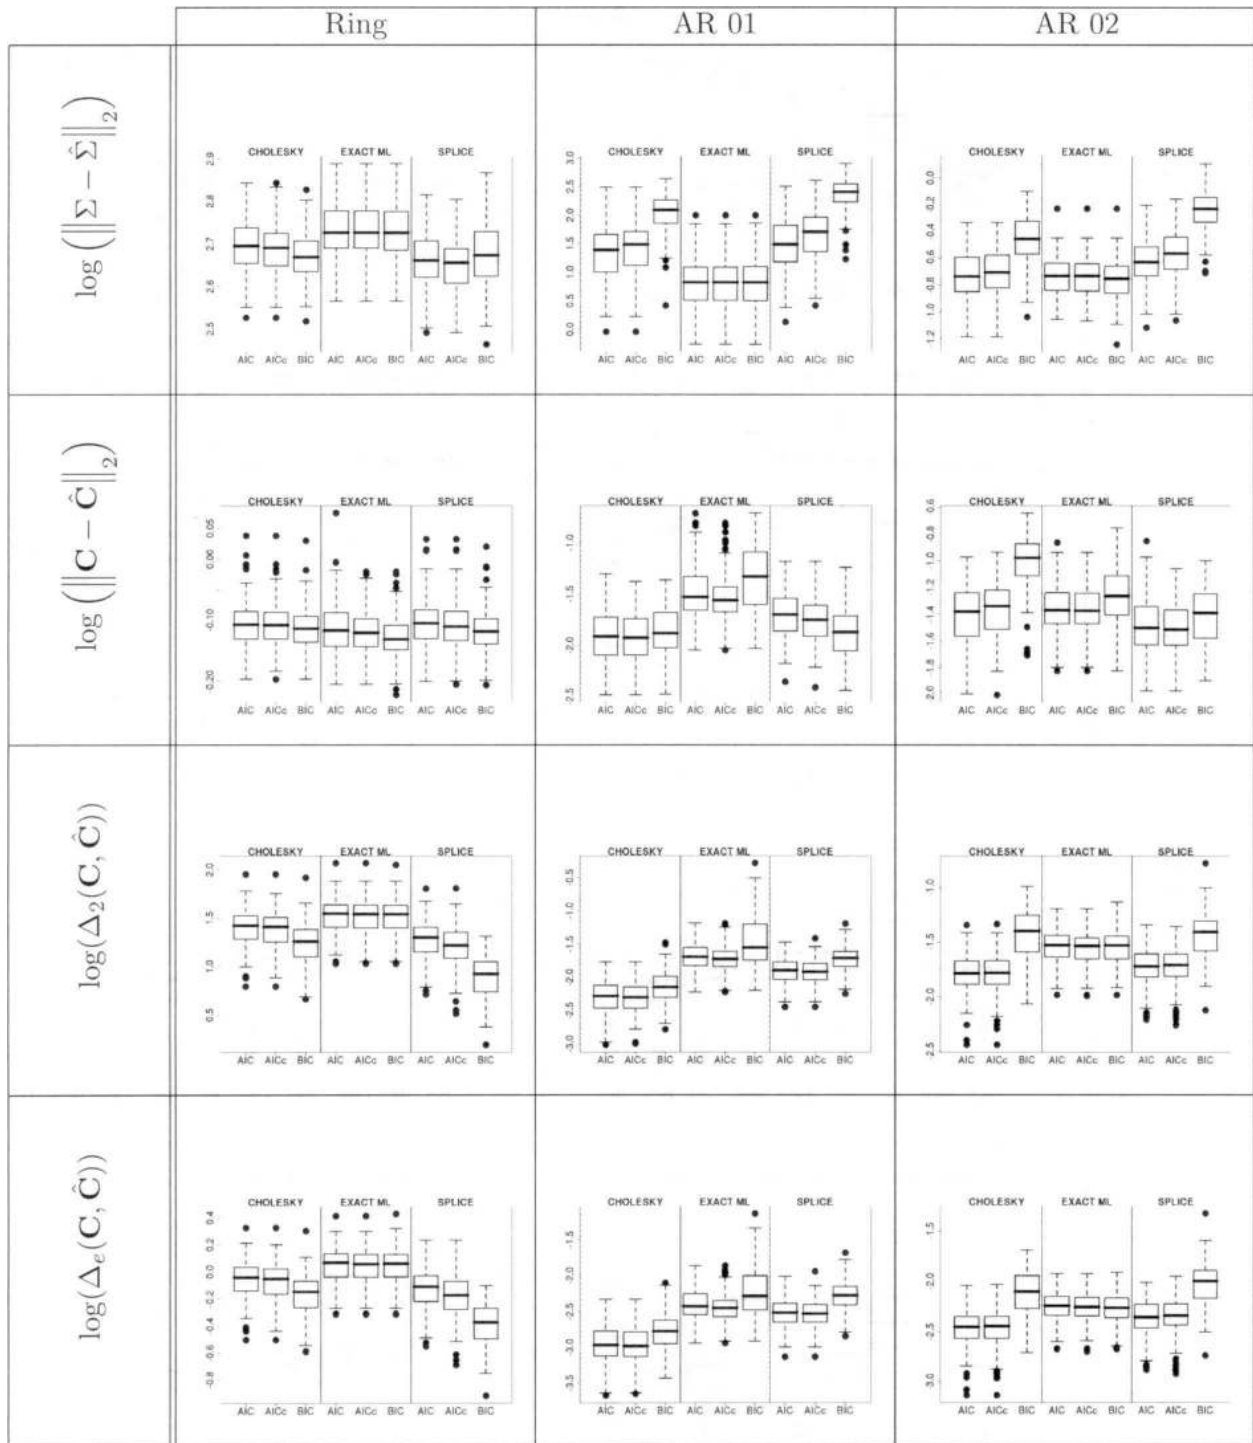


Figure 6: Accuracy metrics for precision matrix estimates in the AR-like cases for $p = 15$ and $n = 1,000$

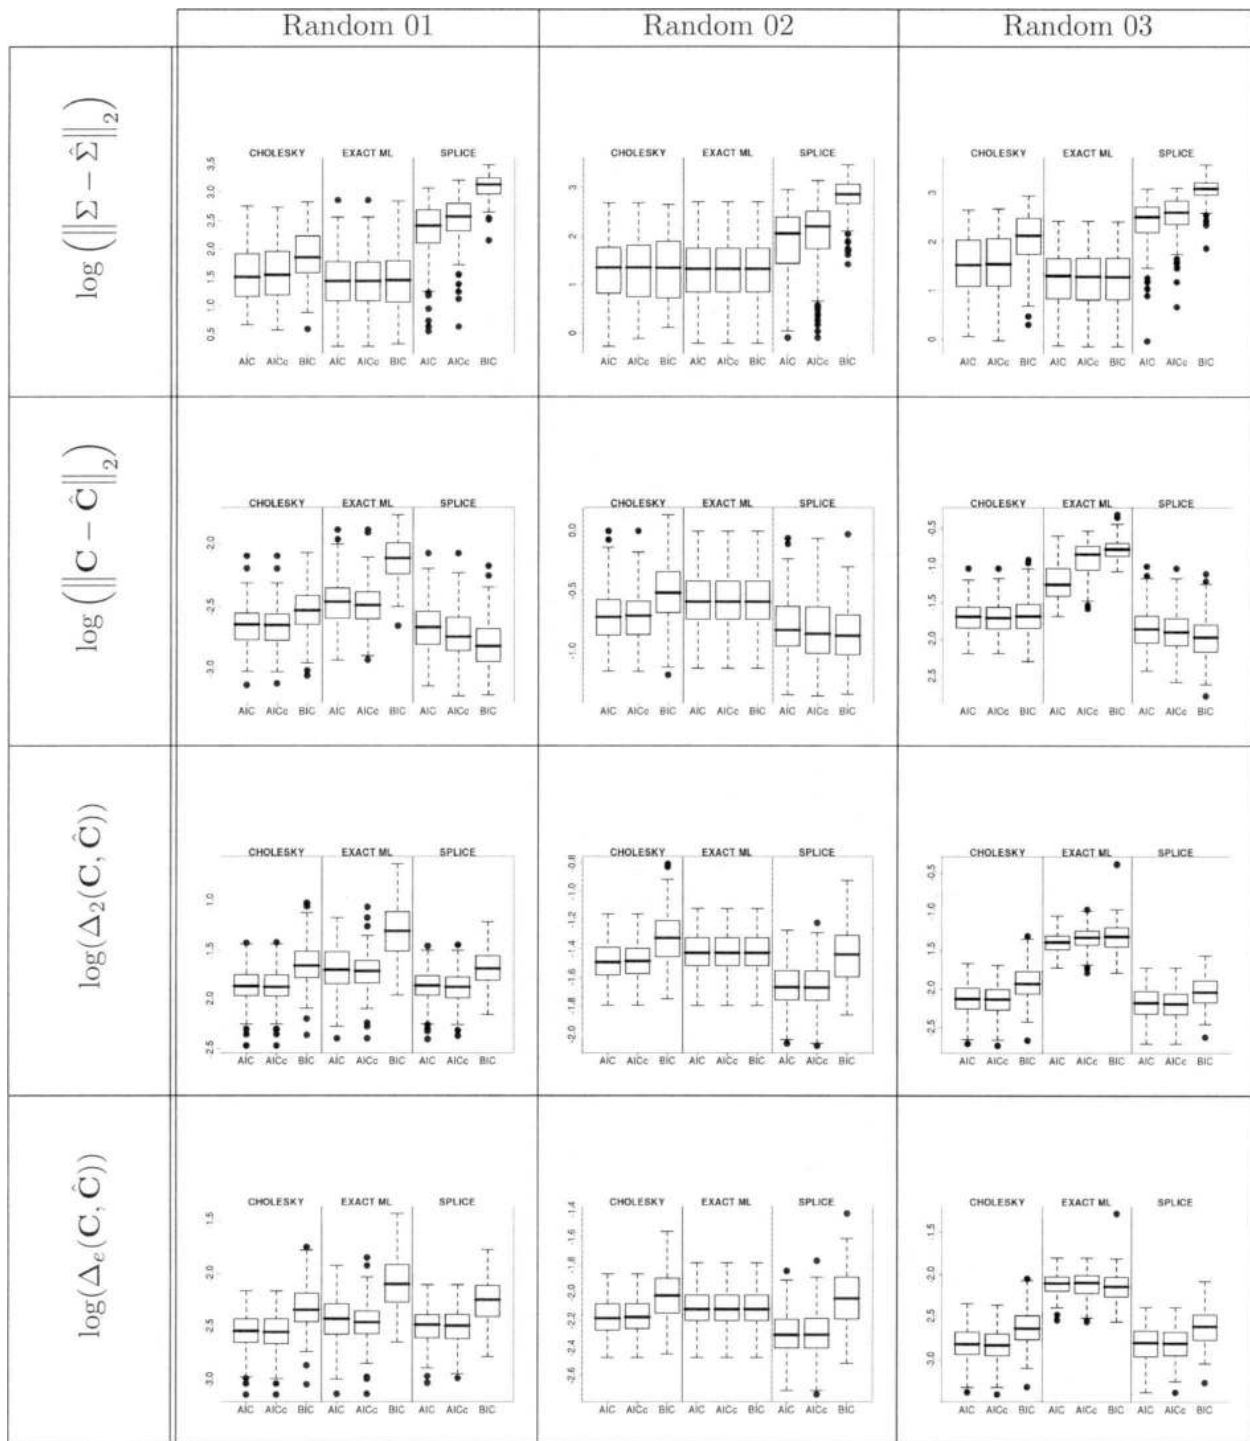


Figure 7: Accuracy metrics for precision matrix estimates in the randomly generated designs for $p = 15$ and $n = 1,000$

For a fixed number of true positives on the vertical axis, we want the expected minimal number of false positives to be as low as possible. A covariance selection method is thus better the closer its ROC curve is to the upper left side of the plot.

Figures 8 and 9 compare the ROC curves for the Cholesky and SPLICE covariance selection procedures for sample sizes $n = 20$ and $n = 1,000$ respectively. The Exact ML does not have its ROC curve shown: the grid used to approximate its regularization path often did not include estimates with a specific number of true positives. A finer grid can ameliorate the problem, but would be prohibitively expensive to compute (recall we used a grid with 500 equally spaced values of λ). This illustrates an advantage of path following algorithms over using grids: path following performs a more thorough search on the space of models.

The mean ROC curves on Figures 8 and 9 show that SPLICE had a better performance in terms of model selection in comparison to the Cholesky method over all cases considered. In addition, for a given number of true positives, the number of false positives incurred by SPLICE decreases significantly as the sample size increases. Our results also suggest that, with the exception of the Random Design 02, the chance that the SPLICE path contains the true model approaches one as the sample size increases for all simulated cases.

Finally, we consider the effect of ordering the variables on the selection performance of SPLICE and Cholesky. To do this, we restrict attention to the “star” cases and compare the performance of SPLICE and Cholesky when the variables are presented in the “correct” order (X_1, X_2, \dots, X_p) and in the inverted order $(X_p, X_{p-1}, \dots, X_1)$. Figure 10 shows the boxplot of the minimal number of false positives on 200 replications of the Cholesky and SPLICE paths for selected number of true positives in the small sample case ($n = 20$). In addition to outperforming Cholesky by a wide margin, SPLICE is not sensitive to the order in which the variables are presented. Cholesky, on the other hand, suffers some further degradation in terms of model selection when the variables are presented in the reverse order.

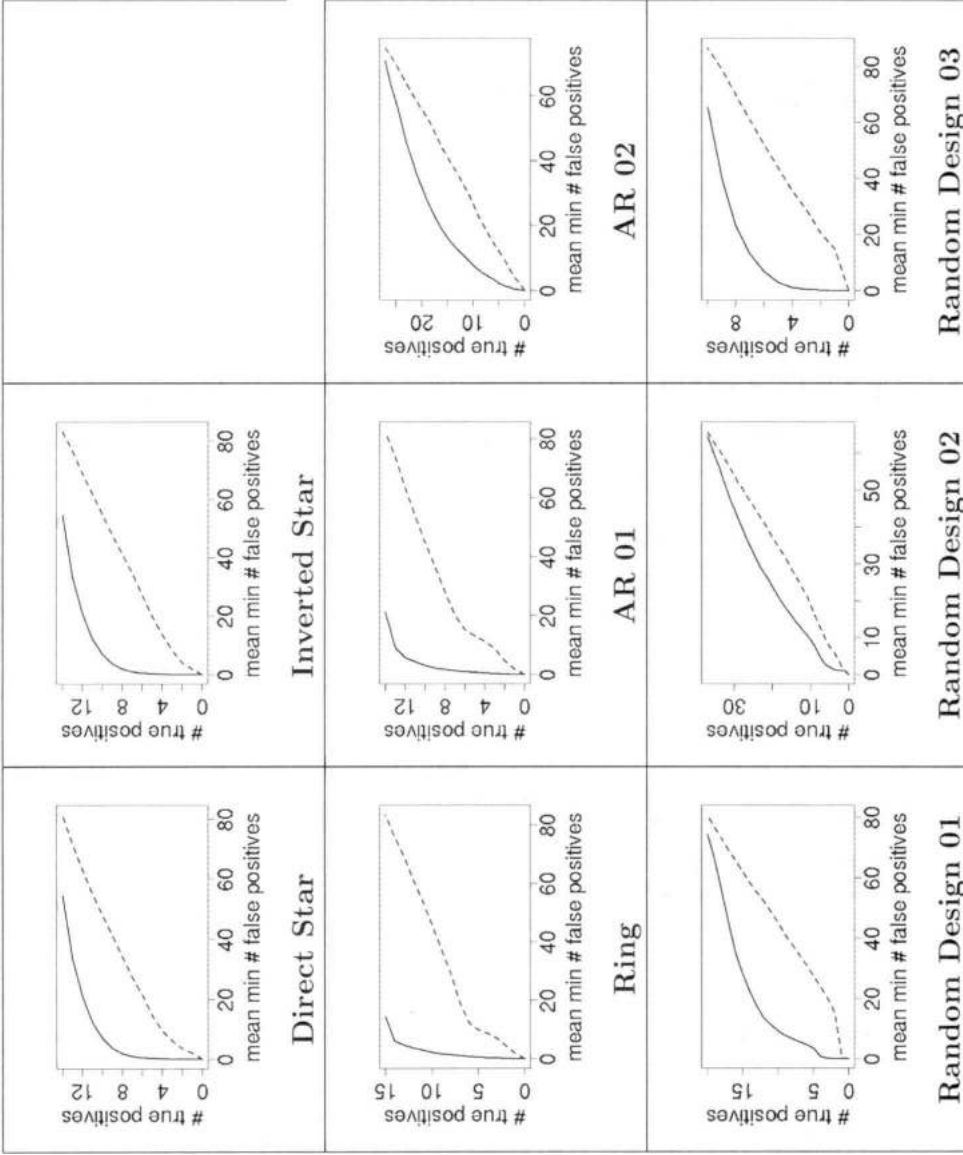


Figure 8: “Mean ROC curves” for the Cholesky and SPLICE covariance selection for $p = 15$ and $n = 20$: Within each panel, the relative operating characteristic (ROC) curve shows the mean minimal number of false positives (horizontal axis) needed to achieve a given number of true positives (vertical axis) for both the SPLICE (solid lines) and Cholesky (dashed lines). A selection procedure is better the more its curve approaches the upper left corner of the plot. Our results suggest that SPLICE trades off better than Cholesky between false and true positives across all cases considered.

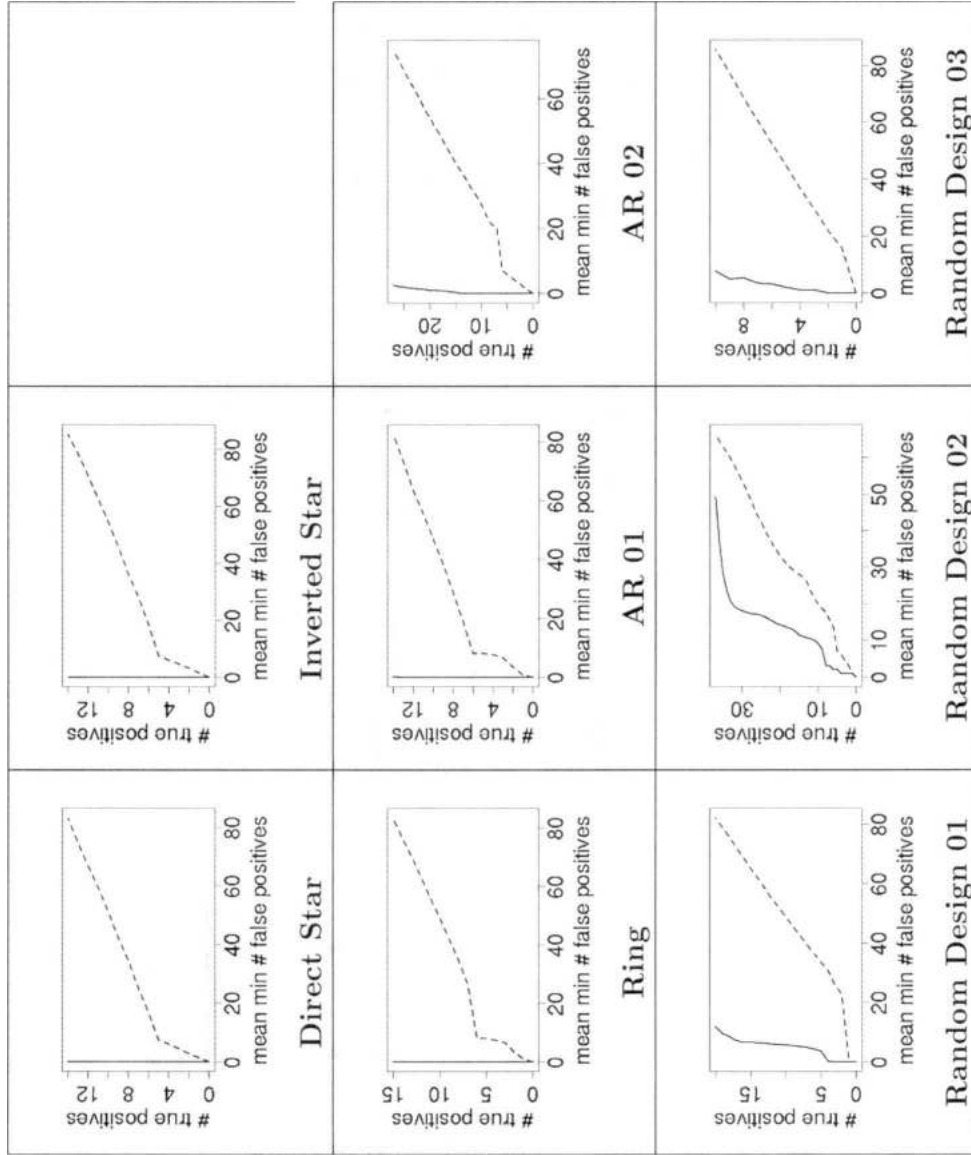


Figure 9: “Mean ROC curves” for the Cholesky and SPLICE covariance selection for $p = 15$ and $n = 1,000$: Within each panel, the relative operating characteristic (ROC) curve shows the mean minimal number of false positives (horizontal axis) needed to achieve a given number of true positives (vertical axis) for both the SPLICE (solid lines) and Cholesky (dashed lines). A selection procedure is better the more its curve approaches the upper left corner of the plot. Our results suggest that SPLICE picks correct covariance terms with a larger sample size for large enough sample sizes. The Cholesky estimates, however, do not seem much more likely to select correct covariance terms in this larger sample size in the comparison with the case $n = 20$.

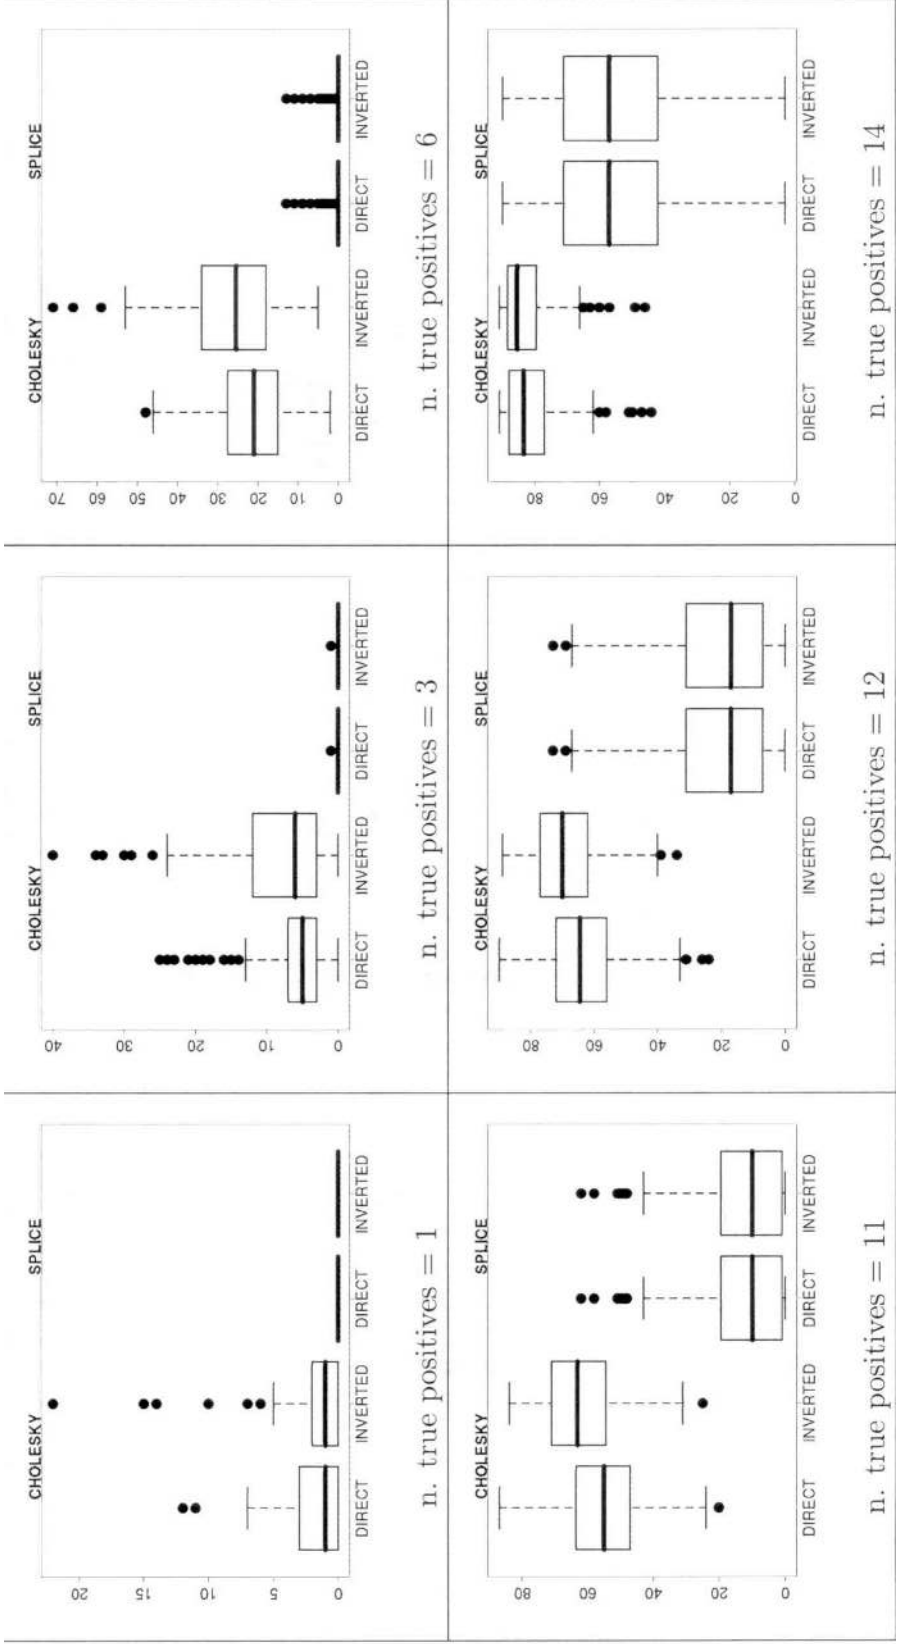


Figure 10: **Horizontal detail for ROC curves in the “star” cases:** Within each panel, we show a boxplot of the minimal number of false positives on the path for the indicated number of true positives. Each panel can be thought of as the boxplot corresponding to a horizontal layer on the graph shown in Figure 8. SPLICE is insensitive to the permutation of the variables (compare inverted and direct). Cholesky performs worse than SPLICE in both the direct and inverted cases: its performance deteriorates further if the variables are presented in inverted order.

4.2 Positive Semi-Definiteness along the regularization path

As noted in Section 2 above, there is no theoretical guarantee that the SPLICE estimates be positive semi-definite. In the somewhat well-behaved cases studied in the experiments of Section 4.1 (see Figure 1) all of the estimates selected by either AIC_C and BIC were positive semi-definite cases. In only 6 out of the 1,600 simulated cases, did AIC choose a slightly negative SPLICE estimate. This, however, tells little about the positive-definiteness of SPLICE estimates in badly behaved cases. We now provide some experimental evidence that the SPLICE estimates can be positive semi-definite for most of the regularization path even when the true covariance matrix is nearly singular.

The results reported for this experiment are based on 200 replications of SPLICE applied to a data matrix \mathbf{X} sampled from a Gaussian distribution with near singular covariance matrix. The number of observations ($n = 40$) and the dimension of the problem ($p = 30$) are kept fixed throughout. To obtain a variety of near singular covariance matrices, the sample covariance $\Sigma \in \mathbb{R}^{p \times p}$ of each of the 200 replications is sampled from:

$$\Sigma \sim \text{Wishart}(n, \bar{\Sigma}), \text{ with } [\bar{\Sigma}]_{ij} = 0.99^{|i-j|}.$$

The covariance matrices sampled from distribution have expected value $\bar{\Sigma}$, which is itself close to singular. We let the number of degrees of freedom of the Wishart distribution be small (equal to the sample size $n = 30$) to make designs close to singular more likely to happen. Once Σ is sampled, the data matrix \mathbf{X} is then formed from an i.i.d. sample of a $N(0, \Sigma)$ distribution.

To align the results along the path of different replications, we create an index $\bar{\lambda}$ formed by dividing a λ on the path by the the maximum λ at that path. This index varies over $[0, 1]$ and lower values of $\bar{\lambda}$ correspond to less regularized estimates. Figure 11 shows the minimum eigenvalue of $\hat{\mathbf{C}}(\lambda)$ versus the index $\bar{\lambda}$ for the 200 simulated replications. We can see that non-positive estimates only occur near the very end of the path (small values of λ) even in such an extreme design.

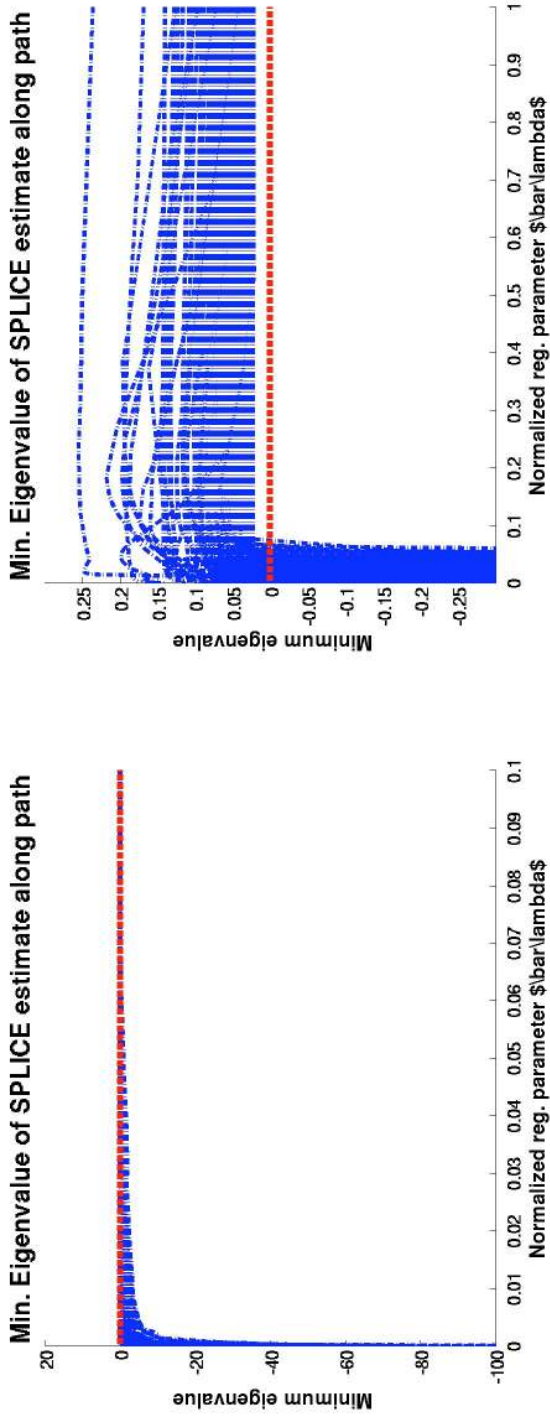


Figure 11: **Positive definiteness along SPLICE path:** At different vertical and horizontal scales, the two panels show the the minimum eigenvalue of the SPLICE estimate $\hat{C}(\lambda)$ along the regularization path as a function of the normalized regularization parameter λ . At the unregularized extreme (small λ), \hat{C} can have very negative eigenvalues (left panel). However, for a long stretch of the regularization path – over 90% of the path as measured by the index λ – the SPLICE estimate \hat{C} is positive definite (right panel).

5 Discussion

In this paper, we have defined Sparse Pseudo-Likelihood Inverse Covariance Estimates (SPLICES) as a ℓ_1 -penalized pseudo-likelihood estimate for precision matrices. The SPLICE loss function (6) is obtained from extending previous work by Meinshausen and Bühlmann (2006) to obtain estimates of precision matrix that are symmetric. The SPLICE estimates are formed from estimates of the coefficients and variance of the residuals of linear regression models.

The main advantage of the estimates proposed here is algorithmic. The regularization path for SPLICE estimates can be efficiently computed by alternating the estimation of coefficients and variance of the residuals of linear regressions. For fixed estimates of the variance of residuals, the complete path of the regression coefficients can be traced efficiently using an adaptation of the homotopy/LARS-LASSO algorithm (Osborne et al., 2000; Efron et al., 2004) that enforces the symmetry constraints along the path. Given the path of regression coefficients, the variance of the residuals can be estimated by means of closed form solutions. An analysis of the complexity of the algorithm suggests that early stopping can reduce its computational cost further. A comparison of the pseudo-likelihood approximation to the exact likelihood function provides another argument in favor of early stopping: the pseudo-likelihood approximation is better the sparser the estimated model. Thus moving on to the lesser sparse stretches of the regularization path can be not only computationally costly but also counterproductive to the quality of the estimates.

We have compared SPLICE with ℓ_1 -penalized covariance estimates based on Cholesky decomposition (Huang et al., 2006) and the exact likelihood expression (Banerjee et al., 2005, 2007; Yuan and Lin, 2007; Friedman et al., 2008) for a variety of sparse precision matrix cases and in terms of four different metrics, namely: quadratic loss, entropy loss, spectral norm of $\mathbf{C} - \hat{\mathbf{C}}$ and spectral norm of $\Sigma - \hat{\Sigma}$. SPLICE estimates had the best performance in all metrics with the exception of the spectral norm of $\Sigma - \hat{\Sigma}$. For this last metric, the best results were achieved by using the ℓ_1 -penalized exact likelihood.

In terms of selecting the right terms of the precision matrix, SPLICE was able to pick a given number of correct covariance terms while incurring in less false positives than the ℓ_1 -penalized Cholesky estimates of Huang et al. (2006) in various simulated cases. Using an uniform grid

for the exact penalized maximum likelihood provided few estimates with a mid-range number of correctly picked covariance terms. This reveals that path following methods perform a more thorough exploration of the model space than penalized estimates computed on (pre-determined) grids of values for the regularization parameter.

While SPLICE estimates are not guaranteed to be positive semi-definite along the entire regularization path, they have been observed to be such for most of the path even in a almost singular problem. Over tamer cases, the estimates selected by AIC, BIC and AIC_C were positive semi-definite in the overwhelming majority of cases (1,594 out of 1,600).

A Proofs

A.1 Positive Semi-definiteness of ℓ_2 -penalized Pseudo-likelihood estimate

We now prove Theorem 1. First, we rewrite the ℓ_2 -norm penalty in a more convenient form:

$$p + \text{tr} \left[\tilde{\mathbf{B}}' \tilde{\mathbf{B}} \right] = p + \sum_{j=1}^p \sum_{k=1}^p (-b_{jk})^2 = \sum_{j=1}^p (1 - b_{jj})^2 + \sum_{j=1}^p \sum_{k=1}^p (-b_{jk})^2 = \text{tr} \left[(\mathbf{I}_p - \tilde{\mathbf{B}})' (\mathbf{I}_p - \tilde{\mathbf{B}}) \right]$$

Hence, the ℓ_2 -penalized estimate defined in (15), can be rewritten as:

$$\begin{aligned} \hat{\mathbf{B}}_2(\lambda_2) &= \arg \min_{\tilde{\mathbf{B}}} \text{tr} \left[(\mathbf{I}_p - \tilde{\mathbf{B}}) \left(\tilde{\mathbf{X}}' \tilde{\mathbf{X}} + \lambda_2 \mathbf{I}_p \right) (\mathbf{I}_p - \tilde{\mathbf{B}}') \right] \\ \text{s.t.} & \begin{cases} \tilde{b}_{jj} &= 0, \text{ for } j = 1, \dots, p \\ \tilde{b}_{jk} &= \tilde{b}_{kj} \text{ for } j = 1, \dots, p, k = j + 1, \dots, p \end{cases} \end{aligned} \quad (19)$$

Using convexity, The Karush-Kuhn-Tucker (KKT) conditions are necessary and sufficient to characterize a solution of problem (19):

$$\left(\tilde{\mathbf{X}}' \tilde{\mathbf{X}} + \lambda_2 \mathbf{I}_p \right) (\mathbf{I}_p - \hat{\mathbf{B}}_2(\lambda_2)) + \Theta + \Omega = 0, \quad (20)$$

where Θ is a diagonal matrix and Ω is an anti-symmetric matrix. Given that $\tilde{b}_{jj} = 0$ and $\omega_{jj} = 0$

(anti-symmetry), it follows that, for $\lambda_2 > 0$:

$$\theta_{jj} = -\left(\tilde{\mathbf{X}}_j' \tilde{\mathbf{X}}_j + \lambda_2\right) < 0, \quad (21)$$

that is, $-\Theta$ is a positive definite matrix.

From (20), we can conclude that $\left(\mathbf{I}_p - \hat{\mathbf{B}}_2\right)$ satisfies:

$$\left(\tilde{\mathbf{X}}' \tilde{\mathbf{X}} + \lambda_2 \mathbf{I}_p\right) \left(\mathbf{I}_p - \hat{\mathbf{B}}_2(\lambda_2)\right) + \left(\mathbf{I}_p - \hat{\mathbf{B}}_2(\lambda_2)\right)' \left(\tilde{\mathbf{X}}' \tilde{\mathbf{X}} + \lambda_2 \mathbf{I}_p\right) = -2\Theta.$$

Theorem 1 then follows from setting $\mathbf{U} = \left(\mathbf{I}_p - \hat{\mathbf{B}}_2(\lambda_2)\right)$, $\mathbf{V} = \left(\tilde{\mathbf{X}}' \tilde{\mathbf{X}} + \lambda_2 \mathbf{I}_p\right)$ and $\mathbf{W} = -\Theta$ and applying the following lemma:

Lemma 1. *Let \mathbf{U} , \mathbf{V} and \mathbf{W} be $p \times p$ symmetric matrices. Suppose that \mathbf{V} is strictly positive definite and \mathbf{W} is positive semi-definite and:*

$$\mathbf{UV} + \mathbf{VU} = \mathbf{W}.$$

It follows that \mathbf{U} is positive semi-definite.

Proof. Since \mathbf{V} is symmetric positive semi-definite, we can write it as $\mathbf{V} = \mathbf{A}\mathbf{A}'$. Take a vector $z \in \mathbb{R}^p$ and rewrite it as $z = \sum_{k=1}^p \gamma_k a_k$ where a_k are eigenvectors of V (no need for uniqueness).

From positive semi-definiteness of \mathbf{W} and the assumed identity:

$$\begin{aligned} 0 \leq z' \mathbf{W} z &= z' \mathbf{UV} z + z' \mathbf{VU} z \\ &= \sum_{j=1}^p \sum_{k=1}^p \gamma_j \gamma_k (a_k' \mathbf{UV} a_j + a_k' \mathbf{VU} a_j) \\ &= 2 \sum_{j=1}^p \gamma_j^2 a_k' \mathbf{UV} a_k \end{aligned}$$

where the second follows from the cross products being zero:

$$a_j' \mathbf{UV} a_k = \text{trace}(a_j' \mathbf{UV} a_k) = \text{trace}(a_k a_j' \mathbf{UV}) = 0, \text{ whenever } j \neq k.$$

Now, taking in particular $z = a_k$ we have, for every $k = 1, \dots, p$:

$$2a'_k \mathbf{U} \mathbf{V} a_k = 2\lambda_k a'_k \mathbf{U} a_k = a'_k \mathbf{W} a_k \geq 0$$

We can conclude that for every k having $\lambda_k > 0$, $a'_k \mathbf{U} a_k \geq 0$ and the result follows. \square

B Algorithms

B.1 Appendix: A Path-Following Algorithm for the Cholesky estimate

In this section, we describe a path tracing algorithm for the precision matrix estimate based on Cholesky decomposition introduced in Huang et al. (2006). The algorithm can be understood as a block-wise coordinate optimization in the same spirit as Friedman et al. (2007).

For a fixed diagonal matrix $\mathbf{D}^2 = \text{diag}(d_1^2, \dots, d_p^2)$, the sparse Cholesky estimate of Huang et al. (2006) is:

$$\hat{\mathbf{U}}(\lambda) = \arg \min_{\mathbf{U} \in \text{UUT}} \mathbf{X} \mathbf{U} \mathbf{D}^{-2} \mathbf{U}' \mathbf{X}' + \lambda \|\mathbf{U}\|_1,$$

where UUT denotes the space of upper triangular matrices with unit diagonal. This is equivalent to solving:

$$\hat{\beta}(\lambda) = \arg \min_{\beta} \sum_{j=1}^p \frac{\|\mathbf{X}_j - \sum_{k=1}^{j-1} \mathbf{X}_k \beta_{jk}\|^2}{d_j^2} + \lambda \sum_{j=1}^p \sum_{k=1}^{j-1} |\beta_{jk}|.$$

It is not hard to see that the objective function can be broken into $p - 1$ uncoupled small components. As a result, the optimization problem can be separated into $p - 1$ smaller problems, that is, $\hat{\beta}(\lambda) = (\hat{\beta}_2(\lambda), \dots, \hat{\beta}_p(\lambda))$ with:

$$\hat{\beta}_j(\lambda d_j^2) = \arg \min_{\beta} \|\mathbf{X}_j - \sum_{k=1}^{j-1} \mathbf{X}_k \beta_{jk}\|^2 + \lambda d_j^2 \sum_{k=1}^{j-1} |\beta_{jk}|.$$

Each of these $p-1$ subproblems can have its path regularization traced by means of the homotopy/LARS-LASSO algorithm in Osborne et al. (2000); Efron et al. (2004). The $\hat{\beta}(\lambda)$ parameter estimate is

recovered by $(\hat{\beta}_2(\lambda d_2^2), \dots, \hat{\beta}_p(\lambda d_p^2))$. All is needed is a little care in merging the $p-1$ paths together as the scaling of the regularization parameter changes from one subproblem to the next.

An alternative way to understand how the problem can be broken into these smaller pieces stems from the representation of program in (22) as a linear regression. A little manipulation can be used to show that (22) can be represented as a linear regression of $\tilde{\mathbf{Y}}$ against $\tilde{\mathbf{Z}}$ as below:

$$\tilde{\mathbf{Y}} = \begin{bmatrix} \frac{\mathbf{X}_2}{d_2^2} \\ \frac{\mathbf{X}_3}{d_3^2} \\ \frac{\mathbf{X}_4}{d_4^2} \\ \vdots \\ \frac{\mathbf{X}_{p-1}}{d_{p-1}^2} \\ \frac{\mathbf{X}_p}{d_p^2} \end{bmatrix} \quad \text{and,}$$

$$\tilde{\mathbf{Z}} = \begin{bmatrix} \frac{\tilde{\mathbf{X}}_1}{d_2^2} & 0 & 0 & 0 & 0 & 0 & \cdots & 0 & \cdots & 0 & 0 & \cdots & 0 \\ 0 & \frac{\tilde{\mathbf{X}}_1}{d_3^2} & \frac{\tilde{\mathbf{X}}_2}{d_3^2} & 0 & 0 & 0 & \cdots & 0 & \cdots & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & \frac{\tilde{\mathbf{X}}_1}{d_4^2} & \frac{\tilde{\mathbf{X}}_2}{d_4^2} & \frac{\tilde{\mathbf{X}}_3}{d_4^2} & \cdots & 0 & \cdots & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & 0 & 0 & \cdots & \frac{\tilde{\mathbf{X}}_1}{d_{p-1}^2} & \cdots & \frac{\tilde{\mathbf{X}}_{p-2}}{d_{p-1}^2} & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 & \cdots & 0 & \frac{\tilde{\mathbf{X}}_1}{d_p^2} & \cdots & \frac{\tilde{\mathbf{X}}_{p-1}}{d_p^2} \end{bmatrix}$$

The separability of the program (22) into the subprograms (22) follows from the block diagonal structure of the matrix $\tilde{\mathbf{Z}}'\tilde{\mathbf{Z}}$. The application of the homotopy/LARS-LASSO algorithm to each of the problems and the subsequent merging of the resulting paths into a single path can be seen as a path version of the coordinate wise algorithms described in Friedman et al. (2007).

C Appendix: Sampling sparse precision matrices

In Section 4, we use three randomly selected precision matrices in the simulation studies presented therein. These random precision matrices are sampled as follows.

- A random sample containing 20 observations of X is sampled from a zero-mean Gaussian

distribution with precision matrix containing 2 along its main diagonal and 1 on its off-diagonal entries;

- A random precision matrix is formed by computing $G = (X^T X)^{-1}$;
- The number of off-diagonal terms N is sampled from the a geometric distribution with parameter $\lambda = 0.05$ conditioned to be between 1 and $\frac{15 \times 14}{2} = 105$;
- A new matrix H is formed by setting all off-diagonal of the matrix G are set to zero, except for N randomly selected entries (all entries are equally likely to be picked);
- Since H may not be positive definite, the precision matrix is formed by adding $H + \mathbf{I}_{15} \cdot \max(0, 0.02 - \varphi(H))$ where $\varphi(H)$ is the smallest eigenvalue of H .

References

- AKAIKE, H. 1973. Information theory and an extension of the maximum likelihood principle. In *2nd International Symposium on Information Theory*, B. N. Petrov and F. Csáki, Eds. Akadémia Kiadó, Budapest, 267–281.
- BANERJEE, O., D’ASPREMONT, A., AND ELGHAOUI, L. 2005. Sparse covariance selection via robust maximum likelihood estimation. Tech. rep., arXiv <http://arxiv.org/abs/cs.CE/0506023>.
- BANERJEE, O., GHAOUI, L. E., AND D’ASPREMONT, A. 2007. Model selection through sparse maximum likelihood estimation for multivariate gaussian or binary data. *Journal of Machine Learning Research* 101.
- BESAG, J. 1974. Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society, Series B* 36, 2, 192–236.
- BICKEL, P. AND LEVINA, E. 2004. Some theory for fisher’s linear discriminant function, ‘naive bayes’, and some alternatives when there are many more variables than observations. *Bernoulli* 10, 6, 989–1010.
- BICKEL, P. AND LEVINA, E. 2008. Regularized estimation of large covariance matrices. *The Annals of Statistics* 36, 1, 199–227.
- BILMES, J. A. 2000. Factored sparse inverse covariance matrices. In *IEEE International Conference on Acoustics, Speech and Signal Processing*.
- BOYD, S. AND VANDENBERGHE, L. 2004. *Convex Optimization*. Cambridge University Press, Cambridge, UK ; New York.
- BREIMAN, L. 1995. Better subset regression using the nonnegative garrote. *Technometrics* 37, 4, 373–384.
- CHEN, S. S., DONOHO, D. L., AND SAUNDERS, M. A. 2001. Atomic decomposition by basis pursuit. *SIAM Review* 43, 1, 129–159.
- DANIELS, M. AND KASS, R. E. 1999. Nonconjugate bayesian estimation of covariance matrices and its use in hierarchical models. *Journal of the American Statistical Association* 94, 448, 1254–1263.
- DANIELS, M. J. AND KASS, R. E. 2001. Shrinkage estimators for covariance matrices. *Biometrics* 57, 1173–1184.
- DEMPSTER, A. P. 1972. Covariance selection. *Biometrics* 28, 1 (March), 157–175.

- EFRON, B., HASTIE, T., JOHNSTONE, I., AND TIBSHIRANI, R. 2004. Least angle regression. *The Annals of Statistics* 35, 407–499.
- FAN, J., FAN, Y., AND LV, J. 2006. High dimensional covariance matrix estimation using a factor model. Tech. rep., Princeton University.
- FRIEDMAN, J., HASTIE, T., HÖFLING, H., AND TIBSHIRANI, R. 2007. Pathwise coordinate optimization. *The Annals of Applied Statistics* 1, 2, 302–332.
- FRIEDMAN, J., HASTIE, T., AND TIBSHIRANI, R. 2008. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics* 9, 3, 432–441.
- FURRER, R. AND BENGTTSSON, T. 2007. Estimation of high-dimensional prior and posterior covariance matrices in kalman filter variants. *Journal of Multivariate Analysis* 98, 227–255.
- HOCKING, R. R. 1996. *Methods and Applications of Linear Models: Regression and the Analysis of Variance*. John Wiley and Sons.
- HUANG, J. Z., LIU, N., POURAHMADI, M., AND LIU, L. 2006. Covariance matrix selection and estimation via penalized normal likelihood. *Biometrika* 93, 1, 85–98.
- HURVICH, C. M., SIMONOFF, J. S., AND TSAI, C.-L. 1998. Smoothing parameter selection in nonparametric regression using an improved akaike information criterion. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* 60, 2, 271–293.
- JOHNSTONE, I. M. AND LU, A. Y. 2004. Sparse principal component analysis. Tech. rep., Stanford University Department of Statistics.
- JOLLIFFE, I. T. 2002. *Principal Component Analysis*, 2nd ed. Springer.
- KNIGHT, K. AND FU, W. J. 2000. Asymptotics for lasso type estimators. *The Annals of Statistics* 28, 5, 1356–1378.
- LEDOIT, O. AND WOLF, M. 2004. A well conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis* 88, 2, 365–411.
- MARCHENKO, V. A. AND PASTUR, L. A. 1967. The distribution of eigenvalues in certain sets of random matrices. *Mat. Sb.* 72, 507–536.
- MEINSHAUSEN, N. AND BÜHLMANN, P. 2006. High-dimensional graphs and variable selection with the lasso. *The Annals of Statistics* 34, 3, 1436–1462.
- OSBORNE, M., PRESNELL, B., AND TURLACH, B. A. 2000. On the LASSO and its dual. *Journal of Computational and Graphical Statistics* 9, 2 (June), 319–337.
- PAUL, D., BAIR, E., HASTIE, T., AND TIBSHIRANI, R. 2008. Pre-conditioning for feature selection and regression in high-dimensional problems. *Annals of Statistics* 36, 4, 1595–1618.
- ROTHMAN, A. J., BICKEL, P., LEVINA, E., AND ZHU, J. 2007. Sparse permutation invariant covariance estimation. *Electronic Journal of Statistics* 2, 494–515.
- SCHWARTZ, G. 1978. Estimating the dimension of a model. *The Annals of Statistics* 6, 461–464.
- STEIN, C. 1975. Estimating a covariance matrix. In *39th Annual Meeting IMS*.
- STEIN, C. 1986. Lectures on the theory of estimation of many parameters (reprint). *Journal of Mathematical Sciences* 34, 1.
- TIBSHIRANI, R. 1996. Regression shrinkage and selection via the LASSO. *Journal of the Royal Statistical Society, Series B* 58, 1, 267–288.
- WU, W. B. AND POURAHMADI, M. 2003. Nonparametric estimation of large covariance matrices of longitudinal data. *Biometrika* 90, 831–844.
- YUAN, M. AND LIN, Y. 2007. Model selection and estimation in the gaussian graphical model. *Biometrika* 94, 1, 19–35.

- ZOU, H. AND HASTIE, T. 2005. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society, Series B* 67, 2, 301–320.
- ZOU, H., HASTIE, T., AND TIBSHIRANI, R. 2006. Sparse principal component analysis. *Journal of Computational and Graphical Statistics* 15, 2, 265–286.
- ZOU, H., HASTIE, T., AND TIBSHIRANI, R. 2007. On the “degrees of freedom” of the LASSO. *The Annals of Statistics* 35, 5, 2173–2192.