

Penalized Likelihood Methods for Estimation of Sparse High Dimensional Directed Acyclic Graphs

Ali Shojaie and George Michailidis
 Department of Statistics, University of Michigan, Ann Arbor
 shojaie@umich.edu, gmichail@umich.edu

Abstract

Directed acyclic graphs (DAGs) are commonly used to represent causal relationships among random variables in graphical models. Applications of these models arise in the study of physical, as well as biological systems, where directed edges between nodes represent the influence of components of the system on each other. The general problem of estimating DAGs from observed data is computationally NP-hard. Moreover two directed graphs may be observationally equivalent. When the nodes exhibit a natural ordering, the problem of estimating directed graphs reduces to the problem of estimating the structure of the network. In this paper, we propose a penalized likelihood approach that directly estimates the adjacency matrix of DAGs. Both lasso and adaptive lasso penalties are considered and an efficient algorithm is proposed for estimation of high dimensional DAGs. We study variable selection consistency of the two penalties when the number of variables grows to infinity with the sample size. We show that although lasso can only consistently estimate the true network under stringent assumptions, adaptive lasso achieves this task under mild regularity conditions. The performance of the proposed methods are compared to alternative methods in simulated, as well as real, data examples.

1 Introduction

Graphical models provide efficient tools for the study of statistical models through a compact representation of the joint probability distribution of the underlying random variables. The nodes of the graph represent the random variables, while the edges capture the relationships among them. Both directed and undirected edges are used to represent interactions among random variables. However, there is a conceptual difference between these two types of edges: while undirected edges are used to represent similarity or correlation, directed edges are usually interpreted as causal relationships. The study of directed edges is therefore directly related to the theory of causality, and of main interest in many applications. A special class of directed graphical models (also known as Bayesian Networks) are based on directed acyclic graphs (DAGs), where all the edges of the graph are directed and there are no directed cycles present in the graph. DAGs are used in graphical models and belief networks and have been the focus of research in the computer science literature (see [Pearl \(2000\)](#)). Important applications involving DAGs also arise in the study of biological systems, as many cellular mechanisms are known to include causal relationships. Cell signalling pathways and gene regulatory networks are two examples, where causal relationships play an important role ([Markowitz and Spang, 2007](#)).

The problem of estimating DAGs is an NP-hard problem, and estimation of direction of edges may not be possible due to observational equivalence (see section 2). Most of the earlier methods for estimating DAGs correspond to greedy search algorithms that search through the space of possible DAGs. A number of methods are available for estimating the structure of DAGs for small to moderate number of nodes. The max-min hill climbing algorithm ([Tsamardinos et al., 2006](#)), and the PC-Algorithm ([Spirtes et al., 2000](#)) are two such examples. However, the space of possible DAGs grows super-exponentially with the number of variables (nodes), and estimation of DAGs using these methods, especially in a small n , large p setting, becomes impractical. Bayesian methods of estimating DAGs (e.g [Heckerman et al., 1995](#)) are also computationally very intensive and therefore not particularly appropriate for large graphs. [Kalisch and Bühlmann \(2007\)](#) recently proposed an implementation of the PC-Algorithm with polynomial complexity that can be used for

estimation of high dimensional sparse DAGs. However, when the variables inherit a *natural ordering*, estimation of a DAG is reduced to estimating its structure or skeleton. Applications with natural ordering of variables include estimation of causal relationships from temporal observations, or settings where additional experimental data can determine the ordering of variables, and estimation of transcriptional regulatory networks from gene expression data. Examples of these applications are presented in section 6.

The structure of the graph can be determined from conditional independence relations among random variables. For undirected graphs, this is equivalent to learning the structure of the conditional independence graph (CIG), which in the case of Gaussian random variables, is determined by zeros in the inverse covariance matrix (also known as precision or concentration matrix). Different penalization methods, have been recently proposed to obtain sparse estimates of the concentration matrix. [Meinshausen and Bühlmann \(2006\)](#) considered an approximation to the problem of sparse inverse covariance estimation using the lasso penalty. They showed under a set of assumptions, that their proposed method correctly determines the neighborhood of each node. [Banerjee et al. \(2008\)](#) and [Friedman et al. \(2008b\)](#) explored different aspects of the problem of estimating the concentration matrix using the lasso penalty, while [Yuan and Lin \(2007\)](#) and [Fan et al. \(2007\)](#) considered other choices for the penalty. [Rothman et al. \(2008\)](#) proved consistency in Frobenius norm, as well as in matrix norm, of the ℓ_1 -regularized estimate of the concentration matrix when $p \gg n$, while [Lam and Fan \(2008\)](#) extended their result and considered estimation of matrices related to the precision matrix, including the Cholesky factor of the inverse covariance matrix, using general penalties. Penalization of the Cholesky factor of the inverse covariance matrix has been also considered by [Huang et al. \(2006\)](#), where they used the lasso penalty in order to obtain a sparse estimate of the inverse covariance matrix. This method is based on the regression interpretation of the Cholesky factorization model and therefore requires the variables to be ordered *a priori*.

In this paper, we consider the problem of estimating the skeleton of DAGs, where the variables exhibit a natural ordering. We use graph theoretic properties of DAGs and reformulate the likelihood as a function of the adjacency matrix of the graph. We then exploit the ordering of variables to propose an efficient algorithm for estimation of structure of DAGs, which offers considerable improvement in terms of computational complexity. Both lasso and adaptive lasso penalties are considered and variable selection consistency of estimators is established in the $p \gg n$ setting. In particular, we show that although lasso is only variable selection consistent under stringent conditions, adaptive lasso can consistently estimate the true DAG under the usual regularity assumptions. We also present a data dependent choice of the tuning parameter that controls the probability of errors. Theoretical as well as empirical evidence shows that when the underlying causal mechanism in the network is linear, the proposed method can also be applied to non-Gaussian observations. Finally, additional simulations indicate that although the proposed method is derived based on the ordering of variables, the method is not sensitive to random permutations of the order of variables in high dimensional sparse settings.

2 Representation of Directed Acyclic Graphs

Consider a graph $\mathcal{G} = (V, E)$, where V corresponds to the set of nodes with p elements and $E \subset V \times V$ to the edge set. The nodes of the graph represent random variables X_1, \dots, X_p and the edges capture associations amongst them. An edge is called directed if $(i, j) \in E \Rightarrow (j, i) \notin E$ and undirected if $(i, j) \in E \Rightarrow (j, i) \in E$. The main focus of this paper is a special class of graphs where E consists of only directed edges, and does not include directed cycles. We denote by pa_i the set of parents of node i and for $j \in pa_i$, we denote $j \rightarrow i$. The *skeleton* of a DAG is the undirected graph that is obtained by replacing directed edges in E with undirected ones. Finally, throughout this paper, we represent E using the adjacency matrix A of the graph; i.e. a $p \times p$ matrix whose (j, i) th entry indicates whether there is an edge (and possibly its weight) between nodes j and i .

The estimation of DAGs is a challenging problem due to the so-called *observational equivalence* of DAGs with respect to the same probability distribution. More specifically, regardless of the sample size, it may not be possible to infer the direction of causation among random variables from observational data. As an illustration of the observational equivalence, consider the simple DAG in the right panel of Figure 1. Reversing the direction of all edges of the graph results in a new DAG, which is the same as the original graph, except for changes in the node labels and is therefore *polymorphic* to the original one. It is therefore natural to estimate the *equivalence class* of DAGs corresponding to the same probability distribution \mathcal{P} starting with the skeleton of the network. The second challenge in estimating



Figure 1: Left: A simple DAG, Right: Illustration of observational equivalence in DAGs

DAGs is that conditional independence among random variables may not reveal the skeleton. The notion of conditional independence in DAGs is either represented using the concept of *d-separation* (Pearl, 2000) or the *moral graph* of a DAG (Lauritzen, 1996). The moral graph is obtained by removing the directions of the graph and “marrying” the parents of each node. Therefore estimation of the conditional independence structure reveals the structure of the moral graph of the DAG, which includes additional edges between parents of each node. This can also be illustrated using the simple graph in the left panel of Figure 1. Suppose $X_i, i = 1, \dots, 4$ are normally distributed with covariance matrix Σ . The only zero elements of the inverse covariance matrix are $\Sigma_{14}^{-1} = \Sigma_{41}^{-1}$, as X_2 and X_3 are connected in the moral graph of \mathcal{G} .

2.1 The Latent Variable Model

The causal effect of random variables in a directed acyclic graph can be explained using *structural equation models*, where each variable is modeled as a (nonlinear) function of its parents. The general form of these models is given by (Pearl, 2000):

$$X_i = f_i(pa_i, Z_i), \quad i = 1, \dots, p \quad (2.1)$$

The random variables Z_i are the latent variables representing the unexplained variation in each node. To model the association among nodes of a DAG, we consider a simplification of (2.1) with f_i being linear. More specifically, let ρ_{ij} represent the *effect* of node j on i for $j \in pa_i$, then

$$X_i = \sum_{j \in pa_i} \rho_{ij} X_j + Z_i, \quad i = 1, \dots, p \quad (2.2)$$

In the special case where the random variables are Gaussian, equations (2.1) and (2.2) are equivalent in the sense that ρ_{ij} are coefficients of the linear regression model of X_i 's on $X_j, j \in pa_i$. It is known in the normal case that $\rho_{ij} = 0, j \notin pa_i$.

Consider the simple DAG in the left panel of Figure 1; denoting the *influence matrix* of the graph by Λ , (2.2) can be written in compact form as $X = \Lambda Z$, where for the simple example above, we have

$$\Lambda = \begin{pmatrix} 1 & 0 & 0 \\ \rho_{12} & 1 & 0 \\ \rho_{12}\rho_{23} & \rho_{23} & 1 \end{pmatrix}$$

Let the latent variables Z_i be independent with mean μ_i and variance σ_i^2 . Then $E(X) = \Lambda\mu$ and $\Sigma = \text{var}(X) = \Lambda D \Lambda^T$, where $D = \text{diag}(\sigma_i^2)$ and Λ^T denotes the transpose of the matrix Λ .

The following result from Shojaie and Michailidis (2009a) establishes relationships between the influence matrix Λ , and the adjacency matrix of the graph, A . The second part of the lemma establishes a compact relationship between Λ and A in the case of DAGs, which is explored in section 3 to directly formulate the problem of estimating the skeleton of a DAG.

Lemma 2.1. For any graph $\mathcal{G} = (V, A)$,

- (i) $\Lambda = A^0 + A^1 + A^2 + \dots = \sum_{r=0}^{\infty} A^r$, where $A^0 \equiv I$.
- (ii) If \mathcal{G} is a DAG, Λ has full rank and $\Lambda = (I - A)^{-1}$.

Remark 2.2. Part (ii) of Lemma 2.1 and the fact that $\Sigma = \Lambda D \Lambda^\top$ imply that for any DAG, if $D_{ii} > 0$ for all i , then Σ is full rank. More specifically, let $\phi_j(M)$ denote the j th eigenvalue of matrix M . Then, $\phi_{\min}(\Sigma) > 0$ (or $\phi_{\max}(\Sigma^{-1}) < \infty$). Similarly, since $\Sigma^{-1} = \Lambda^{-\top} D^{-1} \Lambda^{-1}$, full rankness of Λ implies that $\phi_{\min}(\Sigma^{-1}) > 0$ (or equivalently $\phi_{\max}(\Sigma) < \infty$). This result also applies to all subnetworks of a DAG.

The properties of the proposed latent variable model established in Lemma 2.1 are independent of the choice of probability distribution \mathcal{P} . In fact, since the latent variables Z_i in (2.2) are assumed independent, given the entries of the adjacency matrix, the distribution of each random variable X_i in the graph only depends on the values of pa_i . Therefore, regardless of the choice of the probability distribution, the joint distribution of the random variables is compatible with \mathcal{G} (see for example Pearl (2000) p. 16). In section 5, we illustrate this result using data generated according to non-Gaussian distributions.

3 Penalized Estimation of DAGs

3.1 Problem Formulation

Consider the latent variable model of section 2.1 and denote by \mathcal{X} the $n \times p$ data matrix. We assume, without loss of generality, that the X_i 's are centered and scaled, so that $\mu_i = 0$ and $\sigma_i^2 = 1, i = 1, \dots, p$. Note that the results in section 2.1 were established independent of the choice of the probability distribution. As mentioned before, under the normality assumption, the latent variable model is equivalent to the general structural equation model. Although we focus on Gaussian random variables in the remainder of this paper, the estimation procedure proposed in this section can be applied to a variety of other distributions, if one is willing to assume the linear structure in (2.2).

Denote by $\Omega \equiv \Sigma^{-1}$ the precision matrix of a p -vector of Gaussian random variables and consider a general penalty function by $J(\Omega)$. The penalized likelihood function is then given by

$$\hat{\Omega} = \underset{\Omega > 0}{\operatorname{argmin}} \{-\log \det(\Omega) + \operatorname{tr}(\Omega S) + \lambda J(\Omega)\} \quad (3.1)$$

where $S = n^{-1} \mathcal{X}^\top \mathcal{X}$ denotes the empirical covariance matrix and λ is the tuning parameter controlling the size of the penalty. Applications in biological and social networks often involve sparse networks. It is therefore desirable to find a sparse solution for (3.1). This becomes more important in the small n , large p setting, where the unpenalized solution includes many additional edges. The lasso penalty of Tibshirani (1996) and the adaptive lasso penalty proposed by Zou (2006) are singular at zero and therefore result in sparse solutions. We consider these two penalties in order to find a sparse estimate of the adjacency matrix. Other choices of the penalty function are briefly discussed in the conclusions section.

Using the latent variable model of section 2.1, and the relationship between the covariance matrix and the adjacency matrix of DAGs established in Lemma 2.1, the problem of estimating the adjacency matrix of the graph can be directly formulated as an optimization problem based on A . As noted in Lemma 2.1, if the underlying graph is a DAG and the ordering of the variables is known, then A is a lower triangular matrix with zeros on the diagonal. Let $\mathcal{A} = \{A : A_{ij} = 0, j \geq i\}$. Then using the facts that $\det(A) = 1$ and $\sigma_i^2 = 1$, A can be estimated as the solution of the following optimization problem

$$\hat{A} = \underset{A \in \mathcal{A}}{\operatorname{argmin}} \{\operatorname{tr}[(I - A)^\top (I - A)S] + \lambda J(A)\} \quad (3.2)$$

In this paper, we consider the general weighted lasso problem, where

$$J(A) = \lambda \sum_{i,j=1:p, j < i} w_{ij} |A_{ij}| \quad (3.3)$$

Lasso and adaptive lasso problems are special cases of this general penalty. In the case of lasso, $w_{ij} = 1$. The original weights in adaptive lasso, proposed by Zou (2006) are obtained by setting $w_{ij} = |\tilde{A}_{ij}|^{-\gamma}$, for some initial estimate of the adjacency matrix \tilde{A} and some power γ . To facilitate the study of asymptotic properties of adaptive lasso estimates we consider the following modification of the original weights

$$w_{ij} = 1 \vee |\tilde{A}_{ij}|^{-\gamma} \quad (3.4)$$

where the original estimates \tilde{A} are obtained from the regular lasso estimates.

The objective function for both lasso and adaptive lasso problems is convex. However, since the ℓ_1 penalty is non-differentiable, these problems can be reformulated using matrices $A_+ = \max(A, 0)$ and $A_- = -\min(A, 0)$. To that end, let W be the $p \times p$ matrix of weights for adaptive lasso, or the matrix of ones for the lasso estimation problem. Problem (3.2) can then be formulated as:

$$\min_{A_+, A_- \succeq 0} \text{tr} \left\{ S(I - A_+ + A_-)^\top (I - A_+ + A_-) + \lambda(A_+ + A_-)W + \Delta(A_+ + A_-)\mathbf{1}_{l^+} \right\} \quad (3.5)$$

where $\succeq 0$ is interpreted componentwise, Δ is a large positive number and $\mathbf{1}_{l^+}$ is the indicator matrix for lower triangular elements of a $p \times p$ matrix, including the diagonal elements. The last term of the objective function ($\text{tr}[\Delta(A_+ + A_-)\mathbf{1}_{l^+}]$) prevents the upper triangular elements of the matrices A_+ and A_- to be nonzero.

Problem (3.5) is a quadratic optimization problem with non-negativity constraints and can be solved using standard interior point algorithms. However, such algorithms do not scale well with dimension and are only applicable if p ranges in the hundreds. In section 3.2, we present an alternative formulation of the problem, which leads to considerably more efficient algorithms.

3.2 Optimization Algorithm

Consider again the problem of estimating the adjacency matrix of DAGs with either lasso or adaptive lasso penalties. Denoting the i th row of matrix A as A_i we can write (3.2) as:

$$\hat{A} = \underset{A \in \mathcal{A}}{\text{argmin}} \left\{ \sum_{i=1}^p A_i^\top S A_i - 2A_i S_i^\top + \lambda W_i^\top |A_i| \right\} \quad (3.6)$$

It can be seen that the objective function in (3.6) is separable and therefore it suffices to solve the optimization problem over each row of matrix A . Denote by \underline{l} the set of indices up to l , i.e. $\underline{l} = j : 1 \leq j \leq l$.

Then, taking advantage of the lower triangular structure of A , solving (3.6) is equivalent to solving the following $p - 1$ optimization problems ($A_{11} = 0$)

$$\hat{A}_{i, \underline{i-1}} = \underset{\theta \in \mathbb{R}^{i-1}}{\text{argmin}} \left\{ \theta^\top S_{\underline{i-1}, \underline{i-1}} \theta - 2S_{i, \underline{i-1}} \theta + \lambda \sum_{j=1}^{i-1} |\theta_j| w_{ij} \right\}, \quad i = 2, \dots, p \quad (3.7)$$

Using the facts that $S_{\underline{i-1}, \underline{i-1}} = n^{-1}(\mathcal{X}_{\underline{n}, \underline{i-1}})^\top \mathcal{X}_{\underline{n}, \underline{i-1}}$ and $S_{i, \underline{i-1}} = n^{-1}(\mathcal{X}_{\underline{n}, i})^\top \mathcal{X}_{\underline{n}, \underline{i-1}}$, the problem in (3.7) can be reformulated as following ℓ_1 -regularized least squares problems

$$\hat{A}_{i, \underline{i-1}} = \underset{\theta \in \mathbb{R}^{i-1}}{\text{argmin}} \left\{ n^{-1} \|\mathcal{X}_{\underline{n}, \underline{i-1}} \theta - \mathcal{X}_{\underline{n}, i}\|_2^2 + \lambda_i \sum_{j=1}^{i-1} |\theta_j| w_{ij} \right\}, \quad i = 2, \dots, p \quad (3.8)$$

The formulation in (3.8) indicates that the i th row of matrix A includes the coefficient of projecting X_i on X_j , $j = 1, \dots, i - 1$, which is in agreement with the discussion in section 2.1. It also reveals a connection between covariance selection methods and the neighborhood selection approach of Meinshausen and Bühlmann (2006); namely, when the underlying graph is a DAG, the approximate solution of the neighborhood selection problem is exact, if the regression model is fitted on the set of parents of each node instead of all other nodes in the graph.

Using (3.8), the problem of estimating DAGs can be solved very efficiently. In fact, it suffices to solve $p - 1$ lasso problems for estimation of least squares coefficients, with dimensions ranging from 1 to $p - 1$. To solve these problems, we use the efficient pathwise coordinate optimization algorithm of Friedman et al. (2008a), implemented in the R-package `glmnet`. The proposed algorithm is summarized in Algorithm 1.

3.3 Analysis of Computational Complexity

In this section, we provide a comparison of the computational complexity of the algorithm proposed in section 3.2 and the PC-Algorithm. As mentioned in the introduction, the space of all possible DAGs is super-exponential in the

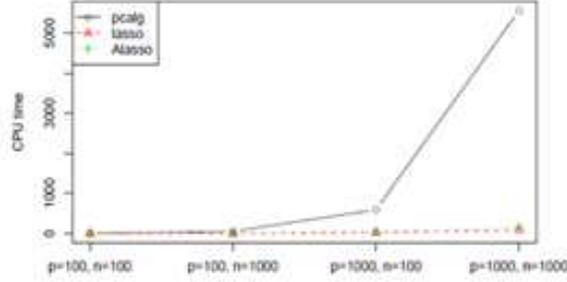


Figure 2: CPU time for analysis of simulated DAGs with different number of nodes and sample size. The results are average times of 10 repetitions on an Intel 2 · 0GH processor with 2 · 0GB of RAM.

number of nodes and hence it is not surprising that the PC-Algorithm, without any restriction on the space of DAGs has exponential complexity. [Kalisch and Bühlmann \(2007\)](#) propose an efficient implementation of the PC-Algorithm for sparse DAGs; its complexity where the maximal neighborhood size q is small, is bounded with high probability by $O(p^q)$. Although this is a considerable improvement over other methods of estimating DAGs, in many applications it can become fairly expensive. For example, gene regulatory networks and signaling pathways exhibit a “hub” structure, which leads to large values for q . The graphical lasso algorithm, proposed by [Friedman et al. \(2008b\)](#), uses an iterative algorithm for estimation of the inverse covariance matrix and has computational complexity $O(p^3)$.

The reformulation of the DAG estimation problem in (3.8) requires solving $p-1$ lasso regression problems. The cost of solving a lasso problem comprised of k covariates and n observations using the pathwise coordinate optimization (shooting) algorithm of [Friedman et al. \(2008a\)](#) is $O(nk)$; hence, the total cost of estimating the adjacency matrix of the graph is $O(np^2)$, which is the same to the cost of calculating the (full) empirical covariance matrix S_n . Moreover, the formulation in (3.8) includes a set of non-overlapping sub-problems. Therefore, for problems with very large number of nodes and/or observations, the performance of the algorithm can be further improved by parallelizing the estimation of these sub-problems. The adaptive lasso version of the problem is similarly solved using the modification of the regular lasso problem proposed in [Zou \(2006\)](#), which results in the same computational cost as the regular lasso problem.

Figure 2, compares the CPU time required for estimation of DAGs using both the PC-Algorithm, as well as our proposed algorithm for a range of values of p and n . To control the complexity of the PC algorithm, the average neighborhood size is set to 5 and the significance level for the PC-Algorithm, as well as the tuning parameter for lasso and adaptive lasso penalties are set according to the optimal values discussed in Section 5. The time reported for the PC-Algorithm is the CPU time required for estimation of the skeleton of the graph. The plot demonstrates the higher order of complexity of the PC-Algorithm, as well as the dependency of the algorithm on the sample size.

Algorithm 1 Penalized Likelihood Estimation of DAGs

1. Given the ordering \mathcal{O} , order the columns of observation matrix \mathcal{X} in increasing order
 2. For $i = 2, 3, \dots, p$
 - 2.1. Let $y = \mathcal{X}_{\underline{n},i}$, $X = \mathcal{X}_{\underline{n},i-1}$ and $w = W_{i,i-1}$
 - 2.2. Given the weight matrix W , solve $\hat{A}_{i,i-1} = \operatorname{argmin} \left\{ n^{-1} \|X\theta - y\|_2^2 + \lambda_i \sum_{j=1}^{i-1} |\theta_j| w_j \right\}$
-

4 Asymptotic Properties

4.1 Preliminaries

We next establish theoretical properties for both lasso, as well as adaptive lasso estimators of the adjacency matrix of DAGs. Asymptotic properties of lasso-type estimates with fixed design matrices have been studied by a number of researchers (see e.g. [Knight and Fu, 2000](#); [Zou, 2006](#); [Huang et al., 2008](#)). Lasso estimates with random design matrices have been also considered by [Meinshausen and Bühlmann \(2006\)](#). On the other hand, [Rothman et al. \(2008\)](#) and [Lam and Fan \(2008\)](#) among others have studied asymptotic properties of estimates of covariance and precision matrices.

As discussed in section [3.1](#), the problem of estimating the adjacency matrix of a DAG is equivalent to solving $p - 1$ non-overlapping penalized least square problems described in [\(3.7\)](#). In order to study the asymptotic properties of the proposed estimators, we focus on the asymptotic consistency of network estimation, i.e. the probability of correctly estimating the network structure, in terms of type I and type II errors. We allow the total number of nodes in the graph to grow as an arbitrary polynomial function of the sample size, while assuming that the true underlying network is sparse. The assumptions required for establishing asymptotic properties of lasso and adaptive lasso estimates of DAGs are presented in section [4.2](#). We then study variable selection consistency of both lasso and adaptive lasso estimates in section [4.3](#). The choice of penalty parameter is studied in section [4.4](#). Technical proofs are given in the Appendix.

4.2 Assumptions

Let $X = (X_1, \dots, X_p)$ be a collection of p zero-mean Gaussian random variables with covariance matrix Σ . Denote by \mathcal{X} the $n \times p$ matrix of observations, and by S the empirical covariance matrix. To simplify the notation, denote by $\theta^i = A_{i, \underline{i-1}}$ the entries of the i th row of A to the left of the diagonal. Further, let $\theta^{i, \mathcal{I}}$ be the estimate for the i th row, with values outside the set of indices \mathcal{I} set to zero; i.e., $\theta^{i, \mathcal{I}} \equiv A_{i, \underline{i-1}}$ and $A_{i, j} = 0$, $j \notin \mathcal{I}$. Consider the following assumptions:

- (A-0) For some $a > 0$, $p = p(n) = O(n^a)$ as $n \rightarrow \infty$, and there exists a $0 \leq b < 1$ such that $\max_{i \in V} \text{card}(pa_i) = O(n^b)$ as $n \rightarrow \infty$.
- (A-1) There exists $\nu > 0$ such that for all $n \in \mathbb{N}$ and all $i \in V$, $\text{var}(X_i | X_{\underline{i-1}}) \geq \nu$.
- (A-2) There exists $\delta > 0$ and some $\xi > b$ (with b defined above) such that for all $i \in V$ and for every $j \in pa_i$, $|\pi_{ij}| \geq \delta n^{-(1-\xi)/2}$, where π_{ij} is the partial correlation between X_i and X_j after removing the effect of the remaining variables.
- (A-3) There exists $\Psi < \infty$ such that for all $n \in \mathbb{N}$ and every $i \in V$ and for every $j \in pa_i$, $\|\theta^{j, pa_i}\| \leq \Psi$.
- (A-4) There exists $\kappa < 1$ such that for all $i \in V$ and for every $j \notin pa_i$, $\left| \sum_{k \in pa_i} \text{sign}(\theta_k^{i, pa_i}) \theta_k^{j, pa_i} \right| < \kappa$.

Assumption (A-3) limits the magnitude of the effects that each node in the network receives from its parents. This is less restrictive than the neighborhood selection criterion, where the effects over all neighboring nodes are assumed to be bounded. In fact, empirical data indicate that the average number of upstream-regulators per gene in regulatory networks is less than 2 ([Leclerc, 2008](#)). Thus, the size of parents of each node is small, but each hub node can affect many downstream nodes.

Assumption (A-4) is referred to as *neighborhood stability* and is equivalent to the *irrepresentability* assumption proposed by [Huang et al. \(2008\)](#). It has been shown that the lasso estimates are not in general variable selection consistent if this assumption is violated. [Huang et al. \(2008\)](#) considered adaptive lasso estimates with general initial weights and showed their variable selection consistency under a weaker form of irrepresentability assumption, referred to as *adaptive irrepresentability*. We will show that when the initial weights for adaptive lasso are derived from the regular lasso estimates (as in [\(3.4\)](#)), the assumption of neighborhood stability, as well as the less stringent assumption (A-3) are not required for establishing variable selection consistency of adaptive lasso. This relaxation in assumptions required for variable selection consistency, is a result of the consistency of regular lasso estimates, as well as the special structure of DAGs. However, the results of this section can be extended to adaptive lasso estimates of the precision matrix, as well as regression models with fixed and random design matrices, under additional mild assumptions.

4.3 Asymptotic Consistency of DAG Estimation

Our first result studies the variable selection consistency of the lasso penalty.

Theorem 4.1 (Variable Selection Consistency of Lasso). *Suppose that (A-1)-(A-4) hold and $\lambda \asymp dn^{-(1-\zeta)/2}$ for some $b < \zeta < \xi$ and $d > 0$. Then there exist constants $c_{(i)}, \dots, c_{(iv)} > 0$ for the lasso estimation problem, such that for all $i \in V$, as $n \rightarrow \infty$*

1. (i)] *Estimation of the direction of influence:*

$$\text{pr} \left\{ \text{sign}(\hat{\theta}_j^{i, pa_i}) = \text{sign}(\theta_j^{i, pa_i}) \text{ for all } j \in pa_i \right\} = 1 - O \left\{ \exp(-c_{(i)} n^\zeta) \right\}$$

(ii) *Control of type I error:* $\text{pr}(\hat{pa}_i \subseteq pa_i) = 1 - O \left\{ \exp(-c_{(ii)} n^\zeta) \right\}$.

(iii) *Control of type II error:* $\text{pr}(pa_i \subseteq \hat{pa}_i) = 1 - O \left\{ \exp(-c_{(iii)} n^\zeta) \right\}$.

(iv) *Let \hat{E} be the lasso estimate for the set of edges in the network. Then*

$$\text{pr}(\hat{E} = E) = 1 - O \left\{ \exp(-c_{(iv)} n^\zeta) \right\}.$$

Proof. The proof of this theorem follows from arguments similar to those presented in [Meinshausen and Bühlmann \(2006\)](#), with minor modifications and replacing conditional independence in undirected graphs with d-separation in DAGs. \square

The next result establishes similar properties for adaptive lasso estimates, without the assumptions of neighborhood stability. The proof of [Theorem 4.3](#) makes use of consistency of a class of sparse estimates of the Cholesky factors of covariance matrices, established in [Theorem 9 of Lam and Fan \(2008\)](#). For completeness, we restate a simplified version of the theorem for our lasso problem, for which $\sigma_i = 1, i = 1, \dots, p$ and the eigenvalues of the covariance matrix are bounded (see [Remark 2.2](#)). Throughout this section, we denote by s the total number of nonzero element of the true adjacency matrix, A of the DAG.

Theorem 4.2 ([Lam and Fan \(2008\)](#)). *If $n^{-1}(s+1) \log p = o(1)$ and $\lambda = O \left\{ (\log p/n)^{1/2} \right\}$, then $\|\hat{A} - A\|_F = O_p \left\{ (n^{-1} s \log p)^{1/2} \right\}$.*

It can be seen from [Theorem 4.2](#) that lasso estimates are consistent as long as $n^{-1}(s+1) \log p = o(1)$. To take advantage of this result, we replace (A-0) with the following assumption

1. (A-0')] For some $a > 0$, $p = p(n) = O(n^a)$ as $n \rightarrow \infty$. Also, $\max_{i \in V} \text{card}(pa_i) = O(n^b)$ as $n \rightarrow \infty$, where $sn^{2b-1} \log n = o(1)$ as $n \rightarrow \infty$.

Assumption (A-0') further restricts the number of parents of each node and also enforces a restriction on the total number of nonzero elements of the adjacency matrix. Condition $sn^{2b-1} \log n = o(1)$, implies that $b < 1/2$. Therefore, although the consistency of adaptive lasso in [Theorem 4.3](#) is established without making any further assumptions on the structure of the network (compared to [Theorem 4.1](#)), it is achieved at the price of requiring higher degree of sparsity in the network. We now state the main result regarding variable selection consistency of adaptive lasso. Note that the theorem only requires assumptions (A-0'), (A-1) and (A-2), and assumptions (A-3) and (A-4) are no longer required.

Theorem 4.3 (Variable Selection Consistency of Adaptive Lasso). *Consider the adaptive lasso estimation problem, where the initial weights are calculated using regular lasso estimates of the adjacency matrix of the graph in (3.7). Suppose (A-0') and (A-1)-(A-2) hold and $\lambda \asymp dn^{-(1-\zeta)/2}$ for some $b < \zeta < \xi$ and $d > 0$. Also suppose that the initial lasso estimates are found using a penalty parameter λ^0 that satisfies $\lambda^0 = O \left\{ (\log p/n)^{1/2} \right\}$. Then there exist constants $c_{(i)}, \dots, c_{(iv)} > 0$ such that for all $i \in V$, as $n \rightarrow \infty$ (i)-(iv) in [Theorem 4.1](#) hold.*

Proof. A proof is given in the Appendix. \square

4.4 Choice of the Tuning Parameter

Both lasso, as well as adaptive lasso estimates of the adjacency matrix, depend on the choice of the tuning parameter λ . Different methods have been proposed for selecting the value of tuning parameter, including cross validation (Rothman et al., 2008; Fan et al., 2007) and Bayesian Information Criteria (BIC) (Yuan and Lin, 2007). In both of these methods, the value of λ is chosen to minimize the (penalized) likelihood function. However, choices of λ that result in the optimal classification error do not guarantee a small error for the network reconstruction. We propose next a choice of λ for DAGs. Consider the following choice of the tuning parameter for the general weighted lasso problem with weights w_{ij} . Let Z_q^* denote the $(1 - q)$ th quantile of standard normal distribution, and define

$$\lambda_i(\alpha) = 2n^{-1/2}Z_{\frac{\alpha}{2p(i-1)}}^* \quad (4.1)$$

The following theorem, states that such a choice controls the probability of falsely joining two distinct ancestral sets, defined next.

Definition 4.4. For every node $i \in V$, the ancestral set of node i , AN_i consists of all nodes j such that j is an ancestor of i or i is an ancestor of j or i and j have a common ancestor k .

Theorem 4.5 (Choice of the Tuning Parameter). *Under the assumptions of Theorems 4.1 and 4.3 above (for lasso and adaptive lasso, respectively), for all $n \in \mathbb{N}$ the solution of the general weighted lasso estimation problem with tuning parameter determined in (4.1) satisfies*

$$\text{pr}(\exists i \in V : \hat{AN}_i \not\subseteq AN_i) \leq \alpha$$

Proof. A proof is given in the Appendix. □

Note that as in Meinshausen and Bühlmann (2006), Theorem 4.5 is true for all values of p and n . However, the theorem does not provide any guarantee on false positive or false negative probabilities for individual edges in the graph. We also need to determine the optimal choice of penalty parameter λ^0 for the first phase of the adaptive lasso, where the weights are estimated using lasso. Since the goal of the first phase is to achieve prediction consistency, cross validation can be used to determine the optimal choice of λ^0 . On the other hand, it is easy to see that the error-based proposal in (4.1) satisfies the requirement of Theorem 4.2 and can therefore be used to define λ^0 . It is however recommended to use a higher value of significance level in estimating the initial weights, in order to prevent an over-sparse solution.

5 Performance Analysis

5.1 Preliminaries

In this section, we consider examples of estimating DAGs of varying number of edges from randomly generated data. To randomly generate data from DAGs, one needs to generate lower-triangular adjacency matrices with sparse nonzero elements, ρ_{ij} . We use the random DAG generator in the R-package `pcalg` (Kalisch and Bühlmann (2007)) which also controls the neighborhood size. The sparsity levels in DAGs with different sizes are set according to the theoretical bounds in section 4, as well as the recommendations of (Kalisch and Bühlmann, 2007), for neighborhood size. More specifically, in simulations throughout this section, we use a maximum neighborhood size of 5, while limiting the total number of true edges to be equal to the sample size n .

Different measures of structural difference can be used to evaluate the performance of estimators. The Structural Hamming Distance (SHD) between the structures of the estimated and true DAGs, represents the number of edges that are not in common between the two graphs and is equal to the sum of false positive and false negative edges in the estimated graph. The main drawback of this measure is its dependency on the number of nodes, as well as the sparsity of the network. The second measure of goodness of estimation considered here is the Matthews Correlation Coefficient (MCC). MCC is commonly used to assess the performance of binary classification methods and is defined as

$$\text{MCC} = \frac{(\text{TP} \times \text{TN}) - (\text{FP} \times \text{FN})}{(\text{TP} + \text{FP})(\text{TP} + \text{FN})(\text{TN} + \text{FP})(\text{TN} + \text{FN})^{1/2}} \quad (5.1)$$

where TP, TN, FP and FN correspond to true positive, true negative, false positive and false negative, respectively. The value of MCC ranges from -1 to 1 with larger values corresponding to better fits (-1 and 1 represent worst and best fits, respectively). Finally, in order to compare the performance of different estimation methods with theoretical bounds established in section 4.3, we also report the values of false positive rates.

The performance of both the PC-Algorithm as well as our proposed estimators based on the choice of tuning parameter in (4.1) vary with different values of significance level α . In the following experiments, we first investigate the appropriate choice of α for each estimator. We then compare the performance of the estimators with an optimal choice of lambda using both numerical measures of performance, as well as gray-scale images of the estimated DAGs with the true structure. Gray-scale images are obtained by calculating the proportion of times that a specific edge is present in the simulations (i.e. $\hat{A}_{ij} \neq 0$). To offset the effect of numerical instability, we consider an edge present, if $|\hat{A}_{ij}| > 10^{-4}$.

5.2 Estimation of DAGs from Normally Distributed Observations

We begin with an example that illustrates the differences between estimation of DAGs and the conditional independence networks (CIG). The first two images in Figure 3 represent a randomly generated DAG of size 50 along with the gray-scale image of the average precision matrix estimated based on 100 observations using the graphical lasso algorithm (implemented in the R package `glasso`). To control the probability of falsely connecting two components of the graph, the value of the tuning parameter for `glasso` is defined based on the error-based proposal of Banerjee et al. (2008) (Theorem 2). It can be seen that the CIG has many more edges (8% false positives compared to 1% for lasso and adaptive lasso), and does not reveal the true structure of the underlying DAG. It can be seen that although methods of estimating CIGs are computationally efficient, they should not be used in applications like estimation of gene regulatory networks, where the underlying graph is directed. In simulations throughout this section, the sample size is fixed at

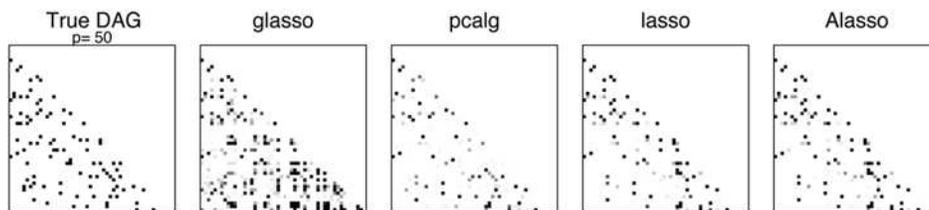


Figure 3: True DAG along with estimates from Gaussian observations using `glasso`, `pcalg`, `lasso` and `Alasso`. Gray scale represents the percentage of inclusion of edges.

$n = 100$, and estimators are evaluated for an increasing number of nodes ($p = 50, 100, 200$). Figure 4 shows the mean and standard deviation of Hamming distances for estimates based on the PC-Algorithm (`pcalg`), as well as the proposed lasso (`lasso`) and adaptive lasso (`Alasso`) methods for different values of the tuning parameter α and different network sizes. For all values of p and α , the adaptive lasso estimate gives the best results, and the proposed penalized likelihood methods outperform the PC-Algorithm. This difference becomes more significant as the size of the network increases.

As mentioned in section 2, it is not always possible to estimate the direction of the edges of a DAG and therefore, the estimate from the PC-Algorithm may include undirected edges. Since our penalized likelihood methods assume knowledge of the ordering of variables and estimate the structure of the network, in the simulations considered here, we only estimate the skeleton (structure) of the network using the PC-Algorithm. We then use the ordering of the variables to determine the direction of the edges. The performance of the the PC-Algorithm for estimation of the partially completed DAG (PCDAG) may therefore be worse than the results reported here.

In the simulation results reported here, observations are generated according to the linear structural equation model (2.2) with standard normal latent variables and $\rho_{ij} = \rho = 0.8$. Additional simulation studies with different values of σ and ρ indicate that changes in σ do not have a considerable affect on the performance of the proposed models. On the other hand, as the magnitude of ρ decreases, the performance of the proposed methods (as well as the `pcalg`) algorithm deteriorates, but the findings of the above comparison remain unchanged.

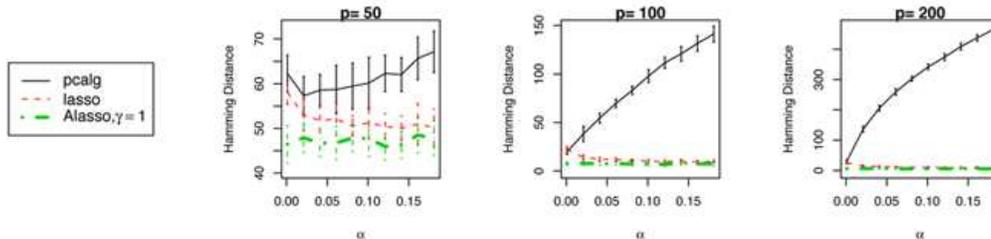


Figure 4: Hamming Distance for estimation of DAG using `pcalg`, `lasso` and `Alasso` from normal observations.

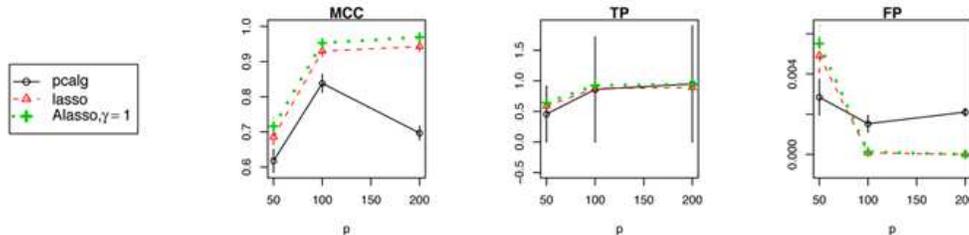


Figure 5: MCC, FP and TP for estimation of DAG using `pcalg`, `lasso` and `Alasso` from normal observations.

The above simulation results suggest that the optimal performance of the PC-Algorithm is achieved when $\alpha = 0.01$. The performance of lasso and adaptive lasso methods is less sensitive to the choice of α ; however, $\alpha = 0.10$ seems to deliver more reliable estimates. Similar results were also observed in simulations with other choices of σ and ρ . In addition, our extended simulations indicate that the performance of adaptive lasso does not vary significantly with the choice of power (γ), and therefore we only present the results for $\gamma = 1$. Figure 3 represents images of estimated and true DAGs created based on the above considerations for tuning parameters for $p = 50$. Similar results are observed for $p = 100$, $p = 200$, and are excluded to conserve space. Plots in figure 5 compare the performance of the three methods with the optimal settings of tuning parameters, over a range of values of p . It can be seen that the values of MCC confirm the above findings based on SHD. However, false positive and true positive rates only focus on one aspect of estimation at a time and do not provide a clear distinction between the methods. These plots also suggest that estimates from the `pcalg` may vary significantly. This variation is represented more significantly in terms of false positive rates, where standard deviations for estimates based are up to 10 times larger than those of `lasso` and `Alasso` estimates ($\sim 40\%$ for `pcalg` compared to $< 4\%$ for `lasso` and `Alasso`).

As mentioned in section 2, the representation of conditional independence in DAGs adapted in our proposed algorithm, is not restricted to normally distributed random variables. Moreover, if the underlying structural equations are linear, the method proposed in this paper can correctly estimate the underlying DAG. In order to assess the sensitivity of the estimates to the underlying distribution, we performed two simulation studies with non-Normal observations. In both simulations, observations were generated according to a linear structural model. However, in the first simulation, each latent variable was generated from a mixture of a standard normal and a t-distribution with 3 degrees of freedom, while in the second simulation, a t-distribution with 4 degrees of freedom was used. The performance of the proposed algorithm for non-normal observations was similar to the case of Gaussian observations, with `Alasso` providing the best estimates, and the performance of penalized methods improving as the dimension and sparsity increase.

5.3 Sensitivity to Perturbations in the Ordering of the Variables

Algorithm 3.2 assumes a known ordering of the variables. The superior performance of the proposed penalized likelihood methods in comparison to the PC-Algorithm may be explained by the fact that additional information about the order of the variables significantly simplifies the problem of estimating DAGs. Therefore, when such additional information is available, estimates using the PC-Algorithm suffer from a natural disadvantage. However, as the underlying

network becomes more sparse, the network includes fewer complex structures and it is expected that the ordering of variables should play a less significant role.

Next, we study the performance of the proposed penalized likelihood methods as well as the PC-Algorithm in problems where the ordering of variables is unknown. To this end, we generate normally distributed observations from the latent variable model of section 2.1. We then randomly permute the order of variables in the observation matrix and use the permuted matrix to estimate the original DAG. Figure 6 illustrates the performance of the three methods for choices of α described in section 5.2. It can be seen that for small, dense networks, the PC-Algorithm outperforms the proposed methods. This is expected since the change in the order of variables causes the algorithm to include unnecessary moral edges, while failing to recognize some of the existing associations. On the other hand, as the size of the network and correspondingly the degree of sparsity in the network increase, the local structures become simpler and therefore the ordering of the variables becomes less crucial. Thus, the performance of penalized likelihood algorithms is improved compared to that of the PC-Algorithm. For the high dimensional sparse case, where the computational cost of the PC-Algorithm becomes more significant, both lasso and adaptive lasso methods outperform the PC-Algorithm.

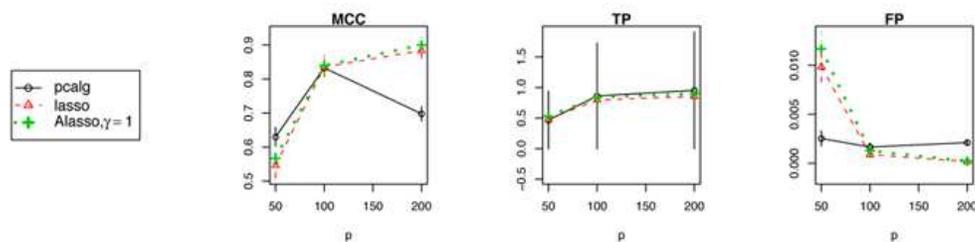


Figure 6: MCC, FP and TP for estimation of DAG using `pcalg`, `lasso` and `Alasso` with random ordering.

6 Real Data Application

6.1 Analysis of Cell Signalling Pathway Data

Sachs et al. (2003) carried out a set of flow cytometry experiments on signaling networks of human immune system cells. The ordering of the connections between pathway components were established based on perturbations in cells using molecular interventions and we consider the ordering to be known a priori. The data set includes $p = 11$ proteins and $n = 7466$ samples.

Friedman et al. (2008b) analyzed this data set using the `glasso` algorithm. They estimated the graph for a range of values of the ℓ_1 penalty and reported moderate agreement (around 50% false positive and false negative rates) between one of the estimates and the findings of Sachs et al. (2003). True and estimated signaling networks using the PC-Algorithm, as well as both lasso and adaptive lasso algorithms, along with performance measures are given in Figure 7. The estimated network using the `pcalg` includes a number of undirected edges. As in the simulation studies, we only estimate the structure of the network using the `pcalg` and determine the direction of edges by enforcing the ordering of nodes in the DAG. It can be seen that the adaptive lasso and lasso provides estimates that are closest to the true structure.

6.2 Transcription Regulatory Network of E-coli

Transcriptional regulatory networks play an important role in controlling the gene expression in cells and incorporating the underlying regulatory network results in more efficient estimation and inference (Shojaie and Michailidis, 2009a,b). Kao et al. (2004) proposed to use Network Component Analysis to infer transcriptional regulatory network of *Escherichia coli* (E-coli). They also provide whole genome expression data over time ($n = 24$), as well as information about the known regulatory network of E-coli.

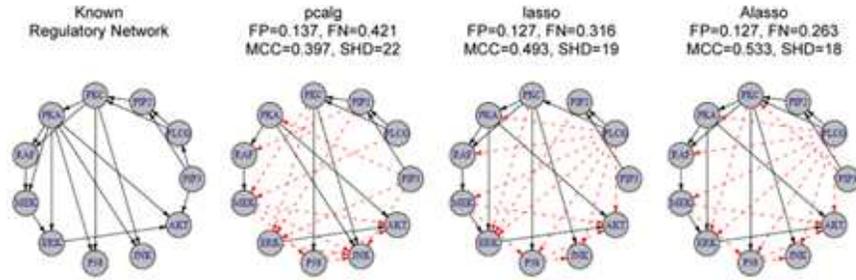


Figure 7: Known and estimated networks for human cell signalling data. True edges (True Positives in estimated networks) are marked with solid blue arrows, while False Positives are indicated by dashed red arrows.

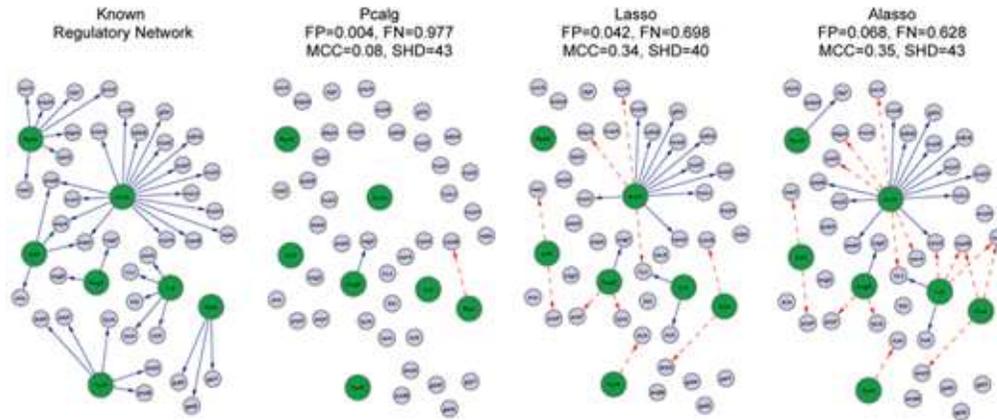


Figure 8: Known and estimated transcription regulatory network of E-coli. Large nodes indicate the transcription factors (TFs) and smaller nodes refer to their regulated genes (ORFs). True edges (True Positives in estimated networks) are marked with solid blue arrows, while False Positives are indicated by dashed red arrows.

In this application, the set of transcription factors (TFs) are known *a priori* and the goal is to find connections among transcription factors and regulated genes through analysis of whole genome transcriptomic data. Therefore, the algorithm proposed in this paper can be used by exploiting the natural hierarchy of TFs and genes. [Kao et al. \(2004\)](#) provide gene expression data for 7 transcription factors and 40 genes regulated by these TFs. Figure 8 presents the known regulatory network of E-coli along with the networks estimated using three different methods as well as different measures of performance. The relatively poor performance of the algorithms in this example can be partially attributed to the small sample size. However, it is also known that no single source of transcriptomic data is expected to successfully reveal the regulatory networks and methods that combine different sources of data are considered to be more efficient. It can be seen that the PC-Algorithm can only detect one of the true regulatory connections, while the proposed algorithm, with both lasso, as well as adaptive lasso penalties, offers significant improvements over the results of the PC-Algorithm, mostly due to the significant drop in proportion of false negatives (from 97% for the PC-Algorithm to 63% for adaptive lasso). Although the estimates based on lasso and adaptive lasso penalties are very similar, the choice of the best estimate depends on the performance evaluation metric.

7 Conclusion

We proposed efficient penalized likelihood methods for estimation of the structure of DAGs when variables inherit a natural ordering. Both lasso and adaptive lasso penalties were considered in this paper. However, the proposed algorithm can also be used for estimation of adjacency matrix of DAGs under other choices of penalty, as long as the

penalty function is applied to individual elements of the adjacency matrix.

There are a number of biological applications where the ordering of the variables is known *a priori*. Estimation of transcriptional regulatory networks from gene expression data and reconstruction of causal networks from temporal observations, based on the concept of Granger causality (Granger, 1969) are areas of potential application for the proposed algorithm. Our simulation studies also indicate that the correct ordering of variables is less crucial for estimation high dimensional sparse DAGs. Thus, even in high dimensional sparse applications where the ordering among variables is not known, the methods proposed in this paper may be an efficient alternative for search based methods of estimating DAGs.

Appendix: Technical Proofs

Recall from section 4.2 that $\theta^i = A_{i, i-1}$ denotes the entries of the i th row of the adjacency matrix to the left of the diagonal. Also, denote by $\theta^{i, \mathcal{I}}$ the estimate for the i th row with values outside the set of indices \mathcal{I} set to zero and let $\theta_j^{i, \mathcal{I}}$ be the j component of $\theta^{i, \mathcal{I}}$.

The following lemma is a consequence of the KarushKuhnTucker conditions for the general weighted lasso problem and is used in the proof of Theorems 4.3 and 4.5.

Lemma 7.1. *Let $\hat{\theta}^{i, \mathcal{I}}$ be the general weighted lasso estimate of $\theta^{i, \mathcal{I}}$, i.e.*

$$\hat{\theta}^{i, \mathcal{I}} = \underset{\theta: \theta_k = 0 \forall k \notin \mathcal{I}}{\operatorname{argmin}} \left\{ n^{-1} \|\mathcal{X}_i - \mathcal{X}\theta\|_2^2 + \lambda \sum_{k=1}^p |\theta_k| w_{ik} \right\} \quad (7.1)$$

Let

$$G_j(\theta) = -2n^{-1} \mathcal{X}_j^\top (\mathcal{X}_i - \mathcal{X}\theta)$$

and w_i be the vector of initial weights in adaptive lasso estimation problem. Then a vector $\hat{\theta}$ with $\hat{\theta}_k = 0, \forall k \notin \mathcal{I}$ is a solution of (7.1) iff $\forall j \in \mathcal{I}, G_j(\theta) = -\operatorname{sign}(\hat{\theta}_j) w_{ij} \lambda$ if $\hat{\theta}_j \neq 0$ and $|G_j(\theta)| \leq w_{ij} \lambda$ if $\hat{\theta}_j = 0$. Moreover, if the solution is not unique and $|G_j(\theta)| < w_{ij} \lambda$ for some solution $\hat{\theta}$, then $\hat{\theta}_j = 0$ for all solutions of (7.1).

Proof. The proof of the lemma is identical to the proof of Lemma (A.1) in Meinshausen and Bühlmann (2006) (except for inclusion of general weights w_{ij}) and is therefore omitted. \square

of Theorem 4.3. To prove (i), note that by Bonferroni's inequality, and the fact that $\operatorname{card}(pa_i) = o(n)$ as $n \rightarrow \infty$, it suffices to show that there exists some $c_{(i)} > 0$ such that for all $i \in V$ and for every $j \in pa_i$,

$$\operatorname{pr} \left\{ \operatorname{sign}(\hat{\theta}_j^{i, pa_i}) = \operatorname{sign}(\theta_j^{i, pa_i}) \right\} = 1 - O \left\{ \exp(-c_{(i)} n^\zeta) \right\} \quad \text{as } n \rightarrow \infty$$

Let $\hat{\theta}^{i, pa_i}(\beta)$ be the estimate of θ^{i, pa_i} in (7.1), with the j th component fixed at a constant value β ,

$$\hat{\theta}^{i, pa_i}(\beta) = \underset{\theta \in \Theta_\beta}{\operatorname{argmin}} \left\{ n^{-1} \|\mathcal{X}_i - \mathcal{X}\theta\|_2^2 + \lambda \sum_{k=1}^p |\theta_k| w_k \right\} \quad (7.2)$$

where $\Theta_\beta \equiv \{\theta \in \mathbb{R}^p : \theta_j = \beta, \theta_k = 0, \forall k \notin pa_i\}$. Note that for $\beta = \hat{\theta}_j^{i, pa_i}$, $\hat{\theta}^{i, pa_i}(\beta)$ is identical to $\hat{\theta}^{i, pa_i}$. Thus, if $\operatorname{sign}(\hat{\theta}_j^{i, pa_i}) \neq \operatorname{sign}(\theta_j^i)$, there would exist some β with $\operatorname{sign}(\beta) \operatorname{sign}(\theta_j^i) \leq 0$ such that $\hat{\theta}^{i, pa_i}(\beta)$ is a solution to (7.2). Since $\theta_j^i \neq 0, \forall j \in pa_i$, it suffices to show that for all β with $\operatorname{sign}(\beta) \operatorname{sign}(\theta_j^i) < 0$, with high probability, $\hat{\theta}^{i, pa_i}(\beta)$ can not be a solution to (7.2). Without loss of generality, we consider the case where $\theta_j^i > 0$ ($\theta_j^i < 0$ can be shown similarly). Then if $\beta \leq 0$, from Lemma 7.1, $\hat{\theta}^{i, pa_i}(\beta)$ can be a solution to (7.2) only if $G_j(\hat{\theta}^i(\beta)) \geq -\lambda w_{ij}$. Hence, it suffices to show that for some $c_{(i)} > 0$ and all $j \in pa_i$ with $\theta_j^i > 0$,

$$\operatorname{pr} \left[\sup_{\beta \leq 0} \left\{ G_j \left(\hat{\theta}^i(\beta) \right) < -\lambda w_{ij} \right\} \right] = 1 - O \left\{ \exp(-c_{(i)} n^\zeta) \right\} \quad \text{as } n \rightarrow \infty \quad (7.3)$$

Define,

$$\mathcal{R}_i(\beta) = \mathcal{X}_i - \mathcal{X}\hat{\theta}^i(\beta) \quad (7.4)$$

For every $j \in pa_i$ we can write,

$$X_j = \sum_{k \in pa_i \setminus \{j\}} \theta_k^{j, pa_i \setminus \{j\}} X_k + Z_j \quad (7.5)$$

where Z_j is independent of $\{X_k; k \in pa_i \setminus \{j\}\}$. Then by (7.5),

$$G_j(\hat{\theta}^i(\beta)) = -2n^{-1} \mathcal{Z}_j^T \mathcal{R}_i(\beta) - \sum_{k \in pa_i \setminus \{j\}} \theta_k^{j, pa_i \setminus \{j\}} 2n^{-1} \mathcal{X}_k^T \mathcal{R}_i(\beta)$$

By Lemma 7.1, it follows that for all $k \in pa_i \setminus \{j\}$, $|G_k(\hat{\theta}^i(\beta))| = |2n^{-1} \mathcal{X}_k^T \mathcal{R}_i(\beta)| \leq \lambda w_{ik}$, thus,

$$G_j(\hat{\theta}^i(\beta)) \leq -2n^{-1} \mathcal{Z}_j^T \mathcal{R}_i(\beta) + \lambda \sum_{k \in pa_i \setminus \{j\}} |\theta_k^{j, pa_i \setminus \{j\}}| w_{ik} \quad (7.6)$$

Using the fact that $|\theta_k^{j, pa_i \setminus \{j\}}| \leq 1$, it suffices to show that

$$\text{pr} \left[\sup_{\beta \leq 0} \{-2n^{-1} \mathcal{Z}_j^T \mathcal{R}_i(\beta)\} < -\lambda \sum_{k \in pa_i} w_{ik} \right] = 1 - O \left\{ \exp(-c_{(i)} n^\zeta) \right\} \quad \text{as } n \rightarrow \infty, \quad (7.7)$$

or equivalently,

$$\text{pr} \left[\inf_{\beta \leq 0} \{2n^{-1} \mathcal{Z}_j^T \mathcal{R}_i(\beta)\} < \lambda \sum_{k \in pa_i} w_{ik} \right] = O \left\{ \exp(-c_{(i)} n^\zeta) \right\} \quad \text{as } n \rightarrow \infty. \quad (7.8)$$

It is shown in Lemma A.2. of [Meinshausen and Bühlmann \(2006\)](#) that for any $q > 0$, there exists $c_{(i)} > 0$ such that for all $j \in pa_i$ with $\theta_j^i > 0$

$$\text{pr} \left[\inf_{\beta \leq 0} \{2n^{-1} \mathcal{Z}_j^T \mathcal{R}_i(\beta)\} \leq q\lambda \right] = O \left\{ \exp(-c_{(i)} n^\zeta) \right\} \quad \text{as } n \rightarrow \infty. \quad (7.9)$$

However, by definition $w_{ik} \leq 1$ and therefore, $\sum_{k \in pa_i} w_{ik} \leq \text{card}(pa_i) \leq 1$, which implies that (i) follows from 7.9.

To prove (ii), note that the event $\hat{pa}_i \not\subseteq pa_i$ is equivalent to the event that there exists a node $j \in \underline{i-1} \setminus pa_i$ such that $\hat{\theta}_j^i \neq 0$, i.e.

$$\text{pr}(\hat{pa}_i \subseteq pa_i) = 1 - \text{pr}(\exists j \in \underline{i-1} \setminus pa_i : \hat{\theta}_j^i \neq 0) \quad (7.10)$$

By Lemma 7.1, and using the fact that by definition $w_{ij} \geq 1$

$$\begin{aligned} \text{pr}(\exists j \in \underline{i-1} \setminus pa_i : \hat{\theta}_j^i \neq 0) &= \text{pr}(\exists j \in \underline{i-1} \setminus pa_i : |G_j(\hat{\theta}^{i, pa_i})| \geq w_{ij} \lambda) \\ &\leq \text{pr}(\exists j \in \underline{i-1} \setminus pa_i : |G_j(\hat{\theta}^{i, pa_i})| \geq q\lambda \text{ and } w_{ij} \lambda \leq q\lambda \text{ for some } q \geq 1) \\ &\leq \text{pr}(\exists j \in \underline{i-1} \setminus pa_i : w_{ij} \leq q \text{ for some } q \geq 1) \end{aligned}$$

Since $w_{ij} = 1 \vee |\tilde{\theta}_j^i|^{-\gamma}$, with $\tilde{\theta}_j^i$ the lasso estimate of the adjacency matrix from (3.8), using Lemma 7.1,

$$\begin{aligned} \text{pr}(\exists j \in \underline{i-1} \setminus pa_i : w_{ij} \leq q \text{ for some } q > 0) &= \text{pr}(\exists j \in \underline{i-1} \setminus pa_i : |\tilde{\theta}_j^i| \geq q^{-1/\gamma} \text{ for some } q \geq 1) \\ &\leq \text{pr}(\exists j \in \underline{i-1} \setminus pa_i : |\tilde{\theta}_j^i| \geq q' \text{ for some } q' > 0) \\ &\leq \text{pr}(\exists j \in \underline{i-1} \setminus pa_i : \tilde{\theta}_j^i \neq 0) \\ &= \text{pr}(\exists j \in \underline{i-1} \setminus pa_i : |G_j(\tilde{\theta}^{i, pa_i})| \geq \lambda^0) \end{aligned}$$

Since $\text{card}(pa_i) = o(n)$, we can assume without loss of generality that $\text{card}(pa_i) < n$, which implies that $\tilde{\theta}^{i,pa_i}$ is an almost sure unique solution to (7.1) with $\mathcal{I} = pa_i$. Let

$$\mathcal{E} = \left\{ \max_{j \in \underline{i-1} \setminus pa_i} |G_j(\tilde{\theta}^{i,pa_i})| < \lambda^0 \right\}.$$

Then conditional on event \mathcal{E} , it follows from the first part of Lemma 7.1 that $\tilde{\theta}^{i,pa_i}$ is also a solution of the unrestricted weighted lasso problem (7.1) with $\mathcal{I} = \underline{i-1}$. Since $\tilde{\theta}_j^{i,pa_i} = 0, \forall j \in \underline{i-1} \setminus pa_i$, it follows from the second part of Lemma 7.1 that $\tilde{\theta}_j^i = 0, \forall j \in \underline{i-1} \setminus pa_i$. Hence

$$\text{pr} \left(\exists j \in \underline{i-1} \setminus pa_i : \tilde{\theta}_j^i \neq 0 \right) \leq 1 - \text{pr}(\mathcal{E}) = \text{pr} \left(\max_{j \in \underline{i-1} \setminus pa_i} |G_j(\tilde{\theta}^{i,pa_i})| \geq \lambda^0 \right) \quad (7.11)$$

where,

$$G_j(\tilde{\theta}^{i,pa_i}) = -2n^{-1} \mathcal{X}_j^\top (\mathcal{X}_i - \mathcal{X} \tilde{\theta}^{i,pa_i}) \quad (7.12)$$

Since $\text{card}(V) = O(n^a)$ for some $a > 0$, Bonferroni's inequality implies that to verify (7.10) it suffices to show that there exists a constant $c_{(ii)} > 0$ such that for all $j \in \underline{i-1} \setminus pa_i$,

$$\text{pr} \left(|G_j(\tilde{\theta}^{i,pa_i})| \geq \lambda^0 \right) = O \left\{ \exp(-c_{(ii)} n^\zeta) \right\} \text{ as } n \rightarrow \infty. \quad (7.13)$$

where $R_j \sim N(0, \sigma_j^2)$, $\sigma_j^2 \leq 1$ and R_j is independent from $X_l, l \in pa_i$. Similarly, with R_i satisfying the same requirements as R_j , we get $X_i = \sum_{k \in pa_i} \theta_k^{i,pa_i} X_k + R_i$.

Denote by \mathcal{X}_{pa_i} the columns of \mathcal{X} corresponding to pa_i and let θ_{pa_i} the column vector of coefficients with dimension $\text{card}(pa_i)$ corresponding to pa_i . Then,

$$\begin{aligned} \text{pr} \left\{ |G_j(\tilde{\theta}^{i,pa_i})| \geq \lambda^0 \right\} &= \text{pr} \left\{ \left| -2n^{-1} \mathcal{X}_j^\top (\mathcal{X}_i - \mathcal{X} \tilde{\theta}^{i,pa_i}) \right| \geq \lambda^0 \right\} \\ &= \text{pr} \left[\left| -2n^{-1} \{ \mathcal{X}_{pa_i} \theta_{pa_i}^{j,pa_i} + \mathcal{R}_j \}^\top \{ \mathcal{X}_{pa_i} (\theta_{pa_i}^{i,pa_i} - \tilde{\theta}_{pa_i}^{i,pa_i}) + \mathcal{R}_i \} \right| \geq \lambda^0 \right]. \end{aligned}$$

Therefore,

$$\begin{aligned} \text{pr} \left\{ |G_j(\tilde{\theta}^{i,pa_i})| \geq \lambda^0 \right\} &\leq \text{pr} \left\{ \left| -2n^{-1} (\theta_{pa_i}^{i,pa_i} - \tilde{\theta}_{pa_i}^{i,pa_i})^\top \mathcal{X}_{pa_i}^\top \mathcal{X}_{pa_i} \theta_{pa_i}^{j,pa_i} \right| \geq \lambda^0/3 \right\} \\ &\quad + \text{pr} \left\{ \left| -2n^{-1} (\theta_{pa_i}^{i,pa_i} - \tilde{\theta}_{pa_i}^{i,pa_i})^\top \mathcal{X}_{pa_i}^\top \mathcal{R}_j \right| \geq \lambda^0/3 \right\} \\ &\quad + \text{pr} \left\{ \left| -2n^{-1} (\mathcal{X}_{pa_i} \theta_{pa_i}^{j,pa_i} + \mathcal{R}_j)^\top \mathcal{R}_i \right| \geq \lambda^0/3 \right\} \\ &\equiv \text{I} + \text{II} + \text{III} \end{aligned}$$

Let $\mathbf{1}_{pa_i}$ denote a vector of 1's of dimension $\text{card}(pa_i)$. Then, using the fact that $|\theta_l^{j,pa_i}| \leq 1$, for all $l \in pa_i$, we can write $\text{I} \leq \text{pr} \left(2 \|\theta_{pa_i}^{i,pa_i} - \tilde{\theta}_{pa_i}^{i,pa_i}\|_\infty n^{-1} (\mathcal{X}_{pa_i} \mathbf{1}_{pa_i})^\top \mathcal{X}_{pa_i} \mathbf{1}_{pa_i} \geq \lambda^0/3 \right)$. Then $\mathcal{X}_{pa_i}^\top \mathcal{X}_{pa_i} \sim \mathbb{W}_{\text{card}(pa_i)}(\Sigma_{pa_i}, n)$ where $\mathbb{W}_m(\Sigma, n)$ denotes a Wishart distribution with mean $n\Sigma$. Hence, from properties of the Wishart distribution, we get $(\mathcal{X}_{pa_i} \mathbf{1}_{pa_i})^\top \mathcal{X}_{pa_i} \mathbf{1}_{pa_i} \sim \mathbb{W}_1(\mathbf{1}_{pa_i}^\top \Sigma_{pa_i} \mathbf{1}_{pa_i}, n)$.

Since pa_i also forms a DAG, the eigenvalues Σ_{pa_i} are bounded (see Remark 2.2), and hence

$$\mathbf{1}_{pa_i}^\top \Sigma_{pa_i} \mathbf{1}_{pa_i} \leq \text{card}(pa_i) \phi_{\max}(\Sigma_{pa_i}) \quad (7.14)$$

Therefore, if $Z \sim \chi_1^2$, $n^{-1} (\mathcal{X}_{pa_i} \mathbf{1}_{pa_i})^\top \mathcal{X}_{pa_i} \mathbf{1}_{pa_i}$ is stochastically smaller than $\text{card}(pa_i) \phi_{\max}(\Sigma_{pa_i}) Z$. On the other hand, by Theorem 4.2,

$$\|A - \tilde{A}\|_F = O_p \left\{ (n^{-1} s \log p)^{1/2} \right\}$$

and hence

$$\|\theta_{pa_i}^{i,pa_i} - \tilde{\theta}_{pa_i}^{i,pa_i}\|_\infty = O_p \left\{ (n^{-1} s \log p)^{1/2} \right\} \quad (7.15)$$

Noting that $\text{card}(pa_i) = O(n^b)$, $b < 1/2$ and $p = O(n^a)$, $a > 0$, (7.14) and (7.15) imply that

$$\|\theta_{pa_i}^{i,pa_i} - \tilde{\theta}_{pa_i}^{i,pa_i}\|_\infty \text{card}(pa_i) \phi_{\max}(\Sigma_{pa_i}) = O_p \left\{ (sn^{2b-1} a \log n)^{1/2} \right\}$$

By (A-0'), $sn^{2b-1} \log n = o(1)$ and hence by Slutsky's Theorem and properties of the χ^2 -distribution, there exists $c_{(I)} > 0$ such that for all $j \in \underline{i-1} \setminus pa_i$,

$$I = O \left\{ \exp(-c_{(I)} n^\zeta) \right\} \quad \text{as } n \rightarrow \infty$$

Using a similar argument, for II we can write,

$$II \leq \text{pr} \left(2n^{-1} \|\theta_{pa_i}^{i,pa_i} - \tilde{\theta}_{pa_i}^{i,pa_i}\|_\infty |\mathbf{1}_{pa_i} \mathcal{X}_{pa_i}^\top \mathcal{R}_j| \geq \lambda^0/3 \right). \quad (7.16)$$

Since columns of \mathcal{X}_{pa_i} correspond to samples from normal random variables with mean zero and are all independent of \mathcal{R}_j , it suffices to show that there exists $c_{(II)} > 0$ such that for all $j \in \underline{i-1} \setminus pa_i$ and for all $k \in pa_i$,

$$\text{pr} \left(2n^{-1} \|\theta_{pa_i}^{i,pa_i} - \tilde{\theta}_{pa_i}^{i,pa_i}\|_\infty \text{card}(pa_i) |\mathcal{X}_k^\top \mathcal{R}_j| \geq \lambda^0/3 \right) = O \left\{ \exp(-c_{(II)} n^\zeta) \right\} \quad \text{as } n \rightarrow \infty \quad (7.17)$$

By (7.15) and (A-0'), the random variable on the left hand side of (7.17) is stochastically smaller than $2n^{-1} |\mathcal{X}_k \mathcal{R}_j|$. By independence of X_k and R_j , $E(X_k R_j) = 0$. Also using Gaussianity of both X_k and R_j , there exists $g < \infty$ such that $E\{\exp(|X_k R_j|)\} \leq g$. Since $\lambda^0 = O\{(\log p/n)^{1/2}\}$, by Bernstein's inequality (Van der Vaart and Wellner, 1996), $\text{pr}(2n^{-1} |\mathcal{X}_k \mathcal{R}_j| > \lambda^0/3) \leq \exp(-c_{(II)} n^\zeta)$ for some $c_{(II)} > 0$ and hence (7.17) is satisfied. Finally, for III we have

$$\text{pr} \left\{ | -2n^{-1} (\mathcal{X}_{pa_i} \theta_{pa_i}^{j,pa_i} + \mathcal{R}_i)^\top \mathcal{R}_j | \geq \lambda^0/3 \right\} = \text{pr} \left\{ | -2n^{-1} \mathcal{X}_i^\top \mathcal{R}_j | \geq \lambda^0/3 \right\} \quad (7.18)$$

and using the Bernstein's inequality we conclude that there exists $c_{(III)} > 0$ such that for all $j \in \underline{i-1} \setminus pa_i$ and for all $k \in pa_i$, $III = O \left\{ \exp(-c_{(III)} n^\zeta) \right\}$ as $n \rightarrow \infty$. The proof of (ii) is then complete by taking $c_{(ii)}$ to be the minimum of $c_{(I)}, \dots, c_{(III)}$.

To prove (iii), note that $\text{pr}(pa_i \subseteq \hat{p}a_i) = 1 - \text{pr}(\exists j \in pa_i : \hat{\theta}_j^i = 0)$, and let $\mathcal{E} = \left\{ \max_{k \in \underline{i-1} \setminus pa_i} |G_j(\hat{\theta}^{i,pa_i})| < \lambda w_{ij} \right\}$.

It follows from an argument similar to the proof of (ii) that conditional on \mathcal{E} , $\hat{\theta}^{i,pa_i}$ is an almost sure unique solution of the unrestricted adaptive lasso problem (7.1) with $\mathcal{I} = \underline{i-1}$. Therefore,

$$\text{pr} \left(\exists j \in pa_i : \hat{\theta}_j^i = 0 \right) \leq \text{pr} \left(\exists j \in pa_i : \hat{\theta}_j^i = 0 \right) + \text{pr}(\mathcal{E}^c).$$

From (i), there exists a $c_1 > 0$ such that $\text{pr}(\exists j \in pa_i : \hat{\theta}_j^i = 0) = O(\exp(-c_1 n^\zeta))$ and it was shown in (ii) that $\text{pr}(\mathcal{E}^c) = O\{\exp(-c_2 n^\zeta)\}$ for some $c_2 > 0$. Thus (iii) follows from Bonferroni's inequality.

Finally, it is easy to see that since $p = O(n^a)$, the claim in (iv) also follows from (ii) and (iii) and Bonferroni's inequality. \square

of Theorem 4.5. We first show that if $AN_i \cap AN_j = \emptyset$, then i and j are independent. Since $\Sigma = \Lambda \Lambda^\top$ and Λ is lower triangular,

$$\Sigma_{ij} = \sum_{k=1}^{\min(i,j)} \Lambda_{ik} \Lambda_{jk} \quad (7.19)$$

We assume without loss of generality that $i < j$. The argument for the case $j > i$ is similar. Suppose for all $k = 1, \dots, i$ that $\Lambda_{ik} = 0$ or $\Lambda_{jk} = 0$, then by (7.19) i and j are independent. However, by Lemma 2.1, Λ_{jk} is the influence of k th node on j , and this is zero only if there is no path from k to j . Clearly, if i is an ancestor of j , we have $\Sigma_{ij} \neq 0$. On the other hand, if there is no node $k \in \underline{i-1}$ such that k influences both i and j , (i.e. k is a common ancestor of i and j) then for all $k = 1, \dots, i$ we have $\Lambda_{ik} \Lambda_{jk} = 0$ and the claim follows.

Using Bonferroni's inequality twice and Lemma 7.1, we get

$$\begin{aligned}
\Pr(\exists i \in V : \hat{AN}_i \not\subseteq AN_i) &\leq p \max_{i \in V} \Pr(\exists j \in \underline{i-1} \setminus AN_i : j \in \hat{p}a_i) \\
&\leq p(i-1) \max_{i \in V, j \in \underline{i-1} \setminus AN_i} \Pr(j \in \hat{p}a_i) \\
&\leq p(i-1) \max_{i \in V, j \in \underline{i-1} \setminus AN_i} \Pr\{|G_j(\hat{\theta}^{i, AN_i})| \geq \lambda w_{ij}\}.
\end{aligned}$$

However, by definition $w_{ij} \geq 1$, and hence it suffices to show that,

$$(i-1)p \max_{i \in V, j \in \underline{i-1} \setminus AN_i} \Pr\{|G_j(\hat{\theta}^{i, AN_i})| \geq \lambda\} \leq \alpha. \quad (7.20)$$

Note that $G_j(\hat{\theta}^{i, AN_i}) = -2n^{-1} \mathcal{X}_j^\top (\mathcal{X}_i - \mathcal{X} \hat{\theta}^{i, AN_i})$ and X_j is independent of X_k for all $k \in AN_i$. Therefore, conditional on \mathcal{X}_{AN_i} , $G_j(\hat{\theta}^{i, AN_i}) \sim (0, 4R^2/n)$, where $R^2 = n^{-1} \|\mathcal{X}_i - \mathcal{X} \hat{\theta}^{i, AN_i}\|_2^2 \leq n^{-1} \|\mathcal{X}_i\|_2^2 = 1$, by definition of $\hat{\theta}^{i, AN_i}$ and the fact that columns of the data matrix are scaled.

It follows that for all $j \in \underline{i-1} \setminus AN_i$, $\Pr\{|G_j(\hat{\theta}^{i, AN_i})| \geq \lambda \mid \mathcal{X}_{AN_i}\} \leq 2\{1 - \Phi(n^{1/2}\lambda/2)\}$, where Φ is the cumulative distribution function for standard normal random variable. Using the choice of λ proposed in (4.1), we get $\Pr\{|G_j(\hat{\theta}^{i, AN_i})| \geq \lambda \mid \mathcal{X}_{AN_i}\} \leq \frac{\alpha}{(i-1)p}$, and the result follows. \square

References

- Banerjee, O., El Ghaoui, L., and dAspremont Alexandre (2008). Model selection through sparse maximum likelihood estimation for multivariate Gaussian or binary data. *Journal of Machine Learning Research*, 9:485–516.
- Fan, J., Feng, Y., and Wu, Y. (2007). Network exploration via adaptive LASSO and SCAD penalties. *Manuscript*.
- Friedman, J., Hastie, T., and Tibshirani, R. (2008a). Regularization paths for generalized linear models via coordinate descent. *Technical report*.
- Friedman, J., Hastie, T., and Tibshirani, R. (2008b). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432.
- Granger, C. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, pages 424–438.
- Heckerman, D., Geiger, D., and Chickering, D. (1995). Learning Bayesian networks: The combination of knowledge and statistical data. *Machine Learning*, 20(3):197–243.
- Huang, J., Liu, N., Pourahmadi, M., and Liu, L. (2006). Covariance matrix selection and estimation via penalised normal likelihood. *Biometrika*, 93(1):85–98.
- Huang, J., Ma, S., and Zhang, C. (2008). Adaptive Lasso For Sparse High-Dimensional Regression Models. *Statistica Sinica*, 18:1603–1618.
- Kalisch, M. and Bühlmann, P. (2007). Estimating High-Dimensional Directed Acyclic Graphs with the PC-Algorithm. *Journal of Machine Learning Research*, 8:613–636.
- Kao, K., Yang, Y., Boscolo, R., Sabatti, C., Roychowdhury, V., and Liao, J. (2004). Transcriptome-based determination of multiple transcription regulator activities in Escherichia coli by using network component analysis. *Proceedings of the National Academy of Sciences*, 101(2):641–646.
- Knight, K. and Fu, W. (2000). Asymptotics for lasso-type estimators. *Annals Of Statistics*, 28(5):1356–1378.
- Lam, C. and Fan, J. (2008). Sparsity and Rate of Convergence in Large Covariance Matrix Estimation. *Manuscript*.

- Lauritzen, S. (1996). *Graphical models*. Oxford Univ Press.
- Leclerc, R. (2008). Survival of the sparsest: robust gene networks are parsimonious. *Mol Systems Biology*, 4(1).
- Markowitz, F. and Spang, R. (2007). Inferring cellular networks—a review. *BMC Bioinformatics*, 8(Suppl 6):S5.
- Meinshausen, N. and Bühlmann, P. (2006). High-dimensional graphs and variable selection with the Lasso. *Annals Of Statistics*, 34(3):1436.
- Pearl, J. (2000). *Causality: Models, Reasoning, and Inference*. Cambridge Univ Press.
- Rothman, A., Bickel, P., Levina, E., and Zhu, J. (2008). Sparse permutation invariant covariance estimation. *Electronic Journal of Statistics*, 2:494–515.
- Sachs, K., Perez, O., Pe’er, D., Lauffenburger, D., and Nolan, G. (2003). Causal protein-signaling networks derived from multiparameter single-cell data. *Science*, 308(5721):504–6.
- Shojaie, A. and Michailidis, G. (2009a). Analysis of Gene Sets Based on the Underlying Regulatory Network. *Journal of Computational Biology*, 16(3):407–426.
- Shojaie, A. and Michailidis, G. (2009b). Network Enrichment Analysis in Complex Experiments. *Statistical Applications in Genetics and Molecular Biology*, tentatively accepted.
- Spirtes, P., Glymour, C., and Scheines, R. (2000). *Causation, Prediction, and Search*. MIT Press.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B. Methodological*, 58(1):267–288.
- Tsamardinos, I., Brown, L., and Aliferis, C. (2006). The max-min hill-climbing Bayesian network structure learning algorithm. *Machine Learning*, 65(1):31–78.
- Van der Vaart, A. and Wellner, J. (1996). *Weak convergence and empirical processes*. Springer.
- Yuan, M. and Lin, Y. (2007). Model selection and estimation in the Gaussian graphical model. *Biometrika*, 94(1):19.
- Zou, H. (2006). The Adaptive Lasso and Its Oracle Properties. *Journal of the American Statistical Association*, 101(476):1418–1429.