

An algorithm for the multivariate group lasso with covariance estimation

Wilms I, Croux C.



An algorithm for the multivariate group lasso with covariance estimation

Ines Wilms* and Christophe Croux

Leuven Statistics Research Centre, KU Leuven, Belgium

Abstract We study a group lasso estimator for the multivariate linear regression model that accounts for correlated error terms. A block coordinate descent algorithm is used to compute this estimator. We perform a simulation study with categorical data and multivariate time series data, typical settings with a natural grouping among the predictor variables. Our simulation studies show the good performance of the proposed group lasso estimator compared to alternative estimators. We illustrate the method on a time series data set of gene expressions.

Keywords: Categorical variables, Group Lasso, Multivariate Regression, Penalized Maximum Likelihood, Sparsity, Time Series.

*Corresponding author: ines.wilms@kuleuven.be

1 Introduction

Since its introduction by Yuan and Lin (2006), the group least absolute shrinkage and selection operator (group lasso) has received considerable interest in the statistical literature (e.g. Meier et al., 2008, Wang and Leng, 2008, Simon et al., 2013, Alfons et al., 2015). In many applications, the parameter vector in the regression model is structured into groups. Typical examples are (i) regression with categorical variables, where a group of dummies represents each categorical variable, or (ii) time series regression where several lagged values of the same time series are included in the model. In settings with such a natural group structure, one wants to select either all or none of the variables belonging to a particular group. The key strength of the group lasso lies in its ability to perform such groupwise selection.

We consider the group lasso for the multivariate linear regression model. The multivariate linear regression model generalizes the classical linear regression model in that it regresses $q > 1$ responses instead of a single response on p predictors. Let $\mathbf{Y} = (\mathbf{y}_1, \dots, \mathbf{y}_q) \in \mathbb{R}^{n \times q}$ be the response matrix, and $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_p) \in \mathbb{R}^{n \times p}$ be the predictor matrix. The error vectors are assumed to follow a normal $N_q(0, \Sigma)$ distribution, with $\Sigma^{-1} = \Omega$, and are collected in the columns of the error matrix \mathbf{E} . The multivariate linear regression model is given by

$$\mathbf{Y} = \mathbf{X}\mathbf{B} + \mathbf{E}, \tag{1}$$

where $\mathbf{B} \in \mathbb{R}^{p \times q}$ is the coefficient matrix. We assume that this coefficient matrix contains K predefined groups. Denote each group as \mathbf{B}_{G_j} where $j \in \{1, \dots, K\}$.

Recently, Li et al. (2015) discussed the group lasso for the multivariate linear regression model. Their multivariate group lasso estimator¹ is given by

$$\hat{\mathbf{B}} = \underset{\mathbf{B}}{\operatorname{argmin}} \frac{1}{2n} \operatorname{tr}((\mathbf{Y} - \mathbf{X}\mathbf{B})^T(\mathbf{Y} - \mathbf{X}\mathbf{B})) + \sum_{j=1}^K \lambda_{G_j} m_j \|\mathbf{B}_{G_j}\|_2, \tag{2}$$

where $\operatorname{tr}(\cdot)$ denotes the trace, $\lambda_{G_j} > 0$, for $1 \leq j \leq K$, are sparsity parameters, and m_j equals the number of elements in group j . A groupwise penalty is used for the regression

¹Note that Li et al. (2015) consider a more general version of the multivariate group lasso that also allows for selection of predictors within the important groups.

coefficients. As such, variables are selected in a grouped manner: either all elements of a certain group are set to zero or none.

However, Li et al. (2015) do not account for correlated errors. We extend the multivariate group lasso from Li et al. (2015) such that the correlation between the error terms of the different equations of the multivariate regression model is taken into account. To this end, we simultaneously estimate the regression parameters \mathbf{B} and the inverse covariance matrix of the error terms $\mathbf{\Omega}$ using penalized maximum likelihood:

$$(\hat{\mathbf{B}}, \hat{\mathbf{\Omega}}) = \underset{(\mathbf{B}, \mathbf{\Omega})}{\operatorname{argmin}} \frac{1}{2n} \operatorname{tr}((\mathbf{Y} - \mathbf{XB})^T(\mathbf{Y} - \mathbf{XB})\mathbf{\Omega}) - \frac{1}{2} \log |\mathbf{\Omega}| + \sum_{j=1}^K \lambda_{G_j} m_j \|\mathbf{B}_{G_j}\|_2 + \lambda_\omega \sum_{k \neq k'} |\omega_{kk'}|, \quad (3)$$

where $\lambda_\omega > 0$ is a sparsity parameter, and $\omega_{kk'}$ is the k^{th} element of $\mathbf{\Omega}$. We use an L_1 penalty for the elements of the inverse covariance matrix.

Section 2 describes the algorithm used to approximate the minimizer of the objective function in (3). The main modification in the algorithm compared to the proposal of Li et al. (2015) is that the error covariance structure is taken into account. Simulation studies are performed in Section 3. Our simulations show that the group lasso with covariance estimation considerably outperforms the group lasso without covariance estimation. Section 4 contains a real data example.

2 The algorithm

To find the minimum of the penalized negative log-likelihood in (3), we iteratively solve for \mathbf{B} conditional on $\mathbf{\Omega}$ and for $\mathbf{\Omega}$ conditional on \mathbf{B} .

Solving for \mathbf{B} conditional on $\mathbf{\Omega}$. When $\mathbf{\Omega}$ is fixed, the minimization problem in (3) is equivalent to

$$\hat{\mathbf{B}} = \underset{\mathbf{B}}{\operatorname{argmin}} \frac{1}{2n} \operatorname{tr}((\mathbf{Y} - \mathbf{XB})^T(\mathbf{Y} - \mathbf{XB})\mathbf{\Omega}) + \sum_{j=1}^K \lambda_{G_j} m_j \|\mathbf{B}_{G_j}\|_2. \quad (4)$$

To find a solution to (4), we use a block coordinate descent algorithm, analogously to Friedman et al. (2007) for solving the single response lasso problem, or to Li et al. (2015) for the

multivariate group lasso problem without covariance estimation. Lemma 1 (Lemma 4.2 from Chapter 4 in Bühlmann and van de Geer, 2011) provides a necessary and sufficient condition for $\widehat{\mathbf{B}}$ to be a solution of (4).

Lemma 1. *Denote the loss function by*

$$\rho(\mathbf{B}) = \frac{1}{2n} \text{tr}((\mathbf{Y} - \mathbf{XB})^T(\mathbf{Y} - \mathbf{XB})\Omega).$$

The gradient of the loss function evaluated at \mathbf{B} is

$$\nabla\rho(\mathbf{B}) = -\frac{1}{n}\mathbf{X}^T(\mathbf{Y} - \mathbf{XB})\Omega.$$

A necessary and sufficient condition for \mathbf{B} to be a solution of (4) is

1. $\nabla\rho(\mathbf{B})_{G_j} + \lambda_{G_j}m_j \frac{\mathbf{B}_{G_j}}{\|\mathbf{B}_{G_j}\|_2} = \mathbf{0}$ if $\mathbf{B}_{G_i} \neq \mathbf{0}$
2. $\|\nabla\rho(\mathbf{B})_{G_j}\|_2 \leq \lambda_{G_j}m_j$ if $\mathbf{B}_{G_j} = \mathbf{0}$.

To start up the block coordinate descent algorithm, an initial value for \mathbf{B} is needed. We use the lasso estimator obtained by performing q separate lasso regressions. Assume now that $\widehat{\mathbf{B}}^{(m-1)}$ is given, for $m \geq 1$. In the following iteration step m , we update our estimate from $\widehat{\mathbf{B}}^{(m-1)}$ to $\widehat{\mathbf{B}}^{(m)}$. Note that the ik^{th} element of the gradient of the loss function evaluated at \mathbf{B} is given by

$$\begin{aligned} \nabla\rho(\mathbf{B})_{ik} &= -\frac{1}{n}\mathbf{x}_i^T(\mathbf{Y} - \mathbf{XB})\Omega_k \\ &= \frac{1}{n}(-\mathbf{x}_i^T(\mathbf{Y} - \mathbf{XB}^{-ik})\Omega_k + \omega_{kk}\|\mathbf{x}_i\|_2^2 B_{ik}) \\ &= \frac{1}{n}(-\mathbf{S}_{ik} + \omega_{kk}\|\mathbf{x}_i\|_2^2 B_{ik}), \end{aligned}$$

with \mathbf{x}_i the i^{th} column of \mathbf{X} , Ω_k the k^{th} row of Ω , ω_{kk} the kk^{th} element of Ω , \mathbf{B}^{-ik} is \mathbf{B} with element ik replaced by zero, and $\mathbf{S}_{ik} = \mathbf{x}_i^T(\mathbf{Y} - \mathbf{XB}^{-ik})\Omega_k$.

In iteration step m , we cycle through all groups G_j , with $j = 1, \dots, K$. If, for group G_j

$$\|\nabla\rho(\widehat{\mathbf{B}}^{(m-1)})_{G_j}\|_2 \leq \lambda_{G_j}m_j$$

Table 1: Block Coordinate Descent Algorithm to solve for \mathbf{B} conditional on Ω .

1: **Initialization** Let $\mathbf{B}^{(0)}$ be an initial parameter estimate. We use the lasso estimator obtained by performing q separate lasso regressions. Set $m = 0$.

2: **Repeat**

$m \leftarrow m + 1$

For each block $j = 1, \dots, K$:

If $\|\nabla \rho(\widehat{\mathbf{B}}^{(m-1)})_{G_j}\|_2 \leq \lambda_{G_j} m_j$: set $\widehat{\mathbf{B}}_{G_j}^{(m)} = \mathbf{0}$

Else: Update every ik^{th} element $\widehat{B}_{ik}^{(m)}$ of $\widehat{\mathbf{B}}^{(m)}$ belonging to group G_j by

$$\widehat{B}_{ik}^{(m)} = \frac{\widehat{S}_{ik}^{(m-1)}}{\omega_{kk} \|\mathbf{x}_i\|_2^2 + \frac{n \lambda_{G_j} m_j}{\|\widehat{\mathbf{B}}_{G_j}^{(m-1)}\|_2}}.$$

3: **Until** convergence. We iterate until the relative change in the value of the objective function in (4) in two successive iterations is smaller than the tolerance value $\epsilon = 10^{-2}$.

holds, then according to condition 2 from Lemma 1, all elements of group G_j of $\widehat{\mathbf{B}}^{(m)}$ are set to zero. Otherwise, according to condition 1 from Lemma 1, for every element ik of \mathbf{B} belonging to group G_j it needs to hold that

$$\begin{aligned} 0 &= \nabla \rho(\mathbf{B})_{ik} + \lambda_{G_j} m_j \frac{B_{ik}}{\|\mathbf{B}_{G_j}\|_2} \\ \iff 0 &= \frac{-\mathbf{S}_{ik}}{n} + \frac{\omega_{kk} \|\mathbf{x}_i\|_2^2}{n} B_{ik} + \frac{\lambda_{G_j} m_j}{\|\mathbf{B}_{G_j}\|_2} B_{ik} \\ \iff B_{ik} &= \frac{\mathbf{S}_{ik}}{\omega_{kk} \|\mathbf{x}_i\|_2^2 + \frac{n \lambda_{G_j} m_j}{\|\mathbf{B}_{G_j}\|_2}}. \end{aligned} \quad (5)$$

Note that the right-hand-side from equation (5) involves B_{ik} in the computation of $\|\mathbf{B}_{G_j}\|_2$. For this, we use the estimate from the previous iteration. Table 1 provides a schematic overview of the block coordinate descent algorithm.

Note that the estimator in (4) is a multivariate *adaptive* group lasso estimator since each group has its own sparsity parameter λ_{G_j} . We take $\lambda_{G_j} = \lambda / \|\widehat{\mathbf{B}}_{G_j}^{(0)}\|_2$, for $j = 1, \dots, K$. This way, only one tuning parameter for the regression coefficients needs to be selected instead of K . We use a grid of sparsity parameters and search for the optimal one using the Bayesian

Information Criterion (BIC). The BIC is given by

$$BIC_\lambda = -2\log L_\lambda + k_\lambda \log(n),$$

where $\log L_\lambda$ is the estimated log-likelihood, corresponding to the first term of the objective function in (4), using sparsity parameter λ , and k_λ is the number of non-zero estimated regression coefficients.

Solving for Ω conditional on \mathbf{B} . When \mathbf{B} is fixed, the minimization problem in (3) corresponds to the graphical lasso (Friedman et al., 2008) on the residuals $\mathbf{Y} - \mathbf{XB}$. We use the Bayesian Information Criterion to select the optimal value of the sparsity parameter λ_ω (e.g. Yuan and Lin, 2007).

Starting value and convergence. We start by taking $\Omega = \mathbf{I}$ and then iteratively solve for \mathbf{B} conditional on Ω and for Ω conditional on \mathbf{B} . We iterate until the relative change in the value of the objective function in (3) in two successive iterations is smaller than the tolerance value $\epsilon = 10^{-2}$.

3 Simulation

We compare the performance of the multivariate group lasso with covariance estimation, “GroupLasso+Cov”, to

1. The multivariate group lasso without covariance estimation, “GroupLasso”, i.e. the solution of (2),
2. The multivariate lasso with covariance estimation, “Lasso+Cov”, i.e. the solution of (3) with $m_j = 1$, for $1 \leq j \leq K$, where $K = p \times q$. The resulting estimator is equivalent to the Multivariate Lasso With Covariance Estimator introduced in Rothman et al. (2010).
3. The multivariate lasso without covariance estimation, “Lasso”, i.e. the solution of (2) with $m_j = 1$, for $1 \leq j \leq K$, where $K = p \times q$.

Note that “Lasso+Cov” and “Lasso” do not take the group structure among the predictors into account.

3.1 Predictor groups

The first data configuration corresponds to a regression model with categorical predictors, the second to a time series model.

Categorical data. We consider a design similar to model I from Yuan and Lin (2006) for the univariate regression model. We generate a sample Z_{ij} , for $i = 1, \dots, n$ and $j = 1, \dots, K$, of size n from a centered multivariate normal distribution with covariance matrix Σ^Z where

$$\Sigma_{ij}^Z = 0.5^{|i-j|}.$$

Afterwards, Z_{ij} is trichotomized as

$$C_{ij} = \begin{cases} 0 & \text{if } Z_{ij} < \Phi^{-1}\left(\frac{1}{3}\right) \\ 1 & \text{if } Z_{ij} > \Phi^{-1}\left(\frac{2}{3}\right) \\ 2 & \text{if } \Phi^{-1}\left(\frac{1}{3}\right) < Z_{ij} < \Phi^{-1}\left(\frac{2}{3}\right), \end{cases}$$

for $i = 1, \dots, n = 50$ and $j = 1, \dots, K$, where K denotes the number of groups. We take $K \in \{5, 20, 50\}$. The $(n \times p)$ matrix of predictors \mathbf{X} then contains in its columns the $p = 2K$ dummy variables $D_{ij}^0 = \mathbf{I}(C_{ij} = 0)$ and $D_{ij}^1 = \mathbf{I}(C_{ij} = 1)$, for $j = 1, \dots, K$ and $i = 1, \dots, n$, where $\mathbf{I}(\cdot)$ is the indicator function. Next, the $q = 5$ responses are simulated from

$$\mathbf{Y} = \mathbf{B}\mathbf{X} + \mathbf{E}, \tag{6}$$

where $\mathbf{B} = \mathbf{I}_q \otimes \mathbf{b}$, with $\mathbf{b} = (2, -1, \dots, 2, -1)$ a vector of length p/q . For the error covariance matrix Σ we consider different structures, detailed in the Section 3.2. The group lasso accounts for the grouped predictor variables by selecting either all or none of the dummy variables corresponding to a particular categorical variable in one of the equations of the multivariate regression model.

Time series. We generate the data from a VAR(2) model

$$\mathbf{y}_t = \mathbf{B}_1 \mathbf{y}_{t-1} + \mathbf{B}_2 \mathbf{y}_{t-2} + \mathbf{e}_t, \quad (7)$$

for $t = 1, \dots, T = 50$, where \mathbf{y}_t is a q -dimensional vector, with $q \in \{5, 20, 50\}$. The coefficient matrices \mathbf{B}_1 and \mathbf{B}_2 have the same sparse structure and $\mathbf{e}_t \sim N_q(\mathbf{0}, \boldsymbol{\Sigma})$. For the error covariance matrix $\boldsymbol{\Sigma}$ we consider different structures, detailed in Section 3.2.

The above model is a Vector AutoRegressive (VAR) model of order two since two lagged values of each time series are included as predictors. The group lasso accounts for the grouped predictor variables by selecting either all or none of the lagged values of a particular time series in one of the equations of the VAR. As a result, $\widehat{\mathbf{B}}_1$ and $\widehat{\mathbf{B}}_2$ have their zero elements in exactly the same cells.

We generate the sparse coefficient matrices \mathbf{B}_1 and \mathbf{B}_2 from a network structure (see Fujita et al., 2007). This dimensions of this network are similar to the ones in the real data example to be discussed in Section 4. The adjacency matrix \mathbf{A} represents the network structure where the nodes are the q different time series. Element $A_{ij} = 1$ if a directed edge is drawn from node i to node j , otherwise $A_{ij} = 0$. To construct the adjacency matrix \mathbf{A} , we start (iteration $l = 0$) from a network of two randomly selected nodes that are connected with a bidirectional edge. Next, in iteration $l = 1, \dots, q - 2$, a node that is currently not in the network is randomly selected. This new node is connected to a node that is present in the network via an edge whose direction is randomly chosen. The probability

$$\pi_m^{(l-1)} = \frac{d_m^{(l-1)}}{\sum_n d_n^{(l-1)}},$$

that the new node is connected to node m depends on the degree $d_m^{(l-1)}$ of the node present in the network from iteration $l - 1$. The degree of a node equals the number of edges starting from it. Finally, we set $\mathbf{B}_1 = 0.4\mathbf{A}$ and $\mathbf{B}_2 = 0.2\mathbf{A}$.

3.2 Structure of the error terms

We consider three structures for the error covariance matrix $\boldsymbol{\Sigma}$ and its inverse $\boldsymbol{\Omega}$, see e.g. Rothman et al. (2010):

1. **Sparse Ω** : $\Sigma_{ij} = \rho^{|i-j|}$, with $\rho \in \{0.2, 0.4, 0.6, 0.8\}$. The error covariance matrix Σ is a dense matrix, whereas its inverse Ω is a band matrix.
2. **Diagonal Ω** : $\Sigma = \mathbf{I}_q$. Both the error covariance matrix and its inverse are diagonal.
3. **Dense Ω** : $\Sigma_{ij} = 0.5((|i-j|+1)^{2 \times 0.9} - 2|i-j|^{2 \times 0.9} + (|i-j|-1)^{2 \times 0.9})$. Both the error covariance matrix and its inverse have a dense structure.

3.3 Performance measures

We measure estimation accuracy by looking at the Mean Absolute Estimation Error given by

$$\text{MAEE} = \frac{1}{N} \frac{1}{p \times q} \sum_{m=1}^N \sum_{j=1}^q \sum_{i=1}^p |\widehat{b}_{ij}^{(m)} - b_{ij}|, \quad (8)$$

where $\widehat{b}_{ij}^{(m)}$ is the estimate of the ij^{th} element of \mathbf{B} in simulation run m . We take $N = 1000$ simulation runs.

We measure sparsity recognition by looking at the True Positive Rate and the True Negative Rate given by

$$\begin{aligned} \text{TPR} &= \frac{1}{N} \sum_{m=1}^N \frac{\#\{(i, j) : \widehat{b}_{ij}^{(m)} \neq 0 \text{ and } b_{ij} \neq 0\}}{\#\{(i, j) : b_{ij} \neq 0\}} \\ \text{TNR} &= \frac{1}{N} \sum_{m=1}^N \frac{\#\{(i, j) : \widehat{b}_{ij}^{(m)} = 0 \text{ and } b_{ij} = 0\}}{\#\{(i, j) : b_{ij} = 0\}}. \end{aligned}$$

TPR gives the hit rate of including an important variable, whereas the TNR gives the hit rate of excluding an unimportant variable. Both should be as large as possible for reliable variable selection.

3.4 Results

In this section, we discuss the results for the two data configurations. We show that the `GroupLasso+Cov` considerably improves the `GroupLasso` as soon as the errors are correlated.

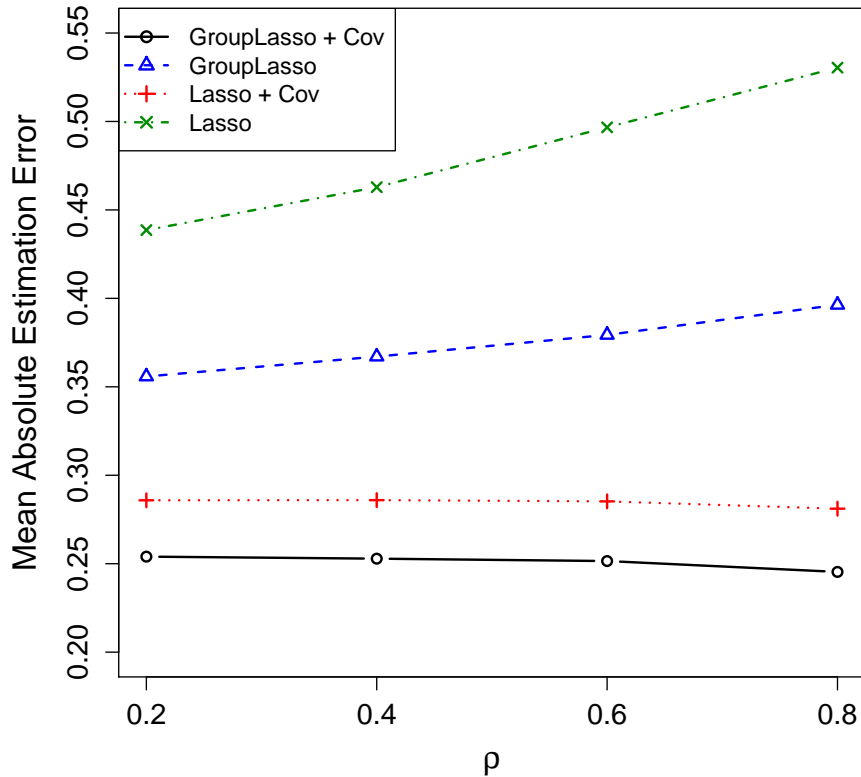


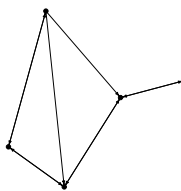
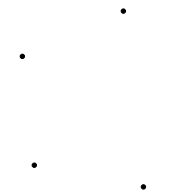
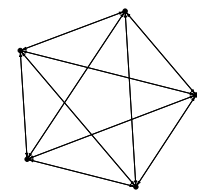
Figure 1: Multivariate regression with $q = 5$ responses, $K = 5$ categorical regressors and $n = 50$: Mean Absolute Estimation Error versus the correlation ρ , for the four considered estimators.

Categorical data. We first discuss the results for the sparse inverse error covariance structure (cfr. Section 3.2). The MAEE with $K = 5$ categorical regressors is displayed in Figure 1 for different values of the correlation ρ . Similar conclusion can be made for $K = 20$ or $K = 50$ categorical regressors and are, hence, omitted.

The **GroupLasso+Cov** substantially outperforms the **GroupLasso** for all values of the correlation ρ . The margin by which the former outperforms the latter increases when ρ increases. The **GroupLasso+Cov** achieves this improved estimation accuracy since it accounts for the error correlation whereas the **GroupLasso** does not. Besides, as expected for grouped predictors, the group lasso estimators outperform the corresponding lasso estimators.

The MAEE for all simulation designs are reported in Table 2. In line with Figure 1,

Table 2: Multivariate regression with $q = 5$ responses, $K \in \{5, 20, 50\}$ categorical regressors and $n = 50$: Mean Absolute Estimation Error, True Positive and True Negative Rate.

| | | Omega sparse | | Omega diagonal | | Omega dense | |
|----------|----------------|---|-----------|--|-----------|---|-----------|
| | |  | |  | |  | |
| | | $\rho = 0.6$ | | | | | |
| | Estimator | MAEE | TPR/TNR | MAEE | TPR/TNR | MAEE | TPR/TNR |
| $K = 5$ | GroupLasso+Cov | 0.251 | 1.00/0.62 | 0.253 | 1.00/0.61 | 0.244 | 1.00/0.67 |
| | GroupLasso | 0.379 | 1.00/0.53 | 0.349 | 1.00/0.54 | 0.394 | 1.00/0.53 |
| | Lasso+Cov | 0.285 | 0.91/0.90 | 0.286 | 0.91/0.88 | 0.282 | 0.91/0.91 |
| | Lasso | 0.497 | 0.96/0.67 | 0.303 | 0.91/0.86 | 0.374 | 0.91/0.85 |
| $K = 20$ | GroupLasso+Cov | 0.156 | 1.00/0.35 | 0.155 | 1.00/0.35 | 0.155 | 1.00/0.35 |
| | GroupLasso | 0.470 | 1.00/0.15 | 0.411 | 1.00/0.36 | 0.503 | 1.00/0.15 |
| | Lasso+Cov | 0.281 | 0.96/0.69 | 0.279 | 0.96/0.69 | 0.281 | 0.96/0.69 |
| | Lasso | 0.547 | 0.96/0.56 | 0.436 | 0.96/0.57 | 0.589 | 0.96/0.56 |
| $K = 50$ | GroupLasso+Cov | 0.196 | 1.00/0.40 | 0.196 | 1.00/0.40 | 0.196 | 1.00/0.40 |
| | GroupLasso | 0.353 | 1.00/0.35 | 0.351 | 1.00/0.35 | 0.353 | 1.00/0.35 |
| | Lasso+Cov | 0.259 | 0.89/0.73 | 0.260 | 0.89/0.73 | 0.259 | 0.89/0.73 |
| | Lasso | 0.526 | 0.89/0.71 | 0.525 | 0.89/0.71 | 0.529 | 0.89/0.71 |

GroupLasso+Cov provides a considerable improvement in MAEE over GroupLasso when the error terms are correlated, see “Omega sparse”, with $\rho = 0.6$. For reasons of brevity, we only report the results for $\rho = 0.6$. The estimation accuracy improves by more than 30%. The improvement of GroupLasso+Cov over GroupLasso becomes even larger when the number of categorical regressors K increases. A paired t -test confirms that this improvement is significant (all p -values < 0.01).

When Ω is diagonal or dense, GroupLasso+Cov also attains the best estimation accuracy. Even though Ω is not sparse in the latter setting, and our proposed estimator provides

a sparse estimate of Ω , it still provides a considerable improvement over the `GroupLasso` by exploiting the correlated error term structure. Furthermore, the `GroupLasso+Cov` also significantly outperforms both lasso estimators.

Table 2 also contains the results on the True Positive Rate and True Negative Rate. The `GroupLasso+Cov` performs very similar to the `GroupLasso`. Accounting for the error correlation mainly affects the estimation accuracy, but only to a lesser extent the sparsity recognition performance. A similar observation is made by Rothman et al. (2010). Furthermore, the group lasso estimators attain, overall, a higher true positive rate than the lasso estimators.

Time series. First consider the settings with a sparse inverse error covariance structure. The MAEE for the VAR(2) model of dimension $q = 5$ is displayed in Figure 2 for different values of ρ . We find that (i) the improvement of `GroupLasso+Cov` over `GroupLasso` is remarkable when the error terms are highly correlated, (ii) `GroupLasso+Cov` and `GroupLasso` perform similarly when the error terms are hardly correlated, (iii) the group lasso estimators perform, overall, better than the corresponding lasso estimators.

The MAEE for all simulation designs are reported in Table 3. For correlated errors (cfr. “Omega sparse” and “Omega dense”), the `GroupLasso+Cov` performs best and attains, in general, a considerably lower MAEE than the `GroupLasso`. For uncorrelated errors (“Omega diagonal”), the differences in estimation accuracy between `GroupLasso+Cov` and `GroupLasso` are less outspoken. Importantly, there is no loss in using the former compared to the latter. By sparsely estimating Ω , the absence of error correlation is accounted for.

Differences in the sparsity recognition between the estimators are less outspoken. While the estimators perform more similarly in terms of sparsity recognition, the considerable improvement in estimation accuracy attained by the `GroupLasso+Cov` gives it a clear advantage over the other estimators.

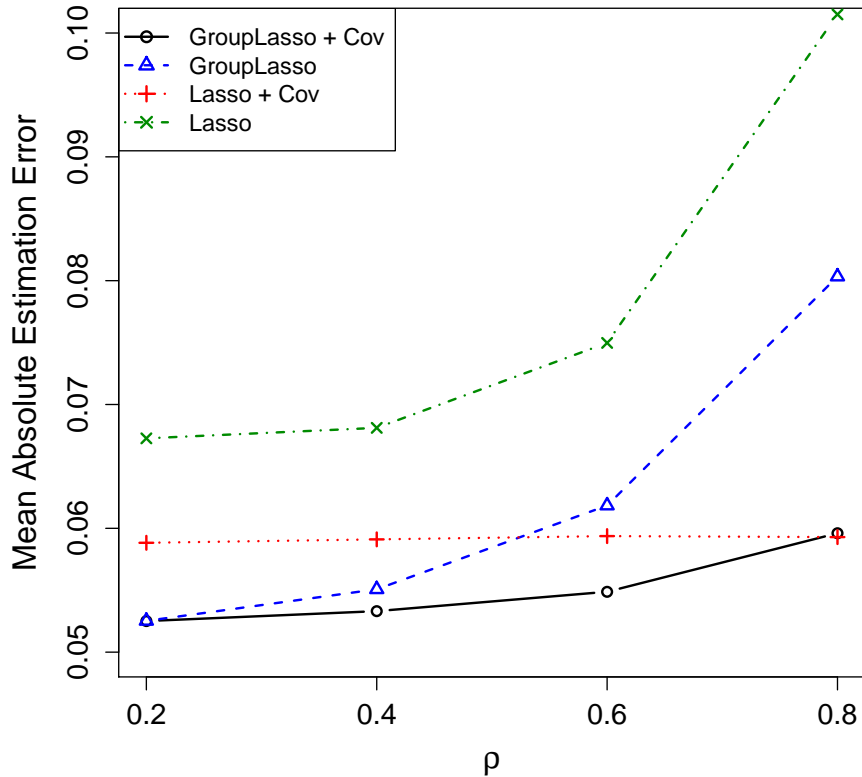


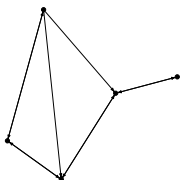

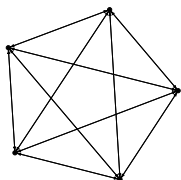
Figure 2: VAR(2) model of dimension $q = 5$ and $T = 50$: Mean Absolute Estimation Error versus the correlation ρ , for the four considered estimators.

4 Application

We consider a data set of 30 mammary gland gene expression variables of mice (Abegaz and Wit, 2013). Data are available for 18 time points, so we estimate a VAR(2) model of dimension $q = 30$, with $T = 18$. Since three samples are available, we estimate the VAR model three times.

We make an out-of-sample forecast comparison between GroupLasso+Cov, GroupLasso, Lasso+Cov, and Lasso. We use an expanding window approach. For $t = 13, \dots, T - 1$, we estimate the VAR(2) model using time points one until t and compute the one-step-ahead forecast. We compare the performance of the different estimators using the Mean Absolute

Table 3: VAR(2) of dimension $q \in \{5, 20, 50\}$ and $T = 50$: Mean Absolute Estimation Error, True Positive and True Negative Rate.

| | | Omega sparse | | Omega diagonal | | Omega dense | |
|----------|----------------|---|-----------|--|-----------|---|-----------|
| | |  | |  | |  | |
| | | $\rho = 0.6$ | | | | | |
| | Estimator | MAEE | TPR/TNR | MAEE | TPR/TNR | MAEE | TPR/TNR |
| $q = 5$ | GroupLasso+Cov | 0.055 | 0.86/0.77 | 0.053 | 0.87/0.89 | 0.058 | 0.85/0.66 |
| | GroupLasso | 0.062 | 0.87/0.75 | 0.051 | 0.87/0.89 | 0.072 | 0.86/0.64 |
| | Lasso+Cov | 0.059 | 0.79/0.62 | 0.059 | 0.80/0.69 | 0.059 | 0.78/0.55 |
| | Lasso | 0.075 | 0.54/0.92 | 0.068 | 0.49/0.97 | 0.090 | 0.55/0.86 |
| $q = 20$ | GroupLasso+Cov | 0.015 | 0.86/0.64 | 0.015 | 0.83/0.76 | 0.017 | 0.89/0.49 |
| | GroupLasso | 0.024 | 0.87/0.54 | 0.018 | 0.84/0.71 | 0.044 | 0.90/0.36 |
| | Lasso+Cov | 0.015 | 0.78/0.51 | 0.015 | 0.76/0.61 | 0.016 | 0.80/0.42 |
| | Lasso | 0.028 | 0.52/0.84 | 0.019 | 0.47/0.91 | 0.069 | 0.58/0.71 |
| $q = 50$ | GroupLasso+Cov | 0.006 | 0.68/0.92 | 0.006 | 0.67/0.95 | 0.008 | 0.73/0.80 |
| | GroupLasso | 0.007 | 0.68/0.92 | 0.006 | 0.67/0.95 | 0.019 | 0.73/0.80 |
| | Lasso+Cov | 0.006 | 0.61/0.36 | 0.006 | 0.62/0.84 | 0.007 | 0.62/0.76 |
| | Lasso | 0.007 | 0.83/0.98 | 0.006 | 0.34/0.98 | 0.027 | 0.43/0.92 |

Forecast Error

$$\text{MAFE} = \frac{1}{5} \sum_{t=13}^{T-1} \frac{1}{q} \sum_{i=1}^q |y_{t+1}^{(i)} - \hat{y}_{t+1}^{(i)}|, \quad (9)$$

where $\hat{y}_{t+1}^{(i)}$ is the estimate of the i^{th} response at time $t + 1$. We repeat this exercise three times, once for each replicate sample. Results are given in Table 4.

The **GroupLasso+Cov** attains the best forecast performance. It is closely followed by the **Lasso+Cov**. An important gain in prediction accuracy is obtained by accounting for the correlation structure of the error terms: the MAFE of the **GroupLasso+Cov** is, on average, 45% lower than the MAFE of the **GroupLasso**. Furthermore, we see from Table 4 that the

Table 4: Mean Absolute Forecast Error for the four considered estimators (rows) and three samples (column). The average MAFE, averaged over the three samples, is provided in the last column.

| Estimator | Sample 1 | Sample 2 | Sample 3 | Average |
|-----------------------|----------|----------|----------|---------|
| GroupLasso+Cov | 0.81 | 0.80 | 0.80 | 0.80 |
| GroupLasso | 1.23 | 1.38 | 1.72 | 1.44 |
| Lasso+Cov | 0.83 | 0.81 | 0.97 | 0.87 |
| Lasso | 1.51 | 1.85 | 2.37 | 1.91 |

group lasso estimators perform better than the corresponding lasso estimators.

We study the interaction between the genes that trigger transitions to the mammary gland’s main development stages. Figure 3 represents the “directed, lagged effects” (Abegaz and Wit, 2013) inferred from $\widehat{\mathbf{B}}$. We discuss the results obtained from the first sample. Results for the other two samples are similar and available from the authors upon request. The nodes in the network are the different genes. A directed edge from gene A to gene B is drawn if the **GroupLasso+Cov** indicates, by giving a non-zero estimate, that gene A has a lagged effect on gene B. The solution is very sparse: 850 out of the possible $900 = 30^2$ effects are estimated as zero. Some genes such as **GTF2A** and **TOR1B**, neither influence any other genes, nor are influenced by other genes. Other genes, such as **HSD17B** and **SAA2** are important hubs in the gene regulatory network. Previous research (Abegaz and Wit, 2013 and references therein) found these genes to play a central role in the mammary gland’s development stages.

Figure 4 represents the “contemporaneous interactions” (Abegaz and Wit, 2013) inferred from $\widehat{\mathbf{\Omega}}$. Again, the genes are the different nodes in the network. The elements of $\widehat{\mathbf{\Omega}}$ have a natural interpretation as partial correlations between the innovations (or error components) of the q equations in the VAR model. An edge is drawn between gene A and gene B if the corresponding element in the inverse error covariance matrix is estimated as non-zero. This means that the innovations of genes A and B are contemporaneously partially correlated: conditional on all other innovations, a shock in the innovation of gene A will lead to an instantaneous shock in the innovation of gene B, and vice versa. As can be seen from Figure

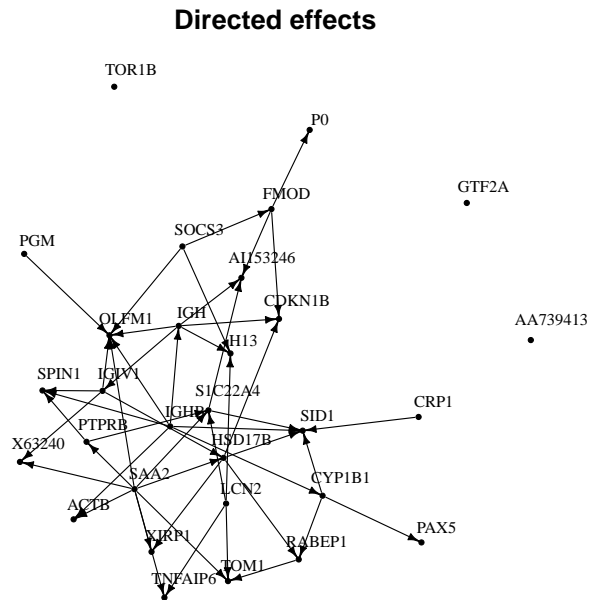


Figure 3: Directed effects: a directed edge is drawn from one gene to another if the **GroupLasso+Cov** estimator indicates, by giving a non-zero regression estimate, that the former influences the latter.

4, contemporaneous interactions are observed only between a subset of 13 gene innovations, indicated by the rectangle. An important advantage of the sparse estimator is that the main interactions in the large gene regulatory network are highlighted. Out of the possible 435 interactions, only 32 are estimated as non-zero. As such, the researcher can concentrate on these results to further deepen our knowledge into the interactions at play in the development stages of the mammary gland.

Acknowledgments. The authors gratefully acknowledge financial support from the FWO (Research Foundation Flanders, contract number 11N9913N).

References

Abegaz, F. and Wit, E. (2013), “Sparse time series chain graphical models for reconstructing genetic networks,” *Biostatistics*, 14(3), 586–599.

Contemporaneous interactions

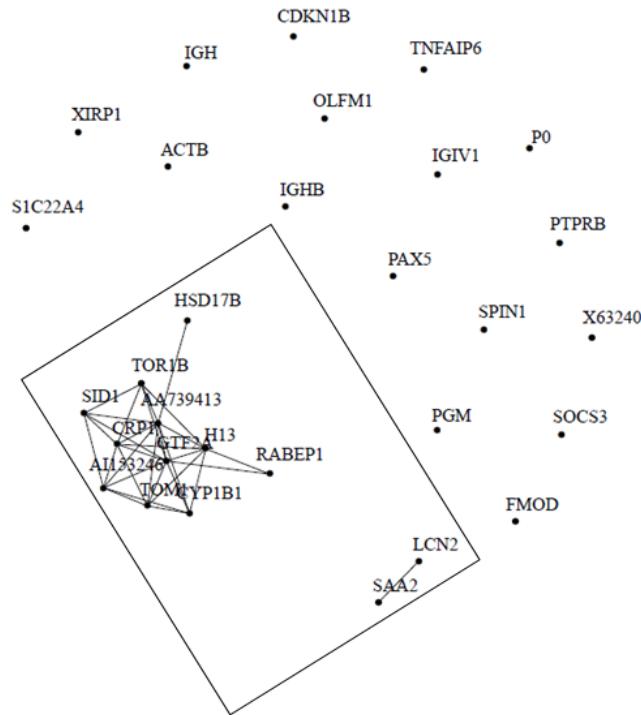


Figure 4: Contemporaneous interactions: an undirected edge is drawn between two genes if the **GroupLasso+Cov** estimator indicates, by giving a non-zero estimate in Ω , that the innovations are partially correlated. Contemporaneous interactions are observed for only a subset of 13 genes, as indicated by the rectangle.

Alfons, A.; Croux, C. and Gelper, S. (2015), “Robust groupwise least angle regression,” *Computational Statistics and Data Analysis*, Available online 17 February 2015.

Bühlmann, P. and van de Geer, S. (2011), *Statistics for high-dimensional data. Methods, Theory and Applications*, Springer.

Friedman, J.; Hastie, T. and Tibshirani, R. (2007), “Pathwise coordinate optimization,” *The Annals of Applied Statistics*, 1(2), 302–332.

— (2008), “Sparse inverse covariance estimation with the graphical lasso,” *Biostatistics*, 9(3), 432–441.

Fujita, A.; Sato, J.; Garay-Malpartida, H.; Yamaguchi, R.; Miyano, S.; Sogayar, M. and Ferreira, C. (2007), “Modeling gene expression regulatory networks with the sparse vector

- autoregressive model,” *BMC Systems Biology*, 1, No. 39.
- Li, Y.; Nan, B. and Zhu, J. (2015), “Multivariate sparse group lasso for the multivariate multiple linear regression with an arbitrary group structure,” *Biometrics*, 71(2), 354–363.
- Meier, L.; van de Geer, S. and Bühlmann, P. (2008), “The group lasso for logistic regression,” *Journal of the Royal Statistical Society, Series B*, 70, 53–71.
- Rothman, A.; Levina, E. and Zhu, J. (2010), “Sparse multivariate regression with covariance estimation,” *Journal of Computational and Graphical Statistics*, 19, 947–962.
- Simon, N.; Friedman, J.; Hastie, T. and Tibshirani, R. (2013), “A sparse-group lasso,” *Journal of Computational and Graphical Statistics*, 22(2), 231–245.
- Wang, H. and Leng, C. (2008), “A note on adaptive group lasso,” *Computational Statistics and Data Analysis*, 52(12), 5277–5286.
- Yuan, M. and Lin, Y. (2006), “Model selection and estimation in regression with grouped variables,” *Journal of the Royal Statistical Society Series B*, 68, 49–67.
- (2007), “Model selection and estimation in the Gaussian graphical model,” *Biometrika*, 94(1), 19–35.

FACULTY OF ECONOMICS AND BUSINESS
Naamsestraat 69 bus 3500
3000 LEUVEN, BELGIË
tel. + 32 16 32 66 12
fax + 32 16 32 67 91
info@econ.kuleuven.be
www.econ.kuleuven.be

