



**HAL**  
open science

# On the robustness of the generalized fused lasso to prior specifications

Vivian Viallon, Sophie Lambert-Lacroix, Holger Höfling, Franck Picard

► **To cite this version:**

Vivian Viallon, Sophie Lambert-Lacroix, Holger Höfling, Franck Picard. On the robustness of the generalized fused lasso to prior specifications. *Statistics and Computing*, 2016, 26 (1), pp.285-301. 10.1007/s11222-014-9497-6 . hal-01067903v2

**HAL Id: hal-01067903**

**<https://hal.science/hal-01067903v2>**

Submitted on 4 May 2017

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Noname manuscript No.  
(will be inserted by the editor)

# On the robustness of the Generalized Fused Lasso to prior specifications

Vivian Viallon · Sophie Lambert-Lacroix · Hölger Hoefling · Franck Picard

Corresponding author :

Vivian Viallon :

IFSTTAR, UMRESTTE, 69675 Bron, France

e-mail: vivian.viallon@univ-lyon1.fr

Received: date / Accepted: date

**Abstract** Using networks as *prior* knowledge to guide model selection is a way to reach structured sparsity. In particular, the fused lasso that was originally designed to penalize differences of coefficients corresponding to successive features has been generalized to handle features whose effects are structured according to a given network. As any *prior* information, the network provided in the penalty may contain misleading edges that connect coefficients whose difference is not zero, and the extent to which the performance of the method depend on the suitability of the graph has never been clearly assessed. In this work we investigate the theoretical and empirical properties of the adaptive generalized fused lasso in the context of generalized linear models. In the fixed  $p$  setting, we show that, asymptotically, adding misleading edges in the graph does not prevent the adaptive generalized fused lasso from enjoying asymptotic oracle properties, while forgetting suitable edges can be more problematic. These theoretical results are complemented by an extensive simulation study that assesses the robustness of the adaptive generalized fused lasso against misspecification of the network as well as its applicability when theoretical coefficients are not exactly equal. Our contribution is also to evaluate the applicability of the generalized fused lasso for the joint modeling of multiple sparse regression functions. Illustrations are provided on two real data examples.

**Keywords** lasso, generalized linear models, joint modeling, model selection

## 1 Introduction

Network data have now become standard in many fields of statistical applications to describe interactions or relationships between observations and/or variables. In molecular biology for instance, protein-protein interaction networks describe physical interactions between proteins (4). In web analysis studies, networks are extracted from structured databases and Web contents to describe hidden structures (10). Network data have become so common that they now constitute some *prior* knowledge for downstream statistical analysis: two proteins of the same biological pathway are likely to share similar effects on the response to a treatment or on disease development. Consequently, statistical methods, like regression and model selection have recently focused on structured sparsity. In addition to the classical sparsity assumption (under which only a small fraction of the variables are relevant), these methods work under the assumption that two connected covariates in the network may share similar effects on the response variable. Consequently the objective of structured sparsity is twofold: improve model selection by using some *prior* knowledge on the structure; increase prediction performance by effective dimensionality reduction based on the *prior* knowledge that several covariates may share the exact same coefficient.

Most methods proposed so far use a penalized version of the log-likelihood based on some structured sparsity-inducing penalty. The fused lasso of Tibshirani (23) is one particular example: in addition to the  $\ell_1$ -norm penalty of the lasso (22), the fused lasso penalizes the  $\ell_1$ -norm of the vector of successive differences. It is therefore especially adapted for smoothing, when covariates are ordered and are likely to

V. Viallon

Université de Lyon, F-69622, Lyon, France; Université Lyon 1, UMRESTTE, F-69373 Lyon; IFSTTAR, UMRESTTE, F-69675 Bron.

E-mail: vivian.viallon@univ-lyon1.fr

S. Lambert-Lacroix

UMR 5525 UJF-Grenoble 1 / CNRS / UPMF / TIMC-IMAG, Grenoble, F-38041, France.

H. Hoefling

Novartis Pharma, Basel, Switzerland.

F. Picard

Laboratoire de Biométrie et Biologie Évolutive, UMR CNRS 5558 Univ. Lyon 1, F-69622 Villeurbanne, France.

share similar effects with their direct neighbor. It has further been generalized to handle more complex structure among feature effects, in particular networks of features (11). The network is modeled as a graph with vertices standing for the  $p$  coefficients of the model, and with a set of edges. It is used in the penalty as *prior* information to penalize the absolute value of the difference of connected coefficients, leading to the generalized fused lasso. Interestingly, structure in the vector of coefficients also naturally arises when jointly estimating multiple regression models. For instance, when data are collected from distinct strata (in epidemiology, these strata can be defined by crossing age, gender, ethnicity), models defined on each stratum are expected to share similarities. Consequently, structured-sparsity and in particular the generalized fused lasso can be used in this context to enforce some structure while estimating the different models (5). In this context, the graph to be used in the penalty is usually provided by the design of the study itself (see Section 2.2 below).

As any *prior* information, the underlying graph can be more or less informative. For instance, the clustered lasso (20) was proposed when only the existence of a structure is assumed but no particular knowledge allows for its precise description. Its main step involves a penalty based on the  $\ell_1$ -norm of the vector of all the  $p(p-1)/2$  differences among the parameter values. This strategy corresponds to the generalized fused lasso with the graph set to a clique that connects all coefficients. When penalizing all differences, it is very likely that some differences are unnecessarily penalized, which raises the question of the method robustness to graph misspecification. Interestingly, any structured-sparsity approach is concerned by this robustness property, but this question has never been thoroughly investigated (1).

In this work we focus on adaptive generalized fused lasso estimates in the context of generalized linear models and show how the framework of joint regressions can be recast as a generalized fused lasso problem. In Section 3, we prove that adaptive generalized fused lasso estimators enjoy asymptotic oracle properties in the fixed  $p$  setting. Our results extend those obtained in the case of clique-based strategies (20; 19), and in the case of joint linear regressions (5). In particular, we observe that only adaptive versions of the generalized fused lasso enjoy asymptotic oracle properties (i.e., such that, as  $n$  grows to infinity, the correct support is recovered with probability tending to one and estimates of non-zero coefficients perform as well as if the true underlying model were given in advance). In a further step we investigate the empirical benefits of using an  $\ell_1$ -based fusion penalty on support recovery and prediction as compared with other penalization strategies, under logistic models. We assess the robustness of the generalized fused lasso to graph misspecification as well as its performance when true coefficients are not exactly equal. We also illustrate the ben-

efits of using adaptive weights and/or relaxation with generalized fused lasso estimates (Section 4). The performance of the generalized fused lasso in the context of joint logistic regression models is also empirically evaluated. Finally the application of the generalized fused lasso is illustrated on two data sets in Sections 5 and 6.

## 2 The adaptive generalized fused lasso in generalized linear models

### 2.1 Models, loss functions and penalty

We consider the generalized linear models framework (13) with  $Y_i$  the *response* variable  $i = 1, \dots, n$  and  $\mathbf{x}_i = (x_{i1}, \dots, x_{ip})^T$  a  $p$ -dimensional vector of features. We further set  $\mathbf{z}_i = (1, \mathbf{x}_i^T)^T$ , and we consider the *fixed design* case with  $\sum_{i=1}^n x_{ij} = 0$ . For generalized linear models the distribution of the response variable is given by

$$f(y_i, \boldsymbol{\beta}^*, \phi) = \exp\left(\frac{y_i \eta_i - b(\eta_i)}{a(\phi)} + c(y_i, \phi)\right),$$

where  $\phi$  is a dispersion parameter and functions  $b(\cdot)$ ,  $a(\cdot)$  and  $c(\cdot, \cdot)$  are known. The linear predictor  $\eta_i$  is given by  $\mathbf{z}_i^T \boldsymbol{\beta}^*$  where  $\boldsymbol{\beta}^* = (\beta_0^*, \boldsymbol{\beta}_{\setminus 0}^*)^T \in \mathbb{R}^{p+1}$  stands for the vector of coefficients, with  $\beta_0^*$  the intercept parameter and  $\boldsymbol{\beta}_{\setminus 0}^* = (\beta_1^*, \dots, \beta_p^*)$ . The mean  $\mu_i = \mathbb{E}(Y_i)$  is related to the linear predictor via the link function  $g: g(\mu_i) = \eta_i$ . Here we consider the canonical link function. Estimation of the parameter vector  $\boldsymbol{\beta}^*$  is usually performed by the maximum likelihood method. It consists in minimizing  $J$ , given by  $J(\boldsymbol{\beta}) = -\sum_{i=1}^n \log f(y_i, \boldsymbol{\beta}, \phi)$ , with respect to  $\boldsymbol{\beta}$ . In the simulation studies and the applications below we focus on the logistic model for which  $Y_i \in \{0, 1\}$ ,  $a(\phi) = 1$ ,  $b(x) = \log(1 + \exp(x))$  and  $c \equiv 0$ . Under logistic models, the mean and the linear predictor are related by  $\mu_i = 1/(1 + \exp(-\mathbf{z}_i^T \boldsymbol{\beta}^*)) = g^{-1}(\eta_i)$ .

As mentioned above, we further focus on the generalized fused lasso (11). Consider a graph  $G = (V, E)$ , with node set  $V = \{1, \dots, p\}$  that corresponds to the coefficient indices in  $\boldsymbol{\beta}_{\setminus 0}$ , and edge set  $E$  that corresponds to pairs of connected coefficient indices  $(j, \ell)$  with  $j > \ell$ . The graph  $G$  that is used in the penalty is fixed and represents some *prior* knowledge, given by an expert. The adaptive generalized fused lasso penalty consists in penalizing all coefficients along with all coefficient differences for which an edge exists in  $G$ :

$$\text{pen}_{\text{Ada}}(\boldsymbol{\beta}; G, \mathbf{w}) = \lambda_n^{(1)} \sum_{j \in V} w_j^{(1)} |\beta_j| + \lambda_n^{(2)} \sum_{(j, \ell) \in E} w_{j\ell}^{(2)} |\beta_j - \beta_\ell|.$$

In the fixed  $p$  case considered here, and following the idea of the adaptive lasso (25), adaptive weights  $w_j^{(1)}$  and  $w_{j\ell}^{(2)}$  are based on initial Maximum-Likelihood estimates  $\tilde{\boldsymbol{\beta}}$

of  $\beta^*$ . More precisely, for some  $\gamma > 0$ , we set  $w_j^{(1)} = |\tilde{\beta}_j|^{-\gamma}$  and  $w_{j\ell}^{(2)} = |\tilde{\beta}_j - \tilde{\beta}_\ell|^{-\gamma}$ . The rationale is to penalize more heavily coefficients (or differences of coefficients) when their initial estimates are small. A typical value (that we use) for  $\gamma$  is 1. The adaptive generalized fused lasso criterion  $Q$  is then simply defined, for given graph  $G$  and weights  $\mathbf{w}$ , as

$$Q(\beta) = J(\beta) + \text{pen}_{\text{Ada}}(\beta; G, \mathbf{w})$$

## 2.2 Application of the generalized fused penalty to joint modeling

Interestingly, in some situations the graph to be used in the penalty is provided by the design of the study itself. This is notably the case when jointly estimating multiple sparse regression models. The joint modeling framework described here has some connections with multi-task learning (12; 3). More specifically we consider the very common case of data collected from distinct *strata*, which often arises in epidemiology where each stratum can be defined by crossing gender, age and ethnicity for instance. The design is structured according to a given **vector of categorical variables**  $(C_1, \dots, C_n)$ , **each variable taking values in  $\{1, \dots, C\}$** , with  $C \geq 1$  the total number of strata. Let  $n_c$  be the number of observations falling into stratum  $c$  (so that  $n = \sum_c n_c$ ). Under generalized linear models, we would have  $g(\mu_i) = \mathbf{z}_i^T \beta_{c_i}^*$ , for  $i = 1, \dots, n$ , where  $\beta_c^*$  denotes the vector of parameters for stratum  $c$ . The purpose of the analysis is to determine whether the distribution of the response varies across strata, *i.e.* to detect which coefficients of  $\beta_c^*$  do vary with  $c$ . Constructing independent (possibly sparse) models for each stratum would not take advantage of the common structure, while constructing a single model for the whole data set would mask the differences. Alternatively, the generalized fused lasso can be used to couple estimations obtained from each stratum, encouraging them to share some common structure (5). More precisely, the following penalty can be used:

$$\sum_{c=1}^C \left\{ \lambda_n^{(1)} \sum_{j=1}^p w_j^{(1)} |\beta_{c,j}| \right\} + \lambda_n^{(2)} \sum_{j=0}^p \sum_{c_1 > c_2} w_{c_1, c_2, j}^{(2)} |\beta_{c_1, j} - \beta_{c_2, j}|,$$

where  $w_j^{(1)}$  and  $w_{c_1, c_2, j}^{(2)}$  are appropriate adaptive weights. Parameter  $\lambda_n^{(2)}$  governs the amount of shrinkage for differences between strata: if null, this penalty resumes to  $C$  independent lasso penalties. If positive, the fused part of the penalty encourages coefficients  $\beta_{c_1, j}$  and  $\beta_{c_2, j}$  to be at least close to each other (that is, the  $j^{\text{th}}$  coefficient in strata  $c_1$

and  $c_2$  respectively). In the general context described here, it can be shown that the problem reduces to an adaptive generalized fused lasso where the graph is composed by  $p + 1$  cliques of size  $C$  and the  $j$ -th clique connects  $\beta_{1, j}, \dots, \beta_{C, j}$  all together (see the Supplementary Material for more details).

## 3 Theoretical results

We study the asymptotic properties of the adaptive generalized fused lasso estimator in generalized linear models. Before stating our results some notations and assumptions are needed. Let  $\mathcal{A} = \{1 \leq j \leq p, \beta_j^* \neq 0\}$  be the *support* of  $\beta_{\setminus 0}^*$  and  $p_0 = |\mathcal{A}|$  its cardinality. Further consider the set

$$\mathcal{B} = \{(j, \ell) \in E, \beta_j^* \neq 0 \text{ and } \beta_j^* = \beta_\ell^*\} \subset \mathcal{A} \times \mathcal{A}.$$

We denote by  $\mathcal{I}(\beta)$  the empirical Fisher's matrix of size  $(p+1) \times (p+1)$ . For future use, observe that  $\mathcal{I}(\beta^*) = \mathbf{Z}^T \mathbf{D} \mathbf{Z}$ , where  $\mathbf{D}$  denotes an  $n \times n$  diagonal matrix. For instance under logistic regression models, we have  $D_{ii} = \mu_i(1 - \mu_i)$ . For any  $\delta \geq 0$ , we further denote by  $N_n(\delta)$  the neighborhood of  $\beta^*$  defined by

$$N_n(\delta) = \left\{ \beta \in \mathbb{R}^{p+1} / \left\| \left[ \mathcal{I}(\beta)^{-\frac{1}{2}} \right]^T (\beta - \beta^*) \right\| \leq \delta \right\}.$$

We will work under the following conditions:

- **AL1**  $\mathcal{I}(\beta^*)/n$  converges to  $\mathbf{C}$  where  $\mathbf{C}$  is a positive definite  $(p+1) \times (p+1)$  matrix.
- **AL2** As  $n$  goes to  $\infty$ ,

$$\max_{\beta \in N_n(\delta)} \left\| \mathcal{I}(\beta)^{-\frac{1}{2}} \mathcal{I}(\beta^*)^T \left[ \mathcal{I}(\beta)^{-\frac{1}{2}} \right]^T - \mathbf{I}_{p+1} \right\| \rightarrow 0.$$

Assumptions **AL1** and **AL2** are standard when working under generalized linear models (13). Assumption **AL1** is similar to the one used to study the fused lasso in the Gaussian context (23). Let us remark that under **AL1**, the minimization of criterion  $Q$  defined in Section 2.1 corresponds, for  $n$  large enough, to a strictly convex optimization problem, and thus is not concerned by the issue of multiple local minima.

In Theorem 1 below, we generalize the first theorem of (23) to the case of generalized fused lasso, in generalized linear models. This result establishes the root- $n$  consistency of non-adaptive generalized fused lasso estimates. However, it also implies that when  $\lambda_n^{(m)} = O(\sqrt{n})$ , for  $m = 1, 2$ , the support of  $\beta^*$  can not be recovered with high probability by non-adaptive fused lasso estimates, as stated in Proposition 1 below. Proofs are given in the Appendix.

**Theorem 1** *Let  $\hat{\beta}$  be the minimizer of criterion  $Q$  defined in Section 2.1 with  $w_j^{(1)} = 1$  and  $w_{j\ell}^{(2)} = 1$  for all  $j, \ell$ . If*

$\lambda_n^{(m)}/\sqrt{n} \rightarrow \lambda_0^{(m)} \geq 0$  ( $m = 1, 2$ ), then under assumptions **AL1-2**,  $\sqrt{n}(\widehat{\beta} - \beta^*) \rightarrow_d \arg \min(\mathcal{V})$ , where  $\mathcal{V}$  is the function defined, for  $\mathbf{u} = (u_0, \dots, u_p) \in \mathbb{R}^{p+1}$ , as

$$\begin{aligned} \mathcal{V}(\mathbf{u}) &= \mathbf{u}^T \mathbf{W} + \frac{1}{2} \mathbf{u}^T \mathbf{C} \mathbf{u} \\ &+ \lambda_0^{(1)} \sum_{j=1}^p \{u_j \text{sign}(\beta_j^*) \mathbb{I}(\beta_j^* \neq 0) + |u_j| \mathbb{I}(\beta_j^* = 0)\} \\ &+ \lambda_0^{(2)} \sum_{(j,\ell) \in E} \{(u_j - u_\ell) \text{sign}(\beta_j^* - \beta_\ell^*) \mathbb{I}(\beta_j^* \neq \beta_\ell^*) \\ &+ |u_j - u_\ell| \mathbb{I}(\beta_j^* = \beta_\ell^*)\}. \end{aligned}$$

Above,  $\mathbf{W}$  has an  $\mathcal{N}(\mathbf{0}_{p+1}, \mathbf{C})$  distribution.

**Proposition 1** Under the assumptions of Theorem 1, and setting  $\widetilde{\mathcal{A}}_n = \{1 \leq j \leq p, \widehat{\beta}_j \neq 0\}$ , we have

$$\limsup_n \mathbb{P}(\widetilde{\mathcal{A}}_n = \mathcal{A}) \leq c < 1,$$

where  $c$  is a constant depending on the true model.

Now we show that for appropriate choices of  $\lambda_n^{(m)} = O(\sqrt{n})$  for  $m = 1, 2$ , the adaptive generalized fused lasso estimator  $\widehat{\beta}^{ad}$ , defined as the minimizer of criterion  $Q$  in Section 2.1, enjoys asymptotic oracle properties, contrasting with its non-adaptive counterpart. Some more notations are needed before stating our result: in particular, the number  $s_0$  of distinct non-zero values in  $\beta_{\setminus 0}^*$  “supported” by  $G$  needs to be precisely defined ( $s_0$  can be seen as the theoretical model complexity “supported” by  $G$ ). To this end, first observe that  $\mathcal{A} \subseteq V$  and  $\mathcal{B} = \{(j, \ell) \in E : \beta_j^* \beta_\ell^* \neq 0, \beta_j^* = \beta_\ell^*\} \subseteq E$ . Then consider the graph  $G_{\mathcal{B}} = (\mathcal{A}, \mathcal{B})$  and denote by  $s_0$  the number of its connected components (e.g., in the particular case where  $G$  is a chain graph,  $s_0$  is the number of segments consisting of non-zero and equal coefficients). Observe that  $d_0 \leq s_0 \leq p_0$ , where  $p_0 = |\mathcal{A}|$  is the number of non-zero coefficients in  $\beta_{\setminus 0}^*$  and  $d_0$  is the number of distinct non-zero values in  $\beta_{\setminus 0}^*$ . We actually have  $s_0 = p_0$  if and only if  $(\beta_j^* = \beta_\ell^* \neq 0 \Rightarrow (j, \ell) \notin E)$ , whereas  $s_0 = d_0$  if and only if for all  $(j, \ell)$  such that  $\beta_j^* = \beta_\ell^*$ ,  $j$  and  $\ell$  belong to the same connected component of  $G_{\mathcal{B}}$ . Now denote by  $\mathcal{A}_1, \dots, \mathcal{A}_{s_0}$  the sets of vertices of each connected components of  $G_{\mathcal{B}}$ . Of course, we have  $\mathcal{A} = \bigcup_{s=1}^{s_0} \mathcal{A}_s$ . Further set  $j_s = \min\{\mathcal{A}_s\}$  for  $s = 1, \dots, s_0$ . Now we can define  $\beta_{\mathcal{B}}^* = (\beta_0^*, \beta_{j_1}^*, \dots, \beta_{j_{s_0}}^*)^T$ , which is composed by the intercept and the  $s_0$  distinct non-zero values of  $\beta_{\setminus 0}^*$  supported by  $G$ ; we further set  $\widehat{\beta}_{\mathcal{B}}^{ad}$  its estimate. Now denote by  $\mathbf{X}_{\mathcal{B}}$  the matrix of size  $n \times s_0$ , whose  $s$ -th column is  $X_{\mathcal{B}_s} = \sum_{j \in \mathcal{A}_s} X_j$ , where  $X_j$  is the  $j$ -th column of  $\mathbf{X}$ . Finally set  $\mathbf{Z}_{\mathcal{B}} = (\mathbf{1}_n, \mathbf{X}_{\mathcal{B}})$  and denote by  $\mathbf{C}_{\mathcal{B}}$  the  $(s_0 + 1) \times (s_0 + 1)$  positive definite matrix that is defined as the limit, as  $n \rightarrow \infty$ , of  $\mathcal{I}(\beta_{\mathcal{B}}^*)/n$ , where  $\mathcal{I}(\beta_{\mathcal{B}}^*) = \mathbf{Z}_{\mathcal{B}}^T \mathbf{D} \mathbf{Z}_{\mathcal{B}}$ . Finally introduce  $\mathcal{A}_n = \{1 \leq j \leq p, \widehat{\beta}_j^{ad} \neq 0\}$  and

$\mathcal{B}_n = \{(j, \ell) \in E, \widehat{\beta}_j^{ad} \neq 0 \text{ and } \widehat{\beta}_j^{ad} = \widehat{\beta}_\ell^{ad}\}$ . We have now all the ingredients to state our main result, whose proof is given in the Appendix (see Section 8.3).

**Theorem 2** If  $\lambda_n^{(m)}/\sqrt{n} \rightarrow 0$  and  $\lambda_n^{(m)} n^{(\gamma-1)/2} \rightarrow \infty$ ,  $m = 1, 2$ , then, under assumptions **AL1-2**, the adaptive generalized fused lasso estimator satisfies the following properties:

1. Consistency in variable selection:  $\mathbb{P}[\mathcal{A}_n = \mathcal{A}] \rightarrow 1$  and  $\mathbb{P}[\mathcal{B}_n = \mathcal{B}] \rightarrow 1$  as  $n \rightarrow +\infty$ .
2. Asymptotic normality:  $\sqrt{n}(\widehat{\beta}_{\mathcal{B}}^{ad} - \beta_{\mathcal{B}}^*) \rightarrow_d \mathcal{N}(\mathbf{0}_{s_0+1}, \mathbf{C}_{\mathcal{B}}^{-1})$ .

Interestingly this result allows us to compare the asymptotic theoretical performance of methods that use a rough structured-sparsity prior (ie. clique-based methods (20; 19)) with graph-based structured sparsity. Observing that  $\mathcal{I}(\beta_{\mathcal{B}}^*)$  is the information matrix of the true submodel as soon as  $s_0 = d_0$ , Theorem 2 states that the estimator  $\widehat{\beta}_{\mathcal{B}}^{ad}$  is asymptotically efficient as soon as  $s_0 = d_0$  ( $p$  is fixed and  $n \rightarrow \infty$ ), which is notably the case for clique-based methods. Our theoretical results state that, asymptotically, adding edges in the graph between coefficients with theoretical different values does not diminish the adaptive generalized fused lasso performance. These results being asymptotic, we evaluate the finite sample properties of the generalized fused lasso in the simulation study, with an emphasis on its robustness to graph mis-specification. In other respect, results presented in the framework of joint linear models (5) readily follow from Theorem 2 above.

## 4 Simulation study

We perform an extensive simulation study to compare the performance of the generalized fused lasso with other penalized strategies in two contexts: (i) when the variables lie on a graph and the support is densely connected on the graph and (ii) in the joint modeling framework. In (i), our main objective is twofold: to study the impact of a graph misspecification on the generalized fused lasso performance and also to study the robustness of the generalized fused lasso when non-null connected coefficients do not share the exact same value. In (ii), the graph being “given” by the design of the study, we will study how the performance of the generalized fused lasso vary with the level of heterogeneity across strata.

### 4.1 Simulation framework, implementation and evaluation

Our results being in the fixed design framework, we set  $N$ , the maximal sample size considered in a given scenario, and we generate  $N$  *i.i.d.* predictors  $\mathbf{x}_i \in \mathbb{R}^p$ ,  $i = 1, \dots, N$ , from a  $\mathcal{N}(\mathbf{0}_p, \mathbf{C}_{\text{AR}(1)})$  distribution, where  $\mathbf{C}_{\text{AR}(1)} = (\rho^{|i-j|})/16$

( $i = 1 \dots p, j = 1 \dots p$ ). Then, given  $n \leq N$ ,  $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^T$  and  $\beta^*$  (with  $\beta_0^* = 0$ ), the vector of labels is generated such that  $Y_i \sim \mathcal{B}(\mu_i)$ , with  $\mu_i$  defined as in Section 2.1. Unless otherwise stated, the  $p_0$  non-null coefficients of  $\beta_{\setminus 0}^*$  are all set to a common value  $\beta^*$ , that we make vary in  $\{\log(1.1), \log(2), \log(4), \log(8), \log(12)\}$ . Fifty replicates are considered for each configuration.

The adaptive generalized fused lasso is solved with the coordinate-wise optimization algorithm (11) with tuning parameters  $(\lambda_1, \lambda_2)$  selected using the BIC, which is standard when  $n \gg p$  (degrees of freedom are set to the number of distinct non-null coefficients in the estimated model). Because using shrunk estimates within the BIC can result in severe over-fitting, we also propose an extension of the relaxed lasso (15) to the generalized (adaptive) fused lasso (see Section 2 in the Supplementary Material).

Prediction accuracy is assessed using an independent test sample of  $N$  observations  $(\mathbf{z}_i^{(0)}, Y_i^{(0)})$  and computing  $\text{Acc. Pred} = (1/N) \sum_{i=1}^N \mathbb{I}(Y_i^{(0)} = \hat{Y}_i^{(0)})$ , with  $\hat{Y}_i^{(0)}$  equal to  $\mathbb{I}(1/[1 + \exp(-\mathbf{z}_i^{(0)T} \hat{\beta})] > 0.5)$ . To assess accuracy on support recovery we define  $\text{Acc.A} = (|\bar{\mathcal{A}} \cap \bar{\mathcal{A}}_n| + |\mathcal{A} \cap \mathcal{A}_n|)/p$  where  $\bar{\mathcal{A}}$  (resp.  $\bar{\mathcal{A}}_n$ ) is the set of null coefficients in  $\beta^*$  (resp.  $\hat{\beta}$ ). Moreover, to evaluate the performance regarding the classification of pairs of coefficients, we focus on  $\text{Acc.B} = |\{(j, \ell) \in \mathcal{B} : \hat{\beta}_j = \hat{\beta}_\ell\}|/|\mathcal{B}|$ . Because the *lasso* does not encourage equality among non-null coefficients, its  $\text{Acc.B}$  only reflects the proportion of pairs in  $\mathcal{B}$  whose elements are both put to zero by the *lasso*. Consequently, the *lasso* should show poor results regarding this criterion (especially when  $\beta^*$  is large). In the joint modeling framework,  $\text{Acc.B}$  assesses the capacity of the methods to detect homogeneity across strata, and could be interpreted as a type-1 error of an interaction test.

In the following simulation study we explore the properties of the “raw” generalized fused lasso (without adaptive weights nor relaxation), the adaptive generalized fused lasso, the relaxed generalized fused lasso, and of the relaxed adaptive generalized fused lasso. Comparisons are made using the relaxed adaptive *lasso* as a reference that does not account for any structure among feature effects. The sparse group lasso (12), which includes the lasso penalty and the group lasso penalty, and an extension of the elastic-net (21) were also tested, but their performance were similar to those achieved by the lasso (see Figures 1 and 6 in the Supplementary Material).

#### 4.2 Performance of adaptive generalized fused lasso estimates and influence of the provided graph

We set  $p = 24$  and sample  $\mathbf{x}_i$  as describe above with  $\rho = -0.39$  (25). Our theoretical results being asymptotic in  $n$  we consider cases where  $n/p \in \{1, 5, 10, 50\}$ . We also ex-

plore different degrees of sparsity with  $p_0 \in \{12, 8, 3\}$ . To study the robustness of selection method to a graph misspecification, we generate graphs with varying suitabilities such that equal (resp. non-equal) coefficients are connected with probability  $\theta$  (resp.  $1 - \theta$ ). When  $\theta$  increases fewer edges connect distinct coefficients, the most [resp. least] favorable configuration for generalized fused estimates being when  $\theta = 1$  [resp.  $\theta = 0$ ]. The most suitable graph in the penalty correspond to the graph that connects equal coefficients. When no prior information is available on graph  $G$ , a strategy can be to use no graph (with the lasso) or to use a graph that connects all coefficients (clique-graph). This latter option was also considered to compare the generalized fused lasso with clique-based methods (19; 20). We mention that edges of the graph are fixed across replicates.

In the following we choose the adaptive lasso as a reference that does not depend on any graph. In the best-case scenario (perfectly suited graph,  $\theta = 1$ ), all graph-based method are more accurate than the lasso for support recovery, which reflects a cooperative effect on support recovery that is characteristic of  $\ell_1$ -based fused penalties (Figure 1). We also start by mentioning that all fused penalties are more accurate than the lasso for prediction, without much difference among them (Figure 1). Then the sensitivity to graph misspecification depends on the penalty: the “raw” fused lasso is very sensitive since its performance drop as soon as  $\theta \leq 0.8$  (or with the clique graph). Interestingly, using adaptive weights and/or relaxation increases robustness to graph misspecifications (compare the first and second lines of Figure 2 for instance), as performance decrease for  $\theta \leq 0.4$  only. Note that the adaptive and the relaxed adaptive generalized fused lasso perform better than (or at least similarly to) the lasso except when  $\theta = 0$ . Averaged accuracies on support recovery show that the *relaxation* of the adaptive generalized fused lasso does not necessarily increase its accuracy (see Table 1, and compare the first and fourth lines of Figure 2 when  $p_0 = 3$  for instance). Then when varying the  $n/p$  ratio (Supplementary Material, Figure 3), adaptive and/or relaxed strategies remain more accurate than the “raw” generalized fused lasso except when the graph is perfectly suited (i.e.,  $\theta = 1$ , which is very unlikely in practice), or in situations where maximum likelihood estimates may lack in precision ( $n/p \simeq 1$ ).

In our simulation design all non-null coefficients were set to a unique value so that the task was to detect a group of  $p_0$  non-null parameters among a set of  $p$ . This was a way to study the effect of varying  $p_0$  on the robustness to a graph misspecification. Then, as expected, all graph-based fused methods perform better when the graph is highly suitable and/or  $p_0$  is high (as compared to the lasso, Figure 2). Taking the example of the adaptive generalized fused lasso (first line in Figure 2), the method performs better than or at least

similarly to the lasso for  $\theta \geq 0.4$  when  $p_0 = 12$ , for  $\theta \geq 0.8$  when  $p_0 = 8$  and only for  $\theta = 1$  when  $p_0 = 3$ .

Next, focusing on the strategy that consists in penalizing all possible differences using the clique-graph (19; 20), we observe that they are close to those of graph-based methods with low suitability ( $\theta = 0; 0.4$ ), whatever the degree of sparsity  $p_0$  of the true model (see Figure 2). With the clique-graph, the number of edges connecting non-null and equal coefficients increases with the number of non-null coefficients  $p_0$ , but so does the number of misleading edges that connect null and non-null coefficients ( $p_0(p - p_0)$  is an increasing function of  $p_0 \in [0, p/2]$ ). As a consequence, generalized fused lasso estimates obtained with the clique-graph never significantly improve upon the lasso for support recovery in our experiments. These results complement those of Theorem 2 above: while using a clique-graph is asymptotically optimal, this strategy is clearly sub-optimal on finite samples.

Lastly, we consider the classification of pairs of non-null (and then equal) coefficients (Acc.B; Figure 1). First observe that  $\mathcal{B} = \emptyset$  when the graph is completely mis-specified ( $\theta = 0$ ) so that Acc.B can not be computed for this value of  $\theta$ . Clearly, no method is robust on graph mis-specification regarding pairs of coefficients, except the "raw" version of the generalized fused lasso. However, what seems to be a good performance for the "raw" fused lasso is actually linked to a poor model selection and estimation precision since the "raw" fused lasso tends to return the same value for all coefficients (Supplementary Material, Figure 4). To explain why the accuracy on pairs drops with the graph mis-specification, we first note that *both* the number of edges connecting non-zero equal coefficients  $|\mathcal{B}|$  and Acc.B decrease as  $\theta$  decreases. When the graph is mis-specified there is an increased discrepancy between the true number of distinct coefficients ( $d_0$ ), the number of distinct coefficients that would be possible to estimate in the asymptotic setting ( $s_0$ ), and  $s_n$  the actual non-asymptotic model complexity. Then as  $\theta$  decreases,  $s_0 - d_0$  obviously tends to increase but so does  $s_n - s_0$ . It means that graph mis-specifications lead to models that are too complex with respect to the theoretical one. This is still true with the clique-graph for which we do have  $s_0 = d_0$ , but Acc.B is still moderate which means that  $s_n > s_0 = d_0$ : even when all the edges that should be present in the graph are indeed present, misleading edges prevent the method from selecting the right model: Acc.B is moderate and, in turn, so is Acc.A.

#### 4.3 Performance of the generalized fused lasso when connected coefficients may not be exactly equal

Here, we adopt the exact same setting as in the previous paragraph, but instead of setting the  $p_0$  non-null coefficients to the same  $\beta^*$  value, we set each of them to  $|\beta^* + \nu|$ ,

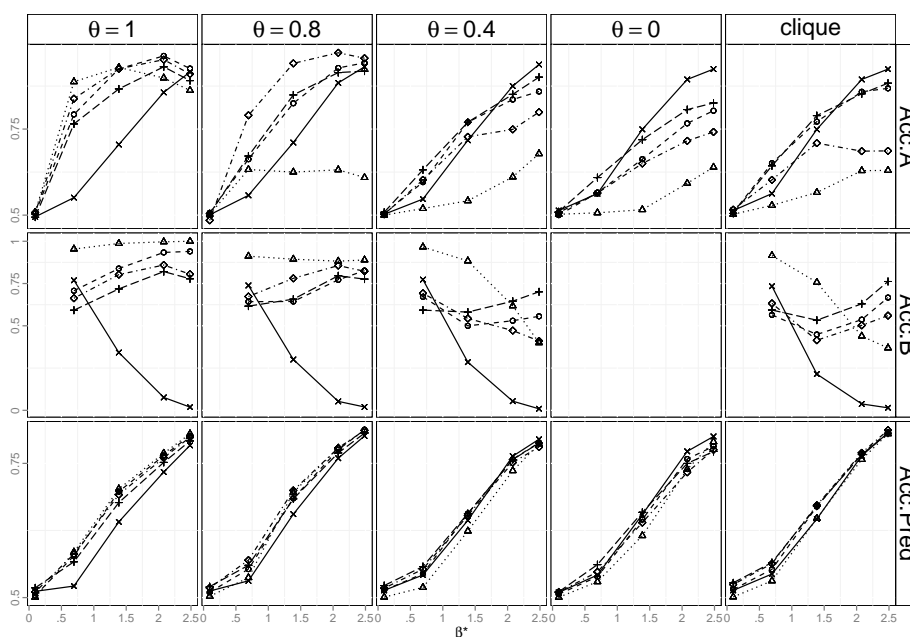
$n/p$		Acc.A		Acc.B		Acc.Pred	
		non-R	R	non-R	R	non-R	R
1	non-adaptive	62	57	75	45	58	58
	adaptive	54	55	41	47	58	58
5	non-adaptive	67	73	87	66	61	62
	adaptive	71	69	66	59	61	61
10	non-adaptive	71	78	82	70	62	63
	adaptive	78	76	72	68	63	63
50	non-adaptive	79	88	75	80	65	65
	adaptive	89	88	87	86	65	65

**Table 1** Average Accuracies ( $\times 100$ ) for support (Acc.A), pairs of non-zero and equal coefficients (Acc.B) and Prediction (Acc.Pred). R, relaxed; non-R, non-relaxed.

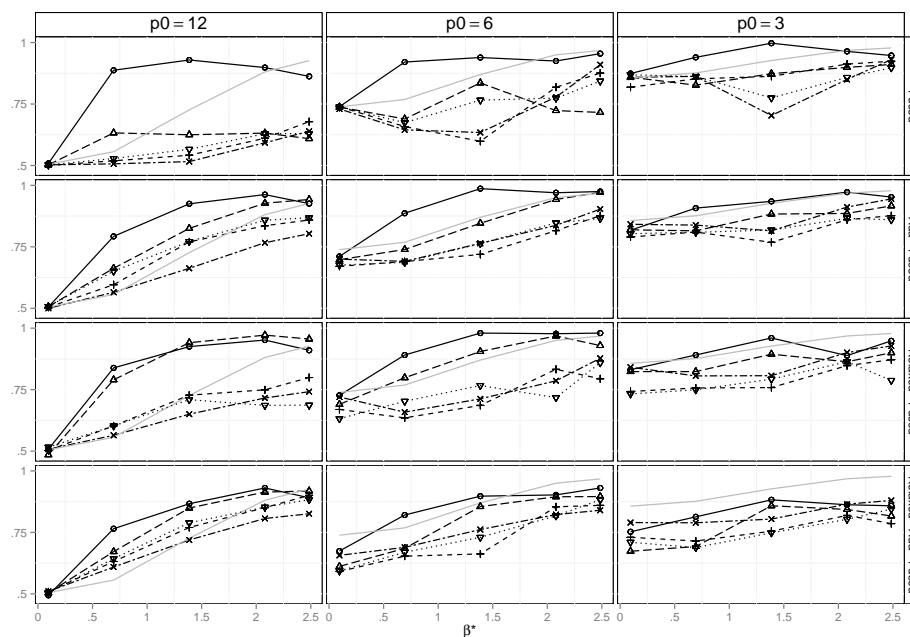
with  $\nu \sim \mathcal{N}(0, \sigma_\nu^2)$ . The variance  $\sigma_\nu^2$  governs the amount of "noise" (variability across the  $p_0$  non-null coefficients), and we make it vary in  $\{0, 0.2, 0.5\}$ . Illustrations of the resulting  $\beta^*$  vector are provided on Figure 5 in the Supplementary Material. In this setting, the reference method is arguably the original group lasso (i.e., with the group lasso penalty only) where the penalty accounts for the true group structure. The relaxed adaptive version of this method was therefore included in this particular setting (using the `grplasso R` package (14)).

We first compare it with generalized fused lasso estimates, for which exact knowledge of the group structure is also assumed; that is we consider generalized fused lasso estimates with  $\theta = 1$  (results for the relaxed adaptive lasso are also provided). Results regarding support recovery (Acc.A) are presented on Figure 3 for  $n = 240$ . Interestingly, most generalized fused lasso estimates are always at least comparable to the group lasso. More importantly, they significantly outperform the group lasso in many situations, especially when the variability across the  $p_0$  non-null coefficients is weak,  $p_0$  is large or signal is weak (low values of  $\beta^*$ ). Comparisons based on prediction accuracy advocate even more for the use of generalized fused lasso estimates (see Figure 6 in the Supplementary Materials).

Figure 4 further shows that, in the case where  $\sigma_\nu^2 = 0.2$ , the performance of the group lasso (with exact knowledge of the group structure) in terms of support recovery is comparable to or worse than that of the generalized fused lasso in the following configurations. (i) when  $p_0 = 12$  and  $\theta \geq 0.8$  or  $\beta^* \leq 1.5$  (irrespective of the  $\theta$  value); (ii) when  $p_0 = 8$  and  $\theta = 1$  or  $\theta = 0.8$  and  $\beta^* \leq 1.5$ ; and (iii) when  $p_0 = 3$  and  $\theta = 1$ . Analogous results were obtained for other values of  $\sigma_\nu^2$ . In other words, even when the true structure is unknown, generalized fused lasso estimates (with  $\theta < 1$  then) can achieve performance similar to those that would achieve the group lasso if the true structure were known, especially if groups of non-null coefficients are large enough. Again, results regarding prediction accuracy support even more the generalized fused lasso (see Figure 7 in the Supplementary Materials).



**Fig. 1** Accuracies for  $\ell_1$ -based fused penalties ( $n/p = 10$ ,  $p_0 = p/2$ ). Raw fused lasso ( $\triangle$ -dotted), adaptive fused lasso ( $\circ$ -dashed), relaxed fused lasso ( $\diamond$ -dotdash), relaxed adaptive fused lasso ( $+$ -longdash), relaxed adaptive lasso ( $\times$ -solid).



**Fig. 2** Accuracy on support recovery  $\mathcal{A}$  for different graphs provided to each method, with  $\theta = 1$  ( $\circ$ -solid),  $\theta = .8$  ( $\triangle$ -long-dash),  $\theta = .4$  ( $+$ -dash),  $\theta = 0$  ( $\times$ -dot-dash), lasso (solid-grey), clique ( $\nabla$ -dot). The lasso strategy corresponds to the case where no graph is provided. The clique strategy corresponds to the case where all differences are penalized.



All these results clearly state the potential benefits of using generalized fused lasso estimates instead of group lasso estimates, in the setting considered in our experiments.

#### 4.4 Simulation Study in the context of joint modeling.

Our objective here is to investigate in more details the performance of the generalized fused lasso in the context of joint logistic regressions, in particular with varying signal-strengths and levels of heterogeneity across strata.

We set  $C = 4$ ,  $n_c = 200$  (so that  $n = \sum_c n_c = 800$ ),  $p = 20$ , and for each stratum we sample covariates  $\mathbf{x}_i$  with  $\rho = 0.5$ . Then we control for the level of heterogeneity across strata by using different repartitions of null and non-null coefficients for a given covariate in the various strata, while keeping  $\beta_{0c}^* = 0$  for every stratum  $c$ . In Configuration 1, all strata share the same  $\beta_c^*$  with 6 non-null elements. In the other two configurations  $\beta_c^*$ s do vary across strata while keeping the number of non-null elements equal to 6 for each stratum. In Configuration 2, 8 coefficients are null on all strata, 6 are null on all but one stratum, and 6 are *non-null* (and equal) on all but one stratum. In Configuration 3, 8 coefficients are null on all strata and 12 are non-null (and equal) on two strata (and null on the other two). See Figure 8 in the Supplementary Material for a graphical description of these three configurations. Then we compare the generalized fused lasso with two versions of the relaxed adaptive lasso. The “*independent*” version consists in computing the relaxed adaptive lasso on each stratum independently. In the “*interaction*” version, a reference stratum is first selected and interaction terms between the remaining strata and the covariates are included as explained in Section 3.2 in the Supplementary Material (this corresponds to the generalized fused lasso with a star-graph and with  $\lambda_1 = \lambda_2$ ). Then one relaxed adaptive lasso is computed on this whole data set. This latter approach can be regarded as the reference method in this context.

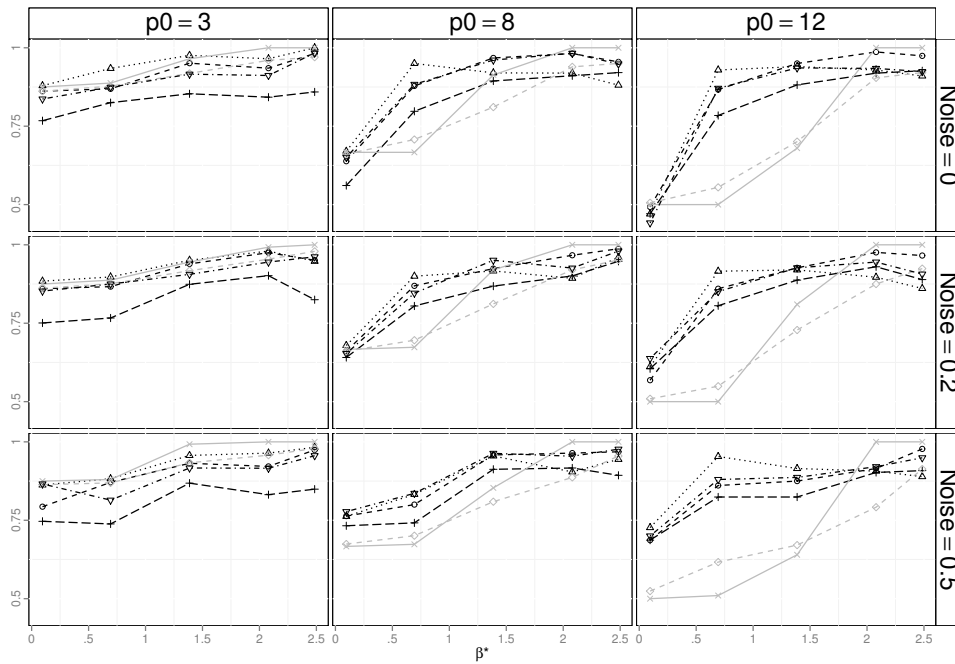
Here we mostly focus on performance for the classification of pairs, which is of primary interest in the context of joint modeling. As before, we first consider Acc.B. Because  $\beta_{c_1,j}^* \beta_{c_2,j}^* \neq 0 \Rightarrow \beta_{c_1,j}^* = \beta_{c_2,j}^* \neq 0$  in our simulations, it assesses the capacity of methods to detect homogeneity across strata, and could be interpreted as a type-1 error of an interaction test, as mentioned above. We also consider the percentage of truly heterogeneous edges (i.e. connecting nodes that correspond to different true coefficients) that are detected as heterogeneous. It can be interpreted as a power of an interaction test, and is denoted by Pow.H. Note that the “*independent*” relaxed adaptive lasso is expected to show good performance regarding this criterion since it does not encourage coefficients to be equal across strata: 1 - Pow.H for this method only represents the

proportion of heterogeneous edges for which both coefficients are set to 0 (so that 1-Pow.H is expected to decrease, and Pow.H to increase, as  $\beta^*$  increases for the “*independent*” relaxed adaptive lasso). Figure 5 shows that the relaxed adaptive generalized fused lasso outperforms the “*interaction*” relaxed adaptive lasso in terms of Pow.H and, in most cases, in terms of Acc.B as well. Moreover, it is almost “perfect” to detect heterogeneity since it shares similar performance with the “*independent*” relaxed adaptive lasso regarding Pow.H. The non-relaxed version is in most cases a little better (resp. worse) than the relaxed one in terms of Acc.B (resp. Pow.H). As for support recovery, (relaxed) adaptive generalized fused lasso also outperforms the “*interaction*” relaxed adaptive lasso, especially as the level of heterogeneity across strata increases (as in Configuration 2 and 3). Interestingly, even under the most heterogeneous configuration (Configuration 3), the relaxed adaptive generalized fused lasso attains performance similar to those achieved by the “*independent*” lasso in terms of Acc.A.

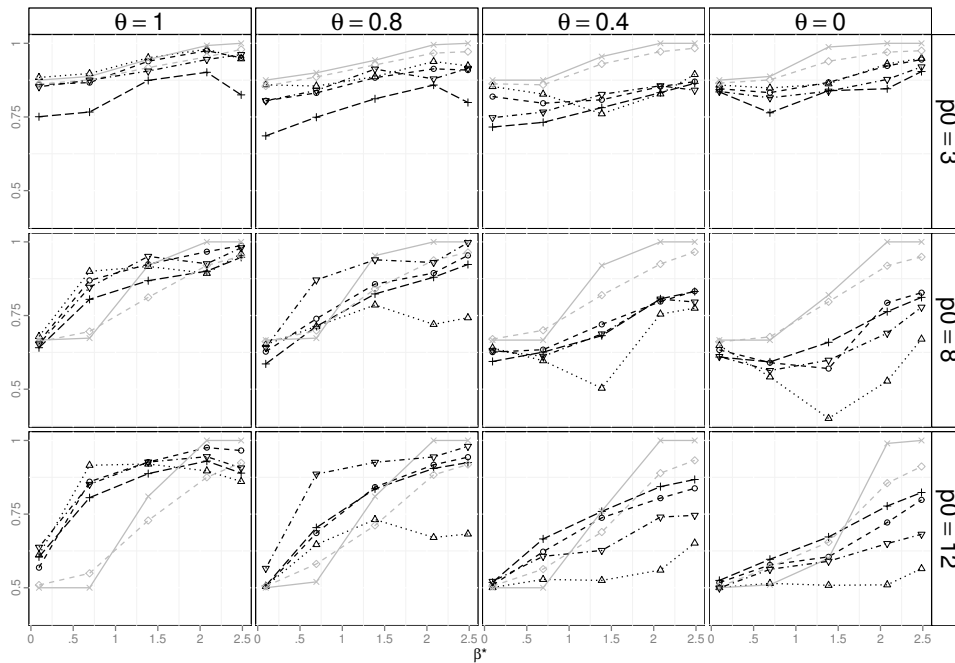
In other respect, Gertheiss and Tutz (5) especially observed that imposing  $\lambda_1 = \lambda_2$  lead to results comparable to those obtained with “free”  $\lambda_1$  and  $\lambda_2$ . As shown on Figure 9 of the Supplementary Material, this was not the case on our experiments for non-adaptive versions of the generalized Fused Lasso, nor for the relaxed adaptive generalized fused lasso, especially under Configuration 3. Therefore, unless computational time is a critical issue, we would not recommend to impose  $\lambda_1 = \lambda_2$  when using generalized fused lasso in the context of joint modeling (especially if non-adaptive versions are to be used).

## 5 Network-based prediction of cancer status based on expression data

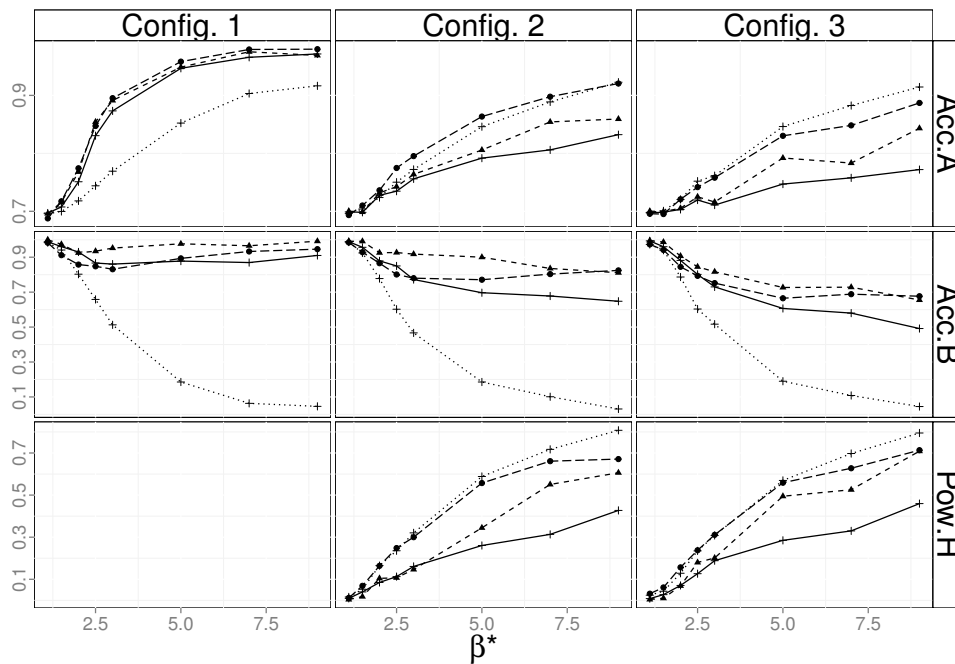
Genomics has faced a flood of network data in the last years, ranging from protein-protein interaction data, pathway data to regulation networks (7; 18). The molecular characterization of cancers has been at the core of many projects, especially to establish molecular subtypes of histologically similar tumors. In particular finding genomic signatures has been the *grail* for many studies to predict patient outcome, survival or relapse (9). Such signatures are determined a penalized logistic regression model based on gene-expression data as covariates, and we consider here a logistic regression to predict the 5-year relapse status of 214 women with breast cancer (80 relapse in the sample) (9). Covariates correspond to the measurement of the  $p = 54,613$  gene expressions reduced to the 248 genes differentially expressed (FDR=0.05), and we use 5-fold Monte Carlo cross validations to assess prediction performance. Interestingly, the expression of different genes is structured according to some unknown regulatory network that can be inferred from the data using Gaussian Graphical models for instance (2). Our hypothesis here



**Fig. 3** Accuracy on support recovery  $\mathcal{A}$  when the graph is perfect ( $\theta = 1$ ) but non-zero (and connected) coefficients may not be exactly equal: fused-like estimates are represented in black ( $\triangle$ -dot: generalized fused lasso,  $\circ$ -solid: adaptive generalized fused lasso,  $\nabla$ -dot: relaxed generalized fused lasso,  $+$ -dash: relaxed adaptive generalized fused lasso). Results for the relaxed adaptive group lasso (grey,  $\times$ -solid) and the relaxed adaptive lasso (grey,  $\circ$ -solid:) are also reported.



**Fig. 4** Accuracy on support recovery  $\mathcal{A}$  when non-zero coefficients are not exactly equal ( $\sigma_v^2 = 0.2$ ) and the graph provided in the generalized fused lasso is imperfect: fused-like estimates are represented in black ( $\triangle$ -dot: generalized fused lasso,  $\circ$ -solid: adaptive generalized fused lasso,  $\nabla$ -dot: relaxed generalized fused lasso,  $+$ -dash: relaxed adaptive generalized fused lasso). Results for the relaxed adaptive group lasso with perfect knowledge of the structure (grey,  $\times$ -solid) and the relaxed adaptive lasso (grey,  $\circ$ -solid:) are also reported.



**Fig. 5** Performance regarding support recovery and the classification of pairs of coefficients in the joint modeling context. Adaptive fused lasso (▲-dashed), relaxed adaptive fused lasso (●-longdash), independent relaxed adaptive lasso (+-solid), interaction relaxed adaptive lasso (+-dotted). Columns correspond to levels of heterogeneity between strata: from homogeneous (Configuration 1) to the most heterogeneous (Configuration 3).

is that using this inferred network can help in the prediction of patients outcome. However, since this regulatory network is not perfectly known, our strategy is based on the hypothesis that there is no "true" regulatory network, and we explore the robustness of the generalized fused lasso to the addition/removal of edges in the penalty, as we did in the simulation study. To proceed we consider the regulatory network that is inferred on the training data by the *SIMONE* package (2). This package is based on sparse Gaussian Graphical models, and by varying the amount of shrinkage, we were able to consider networks with increasing number of edges, and then to assess the impact of changes in the network on prediction performance and estimated model dimensions.

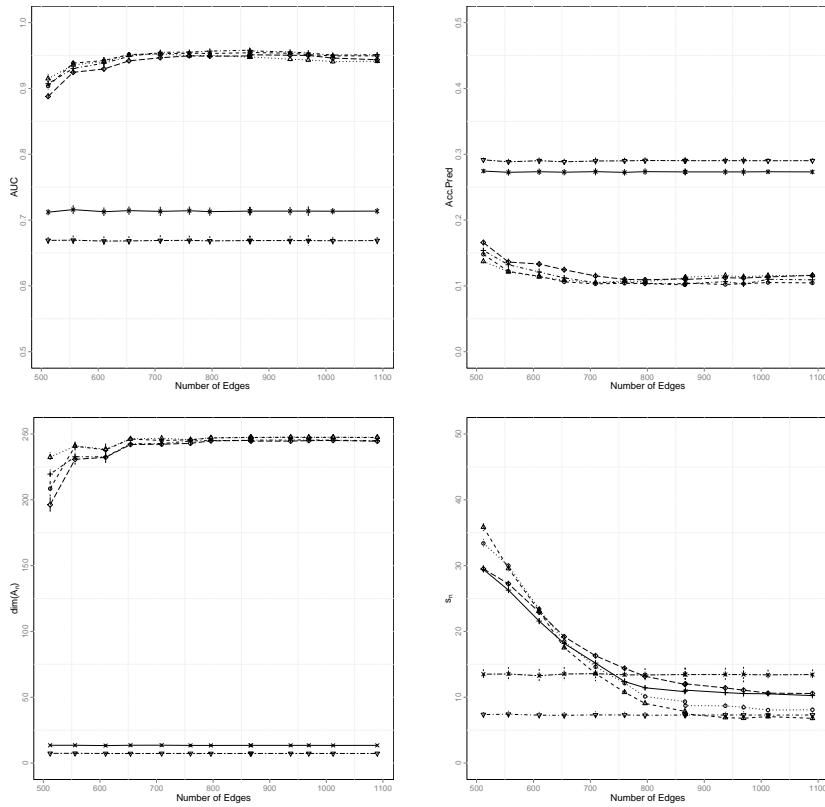
The first conclusion is that the gain in used fused-based strategies is massive: the AUC (Area under the Curve) jumps from  $\sim 0.7$  for the lasso to  $\sim 0.95$  for generalized fused estimates, and the empirical error rate drops from  $\sim 0.3$  to  $\sim 0.1$ , which clearly indicates that the network has helped in the correct classification of samples (Figure 6). Moreover, as previously mentioned in the simulations, fused-based methods are more performant than the lasso, but there is no significant difference among them for prediction. Very interestingly, the classification performance remains with the addition of edges, which indicates that the suitability of this new edges is about  $\theta \simeq 0.5$ . Then the number of non-null estimated coefficients is higher for all fused-penalties, and the estimated number of distinct non-null coefficients (a crude

estimate of quantity  $s_0$ ) converges towards a set of  $\sim 10$  distinct values for estimated parameters.

## 6 Joint modeling to analyze road-safety data

Driving under the influence of alcohol (DUI) is an established risk factor of car accidents. Interestingly, several studies also suggest that DUI increases the risk of dying in an accident, but this result remains controversial: biological evidence supporting this assumption is still lacking and the observed effect of DUI could be due to confounding variables only. We present an application of the generalized fused lasso to the joint modeling framework, where the main objective is to study the effect of alcohol consumption on the risk of death (for drivers involved in a car crash). Our dataset consists of  $n = 21,064$  drivers involved in reported single-car personal injury crashes from 2006 to 2009 in France (16). Current data show 33 covariates including the characteristics of the crash, of the drivers and of crash-involved vehicles, and we focus on the vital status of the driver only. We define 4 strata based on gender and DUI: strata 1-2 (resp. 3-4) for males and females not driving (resp. driving) under the influence when the accident occurred.

We consider logistic models to relate the probability of dying in a car accident to risk factors in each stratum. Since most factors are expected to share similar effects, joint modeling is used to couple the estimations of the four models. Intercept parameters are of particular importance: they should



**Fig. 6** Performance of penalization strategies for AUC (Area Under the Curve) and prediction accuracy (Acc.Pred) on the breast cancer dataset with an increasing number of edges in the network provided in the penalty. Bottom panels display the estimated model size and complexities (number of non-null estimated coefficients, and number of distinct estimated coefficients) according to an increase in the number of edges in the penalty graph. Fused lasso (dotted), adaptive fused lasso (dashed), relaxed fused lasso (dotdash), relaxed adaptive fused lasso (longdash), relaxed adaptive lasso (plain), GE-Net (two-dash).

be homogeneous across strata if neither gender nor alcohol directly modified the risk of death. Then we compare the relaxed adaptive fused lasso, to the “*independent*” and “*interaction*” versions of the relaxed adaptive lasso (with stratum 1, i.e., sober males, as the reference stratum). We also present unpenalized estimates derived from standard logistic regression models independently built on each stratum (see the Supplementary Material, Figure 6 for complete results).

All methods agree on the absence of effect of most factors on the risk of death in a car crash, and on the intensity of the effects of other factors, such as a higher risk for older drivers and a lower risk associated with the use of a seat-belt. Then inspecting the influence of city roads emphasizes interesting differences between the relaxed adaptive fused lasso and “*independent*” relaxed adaptive lassos in the presence of highly correlated variables (see Table 2). Indeed covariates “City Roads” (crash on a road managed by a city) and “City” (crash in a city) were both kept in the analysis despite redundancy. Thus, the global effect should be the sum of the two corresponding coefficients. Based on this global effect, all methods agree on the absence of in-

teraction with gender and DUI (the sum of the two coefficients is roughly constant over the 4 strata according to every method). Accordingly, the relaxed adaptive fused lasso (as well as the “*interactions*” lasso) return equal individual effects for “City Roads” and “City” on every stratum. However, because of these two variables are highly correlated, “*independent*” lassos return widely different individual effect for “City Roads” and “City”, which does not make any sense. There are a few other differences between the three penalized methods, but when relaxed adaptive fused lasso and “*interactions*” relaxed adaptive lasso estimates disagree, the relaxed adaptive fused lasso most often agrees with “*independent*” relaxed adaptive lassos. This is consistent with our conclusions from the simulation study where these latter two methods generally performed the best in terms of support recovery.

Finally, the relaxed adaptive fused lasso agrees with the other methods on the fact that intercepts *do* vary across strata, suggesting an effect of both gender and DUI on the risk of death (see Table 2). More precisely, sober females are at a higher risk than sober males, and, to a lesser extent, females

under the influence are at a higher risk than males under the influence. Moreover, irrespective on gender, drivers under the influence are at a higher risk than sober drivers. However this result should be tempered by potential confounding due to speed, which was not available here. For instance, drivers under the influence are likely to drive faster than sober drivers. The effects of speed may be partly captured by other covariates, but not entirely. Consequently residual “speed effects” could be responsible for detected heterogeneities between intercepts.

## 7 Discussion

In this paper, we investigate theoretical and empirical properties of the generalized fused lasso, in various settings. From the theoretical point of view, we especially show that using adaptive weights leads to estimators enjoying asymptotic oracle properties. However, for the true underlying dimension of the problem  $d_0$  (that is the number of distinct non-null values in  $\beta^*$ ) to be equal to the asymptotic dimension  $s_0$  of the estimator, the graph  $G$  used in the penalty has to enjoy the following property: for all  $(j, \ell)$  such that  $\beta_j^* = \beta_\ell^*$ ,  $j$  and  $\ell$  belong to the same connected component of  $G_{\mathcal{B}}$ , the sub-graph of  $G$  such that  $G_{\mathcal{B}} = (\mathcal{A}, \mathcal{B})$ . In particular, our results indicate that setting  $G$  to the clique connecting all coefficients of  $\beta^*$  together (in which case all the  $p(p-1)/2$  differences are penalized) is asymptotically optimal. Therefore, they confirm those obtained in the linear regression setting in (20), and extend them to generalized linear models. In words, it means that, asymptotically, adding misleading edges in the graph is harmless, while forgetting relevant ones can be harmful.

From the modeling point of view however, we empirically studied the robustness of generalized fused lasso estimates against graph misspecification on finite samples. On our experiments, we observed that adaptive weights and/or relaxation lead to some improved robustness. However, and overall, we demonstrated that the performance of generalized fused lasso estimates on finite samples are deeply related to the suitability of the graph in the penalty, especially for support recovery. In particular, we show that, under the designs considered in our simulations, the clique-based strategy is clearly sub-optimal for support recovery, so that misleading edges are harmful on finite samples. The graph used in the penalty constitutes a formal description of some *prior* knowledge on the problem that is investigated, and has to be determined with caution, especially if support recovery matters. We may stress that this graph does not describe correlations among features but similarity between their effects under the considered model. Of course, correlated features may share similar coefficients, but not necessarily. For instance in epidemiological studies, smoking and alcohol consumption are generally highly correlated.

They further may share similar effects under a logistic model when studying cardiovascular diseases so that it might make sense to penalize the difference of their effects. However, when studying lung cancer, they are not expected to share similar coefficients at all, so that their difference should not be penalized. That being said, in our application of Section 5, using a graph based on (partial) correlation still leads to highly improved prediction accuracy compared to the Lasso for instance (of course, support recovery performance can not be assessed on real data).

We further show that generalized fused lasso can be useful even in situations where theoretical coefficients are not exactly the same. Interestingly, it can improve upon state-of-the-art competing methods, such as the group lasso in this context, even when the graph is not perfectly suited.

A particular situation where the graph is suggested by the design of the study itself is the joint modeling framework, where data come from various strata. When the main question is the detection of heterogeneities across strata, we believe that the graph made of cliques is very appealing as it encourages coefficients to be homogeneous across the strata. This strategy has some connections with the statistical tests theory where tests are generally performed under the null hypothesis (absence of heterogeneity in this case), and data need to be far enough from this assumption for the null hypothesis to be rejected. But even in the joint modeling context, other graphs may be considered: for instance, if strata correspond to various treatments, the control treatment can serve as the reference and star-graphs may be more appropriate than cliques (see Section 1 in the Supplementary Material).

We established the asymptotic oracle properties of the adaptive generalized fused lasso estimates under generalized linear models, for fixed  $p$ . These results are the first established for fused lasso estimates in the setting of generalized linear models and for the generalized fused penalty based on a graph. Even if the fixed  $p$  case is relevant, especially in the joint modeling framework, they should be extended to cover the high-dimensional case. Most published papers on the fused lasso in high-dimension focus on the chain-based fused penalty in the Gaussian sequence model. A notable exception is the work of Vaïter et al. (24) in which results that encompass generalized fused lassos under a bounded noise assumption were recently established. The extension of such results to the random noise case and generalized linear models would be an interesting lead. Moreover, non-asymptotic oracle prediction inequalities still need to be established for fused-like estimates: under the linear model for instance, and denoting by  $\hat{\beta}$  the fused lasso estimator for some appropriate values of the tuning parameters, it is easy to obtain inequalities of the form  $\|\mathbf{X}(\hat{\beta} - \beta^*)\|^2 \leq \kappa p_0 \log(p)/n$  with high probability, for some positive constant  $\kappa$ . Under the Gaussian sequence model, a modification of the Fused

**Table 2** Parameter estimates obtained on the road-safety data for the intercept terms and coefficients of variables ‘‘City’’ and ‘‘City Roads’’. For each method and each parameter, four values are given, corresponding to the estimate obtained on each of the four considered strata.

Variable	relaxed adaptive fused lasso				‘‘independent’’ relaxed adaptive lasso				‘‘interaction’’ relaxed adaptive lasso			
Intercept	-2.86	-3.89	-2.14	-2.14	-2.85	-3.86	-2.01	-2.22	-2.83	-3.67	-2.22	-2.39
City	-0.37	-0.37	-0.37	-0.37	-0.43	0.00	-0.50	-0.70	-0.39	-0.39	-0.39	-0.39
City Roads	-0.21	-0.21	-0.21	-0.21	0.00	-0.54	0.00	0.00	-0.21	-0.21	-0.21	-0.21

Lasso has been shown to enjoy the same kind of bound, but with  $p_0$  replaced by a quantity similar in spirit to  $s_0$  (17). Extending this result to more general contexts is not straightforward and constitutes a promising research direction as well.

## 8 Appendix

### 8.1 Proof of Theorem 1

This proof is an adaptation of the proof given by (23) to account for the generalized linear model loss and the generalized fused lasso penalty. Let us define  $\mathcal{V}_n(\mathbf{u}) = Q(\boldsymbol{\beta}^* + \mathbf{u}/\sqrt{n}) - Q(\boldsymbol{\beta}^*)$  with  $\mathbf{u} = (u_0, \dots, u_p)^T$ , and  $Q$  defined as in Section 2.1. Obviously  $\mathcal{V}_n(\mathbf{u})$  is minimized at  $\sqrt{n}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*)$ . Similarly to (23), we obtain:

$$\begin{aligned} \mathcal{V}_n(\mathbf{u}) &= J\left(\boldsymbol{\beta}^* + \frac{\mathbf{u}}{\sqrt{n}}\right) - J(\boldsymbol{\beta}^*) \\ &+ \lambda_n^{(1)} \sum_{j=1}^p \left\{ \left| \beta_j^* + \frac{u_j}{\sqrt{n}} \right| - |\beta_j^*| \right\} \\ &+ \lambda_n^{(2)} \sum_{(j,\ell) \in E} \left\{ \left| \beta_j^* - \beta_\ell^* + \frac{(u_j - u_\ell)}{\sqrt{n}} \right| - |\beta_j^* - \beta_\ell^*| \right\}. \end{aligned}$$

For any fixed  $\mathbf{u}$ , the last two terms of the right-hand side converge to the last two terms in expression (1) of  $\mathcal{V}(\mathbf{u})$  as  $n$  goes to  $\infty$ . As for the first two terms, a Taylor expansion yields

$$\begin{aligned} J\left(\boldsymbol{\beta}^* + \frac{\mathbf{u}}{\sqrt{n}}\right) - J(\boldsymbol{\beta}^*) &= \nabla J(\boldsymbol{\beta}^*)^T \frac{\mathbf{u}}{\sqrt{n}} \\ &+ \frac{1}{2} \mathbf{u}^T \frac{\mathcal{I}(\boldsymbol{\beta}^*)}{n} \mathbf{u} + o_{\mathbb{P}}(1/n). \end{aligned}$$

Under **AL1**, we have  $\mathbf{u}^T (\mathcal{I}(\boldsymbol{\beta}^*)/n) \mathbf{u} \rightarrow_d \mathbf{u}^T \mathbf{C} \mathbf{u}$ . Moreover, **AL1** implies that the minimum eigenvalue of  $\mathcal{I}(\boldsymbol{\beta}^*)$  goes to  $\infty$  and, under **AL2**, we have  $(\nabla J(\boldsymbol{\beta}^*)/\sqrt{n}) \rightarrow_d \mathbf{W}$ , where  $\mathbf{W}$  has an  $\mathcal{N}(\mathbf{0}_{p+1}, \mathbf{C})$  distribution (8). By Slutsky’s theorem, we therefore have  $\mathcal{V}_n(\mathbf{u}) \rightarrow_d \mathcal{V}(\mathbf{u})$ . Since  $\mathcal{V}_n$  is convex, the epi-convergence results of (6) can finally be used to complete the proof of Theorem 1.

### 8.2 Proof of Proposition 1

If  $\lambda_0^{(2)} = 0$ , the proof is the same as in the lasso case (25). Now assume that  $\lambda_0^{(2)} \neq 0$  (we further assume that  $\lambda_0^{(1)} \neq 0$ ,

the case where  $\lambda_0^{(1)} = 0$  being slightly easier and omitted). First observe that  $\mathbb{P}(\tilde{\mathcal{A}}_n = \mathcal{A}) \leq \mathbb{P}(\sqrt{n}\hat{\beta}_j = 0 \forall j \notin \mathcal{A})$ . Moreover, in virtue of Theorem 1, we have  $\limsup_n \mathbb{P}(\sqrt{n}\hat{\beta}_j = 0 \forall j \notin \mathcal{A}) \leq \mathbb{P}(u_j^* = 0 \forall j \notin \mathcal{A})$ , with  $\mathbf{u}^* = \operatorname{argmin}(\mathcal{V})$ . Therefore, we only need to show that  $c = \mathbb{P}(u_j^* = 0 \forall j \notin \mathcal{A}) < 1$ .

For any  $j \in \{1, \dots, p\}$ , introduce  $E_j^-(\boldsymbol{\beta}^*) = \{\ell : (\ell, j) \in E \text{ or } (j, \ell) \in E \text{ and } \beta_j^* = \beta_\ell^*\}$  and  $E_j^+(\boldsymbol{\beta}^*) = \{\ell : (\ell, j) \in E \text{ or } (j, \ell) \in E \text{ and } \beta_j^* \neq \beta_\ell^*\}$ . Setting  $\mathbf{W} = (W_0, \dots, W_p)^T$  and  $\mathbf{C} \mathbf{u}^* = ((\mathbf{C} \mathbf{u}^*)_0, \dots, (\mathbf{C} \mathbf{u}^*)_p)^T$ , we have, by the KKT conditions,

$$W_0 + (\mathbf{C} \mathbf{u}^*)_0 = 0, \quad (1)$$

and for all  $j \in \mathcal{A}$ ,

$$\begin{aligned} W_j + (\mathbf{C} \mathbf{u}^*)_j + \lambda_0^{(1)} \operatorname{sign}(\beta_j^*) + \lambda_0^{(2)} \left\{ \sum_{\ell \in E_j^+(\boldsymbol{\beta}^*)} \operatorname{sign}(\beta_j^* - \beta_\ell^*) \right. \\ \left. + \sum_{\ell \in E_j^-(\boldsymbol{\beta}^*)} (-1)^{\mathbb{1}(j < \ell)} t_{j\ell} \right\} = 0, \end{aligned}$$

and for all  $j \notin \mathcal{A}$ ,

$$\begin{aligned} W_j + (\mathbf{C} \mathbf{u}^*)_j + \lambda_0^{(1)} r_j + \lambda_0^{(2)} \left\{ \sum_{\ell \in E_j^+(\boldsymbol{\beta}^*)} \operatorname{sign}(\beta_j^* - \beta_\ell^*) \right. \\ \left. + \sum_{\ell \in E_j^-(\boldsymbol{\beta}^*)} (-1)^{\mathbb{1}(j < \ell)} t_{j\ell} \right\} = 0. \end{aligned}$$

Above,  $r_j = \operatorname{sign}(u_j^*)$  if  $u_j^* \neq 0$  and  $r_j$  is some real number in  $[-1, 1]$  otherwise. Similarly,  $t_{j\ell} = \operatorname{sign}(u_j^* - u_\ell^*)$  if  $u_j^* \neq u_\ell^*$  and  $t_{j\ell}$  is some real number in  $[-1, 1]$  otherwise.

For any index  $j \in \mathcal{A}$  there is some  $s = s(j)$  such that  $j \in \mathcal{A}_{s(j)}$ , where  $\mathcal{A}_s$  still denotes the set of vertices of the  $s$ -th connected component of  $G_B$  (see the paragraph before the statement of Theorem 2 for the definitions of these objects). Then summing up the KKT conditions over  $\mathcal{A}_{s(j)}$ , we have

$$\begin{aligned} \sum_{k \in \mathcal{A}_{s(j)}} \left\{ W_k + (\mathbf{C} \mathbf{u}^*)_k + \lambda_0^{(1)} \operatorname{sign}(\beta_k^*) \right. \\ \left. + \lambda_0^{(2)} \sum_{\ell \in E_k^+(\boldsymbol{\beta}^*)} \operatorname{sign}(\beta_k^* - \beta_\ell^*) \right\} = 0. \quad (2) \end{aligned}$$

Similarly, setting  $\tilde{\mathcal{B}} = \{(j, \ell) \in E \cap \bar{\mathcal{A}} \times \bar{\mathcal{A}}\}$  and denoting by  $G_0 = (\bar{\mathcal{A}}, \tilde{\mathcal{B}})$ , the set  $\bar{\mathcal{A}}$  can be decomposed as  $\bar{\mathcal{A}} =$

$\cup_{s=1}^{s_1} \bar{\mathcal{A}}_s$ , where  $1 \leq s_1 \leq p - p_0$  and  $\bar{\mathcal{A}}_s$  is the subset of vertices constituting the  $s$ -th connected component of  $G_0$ . Then, for any  $j \notin \mathcal{A}$ , there exists some  $s = s(j)$  such that  $j \in \bar{\mathcal{A}}_{s(j)}$  and summing up the KKT optimality conditions over  $\bar{\mathcal{A}}_{s(j)}$ , we have

$$\sum_{k \in \bar{\mathcal{A}}_{s(j)}} \left\{ W_k + (\mathbf{C}\mathbf{u}^*)_k + \lambda_0^{(1)} r_k + \lambda_0^{(2)} \sum_{\ell \in E_k^{\neq}(\beta^*)} \text{sign}(\beta_k^* - \beta_\ell^*) \right\} = 0, \quad (3)$$

with  $|s_k| \leq 1$ . If  $u_j^* = 0$  for all  $j \notin \mathcal{A}$ , equations (2) along with equation (1) form a system of  $s_0 + 1$  equations with  $p_0 + 1 \geq s_0 + 1$  variables, that can be written as

$$\mathbf{W}_B + \mathbf{M}_1 \mathbf{u}_{\{0\} \cup \mathcal{A}}^* + \lambda_0^{(1)} \mathbf{r}_B + \lambda_0^{(2)} \mathbf{t}_B = \mathbf{0},$$

where  $\mathbf{M}_1$  is the  $(s_0 + 1) \times (p_0 + 1)$  matrix whose  $(s, j)$  element is  $m_{s,j} = \sum_{k \in \bar{\mathcal{A}}_{s-1}} C_{k,j}$  (with  $C_{k,j}$  the  $(k, j)$  element of  $\mathbf{C}$  and  $\mathcal{A}_0 = \{0\}$ ), and  $\mathbf{W}_B, \mathbf{r}_B$  and  $\mathbf{t}_B$  are vectors in  $\mathbb{R}^{s_0+1}$  whose  $s$ -th elements are  $\sum_{k \in \bar{\mathcal{A}}_{s-1}} W_k, \sum_{k \in \bar{\mathcal{A}}_{s-1}} \text{sign}(\beta_k^*)$  and  $\sum_{k \in \bar{\mathcal{A}}_{s-1}} \sum_{\ell \in E_k^{\neq}(\beta^*)} \text{sign}(\beta_k^* - \beta_\ell^*)$  respectively. Now, denoting by  $\mathbf{M}_1^\dagger$  the pseudo-inverse of  $\mathbf{M}_1$ , there exists some vector  $\boldsymbol{\omega} \in \mathbb{R}^{p_0+1}$  such that  $\mathbf{u}_{\{0\} \cup \mathcal{A}}^* = (\mathbf{I}_{p_0+1} - \mathbf{M}_1^\dagger \mathbf{M}_1) \boldsymbol{\omega} - \mathbf{M}_1^\dagger (\mathbf{W}_B + \lambda_0^{(1)} \mathbf{r}_B + \lambda_0^{(2)} \mathbf{t}_B)$ . Therefore, if  $u_j^* = 0$  for all  $j \notin \mathcal{A}$ , equations (3) form a system of  $s_1$  equations that can be written,

$$\left| \mathbf{W}_{\bar{\mathcal{B}}} + \lambda_0^{(2)} \mathbf{t}_{\bar{\mathcal{B}}} + \mathbf{M}_2 \left\{ (\mathbf{I}_{p_0} - \mathbf{M}_1^\dagger \mathbf{M}_1) \boldsymbol{\omega} - \mathbf{M}_1^\dagger (\mathbf{W}_B + \lambda_0^{(1)} \mathbf{r}_B + \lambda_0^{(2)} \mathbf{t}_B) \right\} \right| \leq \lambda_0^{(1)} \mathbf{r}_{\bar{\mathcal{B}}},$$

where  $\mathbf{W}_{\bar{\mathcal{B}}}, \mathbf{r}_{\bar{\mathcal{B}}}$  and  $\mathbf{t}_{\bar{\mathcal{B}}}$  are the vectors in  $\mathbb{R}^{s_1}$  whose  $s$ -th elements are given by  $\sum_{k \in \bar{\mathcal{A}}_s} W_k, |\bar{\mathcal{A}}_s|$  and

$$\sum_{k \in \bar{\mathcal{A}}_s} \sum_{\ell \in E_k^{\neq}(\beta^*)} \text{sign}(\beta_k^* - \beta_\ell^*)$$

respectively. We can now conclude by observing that

$$c \leq \mathbb{P} \left( \left| \mathbf{W}_{\bar{\mathcal{B}}} + \lambda_0^{(2)} \mathbf{t}_{\bar{\mathcal{B}}} + \mathbf{M}_2 \left\{ (\mathbf{I}_{p_0} - \mathbf{M}_1^\dagger \mathbf{M}_1) \boldsymbol{\omega} - \mathbf{M}_1^\dagger (\mathbf{W}_B + \lambda_0^{(1)} \mathbf{r}_B + \lambda_0^{(2)} \mathbf{t}_B) \right\} \right| \leq \lambda_0^{(1)} \mathbf{r}_{\bar{\mathcal{B}}} \right) < 1.$$

### 8.3 Proof of Theorem 2

The following proof is a modification to the proof given by (25) to account for both the generalized linear model loss and the generalized fused penalty. Let us define  $\mathbf{V}_n(\mathbf{u}) = Q(\beta^* + \mathbf{u}/\sqrt{n}) - Q(\beta^*)$  with  $\mathbf{u} = (u_0, \dots, u_p)^T$  and  $Q$

defined as Section 2.1. Note that  $\mathbf{V}_n(\mathbf{u})$  is minimized at  $\sqrt{n}(\hat{\beta}^{ad} - \beta^*)$ . We have

$$\begin{aligned} \mathbf{V}_n(\mathbf{u}) &= \nabla J(\beta^*)^T \frac{\mathbf{u}}{\sqrt{n}} + \frac{1}{2} \mathbf{u}^T \frac{\mathcal{I}(\beta^*)}{n} \mathbf{u} + o_{\mathbb{P}}(1/n) \\ &+ \frac{\lambda_n^{(1)}}{\sqrt{n}} \sum_{j=1}^p w_j^{(1)} \sqrt{n} \left\{ \left| \beta_j^* + \frac{u_j}{\sqrt{n}} \right| - |\beta_j^*| \right\} \\ &+ \frac{\lambda_n^{(2)}}{\sqrt{n}} \sum_{(j,\ell) \in E} w_{j\ell}^{(2)} \sqrt{n} \left\{ \left| \beta_j^* - \beta_\ell^* + \frac{(u_j - u_\ell)}{\sqrt{n}} \right| - |\beta_j^* - \beta_\ell^*| \right\} \\ &=: \nabla J(\beta^*)^T \frac{\mathbf{u}}{\sqrt{n}} + \frac{1}{2} \mathbf{u}^T \frac{\mathcal{I}(\beta^*)}{n} \mathbf{u} + o_{\mathbb{P}}(1/n) + \sum_{j=1}^p T_j^{(1)} + \sum_{(j,\ell) \in E} T_{j\ell}^{(2)} \end{aligned}$$

We have the two following behaviors :

$$T_j^{(1)} \rightarrow_p \begin{cases} 0 & \text{if } \beta_j^* \neq 0 \text{ or } (\beta_j^* = 0 \text{ and } u_j = 0) \\ \infty & \text{otherwise} \end{cases}$$

and

$$T_{j\ell}^{(2)} \rightarrow_p \begin{cases} 0 & \text{if } \beta_j^* \neq \beta_\ell^* \text{ or } (\beta_j^* = \beta_\ell^* \text{ and } u_j = u_\ell) \\ \infty & \text{otherwise} \end{cases}.$$

Denote by  $\mathbf{C}_{\mathcal{A}}$  the  $(p_0+1) \times (p_0+1)$  sub-matrix of  $\mathbf{C}$  constituted of rows and columns associated with indexes in  $\{0\} \cup \mathcal{A}$  and by  $\mathbf{W}_{\mathcal{A}}$  a random Gaussian vector  $\mathcal{N}(\mathbf{0}_{p_0+1}, \mathbf{C}_{\mathcal{A}})$ . Then, as in the proof of Theorem 1,  $\mathbf{V}_n(\mathbf{u}) \rightarrow_d \mathbf{V}(\mathbf{u})$  for every  $\mathbf{u}$ , with  $\mathbf{V}$  defined for  $\mathbf{u} = (u_0, \dots, u_p) \in \mathbb{R}^{p+1}$ , by

$$\mathbf{V}(\mathbf{u}) = \begin{cases} \frac{1}{2} \mathbf{u}_{\mathcal{A}}^T \mathbf{C}_{\mathcal{A}} \mathbf{u}_{\mathcal{A}} + \mathbf{u}_{\mathcal{A}}^T \mathbf{W}_{\mathcal{A}} & \text{if } u_j = 0 \text{ for } j \notin \mathcal{A} \text{ and } \\ & u_j = u_\ell \text{ for } (j, \ell) \in \mathcal{B}, \\ \infty & \text{otherwise.} \end{cases}$$

Recall the notations introduced just before stating Theorem 2. Any vector  $\mathbf{u} \in \mathbb{R}^{p+1}$  such that  $u_j = 0$  for all  $j \notin \mathcal{A}$  and  $u_j = u_\ell$  for all  $(j, \ell) \in \mathcal{B}$  has  $s_0 + 1$  distinct non-zero values. Denoting by  $u_0, u_{j_1}, \dots, u_{j_{s_0}}$  these values, and setting  $\mathbf{u}_{\mathcal{B}} = (u_0, u_{j_1}, \dots, u_{j_{s_0}})^T$ , we have

$$\mathbf{V}(\mathbf{u}) = \begin{cases} \frac{1}{2} \mathbf{u}_{\mathcal{B}}^T \mathbf{C}_{\mathcal{B}} \mathbf{u}_{\mathcal{B}} + \mathbf{u}_{\mathcal{B}}^T \mathbf{W}_{\mathcal{B}} & \text{if } u_j = 0 \text{ for } j \notin \mathcal{A} \text{ and } \\ & u_j = u_\ell \text{ for } (j, \ell) \in \mathcal{B}, \\ \infty & \text{otherwise,} \end{cases}$$

where  $\mathbf{W}_{\mathcal{B}} \sim \mathcal{N}(\mathbf{0}_{s_0+1}, \mathbf{C}_{\mathcal{B}})$ . Clearly,  $\mathbf{V}$  has a unique minimum for  $\mathbf{u} \in \mathbb{R}^{p+1}$  such that  $u_j = 0$  for all  $j \notin \mathcal{A}$  and  $u_j = u_\ell$  for all  $(j, \ell) \in \mathcal{B}$  and  $\mathbf{u}_{\mathcal{B}} = -\mathbf{C}_{\mathcal{B}}^{-1} \mathbf{W}_{\mathcal{B}}$ . Since  $\mathbf{V}_n$  is convex we can proceed by using the epi-convergence results (6) to prove the asymptotic normality part (25).

Let us now turn our attention to the variable selection consistency. Namely, we have to show that  $\forall j \in \mathcal{A}, \mathbb{P}(j \in \mathcal{A}_n) \rightarrow 1$  and that  $\forall j \notin \mathcal{A}, \mathbb{P}(j \in \mathcal{A}_n) \rightarrow 0$ . The first claim is an easy consequence of the previous asymptotic result (25). To prove the second claim, consider an index

$j \notin \mathcal{A}$  and denote by  $C_j$  the subset of vertices constituting the connected component of  $G$  to which  $j$  belongs. Let  $C_j^0 = \{\ell \in C_j, \beta_\ell^* = 0\}$ ; clearly,  $j \in C_j^0$ . Our aim is to prove that  $\mathbb{P}(\ell \in \mathcal{A}_n) \rightarrow 0$ , for all  $\ell \in C_j^0$ . Observe that the KKT conditions write, for  $k = 1, \dots, p$ :

$$\left[ \nabla J(\hat{\beta}^{ad}) \right]_k = \lambda_n^{(1)} w_k^{(1)} r_k + \lambda_n^{(2)} \left( \sum_{(k,\ell) \in E} w_{k\ell}^{(2)} t_{k\ell} - \sum_{(\ell,k) \in E} w_{k\ell}^{(2)} t_{\ell k} \right)$$

where  $r_k = \text{sign}(\hat{\beta}_k^{ad})$  for  $\hat{\beta}_k^{ad} \neq 0$  and  $s_k$  is some real number in  $[-1, 1]$  if  $\hat{\beta}_k^{ad} = 0$ ; likewise, for any  $(k, \ell) \in E$ ,  $t_{k\ell} = \text{sign}(\hat{\beta}_k^{ad} - \hat{\beta}_\ell^{ad})$  for  $\hat{\beta}_k^{ad} \neq \hat{\beta}_\ell^{ad}$  and  $t_{k\ell}$  is some real number in  $[-1, 1]$  if  $\hat{\beta}_k^{ad} = \hat{\beta}_\ell^{ad}$ . Introducing the set  $\tilde{E} = \{(k, \ell) : (k, \ell) \in E \text{ or } (\ell, k) \in E\}$ , and setting  $t_{k\ell} = -t_{\ell k}$  for  $(\ell, k) \in E$ , we have the following more compact form for the KKT conditions:

$$\left[ \nabla J(\hat{\beta}^{ad}) \right]_k = \lambda_n^{(1)} w_k^{(1)} r_k + \lambda_n^{(2)} \sum_{(k,\ell) \in \tilde{E}} w_{k\ell}^{(2)} t_{k\ell},$$

where, in particular,  $t_{k\ell} = \text{sign}(\hat{\beta}_k^{ad} - \hat{\beta}_\ell^{ad})$  for any  $(k, \ell) \in \tilde{E}$  such that  $\hat{\beta}_k^{ad} \neq \hat{\beta}_\ell^{ad}$ . Next, since  $(\nabla J(\beta^*)/\sqrt{n}) \rightarrow_d \mathbf{W}$  (as shown in the proof of Theorem 1), the assumption **AL1** enables us to show that  $M_n(k) := [\nabla J(\hat{\beta}^{ad})]_k/\sqrt{n} = O_{\mathbb{P}}(1)$  as  $n \rightarrow \infty$  as well.

Let us now suppose that there exist some  $\ell \in C_j^0$  such that  $\hat{\beta}_\ell^{ad} \neq 0$ . In this case, either the set  $\mathcal{S}_{neg} = \{\ell \in C_j^0 : \hat{\beta}_\ell^{ad} < 0\}$  or the set  $\mathcal{S}_{pos} = \{\ell \in C_j^0 : \hat{\beta}_\ell^{ad} > 0\}$  is not empty (or both sets are not empty). If  $\mathcal{S}_{neg} \neq \emptyset$ , let  $b^{\min} = \min_{k \in \mathcal{S}_{neg}} \hat{\beta}_k^{ad}$ . Further denote by  $L$  the subset of  $\mathcal{S}_{neg}$  of connected indices  $\ell$  such that  $\hat{\beta}_\ell^{ad} = b^{\min}$ . Since  $\mathcal{S}_{neg} \neq \emptyset$ ,  $L$  has at least one element. Then, summing up the KKT conditions over  $L$ , we obtain

$$\begin{aligned} \sum_{k \in L} M_n(k) &= \frac{\lambda_n^{(1)}}{\sqrt{n}} n^{\gamma/2} \sum_{k \in L} \frac{r_k}{|\sqrt{n} \tilde{\beta}_k|^\gamma} \\ &+ \frac{\lambda_n^{(2)}}{\sqrt{n}} \sum_{k \in L} \sum_{(k,\ell) \in \tilde{E}, \beta_\ell^* \neq 0} \frac{t_{k\ell}}{|\tilde{\beta}_k - \tilde{\beta}_\ell|^\gamma} \\ &+ \frac{\lambda_n^{(2)}}{\sqrt{n}} n^{\gamma/2} \sum_{k \in L} \sum_{(k,\ell) \in \tilde{E}, \beta_\ell^* = 0 \text{ \& } \hat{\beta}_\ell^{ad} > b^{\min}} \frac{t_{k\ell}}{|\sqrt{n}(\tilde{\beta}_k - \tilde{\beta}_\ell)|^\gamma}. \end{aligned}$$

Since  $L \subset \mathcal{S}_{neg}$ ,  $r_k = -1$ , for all  $k \in L$ , and by definition of  $L$ ,  $t_{k\ell} = -1$  for all  $\ell$  such that  $\hat{\beta}_\ell^{ad} \neq b^{\min}$ . Moreover when  $\beta_\ell^* \neq 0$  then  $\beta_\ell^* \neq \beta_k^*$ , and  $\lambda_n^{(2)} t_{k\ell}/(\sqrt{n}|\tilde{\beta}_k - \tilde{\beta}_\ell|^\gamma) \rightarrow_{\mathbb{P}} 0$ , as  $n$  goes to  $\infty$ . Since  $\lambda_n^{(m)} n^{\gamma/2}/\sqrt{n}$  ( $m = 1, 2$ ) tends to  $\infty$ ,  $\sum_{k \in L} M_n(k)$  tends to  $-\infty$ , which contradicts  $M_n(\ell) = O_{\mathbb{P}}(1)$  for all  $\ell = 1, \dots, p$ . That leads to

$\mathbb{P}(\mathcal{S}_{neg} = \emptyset) \rightarrow 1$ . If  $\mathcal{S}_{neg} = \emptyset$ , then  $\mathcal{S}_{pos} \neq \emptyset$ , and similar arguments can be used (with maxima instead of minima) to get a contradiction. Putting all this together, we conclude that for all  $\ell \in C_j^0$ ,  $\mathbb{P}(\ell \in \mathcal{A}_n) \rightarrow 0$ .

It remains to show the consistency for the set  $\mathcal{B}_n$ . As for  $\mathcal{A}_n$ , we need to prove that  $\forall (j, \ell) \notin \mathcal{B}, \mathbb{P}((j, \ell) \in \mathcal{B}_n^c) \rightarrow 1$  and that  $\forall (j, \ell) \in \mathcal{B}, \mathbb{P}((j, \ell) \in \mathcal{B}_n^c) \rightarrow 0$ . Let us prove the first claim. If  $(j, \ell) \notin \mathcal{B}$  either  $(\beta_j^* = 0 \text{ and/or } \beta_\ell^* = 0)$ , or  $(\beta_j^* \neq 0, \beta_\ell^* \neq 0 \text{ and } \beta_j^* \neq \beta_\ell^*)$ . In the first case, when  $j \in \mathcal{A}^c$ , we have proved previously that  $\mathbb{P}(j \in \mathcal{A}_n^c) \rightarrow 1$ , so  $\mathbb{P}((j, \ell) \in \mathcal{B}_n^c) \rightarrow 1$ . In the second case, if  $(j, \ell) \in \mathcal{A}$ , and  $(j, \ell) \notin \mathcal{B}$ , the asymptotic normality result indicates that  $\hat{\beta}_j^{ad} - \hat{\beta}_\ell^{ad} \rightarrow_{\mathbb{P}} \beta_j^* - \beta_\ell^* \neq 0$ ; thus  $\mathbb{P}((j, \ell) \in \mathcal{B}_n^c) \rightarrow 1$ . Now let us prove the second claim, using KKT conditions as before. Let  $j$  be an index of  $\mathcal{A}$  such that for some  $\ell \in \mathcal{A}$  we have  $(j, \ell) \in \mathcal{B}$ . Then, for some  $1 \leq s(j) \leq s_0$ ,  $j \in \mathcal{A}_{s(j)}$ , where  $\mathcal{A}_s$  still denotes the set of vertices of the  $s$ -th connected component of  $G_{\mathcal{B}}$ . Suppose that there exists some  $\ell \in \mathcal{A}_{s(j)}$  such that  $\hat{\beta}_\ell^{ad} \neq \hat{\beta}_j^{ad}$ . As previously we define  $b^{\min} = \min_{k \in \mathcal{A}_{s(j)}} \hat{\beta}_k^{ad}$  and  $L$  the subset of  $\mathcal{A}_{s(j)}$  of connected indices  $\ell$  such that  $\hat{\beta}_\ell^{ad} = b^{\min}$ . Then, summing up the KKT conditions over  $L$ , we obtain

$$\begin{aligned} \sum_{k \in L} M_n(k) &= \frac{\lambda_n^{(1)}}{\sqrt{n}} \sum_{k \in L} \frac{r_k}{|\tilde{\beta}_k|^\gamma} \\ &+ \frac{\lambda_n^{(2)}}{\sqrt{n}} \sum_{k \in L} \sum_{(k,\ell) \in \tilde{E}, \beta_\ell^* \neq \beta_k^*} \frac{t_{k\ell}}{|\tilde{\beta}_k - \tilde{\beta}_\ell|^\gamma} \\ &+ \frac{\lambda_n^{(2)}}{\sqrt{n}} n^{\gamma/2} \sum_{k \in L} \sum_{(k,\ell) \in \tilde{E}, \beta_\ell^* = \beta_k^* \text{ \& } \hat{\beta}_\ell^{ad} > b^{\min}} \frac{t_{k\ell}}{|\sqrt{n}(\tilde{\beta}_k - \tilde{\beta}_\ell)|^\gamma}. \end{aligned}$$

Since  $L \subset \mathcal{A}$ , the first sum converges to 0 in probability. Moreover, the second sum also converges to 0 in probability, while the third sum tends to  $-\infty$ , which contradicts  $M_n(\ell) = O_{\mathbb{P}}(1)$  for all  $\ell = 1, \dots, p$ . We therefore conclude that  $\mathbb{P}((j, \ell) \in \mathcal{B}_n^c) \rightarrow 0$ , for all  $(j, \ell) \in \mathcal{B}$ , which completes the proof of Theorem 2.

## Acknowledgements

Part of this work was supported by the Interuniversity Attraction Pole (IAP) research network in Statistics P5/24, by the MSTIC project of the Joseph-Fourier University, and by the ABS4NGS ANR project (ANR-11-BINF-0001-06).

## References

1. Azencott, C.A., Grimm, D., Sugiyama, M., Kawahara, Y., Borgwardt, K.M.: Efficient network-guided multi-locus association mapping with graph cuts. *Bioinformatics* **29**(13), i171–179 (2013)



2. Chiquet, J., Smith, A., Grasseau, G., Matias, C., Ambroise, C.: SIMoNe: Statistical Inference for MOdular NETworks. *Bioinformatics* **25**(3), 417–418 (2009)
3. Danaher, P., Wang, P., Witten, D.: The joint graphical lasso for inverse covariance estimation across multiple classes. *Journal of the Royal Statistical Society, Series B (to appear)* (2013)
4. Franceschini, A., Szklarczyk, D., Frankild, S., Kuhn, M., Simonovic, M., Roth, A., Lin, J., Minguez, P., Bork, P., von Mering, C., et al.: String v9. 1: protein-protein interaction networks, with increased coverage and integration. *Nucleic acids research* **41**(D1), D808–D815 (2013)
5. Gertheiss, J., Tutz, G.: Regularization and model selection with categorical effect modifiers. *Statistica Sinica* **22**, 957–982 (2012)
6. Geyer, C.J.: On the Asymptotics of Constrained M-Estimation. *The Annals of Statistics* **22**, 1993–2010 (1994)
7. Girvan, M., Newman, M.E.: Community structure in social and biological networks. *Proc. Natl. Acad. Sci. U.S.A.* **99**(12), 7821–7826 (2002)
8. Gourieroux, C., Monfort, A.: Asymptotic properties of the maximum likelihood estimator in dichotomous logit models. *Journal of Econometrics* **17**, 83–97 (1981)
9. Guedj, M., Marisa, L., de Reynies, A., Orsetti, B., Schiappa, R., Bibeau, F., MacGrogan, G., Lerebours, F., Finetti, P., Longy, M., Bertheau, P., Bertrand, F., Bonnet, F., Martin, A.L., Feugeas, J.P., Bieche, I., Lehmann-Che, J., Lidereau, R., Birnbaum, D., Bertucci, F., de The, H., Theillet, C.: A refined molecular taxonomy of breast cancer. *Oncogene* **31**(9), 1196–1206 (2012)
10. Han, J.: Construction and analysis of web-based computer science information networks. In: *Rough Sets, Fuzzy Sets, Data Mining and Granular Computing*, pp. 1–2. Springer (2011)
11. Höfling, H., Binder, H., Schumacher, M.: A coordinate-wise optimization algorithm for the Fused Lasso. *Arxiv preprint arXiv:1011.6409* (2010)
12. Huang, J., Breheny, P., Ma, S.: A selective review of group selection in high dimensional models. *Statistical Science* **27**(4), 481–499 (2012)
13. McCullagh, P., Nelder, J.A.: *Generalized Linear Models*. 2nd ed. New-York : Chapman & Hall (1989)
14. Meier, L., van de Geer, S., Bühlmann, P.: The group lasso for logistic regression. *Journal of the Royal Statistical Society, Series B* **70**(1), 53–71 (2008)
15. Meinshausen, N.: Relaxed lasso. *Computational Statistics and Data Analysis* pp. 374–393 (2007)
16. Onisr: *La sécurité routière en France, bilan de l'année 2009*. Paris: La documentation Française (2010)
17. Rinaldo, A.: Properties and refinements of the fused lasso. *The Annals of Statistics* **37**(5B), 2922–2952 (2009)
18. Rual, J.F., Venkatesan, K., Hao, T., Hirozane-Kishikawa, T., Dricot, A., Li, N., Berriz, G.F., Gibbons, F.D., Dreze, M., Ayivi-Guedehoussou, N., Klitgord, N., Simon, C., Boxem, M., Milstein, S., Rosenberg, J., Goldberg, D.S., Zhang, L.V., Wong, S.L., Franklin, G., Li, S., Albala, J.S., Lim, J., Fraughton, C., Llamas, E., Cevik, S., Bex, C., Lamesch, P., Sikorski, R.S., Vandenhaute, J., Zoghbi, H.Y., Smolyar, A., Bosak, S., Sequerra, R., Doucette-Stamm, L., Cusick, M.E., Hill, D.E., Roth, F.P., Vidal, M.: Towards a proteome-scale map of the human protein-protein interaction network. *Nature* **437**(7062), 1173–1178 (2005)
19. Sharma, D., Bondell, H., Zhang, H.: Consistent group identification and variable selection in regression with correlated predictors. *Journal of Computational and Graphical Statistics* **22**, 319–340 (2013)
20. She, Y.: Sparse regression with exact clustering. *Electron. J. Statist.* **4**, 1055–1096 (2010)
21. Sun, H., Wang, S.: Penalized logistic regression for high-dimensional DNA methylation data with case-control studies. *Bioinformatics* **28**(10), 1368–1375 (2012)
22. Tibshirani, R.: Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B* **58**, 267–288 (1996)
23. Tibshirani, R., Saunders, M., Rosset, S., Zhu, J., Knight, K.: Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society Series B.* **67**, 91–108 (2005)
24. Vaiter, S., Peyr, G., Dossal, C., Fadili, J.: Robust sparse analysis regularization. *IEEE Transactions on Information Theory* **59**(4), 2001–2016 (2013). DOI 10.1109/TIT.2012.2233859. URL <http://hal.archives-ouvertes.fr/hal-00627452/>
25. Zou, H.: The Adaptive Lasso and Its Oracle Properties. *Journal of the American Statistical Association* **101**(476), 1418–1429 (2006)