



Adaptive Lasso and group-Lasso for functional Poisson regression

Stéphane Ivanoff, Franck Picard, Vincent Rivoirard

► To cite this version:

Stéphane Ivanoff, Franck Picard, Vincent Rivoirard. Adaptive Lasso and group-Lasso for functional Poisson regression. *Journal of Machine Learning Research*, 2016, 17 (55), pp.1–46. hal-01097914v2

HAL Id: hal-01097914

<https://hal.science/hal-01097914v2>

Submitted on 8 Jun 2016

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.

Public Domain

Adaptive Lasso and group-Lasso for functional Poisson regression

Stéphane Ivanoff

*CEREMADE UMR CNRS 7534
Université Paris Dauphine
F-75775 Paris, France*

IVANOFF@CEREMADE.DAUPHINE.FR

Franck Picard

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

FRANCK.PICARD@UNIV-LYON1.FR

Vincent Rivoirard

*CEREMADE UMR CNRS 7534
Université Paris Dauphine
F-75775 Paris, France*

RIVOIRARD@CEREMADE.DAUPHINE.FR

Editor: Zhihua Zhang

Abstract

High dimensional Poisson regression has become a standard framework for the analysis of massive counts datasets. In this work we estimate the intensity function of the Poisson regression model by using a dictionary approach, which generalizes the classical basis approach, combined with a Lasso or a group-Lasso procedure. Selection depends on penalty weights that need to be calibrated. Standard methodologies developed in the Gaussian framework can not be directly applied to Poisson models due to heteroscedasticity. Here we provide data-driven weights for the Lasso and the group-Lasso derived from concentration inequalities adapted to the Poisson case. We show that the associated Lasso and group-Lasso procedures satisfy fast and slow oracle inequalities. Simulations are used to assess the empirical performance of our procedure, and an original application to the analysis of Next Generation Sequencing data is provided.

Keywords: Functional Poisson regression, adaptive lasso, adaptive group-lasso, calibration, concentration

Introduction

Poisson functional regression has become a standard framework for image or spectra analysis, in which case observations are made of n independent couples $(Y_i, X_i)_{i=1,\dots,n}$, and can be modeled as

$$Y_i|X_i \sim \text{Poisson}(f_0(X_i)).$$

The X_i 's (random or fixed) are supposed to lie in a known compact support of \mathbb{R}^d ($d \geq 1$), say $[0, 1]^d$, and the purpose is to estimate the unknown intensity function f_0 assumed to be positive. Wavelets have been used extensively for intensity estimation, and the statistical challenge has been to propose thresholding procedures in the spirit of Donoho and Johnstone (1994), that were adapted to the variance's spatial variability associated with the Poisson framework. An early method to deal with high dimensional count data has been to apply a variance stabilizing-transform (see Anscombe (1948)) and to treat the transformed data as if they were Gaussian. More recently, the same idea has been applied to the data's decomposition in the Haar-wavelet basis, see Fryzlewicz and Nason (2004) and Fryzlewicz (2008), but these methods rely on asymptotic approximations and tend to show lower performance when the level of counts is low (see Besbeas et al. (2004)). Dedicated wavelet thresholding methods were developed in the Poisson setting by Kolaczyk (1999) and Sardy et al. (2004), and a recurrent challenge has been to define an appropriate threshold like the universal threshold for shrinkage and selection, as the heteroscedasticity of the model calls for component-wise thresholding.

In this work we first propose to enrich the standard wavelet approach by considering the so-called dictionary strategy. We assume that $\log f_0$ can be well approximated by a linear combination of p known functions, and we reduce the estimation of f_0 to the estimation of p coefficients. Dictionaries can be built from classical orthonormal systems such as wavelets, histograms or the Fourier basis, which results in a framework that encompasses wavelet methods. Considering overcomplete (*ie* redundant) dictionaries is efficient to capture different features in the signal, by using sparse representations (see Chen et al. (2001) or Tropp (2004)). For example, if $\log f_0$ shows piece-wise constant trends along with some periodicity, combining both Haar and Fourier bases will be more powerful than separate strategies, and the model will be sparse in the coefficients domain. To ensure sparse estimations, we consider the Lasso and the group-Lasso procedures. Group estimators are particularly well adapted to the dictionary framework, especially if we consider dictionaries based on a wavelet system, for which it is well known that coefficients can be grouped scale-wise for instance (see Chicken and Cai (2005)). Finally, even if we do not make any assumption on p itself, it may be larger than n and methodologies based on ℓ_1 -penalties, such as the Lasso and the group-Lasso appear appropriate.

The statistical properties of the Lasso are particularly well understood in the context of regression with *i.i.d.* errors, or for density estimation for which a range of oracle inequalities have been established. These inequalities, now widespread in the literature, provide theoretical error bounds that hold on events with a controllable (large) probability. See for instance Bertin et al. (2011), Bickel et al. (2009), Bunea et al. (2007a,b) and the references therein. For generalized linear models, Park and Hastie (2007) studied ℓ_1 -regularization path algorithms and van de Geer (2008) established non-asymptotic oracle inequalities. The sign consistency of the Lasso has been studied by Jia et al. (2013) for a very specific Poisson model. Finally, we also mention that the Lasso has also been extensively consid-

ered in survival analysis. See for instance Gaïffas and Guilloux (2012), Zou (2008), Kong and Nan (2014), Bradic et al. (2011), Lemler (2013) and Hansen et al. (2014).

Here we consider not only the Lasso estimator but also its extension, the group-Lasso proposed by Yuan and Lin (2006), which is relevant when the set of parameters can be partitioned into groups. The analysis of the group-Lasso has been led in different contexts. For instance, consistency has been studied by Bach (2008), Obozinski et al. (2011) and Wei and Huang (2010). In the linear model, Nardi and Rinaldo (2008) derived conditions ensuring various asymptotic properties such as consistency, oracle properties or persistence. Still for the linear model, Lounici et al. (2011) established oracle inequalities and, in the Gaussian setting, pointed out advantages of the group-Lasso with respect to the Lasso, generalizing the results of Chesneau and Hebiri (2008) and Huang and Zhang (2010). We also mention Meier et al. (2008) who studied the group-Lasso for logistic regression, Blazere et al. (2014) for generalized linear model with Poisson regression as a special case and Dalalyan et al. (2013) for other linear heteroscedastic models.

As pointed out by empirical comparative studies (see Besbeas et al. (2004)), the calibration of any thresholding rule is of central importance. Here we consider Lasso and group-Lasso penalties of the form

$$\text{pen}(\boldsymbol{\beta}) = \sum_{j=1}^p \lambda_j |\beta_j|$$

and

$$\text{pen}^g(\boldsymbol{\beta}) = \sum_{k=1}^K \lambda_k^g \|\boldsymbol{\beta}_{G_k}\|_2,$$

where $G_1 \cup \dots \cup G_K$ is a partition of $\{1, \dots, p\}$ into non-overlapping groups (see Section 1 for more details). By calibration we refer to the definition and to the suitable choice of the weights λ_j and λ_k^g , which is intricate in heteroscedastic models, especially for the group-Lasso. For functional Poissonian regression, the ideal shape of these weights is unknown, even if for the group-Lasso, the λ_k^g 's should of course depend on the groups size. As for the Lasso, most proposed weights in the literature are non-random and constant such that the penalty is proportional to $\|\boldsymbol{\beta}\|_1$, but when facing variable selection and consistency simultaneously, Zou (2006) showed the interest in considering non-constant data-driven ℓ_1 -weights even in the simple case where the noise is Gaussian with constant variance. This issue becomes even more critical in Poisson functional regression in which variance shows spatial heterogeneity. As Zou (2006), our first contribution is to propose here adaptive procedures with weights depending on the data. Weights λ_j for the Lasso are derived by using sharp concentration inequalities, in the same spirit as Bertin et al. (2011), Gaïffas and Guilloux (2012), Lemler (2013) and Hansen et al. (2014), but adapted to the Poissonian setting. To account for heteroscedasticity, weights λ_j are component-specific and depend

on the data (see Theorem 1). We propose a similar procedure for the calibration of the group-Lasso. In most proposed procedures, the analogs of the λ_k^g 's are proportional to the $\sqrt{|G_k|}$'s (see Nardi and Rinaldo (2008), Bühlmann and van de Geer (2011) or Blazere et al. (2014)). But to the best of our knowledge, adaptive group-Lasso procedures (with weights depending on the data) have not been proposed yet. This is the purpose of Theorem 2, which is the main result of this work, generalizing Theorem 1 by using sharp concentration inequalities for infinitely divisible vectors. We show the shape relevance of the data-driven weights λ_k^g by comparing them to the weights proposed by Lounici et al. (2011) in the Gaussian framework. In Theorem 2, we do not impose any condition on the groups size. However, whether $|G_k|$ is smaller than $\log p$ or not highly influences the order of magnitude of λ_k^g .

Our second contribution consists in providing the theoretical validity of our approach by establishing slow and fast oracle inequalities under RE-type conditions in the same spirit as Bickel et al. (2009). Closeness between our estimates and f_0 is measured by using the empirical Kullback-Leibler divergence. We show that classical oracle bounds are achieved. We also show the relevance of considering the group-Lasso instead of the Lasso in some situations. Our results, that are non-asymptotic, are valid under very general conditions on the design $(X_i)_{i=1,\dots,n}$ and on the dictionary. However, to shed some light on our results, we illustrate some of them in the asymptotic setting with classical dictionaries like wavelets, histograms or Fourier bases. Our approach generalizes the classical basis approach and in particular block wavelet thresholding which is equivalent to group-Lasso in that case (see Yuan and Lin (2006)). We refer the reader to Chicken and Cai (2005) for a deep study of block wavelet thresholding in the context of density estimation whose framework shows some similarities with ours in terms of heteroscedasticity. Note that sharp estimation of variance terms proposed in this work can be viewed as an extension of coarse bounds provided by Chicken and Cai (2005). Finally, we emphasize that our procedure differs from Blazere et al. (2014)'s one in several aspects: First, in their Poisson regression setting, they do not consider a dictionary approach. Furthermore, their weights are constant and not data-driven, so are strongly different from ours. Finally, rates of Blazere et al. (2014) are established under much stronger assumptions than ours (see Section 3.1 for more details).

Finally, we explore the empirical properties of our calibration procedures by using simulations. We show that our procedures are very easy to implement, and we compare their performance with variance-stabilizing transforms and cross-validation. The calibrated Lasso and group-Lasso are associated with excellent reconstruction properties, even in the case of low counts. We also propose an original application of functional Poisson regression to the analysis of Next Generation Sequencing data, with the denoising of a Poisson intensity applied to the detection of replication origins in the human genome (see Picard et al. (2014)).

This article is organized as follows. In Section 1, we introduce the Lasso and group-Lasso procedures we propose in the dictionary approach setting. In Section 2, we derive

data-driven weights of our procedures that are extensively commented. Theoretical performance of our estimates are studied in Section 3 in the oracle approach. In Section 4, we investigate the empirical performance of the proposed estimators using simulated data, and an application is provided on next generation sequencing data in Section 5.

1. Penalized log-likelihood estimates for Poisson regression and dictionary approach

We consider the functional Poisson regression model, with n observed counts $Y_i \in \mathbb{N}$ modeled such that:

$$Y_i|X_i \sim \text{Poisson}(f_0(X_i)), \quad (1.1)$$

with the X_i 's (random or fixed) supposed to lie in a known compact support, say $[0, 1]^d$. Since the goal here is to estimate the function f_0 assumed to be positive on $[0, 1]^d$, a natural candidate is a function f of the form $f = \exp(g)$. Then, we consider the so-called dictionary approach which consists in decomposing g as a linear combination of the elements of a given finite dictionary of functions denoted by $\Upsilon = \{\varphi_j\}_{j \in \mathcal{J}}$, with $\|\varphi_j\|_2 = 1$ for all j . Consequently, we choose g of the form:

$$g = \sum_{j \in \mathcal{J}} \beta_j \varphi_j,$$

with $p = \text{card}(\mathcal{J})$ that may depend on n (as well as the elements of Υ). Without loss of generality we will assume in the following that $\mathcal{J} = \{1, \dots, p\}$. In this framework, estimating f_0 is equivalent to estimating the vector of regression coefficients $\beta = (\beta_j)_{j \in \mathcal{J}} \in \mathbb{R}^p$. In the sequel, we write $g_\beta = \sum_{j \in \mathcal{J}} \beta_j \varphi_j$, $f_\beta = \exp(g_\beta)$, for all $\beta \in \mathbb{R}^p$. Note that we do not require the model to be true, that is we do not suppose the existence of β_0 such that $f_0 = f_{\beta_0}$.

The strength of the dictionary approach lies in its ability to capture different features of the function to estimate (smoothness, sparsity, periodicity,...) by sparse combinations of elements of the dictionary so that only few coefficients need to be selected, which limits estimation errors. Obviously, the dictionary approach encompasses the classical basis approach consisting in decomposing g on an orthonormal system. The richer the dictionary, the sparser the decomposition, so p can be larger than n and the model becomes high-dimensional.

We consider a likelihood-based penalized criterion to select β , the coefficients of the dictionary decomposition. We denote by \mathbf{A} the $n \times p$ -design matrix with $A_{ij} = \varphi_j(X_i)$, $\mathbf{Y} = (Y_1, \dots, Y_n)^T$ and the log-likelihood associated with this model is

$$l(\beta) = \sum_{j \in \mathcal{J}} \beta_j (\mathbf{A}^T \mathbf{Y})_j - \sum_{i=1}^n \exp\left(\sum_{j \in \mathcal{J}} \beta_j A_{ij}\right) - \sum_{i=1}^n \log(Y_i!),$$

which is a concave function of β . Next sections propose two different ways to penalize $-l(\beta)$.

1.1 The Lasso estimate

The first penalty we propose is based on the (weighted) ℓ_1 -norm and we obtain a Lasso-type estimate by considering

$$\widehat{\boldsymbol{\beta}}^L \in \underset{\boldsymbol{\beta} \in \mathbb{R}^p}{\operatorname{argmin}} \left\{ -l(\boldsymbol{\beta}) + \sum_{j=1}^p \lambda_j |\beta_j| \right\}. \quad (1.2)$$

The penalty term $\sum_{j=1}^p \lambda_j |\beta_j|$ depends on positive weights $(\lambda_j)_{j \in \mathcal{J}}$ that vary according to the elements of the dictionary and are chosen in Section 2.1. This choice of varying weights instead of a unique λ stems from heteroscedasticity due to the Poisson regression, and a first part of our work consists in providing theoretical data-driven values for these weights, in the same spirit as Bertin et al. (2011) or Hansen et al. (2014) for instance. From the first order optimality conditions (see Bühlmann and van de Geer (2011)), $\widehat{\boldsymbol{\beta}}^L$ satisfies

$$\begin{cases} \mathbf{A}_j^T (\mathbf{Y} - \exp(\mathbf{A} \widehat{\boldsymbol{\beta}}^L)) = \lambda_j \frac{\widehat{\beta}_j^L}{|\widehat{\beta}_j^L|} & \text{if } \widehat{\beta}_j^L \neq 0, \\ |\mathbf{A}_j^T (\mathbf{Y} - \exp(\mathbf{A} \widehat{\boldsymbol{\beta}}^L))| \leq \lambda_j & \text{if } \widehat{\beta}_j^L = 0, \end{cases}$$

where $\exp(\mathbf{A} \boldsymbol{\beta}) = (\exp((\mathbf{A} \boldsymbol{\beta})_1), \dots, \exp((\mathbf{A} \boldsymbol{\beta})_n))^T$ and \mathbf{A}_j is the j -th column of the matrix \mathbf{A} . Note that the larger the λ_j 's, the sparser the estimates. In particular $\widehat{\boldsymbol{\beta}}^L$ belongs to the set of the vectors $\boldsymbol{\beta} \in \mathbb{R}^p$ that satisfies for any $j \in \mathcal{J}$,

$$|\mathbf{A}_j^T (\mathbf{Y} - \exp(\mathbf{A} \boldsymbol{\beta}))| \leq \lambda_j. \quad (1.3)$$

The Lasso estimator of f_0 is now easily derived.

Definition 1 *The Lasso estimator of f_0 is defined as*

$$\widehat{f}^L(x) := \exp(\widehat{g}^L(x)) := \exp\left(\sum_{j=1}^p \widehat{\beta}_j^L \varphi_j(x)\right).$$

We also propose an alternative to \widehat{f}^L by considering the group-Lasso.

1.2 The group-Lasso estimate

We also consider the grouping of coefficients into non-overlapping blocks. Indeed, group estimates may be better adapted than their single counterparts when there is a natural group structure. The procedure keeps or discards all the coefficients within a block and can increase estimation accuracy by using information about coefficients of the same block. In our setting, we partition the set of indices $\mathcal{J} = \{1, \dots, p\}$ into K non-empty groups:

$$\{1, \dots, p\} = G_1 \cup G_2 \cup \dots \cup G_K.$$

For any $\beta \in \mathbb{R}^p$, β_{G_k} stands for the sub-vector of β with elements indexed by the elements of G_k , and we define the block ℓ_1 -norm on \mathbb{R}^p by

$$\|\beta\|_{1,2} = \sum_{k=1}^K \|\beta_{G_k}\|_2.$$

Similarly, \mathbf{A}_{G_k} is the $n \times |G_k|$ submatrix of \mathbf{A} whose columns are indexed by the elements of G_k . Then the group-Lasso $\hat{\beta}^{gL}$ is a solution to the following convex optimization problem:

$$\hat{\beta}^{gL} \in \operatorname{argmin}_{\beta \in \mathbb{R}^p} \left\{ -l(\beta) + \sum_{k=1}^K \lambda_k^g \|\beta_{G_k}\|_2 \right\},$$

where the λ_k^g 's are positive weights for which we also provide a theoretical data-driven expression in Section 2.2. This group-estimator is constructed similarly to the Lasso, with the block ℓ_1 -norm being used instead of the ℓ_1 -norm. In particular, note that if all groups are of size one then we recover the Lasso estimator. Convex analysis states that $\hat{\beta}^{gL}$ is a solution of the above optimization problem if the p -dimensional vector $\mathbf{0}$ is in the subdifferential of the objective function. Therefore, $\hat{\beta}^{gL}$ satisfies:

$$\begin{cases} \mathbf{A}_{G_k}^T (\mathbf{Y} - \exp(\mathbf{A} \hat{\beta}^{gL})) = \lambda_k^g \frac{\hat{\beta}_{G_k}^{gL}}{\|\hat{\beta}_{G_k}^{gL}\|_2} & \text{if } \hat{\beta}_{G_k}^{gL} \neq \mathbf{0}, \\ \|\mathbf{A}_{G_k}^T (\mathbf{Y} - \exp(\mathbf{A} \hat{\beta}^{gL}))\|_2 \leq \lambda_k^g & \text{if } \hat{\beta}_{G_k}^{gL} = \mathbf{0}. \end{cases}$$

This procedure naturally enhances group-sparsity as analyzed by Yuan and Lin (2006), Lounici et al. (2011) and references therein.

Obviously, $\hat{\beta}^{gL}$ belongs to the set of the vectors $\beta \in \mathbb{R}^p$ that satisfy for any $k \in \{1, \dots, K\}$,

$$\|\mathbf{A}_{G_k}^T (\mathbf{Y} - \exp(\mathbf{A} \beta))\|_2 \leq \lambda_k^g. \quad (1.4)$$

Now, we set

Definition 2 *The group Lasso estimator of f_0 is defined as*

$$\hat{f}^{gL}(x) := \exp(\hat{g}^{gL}(x)) := \exp\left(\sum_{j=1}^p \hat{\beta}_j^{gL} \varphi_j(x)\right).$$

In the following our results are given conditionally on the X_i 's, and \mathbb{E} (resp. \mathbb{P}) stands for the expectation (resp. the probability measure) conditionally on X_1, \dots, X_n . In some situations, to give orders of magnitudes of some expressions, we will use the following definition:

Definition 3 *We say that the design $(X_i)_{i=1,\dots,n}$ is regular if either the design is deterministic and the X_i 's are equispaced in $[0, 1]$ or the design is random and the X_i 's are i.i.d. with density h , with*

$$0 < \inf_{x \in [0,1]^d} h(x) \leq \sup_{x \in [0,1]^d} h(x) < \infty.$$

2. Weights calibration using concentration inequalities

Our first contribution is to derive theoretical data-driven values of the weights λ_j 's and λ_k^g 's, specially adapted to the Poisson model. In the classical Gaussian framework with noise variance σ^2 , weights for the Lasso are chosen to be proportional to $\sigma\sqrt{\log p}$ (see Bickel et al. (2009) for instance). The Poisson setting is more involved due to heteroscedasticity and such simple tuning procedures cannot be generalized easily. Sections 2.1 and 2.2 give closed forms of parameters λ_j and λ_k^g . They are based on concentration inequalities specific to the Poisson model. In particular, λ_j is used to control the fluctuations of $\mathbf{A}_j^T \mathbf{Y}$ around its mean, which enhances the key role of V_j , a variance term (the analog of σ^2) defined by

$$V_j = \text{Var}(\mathbf{A}_j^T \mathbf{Y}) = \sum_{i=1}^n f_0(X_i) \varphi_j^2(X_i). \quad (2.1)$$

2.1 Data-driven weights for the Lasso procedure

For any j , we choose a data-driven value for λ_j as small as possible so that with high probability, for any $j \in \mathcal{J}$,

$$|\mathbf{A}_j^T (\mathbf{Y} - \mathbb{E}[\mathbf{Y}])| \leq \lambda_j. \quad (2.2)$$

Such a control is classical for Lasso estimates (see the references above) and is also a key point of the technical arguments of the proofs. Requiring that the weights are as small as possible is justified, from the theoretical point of view, by oracle bounds depending on the λ_j 's (see Corollaries 1 and 2). Furthermore, as discussed in Bertin et al. (2011), choosing theoretical Lasso weights as small as possible is also a suitable guideline for practical purposes. Finally, note that if the model were true, *i.e.* if there existed a true sparse vector β_0 such that $f_0 = f_{\beta_0}$, then $\mathbb{E}[\mathbf{Y}] = \exp(\mathbf{A}\beta_0)$ and β_0 would belong to the set defined by (1.3) with large probability. The smaller the λ_j 's, the smaller the set within selection of $\hat{\beta}^L$ is performed. So, with a sharp control in (2.2), we increase the probability to select β_0 . The following theorem provides the data-driven weights λ_j 's. The main theoretical ingredient we use to choose the weights λ_j 's is a concentration inequality for Poisson processes and to proceed, we link the quantity $\mathbf{A}_j^T \mathbf{Y}$ to a specific Poisson process, as detailed in the proofs Section 7.1.

Theorem 1 *Let j be fixed and $\gamma > 0$ be a constant. Define $\widehat{V}_j = \sum_{i=1}^n \varphi_j^2(X_i)Y_i$ the natural unbiased estimator of V_j and*

$$\widetilde{V}_j = \widehat{V}_j + \sqrt{2\gamma \log p \widehat{V}_j \max_i \varphi_j^2(X_i) + 3\gamma \log p \max_i \varphi_j^2(X_i)}.$$

Set

$$\lambda_j = \sqrt{2\gamma \log p \widetilde{V}_j} + \frac{\gamma \log p}{3} \max_i |\varphi_j(X_i)|, \quad (2.3)$$

then

$$\mathbb{P}\left(|\mathbf{A}_j^T(\mathbf{Y} - \mathbb{E}[\mathbf{Y}])| \geq \lambda_j\right) \leq \frac{3}{p^\gamma}. \quad (2.4)$$

The first term $\sqrt{2\gamma \log p \widetilde{V}_j}$ in λ_j is the main one, and constitutes a variance term depending on \widetilde{V}_j that slightly overestimates V_j (see Section 7.1 for more details about the derivation of \widetilde{V}_j). Its dependence on an estimate of V_j was expected since we aim at controlling fluctuations of $\mathbf{A}_j^T \mathbf{Y}$ around its mean. The second term comes from the heavy tail of the Poisson distribution, and is the price to pay, in the non-asymptotic setting, for the added complexity of the Poisson framework compared to the Gaussian framework.

To shed more lights on the form of the proposed weights from the asymptotic point of view, assume that the design is regular (see Definition 3). In this case, it is easy to see that under mild assumptions on f_0 , V_j is asymptotically of order n . If we further assume that

$$\max_i |\varphi_j(X_i)| = o(\sqrt{n/\log p}), \quad (2.5)$$

then, when p is large, with high probability, \widehat{V}_j (and then \widetilde{V}_j) is also of order n (using Remark 2 in the proofs Section 7.1), and the second term in λ_j is negligible with respect to the first one. In this case, λ_j is of order $\sqrt{n \log p}$. Note that Assumption (2.5) is quite classical in heteroscedastic settings (see Bertin et al. (2011)). By taking the hyperparameter γ larger than 1, then for large values of p , (2.2) is true for any $j \in \mathcal{J}$, with large probability.

2.2 Data-driven weights for the group Lasso procedure

Current group-Lasso procedures are tuned by choosing the analog of λ_k^g proportional to $\sqrt{|G_k|}$ (see Nardi and Rinaldo (2008), Chapter 4 of Bühlmann and van de Geer (2011) or Blazere et al. (2014)). A more refined version of tuning group-Lasso is provided by Lounici et al. (2011) in the Gaussian setting (see below for a detailed discussion). To the best of our knowledge, data-driven weights (with theoretical validation) for the group-Lasso have not been proposed yet. It is the purpose of Theorem 2. Similarly to the previous section, we propose data-driven theoretical derivations for the weights λ_k^g 's that are chosen as small as possible, but satisfying for any $k \in \{1, \dots, K\}$,

$$\|\mathbf{A}_{G_k}^T(\mathbf{Y} - \mathbb{E}[\mathbf{Y}])\|_2 \leq \lambda_k^g \quad (2.6)$$

with high probability (see (1.4)). Choosing the smallest possible weights is also recommended by Lounici et al. (2011) in the Gaussian setting (see in their Section 3 the discussion about weights and comparisons with coarser weights of Nardi and Rinaldo (2008)). Obviously, λ_k^g should depend on sharp estimates of the variance parameters $(V_j)_{j \in G_k}$. The following theorem is the equivalent of Theorem 1 for the group-Lasso. Relying on specific concentration inequalities established for infinitely divisible vectors by Houdré et al. (2008), it requires a known upper bound for f_0 , which can be chosen as $\max_i Y_i$ in practice.

Theorem 2 *Let $k \in \{1, \dots, K\}$ be fixed and $\gamma > 0$ be a constant. Assume that there exists $M > 0$ such that for any x , $|f_0(x)| \leq M$. Let*

$$c_k = \sup_{\mathbf{x} \in \mathbb{R}^n} \frac{\|\mathbf{A}_{G_k} \mathbf{A}_{G_k}^T \mathbf{x}\|_2}{\|\mathbf{A}_{G_k}^T \mathbf{x}\|_2}. \quad (2.7)$$

For all $j \in G_k$, still with $\widehat{V}_j = \sum_{i=1}^n \varphi_j^2(X_i) Y_i$, define

$$\widetilde{V}_j^g = \widehat{V}_j + \sqrt{2(\gamma \log p + \log |G_k|) \widehat{V}_j \max_i \varphi_j^2(X_i)} + 3(\gamma \log p + \log |G_k|) \max_i \varphi_j^2(X_i). \quad (2.8)$$

Let $\gamma > 0$ be fixed. Define $b_k^i = \sqrt{\sum_{j \in G_k} \varphi_j^2(X_i)}$ and $b_k = \max_i b_k^i$. Finally, we set

$$\lambda_k^g = \left(1 + \frac{1}{2\sqrt{2\gamma \log p}}\right) \sqrt{\sum_{j \in G_k} \widetilde{V}_j^g} + 2\sqrt{\gamma \log p} D_k, \quad (2.9)$$

where $D_k = 8Mc_k^2 + 16b_k^2\gamma \log p$. Then,

$$\mathbb{P}\left(\|\mathbf{A}_{G_k}^T(\mathbf{Y} - \mathbb{E}[\mathbf{Y}])\|_2 \geq \lambda_k^g\right) \leq \frac{2}{p^\gamma}. \quad (2.10)$$

Similarly to the weights λ_j 's of the Lasso, each weight λ_k^g is the sum of two terms. The term \widetilde{V}_j^g is an estimate of V_j so it plays the same role as \widehat{V}_j . In particular, \widetilde{V}_j^g and \widehat{V}_j are of the same order since $\log |G_k|$ is not larger than $\log p$. The first term in λ_k^g is a variance term, and the leading constant $1 + 1/(2\sqrt{2\gamma \log p})$ is close to 1 when p is large. So, the first term is close to the square root of the sum of sharp estimates of the $(V_j)_{j \in G_k}$, as expected for a grouping strategy (see Chicken and Cai (2005)).

The second term, namely $2\sqrt{\gamma \log p} D_k$, is more involved. To shed light on it, since b_k and c_k play a key role, we first state the following proposition controlling values of these terms.

Proposition 1 *Let k be fixed. We have*

$$b_k \leq c_k \leq \sqrt{n} b_k. \quad (2.11)$$

Furthermore,

$$c_k^2 \leq \max_{j \in G_k} \sum_{j' \in G_k} \left| \sum_{l=1}^n \varphi_j(X_l) \varphi_{j'}(X_l) \right|. \quad (2.12)$$

The first inequality of Proposition 1 shows that $2\sqrt{\gamma \log p D_k}$ is smaller than $c_k \sqrt{\log p} + b_k \log p \leq 2c_k \log p$ up to a constant depending on γ and M . At first glance, the second inequality of Proposition 1 shows that c_k is controlled by the coherence of the dictionary (see Tropp (2004)) and b_k depends on $(\max_i |\varphi_j(X_i)|)_{j \in G_k}$. In particular, if for a given block G_k , the functions $(\varphi_j)_{j \in G_k}$ are orthonormal, then for fixed $j \neq j'$, if the X_i 's are deterministic and equispaced on $[0, 1]$ or if the X_i 's are i.i.d. with a uniform density on $[0, 1]^d$, then, when n is large

$$\frac{1}{n} \sum_{l=1}^n \varphi_j(X_l) \varphi_{j'}(X_l) \approx \int \varphi_j(x) \varphi_{j'}(x) dx = 0$$

and we expect

$$c_k^2 \lesssim \max_{j \in G_k} \sum_{l=1}^n \varphi_j^2(X_l).$$

In any case, by using the Cauchy-Schwarz Inequality, Condition (2.12) gives

$$c_k^2 \leq \max_{j \in G_k} \sum_{j' \in G_k} \left(\sum_{l=1}^n \varphi_j^2(X_l) \right)^{1/2} \left(\sum_{l=1}^n \varphi_{j'}^2(X_l) \right)^{1/2}. \quad (2.13)$$

To further discuss orders of magnitude for the c_k 's, we consider the following condition

$$\max_{j \in G_k} \sum_{l=1}^n \varphi_j^2(X_l) = O(n), \quad (2.14)$$

which is satisfied for instance for fixed k if the design is regular, since $\|\varphi_j\|_2 = 1$. Under Assumption (2.14), Inequality (2.13) gives

$$c_k^2 = O(|G_k|n).$$

We can say more on b_k and c_k (and then on the order of magnitude of λ_k^g) by considering classical dictionaries of the literature to build the blocks G_k , which is of course realized in practice. In the subsequent discussions, the balance between $|G_k|$ and $\log p$ plays a key role. Note also that $\log p$ is the group size often recommended in the classical setting ($p = n$) for block thresholding (see Theorem 1 of Chicken and Cai (2005)).

2.2.1 ORDER OF MAGNITUDE OF λ_k^g BY CONSIDERING CLASSICAL DICTIONARIES.

Let G_k be a given block and assume that it is built by using only one of the subsequent systems. For each example, we discuss the order of magnitude of the term $D_k = 8Mc_k^2 + 16b_k^2\gamma \log p$. For ease of exposition, we assume that f_0 is supported by $[0, 1]$ but we could easily generalize the following discussion to the multidimensional setting.

Bounded dictionary. Similarly to Blazere et al. (2014), we assume that there exists a constant L not depending on n and p such that for any $j \in G_k$, $\|\varphi_j\|_\infty \leq L$. For instance, atoms of the Fourier basis satisfy this property. We then have

$$b_k^2 \leq L^2 |G_k|.$$

Finally, under Assumption (2.14),

$$D_k = O(|G_k|n + |G_k| \log p). \quad (2.15)$$

Compactly supported wavelets. Consider the one-dimensional Haar dictionary: For $j = (j_1, k_1) \in \mathbb{Z}^2$ we set $\varphi_j(x) = 2^{j_1/2} \psi(2^{j_1}x - k_1)$, $\psi(x) = 1_{[0, 0.5]}(x) - 1_{]0.5, 1]}(x)$. Assume that the block G_k depends on only one resolution level j_1 : $G_k = \{j = (j_1, k_1) : k_1 \in B_{j_1}\}$, where B_{j_1} is a subset of $\{0, 1, \dots, 2^{j_1} - 1\}$. In this case, since for $j, j' \in G_k$ with $j \neq j'$, for any x , $\varphi_j(x)\varphi_{j'}(x) = 0$,

$$b_k^2 = \max_i \sum_{j \in G_k} \varphi_j^2(X_i) = \max_{i, j \in G_k} \varphi_j^2(X_i) = 2^{j_1}$$

and Inequality (2.12) gives

$$c_k^2 \leq \max_{j \in G_k} \sum_{l=1}^n \varphi_j^2(X_l).$$

If, similarly to Condition (2.5), we assume that $\max_{i, j \in G_k} |\varphi_j(X_i)| = o(\sqrt{n/\log p})$, then

$$b_k^2 = o(n/\log p),$$

and under Assumption (2.14),

$$D_k = O(n),$$

which improves (2.15). This property can be easily extended to general compactly supported wavelets ψ , since, in this case, for any $j = (j_1, k_1)$

$$S_j = \{j' = (j_1, k'_1) : k'_1 \in \mathbb{Z}, \varphi_j \times \varphi_{j'} \not\equiv 0\}$$

is finite with cardinal only depending on the support of ψ .

Regular histograms. Consider a regular grid of the interval $[0, 1]$, $\{0, \delta, 2\delta, \dots\}$ with $\delta > 0$. Consider then $(\varphi_j)_{j \in G_k}$ such that for any $j \in G_k$, there exists ℓ such that $\varphi_j = \delta^{-1/2} 1_{(\delta(\ell-1), \delta\ell]}$. We have $\|\varphi_j\|_2 = 1$ and $\|\varphi_j\|_\infty = \delta^{-1/2}$. As for the wavelet case, for $j, j' \in G_k$ with $j \neq j'$, for any x , $\varphi_j(x)\varphi_{j'}(x) = 0$, then

$$b_k^2 = \max_i \sum_{j \in G_k} \varphi_j^2(X_i) = \max_{i, j \in G_k} \varphi_j^2(X_i) = \delta^{-1}.$$

If, similarly to Condition (2.5), we assume that $\max_{i, j \in G_k} |\varphi_j(X_i)| = o(\sqrt{n/\log p})$, then

$$b_k^2 = o(n/\log p),$$

and under Assumption (2.14),

$$D_k = O(n).$$

The previous discussion shows that we can exhibit dictionaries such that c_k^2 and D_k are of order n and the term $b_k^2 \log p$ is negligible with respect to c_k^2 . Then, if similarly to Section 2.1, the terms $(\tilde{V}_j^g)_{j \in G_k}$ are all of order n , λ_k^g is of order $\sqrt{n \times \max(\log p, |G_k|)}$ and the main term in λ_k^g is the first one as soon as $|G_k| \geq \log p$. In this case, λ_k^g is of order $\sqrt{|G_k|n}$.

2.2.2 COMPARISON WITH THE GAUSSIAN FRAMEWORK.

Now, let us compare the λ_k^g 's to the weights proposed by Lounici et al. (2011) in the Gaussian framework. Adapting their notations to ours, Lounici et al. (2011) estimate the vector β_0 in the model $\mathbf{Y} \sim \mathcal{N}(\mathbf{A}\beta_0, \sigma^2 \mathbf{I}_n)$ by using the group-Lasso estimate with weights equal to

$$\tilde{\lambda}_k^g = 2\sqrt{\sigma^2 \left(\text{Tr}(\mathbf{A}_{G_k}^T \mathbf{A}_{G_k}) + 2|||\mathbf{A}_{G_k}^T \mathbf{A}_{G_k}||| (2\gamma \log p + \sqrt{|G_k| \gamma \log p}) \right)},$$

where $|||\mathbf{A}_{G_k}^T \mathbf{A}_{G_k}|||$ denotes the maximal eigenvalue of $\mathbf{A}_{G_k}^T \mathbf{A}_{G_k}$ (see (3.1) in Lounici et al. (2011)). So, if $|G_k| \leq \log p$, the above expression is of the same order as

$$\sqrt{\sigma^2 \text{Tr}(\mathbf{A}_{G_k}^T \mathbf{A}_{G_k})} + \sqrt{\sigma^2 |||\mathbf{A}_{G_k}^T \mathbf{A}_{G_k}||| \gamma \log p}. \quad (2.16)$$

Neglecting the term $16b_k^2 \gamma \log p$ in the definition of D_k (see the discussion in Section 2.2.1), we observe that λ_k^g is of the same order as

$$\sqrt{\sum_{j \in G_k} \tilde{V}_j^g} + \sqrt{M c_k^2 \gamma \log p}. \quad (2.17)$$

Since M is an upper bound of $\text{Var}(Y_i) = f_0(X_i)$ for any i , strong similarities can be highlighted between the forms of the weights in the Poisson and Gaussian settings:

- For the first terms, \tilde{V}_j^g is an estimate of V_j and

$$\sum_{j \in G_k} V_j \leq M \sum_{j \in G_k} \sum_{i=1}^n \varphi_j^2(X_i) = M \times \text{Tr}(\mathbf{A}_{G_k}^T \mathbf{A}_{G_k}).$$

- For the second terms, in view of (2.7), c_k^2 is related to $\|\mathbf{A}_{G_k}^T \mathbf{A}_{G_k}\|$ since we have

$$c_k^2 = \sup_{\mathbf{x} \in \mathbb{R}^n} \frac{\|\mathbf{A}_{G_k} \mathbf{A}_{G_k}^T \mathbf{x}\|_2^2}{\|\mathbf{A}_{G_k}^T \mathbf{x}\|_2^2} \leq \sup_{\mathbf{y} \in \mathbb{R}^{|G_k|}} \frac{\|\mathbf{A}_{G_k} \mathbf{y}\|_2^2}{\|\mathbf{y}\|_2^2} = \|\mathbf{A}_{G_k}^T \mathbf{A}_{G_k}\|.$$

These strong similarities between the Gaussian and the Poissonian settings strongly support the shape relevance of the weights we propose.

2.2.3 SUBOPTIMALITY OF THE NAIVE PROCEDURE

Finally, we show that the naive procedure that considers $\sqrt{\sum_{j \in G_k} \lambda_j^2}$ instead of λ_k^g is suboptimal even if, obviously due to Theorem 1, with high probability,

$$\|\mathbf{A}_{G_k}^T (\mathbf{Y} - \mathbb{E}[\mathbf{Y}])\|_2 \leq \sqrt{\sum_{j \in G_k} \lambda_j^2}.$$

Suboptimality is justified by following heuristic arguments. Assume that for all j and k , the first terms in (2.3) and (2.9) are the main ones and $\tilde{V}_j \approx \tilde{V}_j^g \approx V_j$. Then by considering λ_k^g instead of $\sqrt{\sum_{j \in G_k} \lambda_j^2}$, we improve our weights by the factor $\sqrt{\log p}$, since in this situation,

$$\lambda_k^g \approx \sqrt{\sum_{j \in G_k} V_j}$$

and

$$\sqrt{\sum_{j \in G_k} \lambda_j^2} \approx \sqrt{\log p \sum_{j \in G_k} V_j} \approx \sqrt{\log p} \lambda_k^g.$$

Remember that our previous discussion shows the importance to consider weights as small as possible as soon as (2.6) is satisfied with high probability. The next section will confirm this point.

3. Oracle inequalities

In this section, we establish oracle inequalities to study theoretical properties of our estimation procedures. The X_i 's are still assumption-free, and the performance of our procedures will be only evaluated at the X_i 's. To measure the closeness between f_0 and an estimate,

we use the empirical Kullback-Leibler divergence associated with our model, denoted by $K(\cdot, \cdot)$. Straightforward computations (see for instance Leblanc and Letué (2006)) show that for any positive function f ,

$$\begin{aligned} K(f_0, f) &= \mathbb{E} \left[\log \left(\frac{\mathcal{L}(f_0)}{\mathcal{L}(f)} \right) \right] \\ &= \sum_{i=1}^n [(f_0(X_i) \log f_0(X_i) - f_0(X_i))] - [(f_0(X_i) \log f(X_i) - f(X_i))], \end{aligned}$$

where $\mathcal{L}(f)$ is the likelihood associated with f . We speak about *empirical divergence* to emphasize its dependence on the X_i 's. Note that we can write

$$K(f_0, f) = \sum_{i=1}^n f_0(X_i)(e^{u_i} - u_i - 1), \quad (3.1)$$

where $u_i = \log \frac{f(X_i)}{f_0(X_i)}$. This expression clearly shows that $K(f_0, f)$ is non-negative and $K(f_0, f) = 0$ if and only if for all $i \in \{1, \dots, n\}$, we have $u_i = 0$, that is $f(X_i) = f_0(X_i)$ for all $i \in \{1, \dots, n\}$.

Remark 1 *To weaken the dependence on n in the asymptotic setting, an alternative, not considered here, would consist in considering $n^{-1}K(\cdot, \cdot)$ instead of $K(\cdot, \cdot)$.*

If the classical \mathbb{L}_2 -norm is the natural loss-function for penalized least squares criteria, the empirical Kullback-Leibler divergence is a natural alternative for penalized likelihood criteria. In next sections, oracle inequalities will be expressed by using $K(\cdot, \cdot)$.

3.1 Oracle inequalities for the group-Lasso estimate

In this section, we state oracle inequalities for the group-Lasso. These results can be viewed as generalizations of results by Lounici et al. (2011) to the case of the Poisson regression model. They will be established on the set Ω_g where

$$\Omega_g = \left\{ \|\mathbf{A}_{G_k}^T(\mathbf{Y} - \mathbb{E}[\mathbf{Y}])\|_2 \leq \lambda_k^g \quad \forall k \in \{1, \dots, K\} \right\}. \quad (3.2)$$

Under assumptions of Theorem 2, we have $\mathbb{P}(\Omega_g) \geq 1 - \frac{2K}{p^\gamma} \geq 1 - 2p^{1-\gamma}$. By considering $\gamma > 1$, we have that $\mathbb{P}(\Omega_g)$ goes to 1 at a polynomial rate of convergence when p goes to $+\infty$. Even if our procedure is applied with weights defined in Section 2.2, note that subsequent oracle inequalities would hold for any weights $(\lambda_k^g)_{k=1, \dots, K}$ such that the associated set Ω_g has high probability. For any $\beta \in \mathbb{R}^p$, we denote by

$$f_\beta(x) = \exp \left(\sum_{j=1}^p \beta_j \varphi_j(x) \right),$$

the candidate associated with β to estimate f_0 . We first give a *slow oracle inequality* (see for instance Bunea et al. (2007a), Gaïffas and Guilloux (2012) or Lounici et al. (2011)) that does not require any assumption.

Theorem 3 *On Ω_g ,*

$$K(f_0, \hat{f}^{gL}) \leq \inf_{\beta \in \mathbb{R}^p} \left\{ K(f_0, f_\beta) + 2 \sum_{k=1}^K \lambda_k^g \|\beta_{G_k}\|_2 \right\}. \quad (3.3)$$

Note that

$$\sum_{k=1}^K \lambda_k^g \|\beta_{G_k}\|_2 \leq \max_{k \in \{1, \dots, K\}} \lambda_k^g \times \|\beta\|_{1,2}$$

and (3.3) is then similar to Inequality (3.9) of Lounici et al. (2011). We can improve the rate of (3.3) at the price of stronger assumptions on the matrix \mathbf{A} . We consider the following assumptions:

Assumption 1. We assume that

$$m := \max_{i \in \{1, \dots, n\}} |\log f_0(X_i)| < \infty.$$

This assumption is equivalent to assuming that f_0 is bounded from below and from above by positive constants. We do not assume that m is known in the sequel.

Assumption 2. For some integer $s \in \{1, \dots, K\}$ and some constant r , the following condition holds:

$$0 < \kappa_n(s, r) := \min \left\{ \frac{(\beta^T \mathbf{G} \beta)^{1/2}}{\|\beta_J\|_2} : |J| \leq s, \beta \in \mathbb{R}^p \setminus \{0\}, \sum_{k \in J^c} \lambda_k^g \|\beta_{G_k}\|_2 \leq r \sum_{k \in J} \lambda_k^g \|\beta_{G_k}\|_2 \right\},$$

where \mathbf{G} is the Gram matrix defined by $\mathbf{G} = \mathbf{A}^T \mathbf{C} \mathbf{A}$, where \mathbf{C} is the diagonal matrix with $C_{i,i} = f_0(X_i)$. With a slight abuse, β_J stands for the sub-vector of β with elements indexed by the indices of the groups $(G_k)_{k \in J}$.

This assumption is the natural extension of the classical *Restricted Eigenvalue condition* introduced by Bickel et al. (2009) to study the Lasso estimate where the ℓ_1 -norm is replaced with the weighted $\|\cdot\|_{1,2}$ -norm. In the Gaussian setting, Lounici et al. (2011) considered similar conditions to establish oracle inequalities for their group-Lasso procedure (see their Assumption (3.1)). RE-type assumptions are among the mildest ones to establish oracle inequalities (see van de Geer and Bühlmann (2009)). In particular, if the Gram matrix G has a positive minimal eigenvalue, say $d_{G,min}$, then Assumption 2 is satisfied with $\kappa_n^2(s, r) \geq d_{G,min}$. Assumption 2 can be connected to stronger coherence type conditions involving

the ratio $\frac{\max_k \lambda_k^g}{\min_k \lambda_k^g}$ (see Appendix B.3. of Lounici et al. (2011)). Furthermore, if c_0 is a positive lower bound for f_0 , then for all $\beta \in \mathbb{R}^p$,

$$\beta^T \mathbf{G} \beta = (\mathbf{A} \beta)^T \mathbf{C} (\mathbf{A} \beta) \geq c_0 \|\mathbf{A} \beta\|_2^2 = c_0 \sum_{i=1}^n \left(\sum_{j=1}^p \beta_j \varphi_j(X_i) \right)^2 = c_0 \sum_{i=1}^n g_\beta^2(X_i),$$

with $g_\beta = \sum_{j=1}^p \beta_j \varphi_j$. If $(\varphi_j)_{j \in \mathcal{J}}$ is orthonormal on $[0, 1]^d$ and if the design is regular, then the last term is the same order as

$$n \int g_\beta^2(x) dx = n \|\beta\|_2^2 \geq n \|\beta_J\|_2^2$$

for any subset $J \subset \{1, \dots, K\}$. Under these assumptions, $\kappa_n^{-2}(s, r) = O(n^{-1})$.

We now introduce for any $\mu > 0$

$$\Gamma(\mu) = \left\{ \beta \in \mathbb{R}^p : \max_{i \in \{1, \dots, n\}} \left| \sum_{j=1}^p \beta_j \varphi_j(X_i) \right| \leq \mu \right\}.$$

In the sequel, we restrict our attention to estimates belonging to the convex set $\Gamma(\mu)$. Of course, if m were known we would take $\mu = m$ (or μ a bit larger than m). Note that we do not impose any upper bound on μ so this condition is quite mild. The role of $\Gamma(\mu)$ consists in connecting $K(.,.)$ to some empirical quadratic loss functions (see the proof of Theorem 4). Alternative stronger assumptions have been considered by Lemler (2013) relying on van de Geer (2008) and Kong and Nan (2014). The value of μ only influences constants in subsequent oracle inequalities.

We consider the slightly modified group-Lasso estimate. Let $\alpha > 1$ and let us set

$$\hat{\beta}_{\mu, \alpha}^{gL} \in \operatorname{argmin}_{\beta \in \Gamma(\mu)} \left\{ -l(\beta) + \alpha \sum_{k=1}^K \lambda_k^g \|\beta_{G_k}\|_2 \right\}, \quad \hat{f}_{\mu, \alpha}^{gL}(x) = \exp \left(\sum_{j=1}^p (\hat{\beta}_{\mu, \alpha}^{gL})_j \varphi_j(x) \right)$$

for which we obtain the following fast oracle inequality.

Theorem 4 *Let $\varepsilon > 0$ and s a positive integer. Let Assumption 2 be satisfied with s and*

$$r = \frac{\alpha + 1 + 2\alpha/\varepsilon}{\alpha - 1}.$$

Then, under Assumption 1, there exists a constant $B(\varepsilon, m, \mu)$ depending on ε , m and μ such that, on Ω_g ,

$$K(f_0, \hat{f}_{\mu, \alpha}^{gL}) \leq (1 + \varepsilon) \inf_{\substack{\beta \in \Gamma(\mu) \\ |J(\beta)| \leq s}} \left\{ K(f_0, f_\beta) + B(\varepsilon, m, \mu) \frac{\alpha^2 |J(\beta)|}{\kappa_n^2} \times \left(\max_{k \in \{1, \dots, K\}} \lambda_k^g \right)^2 \right\}, \quad (3.4)$$

where κ_n stands for $\kappa_n(s, r)$, and $J(\beta)$ is the subset of $\{1, \dots, K\}$ such that $\beta_{G_k} = \mathbf{0}$ if and only if $k \notin J(\beta)$.

The estimate $\hat{\beta}_{\mu,\alpha}^{gL}$ studied in Theorem 4 slightly differs from the estimate $\hat{\beta}^{gL}$ introduced in Section 1.2 since it depends on α and minimization is only performed on $\Gamma(\mu)$. Both estimates coincide when $\mu = +\infty$ and when $\alpha = 1$. An examination of the proof of Theorem 4 shows that $\lim_{\mu \rightarrow +\infty} B(\varepsilon, m, \mu) = +\infty$. In practice, we observe that most of the time, with $\alpha = 1$ and μ large enough, $\hat{\beta}_{\mu,\alpha}^{gL}$ and $\hat{\beta}^{gL}$ coincide. So, when α is close to 1 and μ is large, Theorem 4 gives a good flavor of theoretical performances satisfied by our group-Lasso procedure.

Let us comment each term of the right-hand side of (3.4). The first term is an approximation term, which can vanish if f_0 can be decomposed on the dictionary. The second term is a variance term, according to the usual terminology, which is proportional to the size of $J(\beta)$. Its shape is classical in the high dimensional setting. See for instance Theorem 3.2 of Lounici et al. (2011) for the group-Lasso in linear models, or Theorem 6.1 of Bickel et al. (2009) and Theorem 3 of Bertin et al. (2011) for the Lasso. If the order of magnitude of λ_k^g is $\sqrt{n \times \max(\log p; |G_k|)}$ (see Section 2.2.1) and if $\kappa_n^{-2} = O(n^{-1})$, the order of magnitude of this variance term is not larger than $|J(\beta)| \times \max(\log p; |G_k|)$. Finally, if f_0 can be well approximated (for the empirical Kullback-Leibler divergence) by a group-sparse combination of the functions of the dictionary, then the right hand side of (3.4) will take small values. So, the previous result justifies our group-Lasso procedure from the theoretical point of view. Note that (3.3) and (3.4) also show the interest of considering weights as small as possible.

Blazere et al. (2014) established rates of convergence under stronger assumptions, namely all coordinates of the analog of \mathbf{A} are bounded by a quantity L , where L is viewed as a constant. Rates depend on L in an exponential manner and would highly deteriorate if L depended on n and p . So, this assumption is not reasonable if we consider dictionaries such as wavelets or histograms (see Section 2.2.1).

3.2 Oracle inequalities for the Lasso estimate

For the sake of completeness, we provide oracle inequalities for the Lasso. Theorems 3 and 4 that deal with the group-Lasso estimate can be adapted to the non-grouping strategy when we take groups of size 1. Subsequent results are similar to those established by Lemler (2013) who studied the Lasso estimate for the high-dimensional Aalen multiplicative intensity model. The block ℓ_1 -norm $\|\cdot\|_{1,2}$ becomes the usual ℓ_1 -norm and the group support $J(\beta)$ is simply the support of β . As previously, we only work on the probability set Ω defined by

$$\Omega = \left\{ |\mathbf{A}_j^T(\mathbf{Y} - \mathbb{E}[\mathbf{Y}])| \leq \lambda_j \quad \forall j \in \{1, \dots, p\} \right\}. \quad (3.5)$$

Theorem 1 asserts that $\mathbb{P}(\Omega) \geq 1 - \frac{3}{p^{\gamma-1}}$ that goes to 1 as soon as $\gamma > 1$. As previously, subsequent oracle inequalities would hold for any weights $(\lambda_j)_{j=1,\dots,p}$ such that the asso-

ciated probability set Ω has high probability. We obtain a slow oracle inequality for \widehat{f}^L :

Corollary 1 *On Ω ,*

$$K(f_0, \widehat{f}^L) \leq \inf_{\beta \in \mathbb{R}^p} \left\{ K(f_0, f_\beta) + 2 \sum_{j=1}^p \lambda_j |\beta_j| \right\}.$$

Now, let us consider fast oracle inequalities. In this framework, Assumption 2 is replaced with the following:

Assumption 3. For some integer $s \in \{1, \dots, p\}$ and some constant r , the following condition holds:

$$0 < \kappa_n(s, r) := \min \left\{ \frac{(\beta^T \mathbf{G} \beta)^{1/2}}{\|\beta_J\|_2} : |J| \leq s, \beta \in \mathbb{R}^p \setminus \{0\}, \sum_{j \in J^c} \lambda_j |\beta_j| \leq r \sum_{j \in J} \lambda_j |\beta_j| \right\},$$

where \mathbf{G} is the Gram matrix defined by $\mathbf{G} = \mathbf{A}^T \mathbf{C} \mathbf{A}$, where \mathbf{C} is the diagonal matrix with $C_{i,i} = f_0(X_i)$.

As previously, we consider the slightly modified Lasso estimate. Let $\alpha > 1$ and let us set

$$\widehat{\beta}_{\mu, \alpha}^L \in \operatorname{argmin}_{\beta \in \Gamma(\mu)} \left\{ -l(\beta) + \alpha \sum_{j=1}^p \lambda_j |\beta_j| \right\}, \quad \widehat{f}_{\mu, \alpha}^L(x) = \exp \left(\sum_{j=1}^p (\widehat{\beta}_{\mu, \alpha}^L)_j \varphi_j(x) \right)$$

for which we obtain the following fast oracle inequality.

Corollary 2 *Let $\varepsilon > 0$ and s a positive integer. Let Assumption 3 be satisfied with s and*

$$r = \frac{\alpha + 1 + 2\alpha/\varepsilon}{\alpha - 1}.$$

Then, under Assumption 1, there exists a constant $B(\varepsilon, m, \mu)$ depending on ε , m and μ such that, on Ω ,

$$K(f_0, \widehat{f}_{\mu, \alpha}^L) \leq (1 + \varepsilon) \inf_{\substack{\beta \in \Gamma(\mu) \\ |J(\beta)| \leq s}} \left\{ K(f_0, f_\beta) + B(\varepsilon, m, \mu) \frac{\alpha^2 |J(\beta)|}{\kappa_n^2} \left(\max_{j \in \{1, \dots, p\}} \lambda_j^2 \right) \right\},$$

where κ_n stands for $\kappa_n(s, r)$, and $J(\beta)$ is the support of β .

This corollary is derived easily from Theorem 4 by considering all groups of size 1. Comparing Corollary 2 and Theorem 4, we observe that the group-Lasso can improve the Lasso estimate when the function f_0 can be well approximated by a function f_β so that the number of non-zero groups of β is much smaller than the total number of non-zero coefficients. The simulation study of the next section illustrates this comparison from the numerical point of view.

4. Simulation study

Simulation settings. Even if our theoretical results do not concern model selection properties, we propose to start the simulation study by considering a simple toy example such that $\log(f_0) = \mathbf{A}\beta_0$ is generated from a true sparse vector of coefficients of size $p = 2^J$ in the Haar basis, such that $J = 10$ and scales $j = 1, 3, 4$ have all non-null coefficients. More details on this function can be found in our code that is freely available ¹. In this setting, $\log(f_0)$ can be decomposed on the dictionary and we can assess the accuracy of selection of different methods.

We then explore the empirical performance of the Lasso and the group Lasso strategies using simulations. By performance we mean the quality of reconstruction of simulated signals as our theoretical results concern reconstruction properties. We considered different forms for intensity functions by taking the standard functions of Donoho and Johnstone (1994): blocks, bumps, doppler, heavisine, to set g_0 . These functions do not have an exact decomposition on any dictionary considered below. The signal to noise ratio was increased by multiplying the intensity functions by a factor α taking values in $\{1, \dots, 7\}$, $\alpha = 7$ corresponding to the most favorable configuration. Observations Y_i were generated such that $Y_i|X_i \sim \text{Poisson}(f_0(X_i))$, with $f_0 = \alpha \exp(g_0)$, and (X_1, \dots, X_n) was set as the regular grid of length $n = 2^{10}$. Each configuration was repeated 20 times. Our method was implemented using the `grpLasso` R package of Meier et al. (2008) to which we provide our concentration-based weights (the code is fully available ¹).

The basis and the dictionary frameworks. The dictionary we consider is built on the Haar basis, on the Daubechies basis with 6 vanishing moments, and on the Fourier basis, in order to catch piece-wise constant trends, localized peaks and periodicities. Each orthonormal system has n elements, which makes $p = n$ when systems are considered separately, and $p = 2n$ or $3n$ depending on the considered dictionary. For wavelets, the dyadic structure of the decomposition allows us to group the coefficients scale-wise by forming groups of coefficients of size 2^q . As for the Fourier basis, groups (also of size 2^q) are formed by considering successive coefficients (while keeping their natural ordering). When grouping strategies are considered, we set all groups at the same size.

Weights calibration in practice. First for both the Lasso and the group Lasso introduced in Section 1, following the arguments at the end of Section 2.1 that provide some theoretical guarantees, we take γ larger than 1. More precisely, we take $\gamma = 1.01$. Our theoretical results do not provide any information about values smaller than 1, so we conduct an empirical study on a very simple case, namely we estimate f_0 such that $\log(f_0) = 1_{[0,1]}$ on the Haar basis, so that only one parameter should be selected. Fig. 1 shows that when γ is too small, the risk of the selected models explodes (left panel), because too many coefficients are selected (right panel). Taking $\gamma = 1.01$ actually corresponds to a conservative choice avoiding the explosion of the MSE and of the number of selected

1. <http://pbil.univ-lyon1.fr/members/fpicard/software.html>

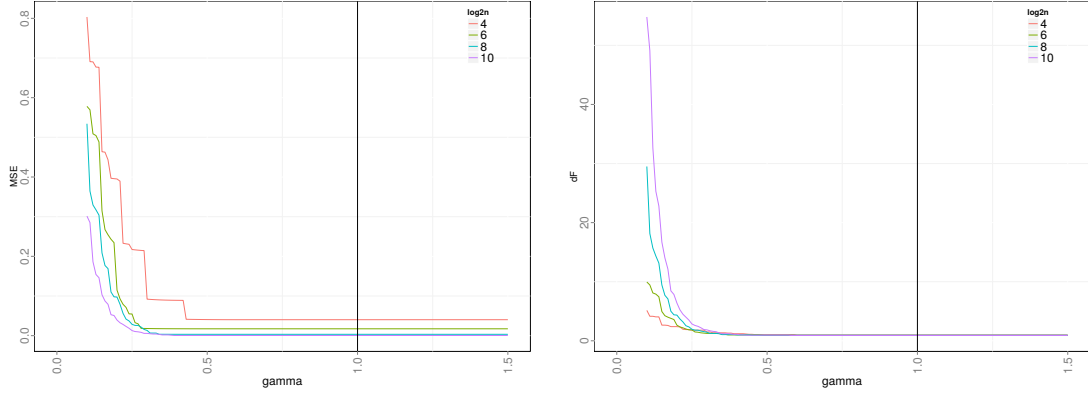


Figure 1: Tuning of the hyperparameter γ for the Lasso penalty. The simulated function is $\log(f_0) = 1_{[0,1]}$, so that only one parameter should be selected on the Haar basis. For different values of n , the left/right panels represent the mean square error / the number of estimated parameters with respect to γ (averaged over 20 simulations). The vertical line corresponds to $\gamma = 1$.

parameters. Furthermore this automatic choice provides good results in practice, as shown by our simulation study that considers many shapes of intensity functions (see below).

Then we estimate V_j (resp V_j^g) by \tilde{V}_j (resp \tilde{V}_j^g). Note that \hat{V}_j (resp \hat{V}_j^g) can also be used: these variances are easier to compute in practice, and this slight modification does not significantly change the performance of the procedures (not shown). The parameter γ being fixed, we use the expression (2.3) for Lasso weights. As for the group Lasso weights (Theorem 2), the first term is replaced by $\sqrt{\sum_{j \in G_k} \tilde{V}_j}$, as it is governed by a quantity that tends to one when p is large. We conducted a preliminary numerical study (not shown) to calibrate the second term, that depends on quantities M , c_k and b_k defined in Theorem 2. The best empirical performance were achieved so that the left- and right-hand terms of (2.9) were approximatively equal. This resumes to group-Lasso weights of the form $2\sqrt{\sum_{j \in G_k} \tilde{V}_j}$.

Competitors. We compare our Lasso procedure (**Lasso.exact** in the sequel), with the Haar-Fisz transform (using the **haarfisz** package), applied to the same data followed by soft-thresholding. Here we mention that we did not perform cycle-spinning (that is often included in denoising procedures) in order to focus on the effects of thresholding only. We also implemented the half-fold cross-validation proposed by Nason (1996) in the Poisson case to set the weights in the penalty, with the proper scaling ($2^{s/2}\lambda$, with s the scale of the wavelet coefficients) as proposed by Sardy et al. (2004). Briefly, this cross-validation procedure is used to calibrate λ by estimating f_0 on a training set made of even positions (for different values of λ), and by measuring the reconstruction error on the test set made

by odd positions (see Nason (1996)). This procedure is referred to cross-validation in the sequel. Then we compare the performance of the group-Lasso with varying group sizes (2,4,8) to the Lasso, to assess the benefits of grouping wavelet coefficients.

Performance measurement. For any estimate \hat{f} , reconstruction performance were measured using the (normalized) mean-squared error $MSE = \|\hat{f} - f_0\|_2^2 / \|f_0\|_2^2$. When there is a true sparse β_0 , model selection performance were measured by the standard indicators: accuracy on support recovery, sensitivity (proportion of true non-null coefficients among selected coefficients) and specificity of detection (proportion of true null coefficients among non-selected coefficients), based on the support of β_0 and on the support of its estimate.

Model selection performance with a true sparse β_0 . The simple toy example described above for model selection perfectly shows the performance of our Lasso procedure (Figure 2). Even if our theoretical results do not concern support recovery, our theoretically calibrated weights provide the best accuracy in support selection, along with the best sensitivity, specificity and reconstruction errors. This example also shows the interest of the theoretical weights calibration compared with the cross-validation procedure that shows bad performance, and also compared with the Haar-Fisz transform. Also, this toy example illustrates the need of scaling for the cross-validated vanilla Lasso: if the procedure proposed by Sardy et al. (2004) is not used, the cross-validated vanilla Lasso lacks of adaptation to heteroscedasticity, which leads to poor results on selection (lack of accuracy and specificity, or too many selected coefficients), and thus to overfitted reconstruction (Fig. 3). Thus, in further simulations we only considered the scaled version of the cross-validated Lasso, in order to compare methods that account for heteroscedasticity.

Performance in the basis setting. Here we focus on wavelet basis (Haar for blocks and Daubechies for bumps, doppler and heavisine) and not on a dictionary approach (considered in a second step) in order to compare our calibrated weights with other methods that rely on penalized strategy. It appears that, except for the **bumps** function, the Lasso with exact weights shows the lowest reconstruction error whatever the shape of the intensity function (Figure 4). Moreover, better performance of the Lasso with exact weights in cases of low intensity emphasize the interest of theoretically calibrated procedures rather than asymptotic approximations (like the Haar-Fisz transform). In the case of **bumps**, cross-validation seems to perform better than the Lasso, but when looking at reconstructed average function (Figure 5a) this lower reconstruction error of cross-validation is associated with higher local variations around the peaks. Compared with Haar-Fisz, the gain of using exact weights is substantial even when the signal to noise ratio is high, which indicates that even in the validity domain of the Haar-Fisz transform (large intensities), the Lasso combined with exact thresholds is more suitable (Figure 5a). As for the group Lasso, its performance highly depend on the group size: while groups of size 2 show similar performance as the Lasso, groups of size 4 and 8 increase the reconstruction error (Figure

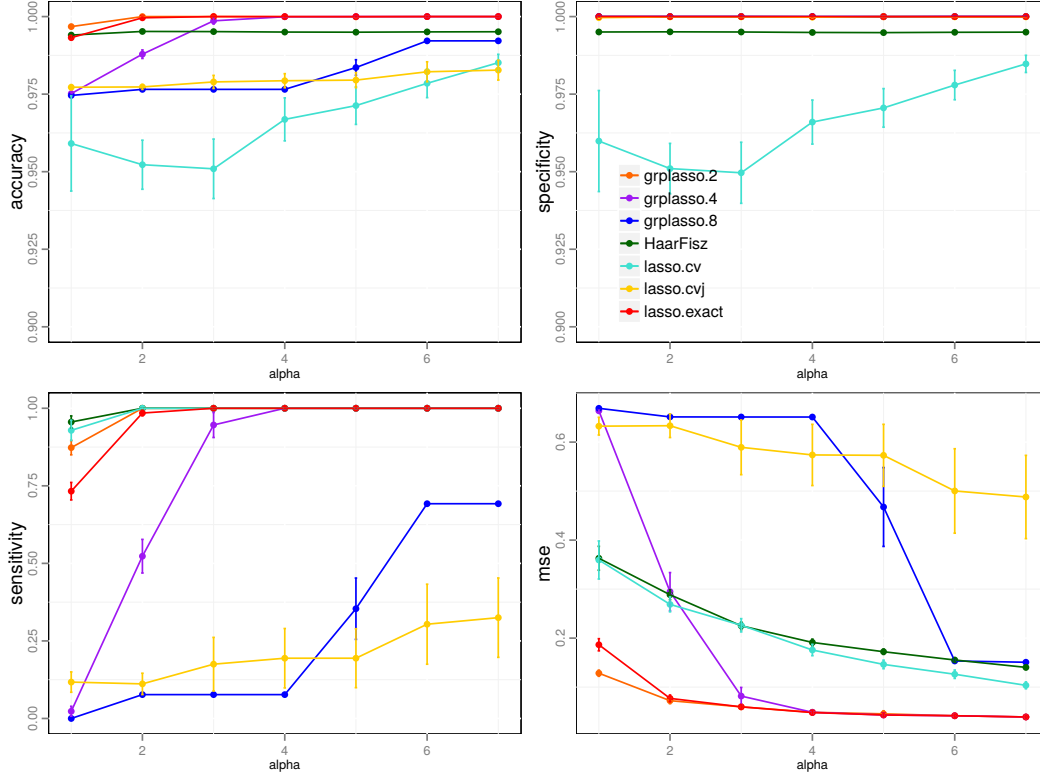


Figure 2: Average (over 20 repetitions) accuracy in support recovery for the toy function f_0 with a true sparse β_0 and the Haar basis, specificity and sensitivity of selection, and Mean Square Error of reconstruction. **Lasso.exact**: Lasso penalty with our data-driven theoretical weights, **Lasso.cv**: Lasso penalty with weights calibrated by cross validation without scaling, **Lasso.cvj**: Lasso penalty with weights calibrated by cross validation with scaling $2^{s/2}\lambda$, **group.Lasso.2/4/8**: group Lasso penalty with our data-driven theoretical weights with group sizes 2/4/8, **HaarFisz**: Haar-Fisz tranform followed by soft-thresholding. Alpha (α) stands for the signal strength.

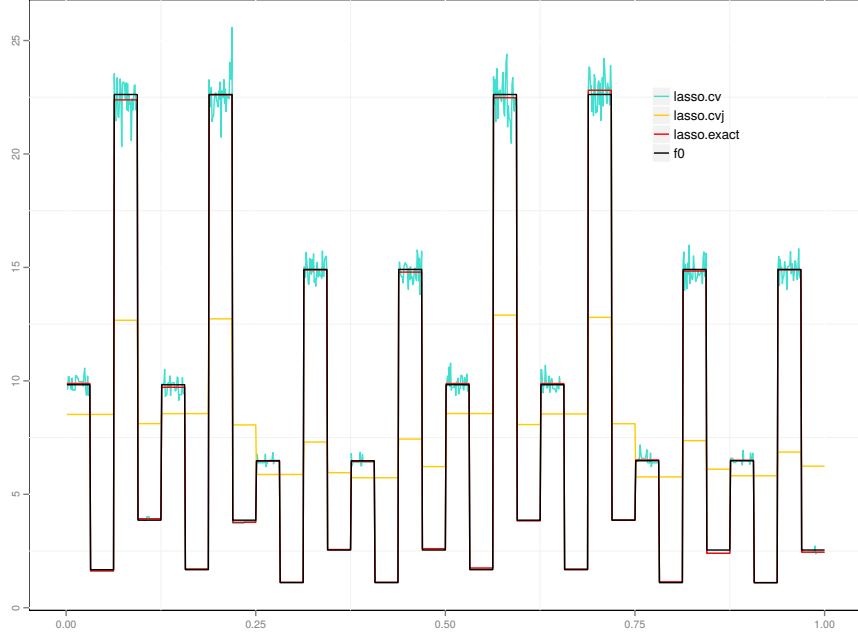


Figure 3: Example of reconstruction for the toy function f_0 with a true sparse β_0 and the Haar basis. The reconstruction with the vanilla-lasso is based on too many coefficients which leads to over-fitting. **Lasso.exact**: Lasso penalty with our data-driven theoretical weights, **Lasso.cv**: Lasso penalty with weights calibrated by cross validation without scaling, **Lasso.cvj**: Lasso penalty with weights calibrated by cross validation with scaling $2^{s/2}\lambda$

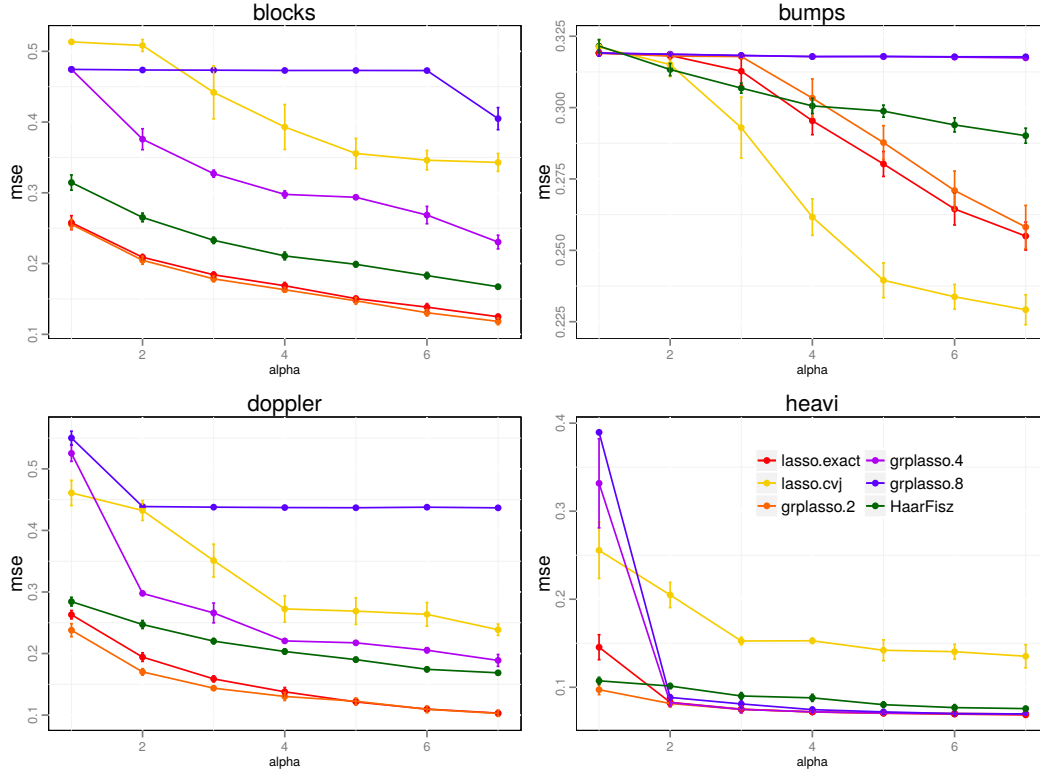


Figure 4: Average (over 20 repetitions) Mean Square Error of reconstruction of different methods for the estimation of simulated intensity functions according to function shapes (blocks, bumps, doppler, heavisine) and signal strength (α). **Lasso.exact**: Lasso penalty with our data-driven theoretical weights, **Lasso.cvj**: Lasso penalty with weights calibrated by cross validation with scaling $2^{s/2}\lambda$, **group.Lasso.2/4/8**: group Lasso penalty with our data-driven theoretical weights with group sizes 2/4/8, **HaarFisz**: Haar-Fisz tranform followed by soft-thresholding.

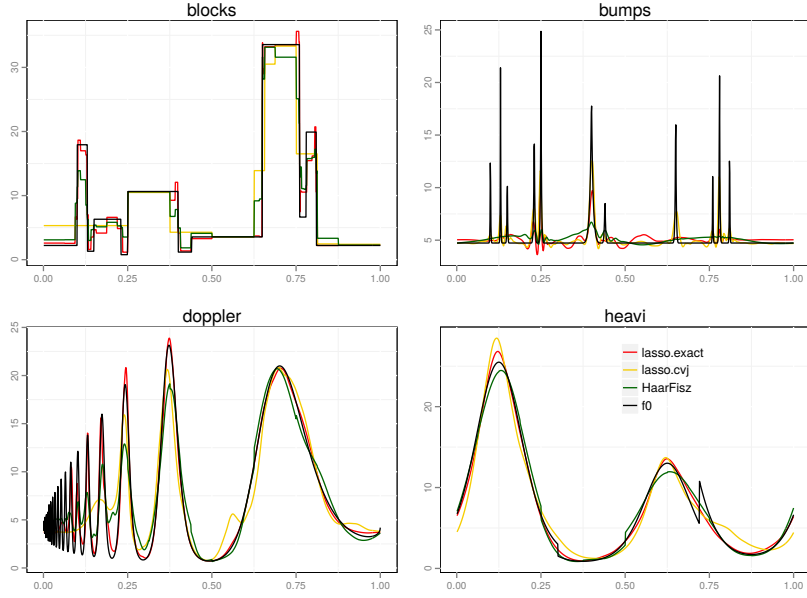
4 and 5b), since they are not scaled to the size of the irregularities in the signal. This trend is not systematic as the group Lasso appears to be adapted to functions that are more regular (Heavisine), and seems to avoid edge effects in some situations.

Performance in the dictionary framework. Lastly, we explored the performance of the dictionary approach, by considering different dictionaries to estimate each function: Daubechies (D), Fourier (F), Haar (H), or their combinations (Figure 6). Rich dictionaries can be very powerful to catch complex shapes in the true intensity function (like the notch in the heavisine case Figure 6b), and the richest dictionary (DFH) often leads to the lowest reconstruction error (MSE) on average. However the richest dictionary (DFH) is not always

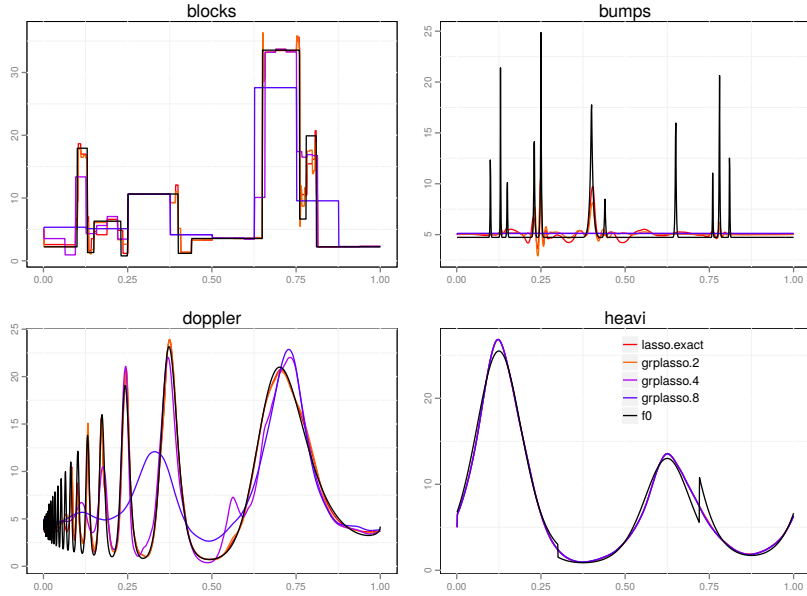
the best choice in terms of reconstruction error, which is striking in the case of the **blocks** function. In this case the Haar system only would be preferable for the Lasso (Figure 6a). For the group-Lasso and the **blocks** intensity function, the combination of the Daubechies and the Haar systems provides the best MSE, but when looking at the reconstructed intensity (Figure 6b-blocks), the Daubechies system introduces wiggles that are not relevant for **blocks**. Also, richer dictionaries do not necessarily lead to more selected parameters (Figure 6a), which illustrates that selection depends on the redundancies between the systems elements of the dictionary. In practice we often do not have any *prior* knowledge concerning the elements that shape the signal, and these simulations suggest that the blind use of the richest dictionary may not be the best strategy in terms of reconstructed functions. Consequently, in the following application, we propose to adapt the half-fold cross validation of Nason (1996) to choose the best combinations of systems.

5. Applications

The analysis of biological data has faced a new challenge with the extensive use of next generation sequencing (NGS) technologies. NGS experiments are based on the massive parallel sequencing of short sequences (reads). The mapping of these reads onto a reference genome (when available) generates counts data (Y_t) spatially organized (in 1D) along the genome (at position X_t). These technologies have revolutionized the perspectives of many fields in molecular biology, and among many applications, one is to get a local quantification of DNA or of a given DNA-related molecule (like transcription factors for instance with chIP-Seq experiments, Furey (2012)). This technology has recently been applied to the identification of replication origins along the human genome. Replication is the process by which a genome is duplicated into two copies. This process is tightly regulated in time and space so that the duplication process takes place in the highly regulated cell cycle. The human genome is replicated at many different starting points called origins of replication, that are loci along the genome at which the replication starts. Until very recently, the number of such origins remained controversial, and thanks to the application of NGS technologies, first estimates of this number could be obtained. The signal is made of counts along the human genome such that reads accumulations indicate an origin activity (see Picard et al. (2014)). Scan statistics were first applied to these data, to detect significant local enrichments reads accumulation, but there is currently no consensus on the best method to analyze such data. Here we propose to use the Poisson functional regression to estimate the intensity function of the data on a portion of the human chromosomes X and 20. Half-fold cross-validation was used to select the appropriate dictionary between Daubechies, Fourier, Haar (and their combinations), and our theoretical weights were used to calibrate the Lasso (Figure 7). Our results are very promising as the sparse dictionary approach is very efficient for denoising (Chromosome X, Figure 7b) and produces null intensities when the signal is low (higher specificity). Another aspect of our method is that it seems to be more powerful in the identification of peaks that are more precise (Chromosome 20,

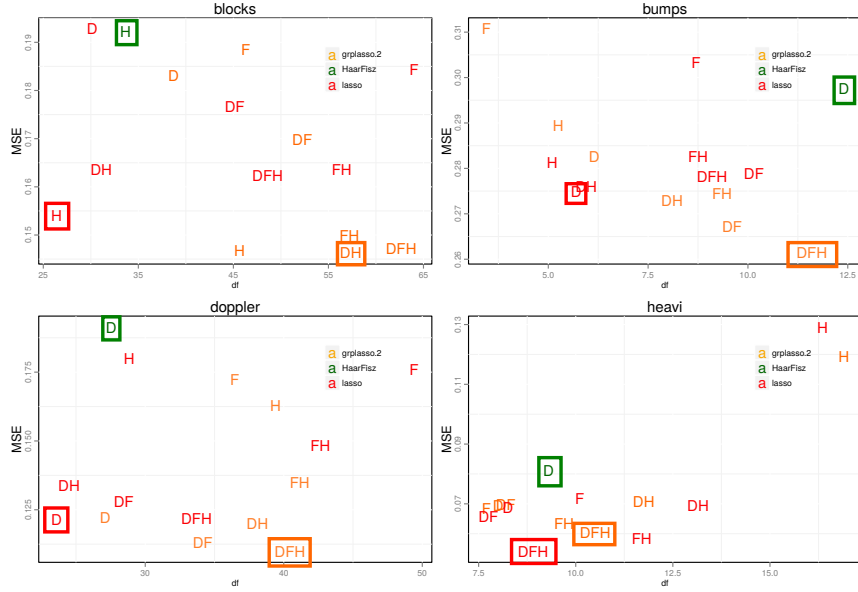


(a) Average reconstructed functions for the Lasso and competitors.

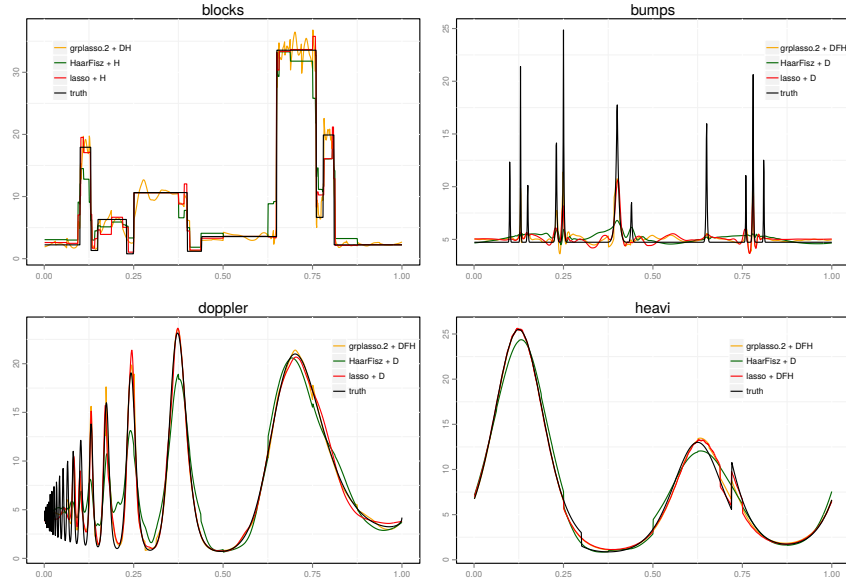


(b) Average reconstructed functions for the group strategies.

Figure 5: Average (over 20 repetitions) reconstructed functions by different methods of estimation according to function shapes (blocks, bumps, doppler, heavisine). Top panel corresponds to non-grouped strategies (5a) and bottom panel compares group-strategies to the Lasso (5b). **Lasso.exact**: Lasso penalty with our data-driven theoretical weights, **Lasso.cvj**: Lasso penalty with weights calibrated by cross validation with scaling $2^{s/2}\lambda$, **group.Lasso.2/4/8**: group Lasso penalty²⁷ with our data-driven theoretical weights with group sizes 2/4/8, **HaarFisz**: Haar-Fisz tranform followed by soft-thresholding, **f0**: simulated intensity function.



(a) Average Mean Square Error for different dictionaries with respect to the average number of selected coefficients (df).



(b) Reconstructed functions for the dictionaries with the smallest MSE.

Figure 6: Average (over 20 repetitions) Mean Square Errors and number of selected coefficients (df) (6a), and reconstructed functions (6b) for different dictionaries: Daubechies (D), Fourier (F), Haar (H) and their combinations. **Lasso.exact**: Lasso penalty with our data-driven theoretical weights, **group.Lasso.2**: group Lasso penalty with our data-driven theoretical weights with group sizes 2, **HaarFisz**: Haar-Fisz tranform followed by soft-thresholding.

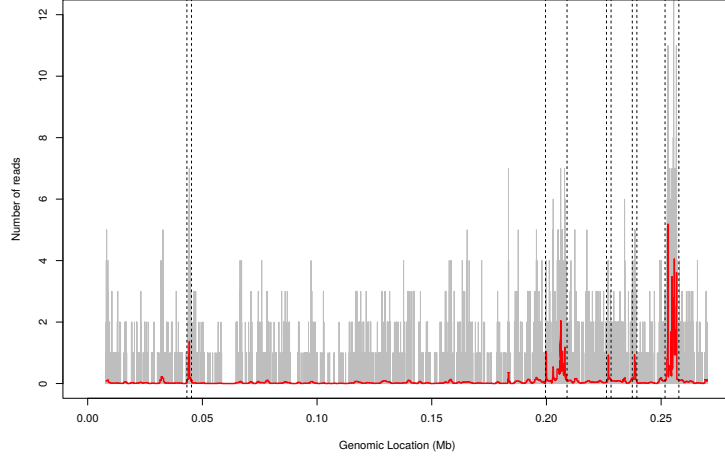
positions 0.20 and 0.25Mb, Figure 7a), which indicates that the dictionary approach may be more sensitive to detect peaks. Given the spread of NGS data and the importance of peak detection in the analysis process, for chIP-Seq Furey (2012), FAIRE-Seq Thurman et al. (2012), OriSeq Picard et al. (2014), our preliminary results suggest that the sparse dictionary approach will be a very promising framework for the analysis of such data.

6. Conclusion

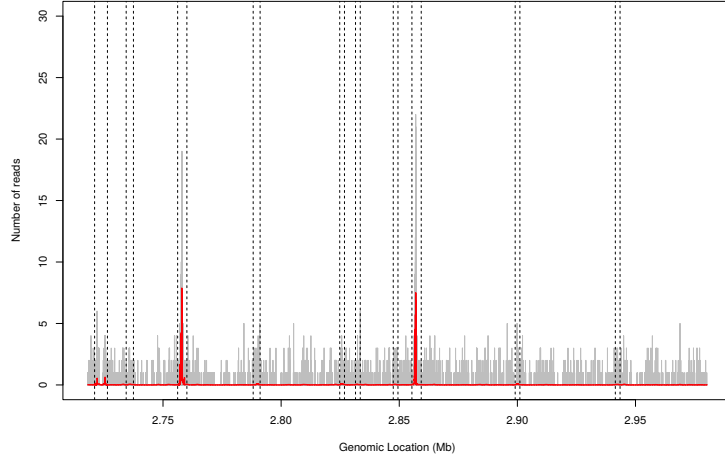
We proposed new adaptive Lasso and group-Lasso procedures to estimate the regression function in high dimensional Poisson regression, with a special focus on the calibration of weights involved in the penalties. Inspired from the adaptive Lasso procedure proposed by Zou (2006) in the context of model selection, we derived data-driven component-specific weights, whose shape appears to be relevant from both theoretical and practical points of view. In the dictionary approach, which extends the classical basis approach, we obtain slow and fast oracle inequalities under RE-type conditions. These theoretical results are enhanced by a numerical study that illustrates the good performance of our procedure on simulated and experimental data. Moreover, even if our main objective is functional reconstruction (that is different from model selection), a numerical toy example shows very promising results of our Lasso and group-Lasso procedures for model selection. The theoretical study of this problem is an exciting challenge we wish to investigate in further works.

The purpose of our data-driven weights is to control the random fluctuations of the normalized observations, namely $(\mathbf{A}_j^T \mathbf{Y})_j$ and $(\mathbf{A}_{G_k}^T \mathbf{Y})_k$, around their expectation. To obtain controls as sharp as possible, we use concentration inequalities for infinitely divisible vectors, which allows us to account for the heteroscedasticity that characterizes the Poisson setting. So, our approach is very different from Zou (2006) who considered weights proportional to the inverse of preliminary estimates (ordinary least squares estimates or maximum likelihood estimates). But both approaches confirm that random weighting schemes can provide suitable procedures for inference in regression models.

Finally, we mention that the constants of our procedure were tuned by using theoretical arguments, and our numerical studies show that they are suitable. Other values for [these constants](#) would probably be appropriate for some signals. Even if our empirical results show that the performance with $\gamma = 1.01$ are excellent, whatever the form of f_0 we considered, our procedure could be enriched by an additional calibration step for γ , with associated extra computational time.



(a) Chromosome 20



(b) Chromosome X

Figure 7: Estimation of the intensity function of Ori-Seq data (chromosomes 20 7a and X 7b). Grey bars indicate the number of reads that match genomic positions (x-axis, in MegaBases). The red line corresponds to the estimated intensity function, and vertical dotted lines stand for the detected origins by scanning statistics.

7. Proofs

7.1 Proof of Theorem 1

We denote by μ the Lebesgue measure on \mathbb{R}^d and we introduce a partition of the set $[0, 1]^d$ denoted $\cup_{i=1}^n S_i$ so that for any $i = 1, \dots, n$, $X_i \in S_i$ and $\mu(S_i) > 0$. Let h the function defined for any $t \in [0, 1]^d$ by

$$h(t) = \sum_{i=1}^n \frac{f_0(X_i)}{\mu(S_i)} 1_{S_i}(t).$$

Finally, we introduce N the Poisson process on $[0, 1]^d$ with intensity h (see Kingman (1993)). Therefore, for any $i = 1, \dots, n$, $N(S_i)$ is a Poisson variable with parameter $\int_{S_i} h(t)dt = f_0(X_i)$ and since $\cup_{i=1}^n S_i$ is a partition of $[0, 1]^d$, $(N(S_1), \dots, N(S_n))$ has the same distribution as (Y_1, \dots, Y_n) . We observe that if for any $j = 1, \dots, p$,

$$\tilde{\varphi}_j(t) = \sum_{i=1}^n \varphi_j(X_i) 1_{S_i}(t),$$

then

$$\int \tilde{\varphi}_j(t) dN(t) \sim \sum_{i=1}^n \varphi_j(X_i) Y_i = \mathbf{A}_j^T \mathbf{Y}.$$

We use the following exponential inequality (see Inequality (5.2) of Reynaud-Bouret (2003)). If g is bounded, for any $u > 0$,

$$\mathbb{P} \left(\int g(x) (dN(x) - h(x)dx) \geq \sqrt{2u \int g^2(x) h(x)dx} + \frac{u}{3} \|g\|_\infty \right) \leq \exp(-u). \quad (7.1)$$

By taking successively $g = \tilde{\varphi}_j$ and $g = -\tilde{\varphi}_j$, we obtain

$$\mathbb{P} \left(|\mathbf{A}_j^T (\mathbf{Y} - \mathbb{E}[\mathbf{Y}])| \geq \sqrt{2u \int \tilde{\varphi}_j^2(x) h(x)dx} + \frac{u}{3} \|\tilde{\varphi}_j\|_\infty \right) \leq 2e^{-u}.$$

Since

$$\int \tilde{\varphi}_j^2(x) h(x)dx = \sum_{i=1}^n \varphi_j^2(X_i) f_0(X_i) = V_j,$$

we obtain

$$\mathbb{P} \left(|\mathbf{A}_j^T (\mathbf{Y} - \mathbb{E}[\mathbf{Y}])| \geq \sqrt{2u V_j} + \frac{u}{3} \|\tilde{\varphi}_j\|_\infty \right) \leq 2e^{-u}. \quad (7.2)$$

To control V_j , we use (7.1) with $g = -\tilde{\varphi}_j^2$ and we have:

$$\mathbb{P} \left(V_j - \hat{V}_j \geq \sqrt{2u \int \tilde{\varphi}_j^4(t) h(t)dt} + \frac{u}{3} \|\tilde{\varphi}_j\|_\infty^2 \right) \leq e^{-u}.$$

We observe that

$$\int \tilde{\varphi}_j^4(t) h(t) dt \leq \|\tilde{\varphi}_j\|_\infty^2 \int \tilde{\varphi}_j^2(t) h(t) dt = \|\tilde{\varphi}_j\|_\infty^2 V_j.$$

Setting $v_j = u\|\tilde{\varphi}_j\|_\infty^2$, we have:

$$\mathbb{P}\left(V_j - \sqrt{2v_j V_j} - \frac{v_j}{3} - \hat{V}_j \geq 0\right) \leq e^{-u}.$$

Let $\alpha_j = \sqrt{\hat{V}_j + \frac{5}{6}v_j} + \sqrt{\frac{v_j}{2}}$, such that α_j is the positive solution to $\alpha_j^2 - \sqrt{2v_j}\alpha_j - (\hat{V}_j + \frac{v_j}{3}) = 0$. Then

$$\mathbb{P}(V_j \geq \alpha_j^2) = \mathbb{P}(\sqrt{V_j} \geq \alpha_j) \leq e^{-u}. \quad (7.3)$$

We choose $u = \gamma \log p$ and observe that $\alpha_j^2 \leq \tilde{V}_j$. Then, by combining (7.2) and (7.3), we have

$$\mathbb{P}\left(|\mathbf{A}_j^T(\mathbf{Y} - \mathbb{E}[\mathbf{Y}])| \geq \sqrt{2\gamma \log p \tilde{V}_j} + \frac{\gamma \log p}{3} \|\tilde{\varphi}_j\|_\infty\right) \leq \frac{3}{p^\gamma}.$$

As $\|\tilde{\varphi}_j\|_\infty = \max_i |\varphi_j(X_i)|$, the theorem follows. \square

Remark 2 *By slightly extending previous computations, we easily show that for $u > 0$,*

$$\mathbb{P}\left(|V_j - \hat{V}_j| \geq \sqrt{2uV_j\|\tilde{\varphi}_j\|_\infty^2} + \frac{u}{3}\|\tilde{\varphi}_j\|_\infty^2\right) \leq 2e^{-u},$$

which leads to

$$\mathbb{P}\left(|V_j - \hat{V}_j| \geq \frac{V_j}{2} + \frac{4\gamma \log p}{3} \|\tilde{\varphi}_j\|_\infty^2\right) \leq \frac{2}{p^\gamma}.$$

7.2 Proof of Theorem 2

For each $k \in \{1, \dots, K\}$, we recall that $b_k^i = \sqrt{\sum_{j \in G_k} \varphi_j^2(X_i)}$, so $b_k^i = \|\mathbf{A}_{G_k}^T \mathbf{e}_i\|_2$, where \mathbf{e}_i is the vector whose i -th coordinate is equal to 1 and all others to 0. We first state the following lemma:

Lemma 1 *Let k be fixed. Assume that there exists some $M > 0$ such that $\forall x, |f_0(x)| \leq M$. Assume further that there exists some $c_k \geq 0$ such that $\forall \mathbf{y} \in \mathbb{R}^n, \|\mathbf{A}_{G_k} \mathbf{A}_{G_k}^T \mathbf{y}\|_2 \leq c_k \|\mathbf{A}_{G_k}^T \mathbf{y}\|_2$. Then, $\forall x > 0, \forall \varepsilon > 0$,*

$$\mathbb{P}\left(\|\mathbf{A}_{G_k}^T(\mathbf{Y} - \mathbb{E}[\mathbf{Y}])\|_2 \geq (1 + \varepsilon) \sqrt{\sum_{j \in G_k} V_j} + x\right) \leq \exp\left(\frac{x}{b_k} - \left(\frac{x}{b_k} + \frac{D_k^\varepsilon}{b_k^2}\right) \log\left(1 + \frac{b_k x}{D_k^\varepsilon}\right)\right),$$

where $D_k^\varepsilon = 8Mc_k^2 + \frac{2}{\varepsilon^2} b_k^2$.

Proof With $k \in \{1, \dots, K\}$ being fixed, we define $f : \mathbb{R}^n \rightarrow \mathbb{R}$ by $f(\mathbf{y}) = \left(\|\mathbf{A}_{G_k}^T \mathbf{y}\|_2 - E \right)_+$, where $E > 0$ is a constant chosen later. We use Corollary 1 from Houdré et al. (2008), applied to the infinitely divisible vector $\mathbf{Y} - \mathbb{E}[\mathbf{Y}] \in \mathbb{R}^n$, whose components are independent, and to f . First note that for any $t > 0$,

$$\begin{aligned} \mathbb{E} e^{tb_k^i |Y_i - \mathbb{E} Y_i|} &\leq \mathbb{E} e^{tb_k^i (Y_i + f_0(X_i))} \\ &= \exp \left(f_0(X_i) (e^{tb_k^i} + tb_k^i - 1) \right) < \infty. \end{aligned}$$

Furthermore, for any $i \in \{1, \dots, n\}$, any $\mathbf{y} \in \mathbb{R}^n$ and any $u \in \mathbb{R}$,

$$\begin{aligned} |f(\mathbf{y} + u\mathbf{e}_i) - f(\mathbf{y})| &\leq \left| \|\mathbf{A}_{G_k}^T(\mathbf{y} + u\mathbf{e}_i)\|_2 - \|\mathbf{A}_{G_k}^T \mathbf{y}\|_2 \right| \\ &\leq \|\mathbf{A}_{G_k}^T(u\mathbf{e}_i)\|_2 \\ &= |u|b_k^i. \end{aligned}$$

Therefore, for all $x > 0$,

$$\mathbb{P} \left(f(\mathbf{Y} - \mathbb{E}[\mathbf{Y}]) - \mathbb{E}[f(\mathbf{Y} - \mathbb{E}[\mathbf{Y}])] \geq x \right) \leq \exp \left(- \int_0^x h_f^{-1}(s) ds \right),$$

where h_f is defined for all $t > 0$ by

$$h_f(t) = \sup_{\mathbf{y} \in \mathbb{R}^n} \sum_{i=1}^n \int_{\mathbb{R}} |f(\mathbf{y} + u\mathbf{e}_i) - f(\mathbf{y})|^2 \frac{e^{tb_k^i|u|} - 1}{b_k^i|u|} \tilde{\nu}_i(du)$$

and $\tilde{\nu}_i$ is the Lévy measure associated with $Y_i - \mathbb{E}[Y_i]$. It is easy to show that $\tilde{\nu}_i = f_0(X_i)\delta_1$, and so

$$h_f(t) = \sup_{\mathbf{y} \in \mathbb{R}^n} \sum_{i=1}^n f_0(X_i) \left(f(\mathbf{y} + \mathbf{e}_i) - f(\mathbf{y}) \right)^2 \frac{e^{tb_k^i} - 1}{b_k^i}.$$

Furthermore, writing $A_i = \left\{ \|\mathbf{A}_{G_k}^T(\mathbf{y} + \mathbf{e}_i)\|_2 \geq E \text{ or } \|\mathbf{A}_{G_k}^T \mathbf{y}\|_2 \geq E \right\}$, we have

$$\begin{aligned} |f(\mathbf{y} + \mathbf{e}_i) - f(\mathbf{y})| &\leq \left| \|\mathbf{A}_{G_k}^T(\mathbf{y} + \mathbf{e}_i)\|_2 - \|\mathbf{A}_{G_k}^T \mathbf{y}\|_2 \right| 1_{A_i} \\ &= \frac{1_{A_i} \left| \|\mathbf{A}_{G_k}^T(\mathbf{y} + \mathbf{e}_i)\|_2^2 - \|\mathbf{A}_{G_k}^T \mathbf{y}\|_2^2 \right|}{\|\mathbf{A}_{G_k}^T(\mathbf{y} + \mathbf{e}_i)\|_2 + \|\mathbf{A}_{G_k}^T \mathbf{y}\|_2} \\ &= \frac{1_{A_i} \left| 2 \langle \mathbf{A}_{G_k}^T \mathbf{e}_i, \mathbf{A}_{G_k}^T \mathbf{y} \rangle + \|\mathbf{A}_{G_k}^T \mathbf{e}_i\|_2^2 \right|}{\|\mathbf{A}_{G_k}^T(\mathbf{y} + \mathbf{e}_i)\|_2 + \|\mathbf{A}_{G_k}^T \mathbf{y}\|_2} \\ &\leq 2 \frac{\left| \langle \mathbf{A}_{G_k}^T \mathbf{e}_i, \mathbf{A}_{G_k}^T \mathbf{y} \rangle \right|}{\|\mathbf{A}_{G_k}^T \mathbf{y}\|_2} + \frac{\|\mathbf{A}_{G_k}^T \mathbf{e}_i\|_2^2}{E}, \end{aligned}$$

with $\langle \cdot, \cdot \rangle$ the usual scalar product. We now have

$$\left(f(\mathbf{y} + \mathbf{e}_i) - f(\mathbf{y})\right)^2 \leq 8 \frac{\langle \mathbf{A}_{G_k}^T \mathbf{e}_i, \mathbf{A}_{G_k}^T \mathbf{y} \rangle^2}{\|\mathbf{A}_{G_k}^T \mathbf{y}\|_2^2} + 2 \frac{\|\mathbf{A}_{G_k}^T \mathbf{e}_i\|_2^4}{E^2}.$$

The first term can be rewritten as $8 \frac{\langle \mathbf{e}_i, \mathbf{A}_{G_k} \mathbf{A}_{G_k}^T \mathbf{y} \rangle^2}{\|\mathbf{A}_{G_k}^T \mathbf{y}\|_2^2}$ and the second one is equal to $2 \frac{b_k^{i4}}{E^2}$, so we can now bound $h_f(t)$ as follows.

$$\begin{aligned} h_f(t) &\leq \sup_{\mathbf{y}} \sum_i f_0(X_i) \frac{e^{tb_k} - 1}{b_k^i} \left(8 \frac{\langle \mathbf{e}_i, \mathbf{A}_{G_k} \mathbf{A}_{G_k}^T \mathbf{y} \rangle^2}{\|\mathbf{A}_{G_k}^T \mathbf{y}\|_2^2} + 2 \frac{b_k^{i4}}{E^2} \right) \\ &\leq \frac{e^{tb_k} - 1}{b_k} \sup_{\mathbf{y}} \left(8M \frac{\|\mathbf{A}_{G_k} \mathbf{A}_{G_k}^T \mathbf{y}\|_2^2}{\|\mathbf{A}_{G_k}^T \mathbf{y}\|_2^2} + \frac{2}{E^2} \sum_i f_0(X_i) b_k^{i4} \right) \\ &\leq \frac{e^{tb_k} - 1}{b_k} \left(8Mc_k^2 + \frac{2}{E^2} \sum_i f_0(X_i) b_k^{i4} \right). \end{aligned}$$

Now, we set

$$E = \varepsilon \sqrt{\sum_{j \in G_k} V_j}.$$

So we have:

$$\begin{aligned} E^2 &= \varepsilon^2 \sum_{j \in G_k} \sum_{i=1}^n f_0(X_i) \varphi_j^2(X_i) \\ &= \varepsilon^2 \sum_{i=1}^n f_0(X_i) \sum_{j \in G_k} \varphi_j^2(X_i) \\ &= \varepsilon^2 \sum_{i=1}^n f_0(X_i) b_k^{i2}. \end{aligned}$$

Thus, we can finally bound the function h_f by the increasing function h defined by

$$h(t) = D_k^\varepsilon \frac{e^{tb_k} - 1}{b_k},$$

with $D_k^\varepsilon = 8Mc_k^2 + \frac{2b_k^2}{\varepsilon^2}$. Therefore,

$$\begin{aligned} \exp \left(- \int_0^x h_f^{-1}(s) ds \right) &\leq \exp \left(- \int_0^x h^{-1}(s) ds \right) \\ &= \exp \left(\frac{x}{b_k} - \left(\frac{x}{b_k} + \frac{D_k^\varepsilon}{b_k^2} \right) \log \left(1 + \frac{b_k x}{D_k^\varepsilon} \right) \right). \end{aligned}$$

Now,

$$\begin{aligned} f(\mathbf{Y} - \mathbb{E}[\mathbf{Y}]) - \mathbb{E}[f(\mathbf{Y} - \mathbb{E}[\mathbf{Y}])] &= \left(\|\mathbf{A}_{G_k}^T(\mathbf{Y} - \mathbb{E}[\mathbf{Y}])\|_2 - E \right)_+ - \mathbb{E} \left(\|\mathbf{A}_{G_k}^T(\mathbf{Y} - \mathbb{E}[\mathbf{Y}])\|_2 - E \right)_+ \\ &\geq \|\mathbf{A}_{G_k}^T(\mathbf{Y} - \mathbb{E}[\mathbf{Y}])\|_2 - E - \mathbb{E} \|\mathbf{A}_{G_k}^T(\mathbf{Y} - \mathbb{E}[\mathbf{Y}])\|_2. \end{aligned}$$

Furthermore, by Jensen's inequality, we have

$$\begin{aligned} \mathbb{E} \|\mathbf{A}_{G_k}^T(\mathbf{Y} - \mathbb{E}[\mathbf{Y}])\|_2 &\leq \sqrt{\mathbb{E} \|\mathbf{A}_{G_k}^T(\mathbf{Y} - \mathbb{E}[\mathbf{Y}])\|_2^2} \\ &= \sqrt{\sum_{j \in G_k} \mathbb{E}[(\mathbf{A}_j^T(\mathbf{Y} - \mathbb{E}[\mathbf{Y}]))^2]} \\ &= \sqrt{\sum_{j \in G_k} \text{Var}(\mathbf{A}_j^T \mathbf{Y})} \\ &= \sqrt{\sum_{j \in G_k} V_j}. \end{aligned}$$

Recalling that $E = \varepsilon \sqrt{\sum_{j \in G_k} V_j}$, we thus have

$$\mathbb{P} \left(f(\mathbf{Y} - \mathbb{E}[\mathbf{Y}]) - \mathbb{E}f(\mathbf{Y} - \mathbb{E}[\mathbf{Y}]) \geq x \right) \geq \mathbb{P} \left(\|\mathbf{A}_{G_k}^T(\mathbf{Y} - \mathbb{E}[\mathbf{Y}])\|_2 - (1 + \varepsilon) \sqrt{\sum_{j \in G_k} V_j} \geq x \right),$$

which concludes the proof. ■

We apply Lemma 1 with

$$\varepsilon = \frac{1}{2\sqrt{2\gamma \log p}} \quad \text{and} \quad x = 2\sqrt{\gamma \log p D_k^\varepsilon}.$$

Then,

$$\begin{aligned} \frac{b_k x}{D_k^\varepsilon} &= \frac{2b_k \sqrt{\gamma \log p}}{\sqrt{D_k^\varepsilon}} \\ &= \frac{2b_k \sqrt{\gamma \log p}}{\sqrt{8M c_k^2 + \frac{2b_k^2}{\varepsilon^2}}} \\ &\leq \varepsilon \sqrt{2\gamma \log p} = \frac{1}{2}. \end{aligned}$$

Finally, using the fact that $\log(1+u) \geq u - \frac{u^2}{2}$, we have:

$$\begin{aligned}
\exp\left(\frac{x}{b_k} - \left(\frac{x}{b_k} + \frac{D_k^\varepsilon}{b_k^2}\right) \log\left(1 + \frac{b_k x}{D_k^\varepsilon}\right)\right) &\leq \exp\left(\frac{x}{b_k} - \left(\frac{x}{b_k} + \frac{D_k^\varepsilon}{b_k^2}\right) \left(\frac{b_k x}{D_k^\varepsilon} - \frac{b_k^2 x^2}{2D_k^{\varepsilon^2}}\right)\right) \\
&= \exp\left(\frac{-x^2}{2D_k^\varepsilon} + \frac{b_k x^3}{2D_k^{\varepsilon^2}}\right) \\
&= \exp\left(\frac{-x^2}{2D_k^\varepsilon} \left(1 - \frac{b_k x}{D_k^\varepsilon}\right)\right) \\
&\leq \exp\left(\frac{-x^2}{4D_k^\varepsilon}\right) = \frac{1}{p^\gamma}.
\end{aligned}$$

We obtain

$$\mathbb{P}\left(\|\mathbf{A}_{G_k}^T(\mathbf{Y} - \mathbb{E}[\mathbf{Y}])\|_2 \geq (1+\varepsilon) \sqrt{\sum_{j \in G_k} V_j} + 2\sqrt{\gamma \log p D_k^\varepsilon}\right) \leq \frac{1}{p^\gamma}.$$

We control V_j as in the proof of Theorem 1, but we take $u = \gamma \log p + \log |G_k|$. The analog of (7.3) is

$$\mathbb{P}(V_j > \tilde{V}_j^g) \leq e^{-u} = \frac{1}{|G_k| p^\gamma}$$

and thus

$$\mathbb{P}(\exists j \in G_k, V_j > \tilde{V}_j^g) \leq \frac{1}{p^\gamma}.$$

This concludes the proof of Theorem 2. □

7.3 Proof of Proposition 1

For the first point, we write:

$$\|\mathbf{A}_{G_k} \mathbf{A}_{G_k}^T \mathbf{x}\|_2^2 = \sum_{l=1}^n \left(\sum_{j \in G_k} \varphi_j(X_l) \sum_{i=1}^n \varphi_j(X_i) x_i \right)^2.$$

Then, we apply the Cauchy-Schwarz inequality:

$$\begin{aligned}
 \|\mathbf{A}_{G_k} \mathbf{A}_{G_k}^T \mathbf{x}\|_2^2 &\leq \sum_{l=1}^n \left(\sum_{j \in G_k} \varphi_j^2(X_l) \right) \left(\sum_{j \in G_k} \left(\sum_{i=1}^n \varphi_j(X_i) x_i \right)^2 \right) \\
 &= \|\mathbf{A}_{G_k}^T \mathbf{x}\|_2^2 \sum_{l=1}^n \left(\sum_{j \in G_k} \varphi_j^2(X_l) \right) \\
 &= \|\mathbf{A}_{G_k}^T \mathbf{x}\|_2^2 \sum_{l=1}^n (b_k^l)^2 \\
 &\leq n b_k^2 \|\mathbf{A}_{G_k}^T \mathbf{x}\|_2^2,
 \end{aligned}$$

which proves the upper bound of (2.11). For the lower bound, we just observe that for any $i = 1, \dots, n$, with \mathbf{e}_i the vector whose i -th coordinate is equal to 1 and all others to 0,

$$\begin{aligned}
 b_k^i{}^2 &= \|\mathbf{A}_{G_k}^T \mathbf{e}_i\|_2^2 \\
 &= \langle \mathbf{A}_{G_k}^T \mathbf{e}_i, \mathbf{A}_{G_k}^T \mathbf{e}_i \rangle \\
 &= \langle \mathbf{e}_i, \mathbf{A}_{G_k} \mathbf{A}_{G_k}^T \mathbf{e}_i \rangle \\
 &\leq \|\mathbf{e}_i\|_2 \|\mathbf{A}_{G_k} \mathbf{A}_{G_k}^T \mathbf{e}_i\|_2 \\
 &\leq c_k \|\mathbf{A}_{G_k}^T \mathbf{e}_i\|_2 \\
 &= c_k b_k^i,
 \end{aligned}$$

which obviously entails $b_k \leq c_k$. For the last point, we observe that

$$\|\mathbf{A}_{G_k}^T \mathbf{x}\|_2^2 = \sum_{j \in G_k} K_j^2,$$

where $K_j = \sum_{i=1}^n \varphi_j(X_i)x_i$. By expressing $\|\mathbf{A}_{G_k}\mathbf{A}_{G_k}^T\mathbf{x}\|_2^2$ with respect to the K_j 's, we obtain:

$$\begin{aligned}
\|\mathbf{A}_{G_k}\mathbf{A}_{G_k}^T\mathbf{x}\|_2^2 &= \sum_{l=1}^n \left(\sum_{j \in G_k} \varphi_j(X_l) \sum_{i=1}^n \varphi_j(X_i)x_i \right)^2 \\
&= \sum_{l=1}^n \sum_{j \in G_k} \varphi_j(X_l) \sum_{i=1}^n \varphi_j(X_i)x_i \sum_{j' \in G_k} \varphi_{j'}(X_l) \sum_{i'=1}^n \varphi_{j'}(X_{i'})x_{i'} \\
&= \sum_{j \in G_k} \sum_{j' \in G_k} \sum_{l=1}^n \varphi_j(X_l)\varphi_{j'}(X_l) \sum_{i=1}^n \varphi_j(X_i)x_i \sum_{i'=1}^n \varphi_{j'}(X_{i'})x_{i'} \\
&= \sum_{j \in G_k} \sum_{j' \in G_k} \sum_{l=1}^n \varphi_j(X_l)\varphi_{j'}(X_l)K_jK_{j'} \\
&\leq \frac{1}{2} \sum_{j \in G_k} \sum_{j' \in G_k} \left| \sum_{l=1}^n \varphi_j(X_l)\varphi_{j'}(X_l) \right| (K_j^2 + K_{j'}^2) \\
&= \sum_{j \in G_k} \sum_{j' \in G_k} \left| \sum_{l=1}^n \varphi_j(X_l)\varphi_{j'}(X_l) \right| K_j^2,
\end{aligned}$$

from which we deduce (2.12). □

7.4 Proof of Theorem 3

For any $\beta \in \mathbb{R}^p$, we have

$$\begin{aligned}
K(f_0, f_\beta) &= \sum_{i=1}^n f_0(X_i) (\log f_0(X_i) - \log f_\beta(X_i)) + f_\beta(X_i) - f_0(X_i) \\
&= \sum_{i=1}^n Y_i (\log f_0(X_i) - \log f_\beta(X_i)) + f_\beta(X_i) - f_0(X_i) \\
&\quad + \sum_{i=1}^n (f_0(X_i) - Y_i) (\log f_0(X_i) - \log f_\beta(X_i)) \\
&= \log \mathcal{L}(f_0) - \log \mathcal{L}(f_\beta) + \sum_{i=1}^n (f_0(X_i) - Y_i) (\log f_0(X_i) - \log f_\beta(X_i)).
\end{aligned}$$

Therefore, for all $\beta \in \mathbb{R}^p$,

$$\begin{aligned}
 K(f_0, \hat{f}^{gL}) - K(f_0, f_\beta) &= l(\beta) - l(\hat{\beta}^{gL}) + \sum_{i=1}^n (f_0(X_i) - Y_i) (\log f_\beta(X_i) - \log \hat{f}^{gL}(X_i)) \\
 &= l(\beta) - l(\hat{\beta}^{gL}) + \sum_{i=1}^n (f_0(X_i) - Y_i) \sum_{j=1}^p (\beta_j - \hat{\beta}_j^{gL}) \varphi_j(X_i) \\
 &= l(\beta) - l(\hat{\beta}^{gL}) + \sum_{j=1}^p (\hat{\beta}_j^{gL} - \beta_j) \sum_{i=1}^n \varphi_j(X_i) (Y_i - f_0(X_i)).
 \end{aligned}$$

Let us write $\eta_j = \sum_{i=1}^n \varphi_j(X_i) (Y_i - f_0(X_i)) = \mathbf{A}_j^T (\mathbf{Y} - \mathbb{E}[\mathbf{Y}])$. We have

$$K(f_0, \hat{f}^{gL}) = K(f_0, f_\beta) + l(\beta) - l(\hat{\beta}^{gL}) + (\hat{\beta}^{gL} - \beta)^T \boldsymbol{\eta}. \quad (7.4)$$

By definition of $\hat{\beta}^{gL}$,

$$-l(\hat{\beta}^{gL}) + \sum_{k=1}^K \lambda_k^g \|\hat{\beta}_{G_k}^{gL}\|_2 \leq -l(\beta) + \sum_{k=1}^K \lambda_k^g \|\beta_{G_k}\|_2.$$

Furthermore, on Ω_g ,

$$\begin{aligned}
 |(\hat{\beta}^{gL} - \beta)^T \boldsymbol{\eta}| &= \left| \sum_{j=1}^p (\hat{\beta}_j^{gL} - \beta_j) (\mathbf{A}_j^T (\mathbf{Y} - \mathbb{E}[\mathbf{Y}])) \right| \\
 &\leq \sum_{k=1}^K \sum_{j \in G_k} |\hat{\beta}_j^{gL} - \beta_j| |\mathbf{A}_j^T (\mathbf{Y} - \mathbb{E}[\mathbf{Y}])| \\
 &\leq \sum_{k=1}^K \left(\sum_{j \in G_k} (\hat{\beta}_j^{gL} - \beta_j)^2 \right)^{1/2} \left(\sum_{j \in G_k} (\mathbf{A}_j^T (\mathbf{Y} - \mathbb{E}[\mathbf{Y}]))^2 \right)^{1/2} \\
 &= \sum_{k=1}^K \|\hat{\beta}_{G_k}^{gL} - \beta_{G_k}\|_2 \|\mathbf{A}_{G_k}^T (\mathbf{Y} - \mathbb{E}[\mathbf{Y}])\|_2 \\
 &\leq \sum_{k=1}^K \lambda_k^g \|\hat{\beta}_{G_k}^{gL} - \beta_{G_k}\|_2.
 \end{aligned} \quad (7.5)$$

Therefore, for all $\beta \in \mathbb{R}^p$,

$$K(f_0, \hat{f}^{gL}) \leq K(f_0, f_\beta) + \sum_{k=1}^K \lambda_k^g \left(\|\hat{\beta}_{G_k}^{gL} - \beta_{G_k}\|_2 - \|\hat{\beta}_{G_k}^{gL}\|_2 + \|\beta_{G_k}\|_2 \right),$$

from which we deduce (3.3). \square

7.5 Proof of Theorem 4

To avoid too tedious notations, we denote \widehat{f}^{gL} instead of $\widehat{f}_{\mu,\alpha}^{gL}$ and $\widehat{\beta}^{gL}$ instead of $\widehat{\beta}_{\mu,\alpha}^{gL}$. We start from Equality (7.4) combined with Inequality (7.5). Then, we have that on Ω_g , for any $\beta \in \Gamma(\mu)$,

$$K(f_0, \widehat{f}^{gL}) + (\alpha - 1) \sum_{k=1}^K \lambda_k^g \|\widehat{\beta}_{G_k}^{gL} - \beta_{G_k}\|_2 \leq K(f_0, f_\beta) + \sum_{k=1}^K \alpha \lambda_k^g \left(\|\widehat{\beta}_{G_k}^{gL} - \beta_{G_k}\|_2 - \|\widehat{\beta}_{G_k}^{gL}\|_2 + \|\beta_{G_k}\|_2 \right).$$

On $J(\beta)^c$, $\|\widehat{\beta}_{G_k}^{gL} - \beta_{G_k}\|_2 - \|\widehat{\beta}_{G_k}^{gL}\|_2 + \|\beta_{G_k}\|_2 = 0$ and

$$K(f_0, \widehat{f}^{gL}) + (\alpha - 1) \sum_{k=1}^K \lambda_k^g \|\widehat{\beta}_{G_k}^{gL} - \beta_{G_k}\|_2 \leq K(f_0, f_\beta) + 2\alpha \sum_{k \in J(\beta)} \lambda_k^g \|\widehat{\beta}_{G_k}^{gL} - \beta_{G_k}\|_2. \quad (7.6)$$

By applying the Cauchy-Schwarz inequality we also have

$$K(f_0, \widehat{f}^{gL}) + (\alpha - 1) \sum_{k=1}^K \lambda_k^g \|\widehat{\beta}_{G_k}^{gL} - \beta_{G_k}\|_2 \leq K(f_0, f_\beta) + 2\alpha |J(\beta)|^{1/2} \left(\sum_{k \in J(\beta)} (\lambda_k^g)^2 \|\widehat{\beta}_{G_k}^{gL} - \beta_{G_k}\|_2^2 \right)^{1/2}. \quad (7.7)$$

If we write $\Delta = \mathbf{D}(\widehat{\beta}^{gL} - \beta)$, where \mathbf{D} is a diagonal matrix with $D_{j,j} = \lambda_k^g$ if $j \in G_k$, then we can rewrite (7.6) as

$$K(f_0, \widehat{f}^{gL}) + (\alpha - 1) \|\Delta\|_{1,2} \leq K(f_0, f_\beta) + 2\alpha \|\Delta_{J(\beta)}\|_{1,2} \quad (7.8)$$

and we deduce from (7.7)

$$K(f_0, \widehat{f}^{gL}) \leq K(f_0, f_\beta) + 2\alpha (|J(\beta)|)^{1/2} \|\Delta_{J(\beta)}\|_{1,2}. \quad (7.9)$$

Now, on the event $\{2\alpha \|\Delta_{J(\beta)}\|_{1,2} \leq \varepsilon K(f_0, f_\beta)\}$, the theorem follows immediately from (7.8). We now assume that $\varepsilon K(f_0, f_\beta) \leq 2\alpha \|\Delta_{J(\beta)}\|_{1,2}$. Since K is non-negative, we deduce from (7.8) that

$$(\alpha - 1) \|\Delta\|_{1,2} \leq 2\alpha \left(1 + \frac{1}{\varepsilon}\right) \|\Delta_{J(\beta)}\|_{1,2},$$

$$(\alpha - 1) \|\Delta_{J(\beta)^c}\|_{1,2} \leq \left(2\alpha \left(1 + \frac{1}{\varepsilon}\right) - (\alpha - 1)\right) \|\Delta_{J(\beta)}\|_{1,2}$$

and

$$\|\Delta_{J(\beta)^c}\|_{1,2} \leq \left(\frac{\alpha + 1 + 2\alpha/\varepsilon}{\alpha - 1}\right) \|\Delta_{J(\beta)}\|_{1,2}.$$

This is equivalent to

$$\sum_{k \in J(\beta)^c} \lambda_k^g \|(\hat{\beta}^{gL} - \beta)_{G_k}\|_2 \leq \left(\frac{\alpha + 1 + 2\alpha/\varepsilon}{\alpha - 1} \right) \sum_{k \in J(\beta)} \lambda_k^g \|(\hat{\beta}^{gL} - \beta)_{G_k}\|_2.$$

From Assumption 2 we have that, if β is such that $|J(\beta)| \leq s$, then

$$\|(\hat{\beta}^{gL} - \beta)_{J(\beta)}\|_2 \leq \frac{1}{\kappa_n} ((\hat{\beta}^{gL} - \beta)^T \mathbf{G}(\hat{\beta}^{gL} - \beta))^{1/2}.$$

Since

$$\mathbf{G}_{j,j'} = \sum_{i=1}^n \varphi_j(X_i) \varphi_{j'}(X_i) f_0(X_i),$$

by setting

$$u_i = \log f_\beta(X_i) - \log f_0(X_i) \quad \text{and} \quad \hat{u}_i^{gL} = \log f_{\hat{\beta}^{gL}}(X_i) - \log f_0(X_i),$$

we have

$$\begin{aligned} (\hat{\beta}^{gL} - \beta)^T \mathbf{G}(\hat{\beta}^{gL} - \beta) &= \sum_{j=1}^p \sum_{j'=1}^p (\hat{\beta}_j^{gL} - \beta_j)(\hat{\beta}_{j'}^{gL} - \beta_{j'}) \mathbf{G}_{j,j'} \\ &= \sum_{i=1}^n f_0(X_i) \left(\sum_{j=1}^p (\hat{\beta}_j^{gL} - \beta_j) \varphi_j(X_i) \right)^2 \\ &= \sum_{i=1}^n f_0(X_i) (\hat{u}_i^{gL} - u_i)^2. \end{aligned}$$

We set $h(f_0, f_\beta) = \sum_{i=1}^n f_0(X_i) u_i^2$ and $h(f_0, \hat{f}^{gL}) = \sum_{i=1}^n f_0(X_i) (\hat{u}_i^{gL})^2$. From (7.9) and since

$$\begin{aligned} \|\Delta_{J(\beta)}\|_2 &\leq (\max_k \lambda_k^g) \|(\hat{\beta}^{gL} - \beta)_{J(\beta)}\|_2 \\ &\leq \frac{\max_k \lambda_k^g}{\kappa_n} ((\hat{\beta}^{gL} - \beta)^T \mathbf{G}(\hat{\beta}^{gL} - \beta))^{1/2}, \end{aligned}$$

we have

$$K(f_0, \hat{f}^{gL}) \leq K(f_0, f_\beta) + \frac{2\alpha}{\kappa_n} |J(\beta)|^{1/2} (\max_k \lambda_k^g) \left(\sqrt{h(f_0, \hat{f}^{gL})} + \sqrt{h(f_0, f_\beta)} \right).$$

To conclude, we use arguments similar to Lemler (2013). We recall them for the safe of completeness. To connect $h(f_0, f_\beta)$ to $K(f_0, f_\beta)$, we use Lemma 1 of Bach (2010) that is recalled now.

Lemma 2 *Let g be a convex three times differentiable function $g : \mathbb{R} \rightarrow \mathbb{R}$ such that for all $t \in \mathbb{R}$, $|g'''(t)| \leq Sg''(t)$ for some $S \geq 0$. Then, for all $t \geq 0$,*

$$\frac{g''(0)}{S^2} \phi(-St) \leq g(t) - g(0) - g'(0)t \leq \frac{g''(0)}{S^2} \phi(St),$$

where $\phi(x) = e^x - x - 1$.

Let h be a real function. We set

$$G(h) = \sum_{i=1}^n (e^{h(X_i)} - f_0(X_i)h(X_i))$$

and

$$g(t) = G(h + tk),$$

where h and k are functions and $t \in \mathbb{R}$. We have :

$$g'(t) = \sum_{i=1}^n (k(X_i)e^{h(X_i)+tk(X_i)} - f_0(X_i)k(X_i)),$$

$$g''(t) = \sum_{i=1}^n (k^2(X_i)e^{h(X_i)+tk(X_i)})$$

and

$$g'''(t) = \sum_{i=1}^n (k^3(X_i)e^{h(X_i)+tk(X_i)}).$$

Therefore $|g'''(t)| \leq Sg''(t)$ with $S = \max_i |k(X_i)|$. We choose $h(X_i) = \log f_0(X_i)$ and $k(X_i) = u_i = \log f_\beta(X_i) - \log f_0(X_i)$ and we apply Lemma 2 to g with $t = 1$. Computations yield that $g(1) - g(0) = K(f_0, f_\beta)$, $g'(0) = 0$ and $g''(0) = \sum_{i=1}^n f_0(X_i)u_i^2 = h(f_0, f_\beta)$. Therefore

$$\frac{\phi(-S)}{S^2} h(f_0, f_\beta) \leq K(f_0, f_\beta) \leq \frac{\phi(S)}{S^2} h(f_0, f_\beta).$$

Finally, for $\beta \in \Gamma(\mu)$, $S = \max_i |u_i| \leq \mu + m$. Furthermore, $x \rightarrow \frac{\phi(x)}{x^2}$ is a nonnegative increasing function and therefore we have

$$\mu' h(f_0, f_\beta) \leq K(f_0, f_\beta) \leq \mu'' h(f_0, f_\beta),$$

where $\mu' = \frac{\phi(-\mu-m)}{(\mu+m)^2}$ and $\mu'' = \frac{\phi(\mu+m)}{(\mu+m)^2}$. It follows that, for $\beta \in \Gamma(\mu)$ such that $|J(\beta)| \leq s$,

$$K(f_0, \widehat{f}^{gL}) \leq K(f_0, f_\beta) + \frac{2\alpha}{\kappa_n \sqrt{\mu'}} |J(\beta)|^{1/2} (\max_k \lambda_k^g) \left(\sqrt{K(f_0, \widehat{f}^{gL})} + \sqrt{K(f_0, f_\beta)} \right).$$

We use twice the inequality $2uv \leq bu^2 + \frac{v^2}{b}$ for any $b > 0$, applied to $u = \frac{\alpha}{\kappa_n} \sqrt{|J(\beta)|} (\max_k \lambda_k^g)$ and v being either $\sqrt{\frac{1}{\mu'} K(f_0, \hat{f}^{gL})}$ or $\sqrt{\frac{1}{\mu'} K(f_0, f_\beta)}$. We have

$$\left(1 - \frac{1}{\mu'b}\right) K(f_0, \hat{f}^{gL}) \leq \left(1 + \frac{1}{\mu'b}\right) K(f_0, f_\beta) + 2b \frac{\alpha^2 |J(\beta)|}{\kappa_n^2} (\max_k \lambda_k^g)^2.$$

Finally,

$$K(f_0, \hat{f}^{gL}) \leq \left(\frac{\mu'b + 1}{\mu'b - 1}\right) K(f_0, f_\beta) + 2 \frac{\mu'b^2}{\mu'b - 1} \frac{\alpha^2 |J(\beta)|}{\kappa_n^2} (\max_k \lambda_k^g)^2.$$

We choose $b > 1/\mu'$ such that $\frac{\mu'b + 1}{\mu'b - 1} = 1 + \varepsilon$ and we set $B(\varepsilon, m, \mu) = 2(1 + \varepsilon)^{-1} \frac{\mu'b^2}{\mu'b - 1}$. Finally, we have, for any $\beta \in \Gamma(\mu)$ such that $|J(\beta)| \leq s$,

$$K(f_0, \hat{f}^{gL}) \leq (1 + \varepsilon) \left(K(f_0, f_\beta) + B(\varepsilon, m, \mu) \frac{\alpha^2 |J(\beta)|}{\kappa_n^2} (\max_k \lambda_k^g)^2 \right).$$

This completes the proof of Theorem 4. □

Acknowledgments

The research of Stéphane Ivanoff and Vincent Rivoirard is partly supported by the french Agence Nationale de la Recherche (ANR 2011 BS01 010 01 projet Calibration) and by the Chaire Economie et Gestion des Nouvelles Données, under the auspices of Institut Louis Bachelier, Havas-Media and Paris-Dauphine. The research of Franck Picard is partly supported by the ABS4NGS ANR project ANR-11-BINF-0001-06. The authors wish to thank two anonymous referees who each made helpful suggestions that improved the presentation of the paper.

References

- Anscombe, F. J. (1948). The transformation of Poisson, binomial and negative-binomial data. *Biometrika*, 35:246–254.
- Bach, F. (2010). Self-concordant analysis for logistic regression. *Electron. J. Stat.*, 4:384–414.
- Bach, F. R. (2008). Consistency of the group lasso and multiple kernel learning. *J. Mach. Learn. Res.*, 9:1179–1225.
- Bertin, K., Le Pennec, E., and Rivoirard, V. (2011). Adaptive Dantzig density estimation. *Ann. Inst. Henri Poincaré Probab. Stat.*, 47(1):43–74.

- Besbeas, P., De Feis, I., and Sapatinas, T. (2004). A comparative simulation study of wavelet shrinkage estimators for Poisson counts. *Intern. Statist. Review*, 72(2):209–237.
- Bickel, P. J., Ritov, Y., and Tsybakov, A. B. (2009). Simultaneous analysis of lasso and Dantzig selector. *Ann. Statist.*, 37(4):1705–1732.
- Blazere, M., Loubes, J.-M., and Gamboa, F. (2014). Oracle inequalities for a group lasso procedure applied to generalized linear models in high dimension. *IEEE Transactions on Information Theory*, 4(12):2303–2318.
- Bradic, J., Fan, J., and Jiang, J. (2011). Regularization for Cox’s proportional hazards model with NP-dimensionality. *Ann. Statist.*, 39(6):3092–3120.
- Bühlmann, P. and van de Geer, S. (2011). *Statistics for high-dimensional data*. Springer Series in Statistics. Springer, Heidelberg. Methods, theory and applications.
- Bunea, F., Tsybakov, A., and Wegkamp, M. (2007a). Sparsity oracle inequalities for the Lasso. *Electron. J. Stat.*, 1:169–194.
- Bunea, F., Tsybakov, A. B., and Wegkamp, M. H. (2007b). Aggregation for Gaussian regression. *Ann. Statist.*, 35(4):1674–1697.
- Chen, S. S., Donoho, D. L., and Saunders, M. A. (2001). Atomic decomposition by basis pursuit. *SIAM Rev.*, 43(1):129–159. Reprinted from *SIAM J. Sci. Comput.* **20** (1998), no. 1, 33–61 (electronic) [MR1639094 (99h:94013)].
- Chesneau, C. and Hebiri, M. (2008). Some theoretical results on the grouped variables Lasso. *Math. Methods Statist.*, 17(4):317–326.
- Chicken, E. and Cai, T. (2005). Block thresholding for density estimation: Local and global adaptivity. *Journal of Multivariate Analysis*, 95:76–106.
- Dalalyan, A. S., Hebiri, M., Meziani, K., and Salmon, J. (2013). Learning heteroscedastic models by convex programming under group sparsity. In *ICML*.
- Donoho, D. and Johnstone, I. (1994). Ideal spatial adaptation by wavelet shrinkage. *Biometrika*, 81:425–455.
- Fryzlewicz, P. (2008). Data-driven wavelet-Fisz methodology for nonparametric function estimation. *Electron. J. Stat.*, 2:863–896.
- Fryzlewicz, P. and Nason, G. P. (2004). A Haar-Fisz algorithm for Poisson intensity estimation. *J. Comput. Graph. Statist.*, 13(3):621–638.
- Furey, T. S. (2012). ChIP-seq and beyond: new and improved methodologies to detect and characterize protein-DNA interactions. *Nat. Rev. Genet.*, 13(12):840–852.

- Gaïffas, S. and Guillaoux, A. (2012). High-dimensional additive hazards models and the Lasso. *Electron. J. Stat.*, 6:522–546.
- Hansen, N., Reynaud-Bouret, P., and Rivoirard, V. (2014). Lasso and probabilistic inequalities for multivariate point processes. *To appear in Bernoulli*.
- Houdré, C., Marchal, P., and Reynaud-Bouret, P. (2008). Concentration for norms of infinitely divisible vectors with independent components. *Bernoulli*, 14(4):926–948.
- Huang, J. and Zhang, T. (2010). The benefit of group sparsity. *Ann. Statist.*, 38(4):1978–2004.
- Jia, J., Rohe, K., and Yu, B. (2013). The lasso under Poisson-like heteroscedasticity. *Statist. Sinica*, 23(1):99–118.
- Kingman, J. F. C. (1993). *Poisson processes*, volume 3 of *Oxford Studies in Probability*. The Clarendon Press, Oxford University Press, New York. Oxford Science Publications.
- Kolaczyk, E. D. (1999). Wavelet shrinkage estimation of certain Poisson intensity signals using corrected thresholds. *Statist. Sinica*, 9(1):119–135.
- Kong, S. and Nan, B. (2014). Non-asymptotic oracle inequalities for the high-dimensional cox regression via lasso. *Statistica Sinica*, 24:2542.
- Leblanc, F. and Letué, F. (2006). Maximum likelihood estimation in poisson regression via wavelet model selection. Technical report, hal-00079298.
- Lemler, S. (2013). Oracle inequalities for the lasso in the high-dimensional multiplicative Aalen intensity model. *Submitted*.
- Lounici, K., Pontil, M., van de Geer, S., and Tsybakov, A. B. (2011). Oracle inequalities and optimal inference under group sparsity. *Ann. Statist.*, 39(4):2164–2204.
- Meier, L., van de Geer, S., and Bühlmann, P. (2008). The group Lasso for logistic regression. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 70(1):53–71.
- Nardi, Y. and Rinaldo, A. (2008). On the asymptotic properties of the group lasso estimator for linear models. *Electron. J. Stat.*, 2:605–633.
- Nason, G. P. (1996). Wavelet shrinkage using cross-validation. *J. Roy. Statist. Soc. Ser. B*, 58(2):463–479.
- Obozinski, G., Wainwright, M. J., and Jordan, M. I. (2011). Support union recovery in high-dimensional multivariate regression. *Ann. Statist.*, 39(1):1–47.
- Park, M. Y. and Hastie, T. (2007). L_1 -regularization path algorithm for generalized linear models. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 69(4):659–677.

- Picard, F., Cadoret, J. C., Audit, B., Arneodo, A., Alberti, A., Battail, C., Duret, L., and Prioleau, M. N. (2014). The spatiotemporal program of DNA replication is associated with specific combinations of chromatin marks in human cells. *PLoS Genet.*, 10(5):e1004282.
- Reynaud-Bouret, P. (2003). Adaptive estimation of the intensity of inhomogeneous Poisson processes via concentration inequalities. *Probab. Theory Related Fields*, 126(1):103–153.
- Sardy, S., Antoniadis, A., and Tseng, P. (2004). Automatic smoothing with wavelets for a wide class of distributions. *J. Comput. Graph. Statist.*, 13(2):399–421.
- Thurman, R. E., Rynes, E., Humbert, R., Vierstra, J., Maurano, M. T., Haugen, E., Sheffield, N. C., Stergachis, A. B., Wang, H., Vernot, B., Garg, K., John, S., Sandstrom, R., Bates, D., Boatman, L., Canfield, T. K., Diegel, M., Dunn, D., Ebersol, A. K., Frum, T., Giste, E., Johnson, A. K., Johnson, E. M., Kutayavin, T., Lajoie, B., Lee, B. K., Lee, K., London, D., Lotakis, D., Neph, S., Neri, F., Nguyen, E. D., Qu, H., Reynolds, A. P., Roach, V., Safi, A., Sanchez, M. E., Sanyal, A., Shafer, A., Simon, J. M., Song, L., Vong, S., Weaver, M., Yan, Y., Zhang, Z., Zhang, Z., Lenhard, B., Tewari, M., Dorschner, M. O., Hansen, R. S., Navas, P. A., Stamatoyannopoulos, G., Iyer, V. R., Lieb, J. D., Sunyaev, S. R., Akey, J. M., Sabo, P. J., Kaul, R., Furey, T. S., Dekker, J., Crawford, G. E., and Stamatoyannopoulos, J. A. (2012). The accessible chromatin landscape of the human genome. *Nature*, 489(7414):75–82.
- Tropp, J. A. (2004). Greed is good: algorithmic results for sparse approximation. *IEEE Trans. Inform. Theory*, 50(10):2231–2242.
- van de Geer, S. A. (2008). High-dimensional generalized linear models and the lasso. *Ann. Statist.*, 36(2):614–645.
- van de Geer, S. A. and Bühlmann, P. (2009). On the conditions used to prove oracle results for the Lasso. *Electron. J. Stat.*, 3:1360–1392.
- Wei, F. and Huang, J. (2010). Consistent group selection in high-dimensional linear regression. *Bernoulli*, 16(4):1369–1384.
- Yuan, M. and Lin, Y. (2006). Model selection and estimation in regression with grouped variables. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 68(1):49–67.
- Zou, H. (2006). The adaptive lasso and its oracle properties. *J. Amer. Statist. Assoc.*, 101(476):1418–1429.
- Zou, H. (2008). A note on path-based variable selection in the penalized proportional hazards model. *Biometrika*, 95(1):241–247.