

Structured regularization for conditional Gaussian graphical models

Julien Chiquet¹ · Tristan Mary-Huard² · Stéphane Robin³

Received: 27 August 2015 / Accepted: 13 April 2016 / Published online: 4 May 2016
© Springer Science+Business Media New York 2016

Abstract Conditional Gaussian graphical models are a reparametrization of the multivariate linear regression model which explicitly exhibits (i) the partial covariances between the predictors and the responses, and (ii) the partial covariances between the responses themselves. Such models are particularly suitable for interpretability since partial covariances describe direct relationships between variables. In this framework, we propose a regularization scheme to enhance the learning strategy of the model by driving the selection of the relevant input features by prior structural information. It comes with an efficient alternating optimization procedure which is guaranteed to converge to the global minimum. On top of showing competitive performance on artificial and real datasets, our method demonstrates capabilities for fine interpretation, as illustrated on three high-dimensional datasets from spectroscopy, genetics, and genomics.

Keywords Multivariate regression · Regularization · Sparsity · Conditional Gaussian graphical model · Structured elastic net · Regulatory motif · QTL study · Spectroscopy

1 Introduction

Multivariate regression, i.e. regression with multiple response variables, is increasingly used to model high dimen-

sional problems. Compared to its univariate counterpart, the general linear model aims to predict several—say q —responses from a set of p predictors, relying on a training data set $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1, \dots, n}$:

$$\mathbf{y}_i = \mathbf{B}^T \mathbf{x}_i + \boldsymbol{\varepsilon}_i, \quad \boldsymbol{\varepsilon}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{R}), \quad \forall i = 1, \dots, n. \quad (1)$$

The $p \times q$ matrix of regression coefficients \mathbf{B} and the $q \times q$ covariance matrix \mathbf{R} of the Gaussian noise $\boldsymbol{\varepsilon}_i$ are unknown. Note that we omit the intercept term and center the data for clarity. Model (1) has been studied by [Mardia et al. \(1979\)](#) in the low dimensional case where both ordinary and generalized least squares estimators of \mathbf{B} coincide and do not depend on \mathbf{R} . These approaches boil down to performing q independent regressions, each column \mathbf{B}_j describing the weights associating the p predictors to the j th response. In the $n < p$ setup however, these estimators are not defined.

Mimicking the univariate-output case, multivariate penalized methods aim to regularize the problem by biasing the regression coefficients toward a given feasible set. Sparsity within the set of predictors is usually the most wanted feature in the high-dimensional setting, which can be met in the multivariate framework by a straightforward application of the most popular penalty-based methods from the univariate world involving ℓ_1 -regularization. In the context of multitask learning, more sophisticated penalties encouraging similarities between parameters across outputs have been proposed. For example, many authors have suggested the use of group-norms penalties to set to zero full rows of \mathbf{B} (see, e.g., [Obozinski et al. 2011](#); [Chiquet et al. 2011](#)). Refinements exist to cope with complex dependency structures, using graph or tree structures ([Kim and Xing 2009, 2010](#)).

Importantly, all aforementioned methods are based on the regularization of parameter matrix \mathbf{B} , that accounts for both

✉ Julien Chiquet
julien.chiquet@genopole.cnrs.fr

¹ LaMME - UMR 8071 CNRS/Université d'Évry-Val-d'Essonne, Boulevard de France 91000, Évry, France

² UMR de Génétique Végétale du Moulon, INRA/Univ. Paris Sud/CNRS, Ferme du Moulon 91190, Gif-sur-Yvette, France

³ UMR 518 AgroParisTech/INRA, Rue Claude Bernard, 70005 Paris, France

direct and undirect links between the inputs and the outputs. Alternatively, one may ask for sparsity in the direct links only. The distinction between direct and undirect links can be formalized when $(\mathbf{x}_i, \mathbf{y}_i)$ are jointly Gaussian. In this context, by conditioning \mathbf{y}_i over \mathbf{x}_i , one obtains an expression of \mathbf{B} that depends on $\mathbf{\Omega}_{xy}$ and $\mathbf{\Omega}_{yy} = \mathbf{R}^{-1}$, where $\mathbf{\Omega}_{xy}$ is the partial covariance matrix between \mathbf{x}_i and \mathbf{y}_i (see, e.g. [Mardia et al. 1979](#)). Note that imposing sparsity on the regression coefficients \mathbf{B} or on $\mathbf{\Omega}_{xy}$ is not equivalent and may therefore lead to different results. Here we chose to impose sparsity on $\mathbf{\Omega}_{xy}$. Indeed, these partial covariances account for the *direct* relationships that exist between the predictors \mathbf{x}_i and the responses \mathbf{y}_i . Approaches encouraging sparsity on direct links $\mathbf{\Omega}_{xy}$ have been referred to as ‘*conditional Gaussian Graphical Model*’ (cGGM) or ‘*partial Gaussian Graphical Model*’ in the recent literature ([Sohn and Kim 2012](#); [Yuan and Zhang 2014](#)). In these papers, the authors propose to regularize the cGGM log-likelihood by two ℓ_1 -norms respectively acting on the partial covariances between the features and the responses $\mathbf{\Omega}_{xy}$, and on the partial covariances between the responses themselves via \mathbf{R}^{-1} .

In the present paper we also consider the direct link regularization approach for multivariate regression with the following specifications. First, we typically consider situations where the number of outputs is small compared to the number of predictors and sample size (while we insist on the fact that the number of predictors may still exceed the sample size). Consequently no sparsity assumption is required for the partial covariances between outputs. Second, we consider applications where structural information about the effect of the inputs is available. Here the structural information will be embedded in the regularization scheme via an additional regularization term using an application-specific metrics, in the same manner as in the ‘structured’ versions of the Elastic-net ([Slawski et al. 2010](#); [Hebiri and Geer 2011](#); [Lorbert et al. 2010](#)), or quadratic penalty function using the Laplacian graph ([Rapaport et al. 2007](#); [Li and Li 2010](#)) proposed in the univariate-output case. Adding a structured regularization term is of importance when interpretability of the estimated model is as much important as its predictive performance. These two specifications (small number of outputs, availability of prior structural information) arise in application fields such as biology, agronomy or health and 3 examples will be investigated hereafter.

We show that the resulting penalized likelihood optimization problem is jointly convex and can be solved efficiently using a two-step procedure for which algorithmic convergence guarantees are provided. Penalized criteria for the choice of the regularization parameters are also provided. The importance of embedding for structural prior information will be exemplified in the various contexts of spectroscopy, genomic selection and regulatory motif discovery, illustrating how accounting for application-specific improves both

performance and interpretability. The procedure is available as an R-package called **spring**, available on the R-`FORGE`.

The outline of the paper is as follows. In Sect. 2 we provide background on cGGM and present our regularization scheme. In Sect. 3, we develop an efficient optimization strategy in order to minimize the associated criterion. A paragraph also addresses the model selection issue. Section 4 is dedicated to illustrative simulation studies. In Sect. 5, we investigate three multivariate data sets: first, we consider an example in spectrometry with the analysis of cookie dough samples; second, the relationships between genetic markers and a series of phenotypes of the plant *Brassica napus* is addressed; and third, we investigate the discovery of regulatory motifs of yeast from time-course microarray experiments. In these applications, some specific underlying structuring priors arise, the integration of which within our model is detailed as it is one of the main contributions of our work.

2 Model setup

2.1 Background on cGGM

The statistical framework of cGGM arises from a different parametrization of (1) which fills the gap between multivariate regression and GGM. It extends to the multivariate case the links existing between the linear model, partial correlations and GGM, as depicted for instance in [Mardia et al. \(1979\)](#). It amounts to investigating the joint probability distribution of $(\mathbf{x}_i, \mathbf{y}_i)$ in the Gaussian case, with the following block-wise decomposition of the covariance matrix $\mathbf{\Sigma}$ and its inverse $\mathbf{\Omega} = \mathbf{\Sigma}^{-1}$:

$$\mathbf{\Sigma} = \begin{pmatrix} \mathbf{\Sigma}_{xx} & \mathbf{\Sigma}_{xy} \\ \mathbf{\Sigma}_{yx} & \mathbf{\Sigma}_{yy} \end{pmatrix}, \quad \mathbf{\Omega} = \begin{pmatrix} \mathbf{\Omega}_{xx} & \mathbf{\Omega}_{xy} \\ \mathbf{\Omega}_{yx} & \mathbf{\Omega}_{yy} \end{pmatrix}. \tag{2}$$

Back to the distribution of \mathbf{y}_i conditional on \mathbf{x}_i , multivariate Gaussian analysis shows that, for centered \mathbf{x}_i and \mathbf{y}_i ,

$$\mathbf{y}_i | \mathbf{x}_i \sim \mathcal{N} \left(-\mathbf{\Omega}_{yy}^{-1} \mathbf{\Omega}_{yx} \mathbf{x}_i, \mathbf{\Omega}_{yy}^{-1} \right). \tag{3}$$

This model, associated with the full sample $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^n$, can be written in a matrix form by stacking in rows first the observations of the responses, and then of the predictors, in two data matrices \mathbf{Y} and \mathbf{X} with respective sizes $n \times q$ and $n \times p$, such that

$$\mathbf{Y} = -\mathbf{X} \mathbf{\Omega}_{xy} \mathbf{\Omega}_{yy}^{-1} + \boldsymbol{\varepsilon}, \quad \text{vec}(\boldsymbol{\varepsilon}) \sim \mathcal{N} \left(\mathbf{0}_{nq}, \mathbf{I}_n \otimes \mathbf{\Omega}_{yy}^{-1} \right), \tag{4}$$

where, for a $n \times p$ matrix \mathbf{A} , denoting \mathbf{A}_j its j th column, the `vec` operator is defined as $\text{vec}(\mathbf{A}) = (\mathbf{A}_1^T \dots \mathbf{A}_p^T)^T$.

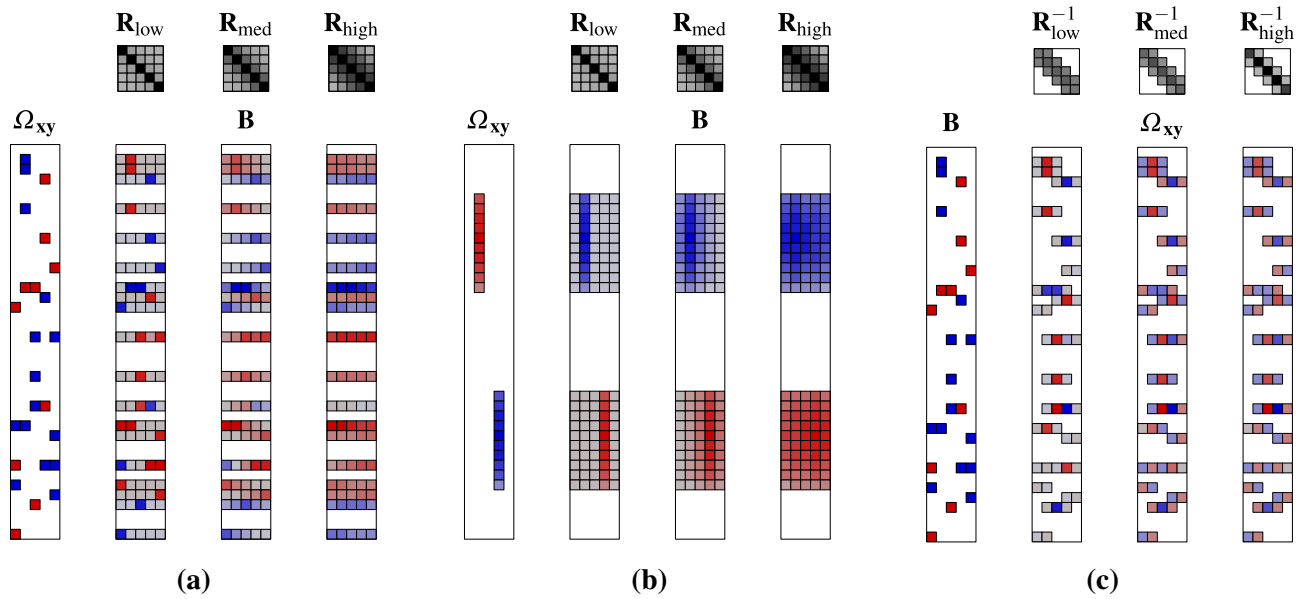


Fig. 1 Toy examples to illustrate the relationships between \mathbf{B} , Ω_{xy} and \mathbf{R} in the cGMM (better seen in *color*): on panel (a), a situation with no particular structure among the predictors; on panel (b), a strong neighborhood structure. For panels (a) and (b), we represent the effect of

stronger correlations in \mathbf{R} on masking the direct links in \mathbf{B} ; panel (c) illustrates the reverse situation of (a) where \mathbf{B} would be sparse and Ω_{xy} would not. (Color figure online)

The log-likelihood of (4)—which is a conditional likelihood regarding the joint model (2)—is written

$$\begin{aligned} \log L(\Omega_{xy}, \Omega_{yy}) &= \frac{n}{2} \log |\Omega_{yy}| \\ &+ \frac{n}{2} \text{tr} \left(\left(\mathbf{Y} + \mathbf{X} \Omega_{xy} \Omega_{yy}^{-1} \right) \Omega_{yy} \right. \\ &\quad \left. \times \left(\mathbf{Y} + \mathbf{X} \Omega_{xy} \Omega_{yy}^{-1} \right)^T \right) + \text{cst.} \end{aligned} \quad (5)$$

Introducing the empirical matrices of covariance $\mathbf{S}_{yy} = n^{-1} \sum_{i=1}^n y_i y_i^T$, $\mathbf{S}_{xx} = n^{-1} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i^T$, and $\mathbf{S}_{yx} = n^{-1} \sum_{i=1}^n y_i \mathbf{x}_i^T$, one has

$$\begin{aligned} -\frac{2}{n} \log L(\Omega_{xy}, \Omega_{yy}) &= -\log |\Omega_{yy}| + \text{tr}(\mathbf{S}_{yy} \Omega_{yy}) \\ &+ 2 \text{tr}(\mathbf{S}_{yx} \Omega_{yx}) \\ &+ \text{tr}(\Omega_{yx} \mathbf{S}_{xx} \Omega_{xy} \Omega_{yy}^{-1}) + \text{cst.} \end{aligned} \quad (6)$$

We notice by comparing the cGGM (3) to the multivariate regression model (1) that $\Omega_{yy}^{-1} = \mathbf{R}$ and $\mathbf{B} = -\Omega_{xy} \Omega_{yy}^{-1}$. Although equivalent to (1), the alternative parametrization (3) shows two important differences with several implications. First, the negative log-likelihood (5) can be shown to be jointly convex in $(\Omega_{xy}, \Omega_{yy})$ Yuan and Zhang (2014). Minimization problems involving (5) will thus be amenable to a global solution, which facilitates both optimization and theoretical analysis. As such, the conditional negative log-likelihood (5) will serve as a building block for our learning criterion. Second, it unveils an easier interpretation of the

relationships between input and output variables, which is especially appealing in modern data settings when the number of potential candidates is huge. Indeed, Ω_{xy} describes the *direct* relationships between predictors and responses, the support of which we are looking for to select relevant interactions. On the other hand, \mathbf{B} entails *both direct and indirect* influences, possibly due to some strong correlations between the responses, described by the covariance matrix \mathbf{R} (or equivalently its inverse Ω_{yy}). To provide additional insights on cGGM, Fig. 1 illustrates the relationships between \mathbf{B} , Ω_{xy} and \mathbf{R} in two simple sparse scenarios where $p = 40$ and $q = 5$. Scenarios (a) and (b) are discriminated by the presence of a strong structure among the predictors. An important point to grasp in both scenarios is how strong correlations between outcomes can completely “mask” the direct links in the regression coefficients: the stronger the correlation in \mathbf{R} , the less it is possible to distinguish in \mathbf{B} the non-zero coefficients of Ω_{xy} . Scenario (c) illustrates the reverse situation of (a), where \mathbf{B} would be sparse and Ω_{xy} would not. Choosing between the two is a matter of modeling. Our approach relies on the assumptions that direct links are sparse.

In this context, sparse inference strategy have been recently proposed in Sohn and Kim (2012), Yuan and Zhang (2014) when both Ω_{xy} and \mathbf{R}^{-1} are assumed to be sparse. In contrast, we discuss in the following a regularization scheme tailored to recover Ω_{xy} and \mathbf{R}^{-1} in the case when the later is sparse and structured and the former is relatively small and dense, as in scenario b of Fig. 1.

2.2 Structured regularization with underlying sparsity

Our regularization scheme starts by considering some structural prior information about the relationships between the coefficients. We are typically thinking of a situation where similar inputs are expected to have similar direct relationships with the outputs. The right panel of Fig. 1 represents an illustration of such a situation, where there exists an extreme neighborhood structure between the predictors. This depicts a pattern that acts along the rows of \mathbf{B} or $\boldsymbol{\Omega}_{xy}$ as substantiated by the following Bayesian point of view.

2.2.1 Bayesian interpretation

Suppose that the similarities between the covariates can be encoded into a $p \times p$ matrix \mathbf{L} . Our aim is to account for this information when learning the coefficients. The Bayesian framework provides a convenient setup for defining the way the structural information should be accounted for. In the single output case (see, e.g. Marin and Robert 2007), a conjugate prior for β (the univariate version of \mathbf{B}) would be $\mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{L}^{-1})$. Combined with the covariance between the outputs, this gives

$$\text{vec}(\mathbf{B}) \sim \mathcal{N}(\mathbf{0}, \mathbf{R} \otimes \mathbf{L}^{-1}),$$

where \otimes is the Kronecker product. This framework is similar to the one considered in Brown et al. (1998) (where the prior for \mathbf{B} would be written as $\mathcal{N}(\mathbf{L}^{-1}, \mathbf{R})$). Two classical choices for \mathbf{L} would be $\mathbf{L} = \lambda \mathbf{I}$, where no prior structure is encoded, or $\mathbf{L} = \mathbf{X}^T \mathbf{X}$, where the prior structure is borrowed from the empirical covariance matrix between the covariates. The later structure is known as the g -prior (see e.g. Brown et al. 1998; Krishna et al. 2009). In following we consider the case where \mathbf{L} is given, is built upon some exogenous prior information and has an arbitrary form. In Sect. 5, we will show how such a matrix can be constructed for a specific application.

As explained in the previous section, we are primarily interested in the (sparse) estimation of $\boldsymbol{\Omega}_{xy}$. By properties of the vec operator, the prior distribution of \mathbf{B} can be restated as

$$\text{vec}(\boldsymbol{\Omega}_{xy}) \sim \mathcal{N}(\mathbf{0}, \mathbf{R}^{-1} \otimes \mathbf{L}^{-1})$$

for the direct links. Choosing such a prior results in

$$\log \mathbb{P}(\boldsymbol{\Omega}_{xy} | \mathbf{L}, \mathbf{R}) = \frac{1}{2} \text{tr} \left(\boldsymbol{\Omega}_{xy}^T \mathbf{L} \boldsymbol{\Omega}_{xy} \mathbf{R} \right) + \text{cst.}$$

2.2.2 Criterion

Through this argument, we propose a criterion with two terms to regularize the conditional negative log-likelihood (5): first, a smooth trace term relying on the available structural information \mathbf{L} ; and second, an ℓ_1 norm that encourages sparsity

among the direct links. We write the criterion as a function of $(\boldsymbol{\Omega}_{xy}, \boldsymbol{\Omega}_{yy})$ rather than $(\boldsymbol{\Omega}_{xy}, \mathbf{R})$, although equivalent in terms of estimation, since the former leads to a convex formulation. The optimization problem turns to the minimization of

$$J(\boldsymbol{\Omega}_{xy}, \boldsymbol{\Omega}_{yy}) = -\frac{1}{n} \log L(\boldsymbol{\Omega}_{xy}, \boldsymbol{\Omega}_{yy}) + \frac{\lambda_2}{2} \text{tr} \left(\boldsymbol{\Omega}_{yx} \mathbf{L} \boldsymbol{\Omega}_{xy} \boldsymbol{\Omega}_{yy}^{-1} \right) + \lambda_1 \|\boldsymbol{\Omega}_{xy}\|_1. \quad (7)$$

2.3 Connection to other sparse methods

To get more insight into our model and to facilitate connections with existing approaches, we shall write the objective function (7) as a penalized *univariate* regression problem. This amounts to “vectorizing” model (4) with respect to $\boldsymbol{\Omega}_{xy}$, i.e., to writing the objective as a function of $(\omega, \boldsymbol{\Omega}_{yy})$ where $\omega = \text{vec}(\boldsymbol{\Omega}_{xy})$. This is stated in the following proposition, which can be derived from straightforward matrix algebra, as proved in “Appendix”. The main interest of this proposition will become clear when deriving the optimization procedure that aims at minimizing (7), as the optimization problem when $\boldsymbol{\Omega}_{yy}$ is fixed turns to a generalized Elastic-Net problem. Note that we use $\mathbf{A}^{1/2}$ to denote the square root of a matrix, obtained for instance by a Cholesky factorization in the case of a symmetric positive definite matrix.

Proposition 1 *Let $\omega = \text{vec}(\boldsymbol{\Omega}_{xy})$. An equivalent vector form of (7) is*

$$J(\omega, \boldsymbol{\Omega}_{yy}) = -\frac{1}{2} \log |\boldsymbol{\Omega}_{yy}| + \frac{1}{2n} \left\| \tilde{\mathbf{y}} + \tilde{\mathbf{X}} \omega \right\|_2^2 + \frac{\lambda_2}{2} \omega^T \left(\boldsymbol{\Omega}_{yy}^{-1} \otimes \mathbf{L} \right) \omega + \lambda_1 \|\omega\|_1,$$

where the $n \times q$ dimensional vector $\tilde{\mathbf{y}}$ and the $nq \times pq$ dimensional matrix $\tilde{\mathbf{X}}$ depends on $\boldsymbol{\Omega}_{yy}$ such that

$$\tilde{\mathbf{y}} = \text{vec}(\mathbf{Y} \boldsymbol{\Omega}_{yy}^{1/2}), \quad \tilde{\mathbf{X}} = \left(\boldsymbol{\Omega}_{yy}^{-1/2} \otimes \mathbf{X} \right).$$

When there is only one response ($q = 1$) the variance-covariance matrix \mathbf{R} and its inverse $\boldsymbol{\Omega}_{yy}$ turns to a scalar denoted by σ^2 , the matrix \mathbf{B} to a p -vector β , and the objective (7) can be rewritten

$$\log(\sigma) + \frac{1}{2n\sigma^2} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + \frac{\lambda_2}{2\sigma^2} \beta^T \mathbf{L} \beta + \frac{\lambda_1}{\sigma^2} \|\beta\|_1. \quad (8)$$

We recognize in (8) the log-likelihood of a “usual” linear model with Gaussian noise, unknown variance σ^2 and regression parameters β , plus a mixture of ℓ_1 and ℓ_2 norms. Assuming that σ is known, we recognize the general structured Elastic-net regression:

$$\frac{1}{2} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + \frac{\lambda_2}{2} \beta^T \mathbf{L} \beta + \lambda_1 \|\beta\|_1.$$

If σ is unknown, letting $\lambda_2 = 0$ in (8) leads to the ℓ_1 -penalized mixture regression model proposed by Stadler et al. (2010), where both β and σ^2 are inferred. In Stadler et al. (2010), the authors insist on the importance of penalizing the vector of coefficients by an amount that is inversely proportional to σ : in the closely related Bayesian Lasso framework of Park and Casella (2008), the prior on β is defined conditionally on σ to guarantee an unimodal posterior.

We also mention the multitask group-lasso that has been proposed to deal with multivariate outcomes. The purpose is to select a subset of covariates that jointly influence all the outcomes. Denoting $\check{y} = \text{vec}(\mathbf{Y})$, $\check{\beta} = \text{vec}(\beta)$ and $\check{\mathbf{X}} = \mathbf{I} \otimes \mathbf{X}$, the objective function is

$$\|\check{y} - \check{\mathbf{X}}\check{\beta}\|_2^2 + \sum_g \lambda_g \|\check{\beta}_g\|_2.$$

where $\check{\beta}_g = (\beta_{g1}, \dots, \beta_{gq})$ stands for the g th row of matrix \mathbf{B} .

3 Learning

3.1 Optimization

In the classical framework of parametrization (1), alternate strategies where optimization is successively performed over Ω_{xy} and Ω_{yy} have been proposed Rothman et al. (2010) and Yin and Li (2011). They come with no guarantee of convergence to the global optimum since the objective is only bi-convex. In the cGGM framework of Sohn and Kim (2012) and Yuan and Zhang (2014), the optimized criterion is jointly convex yet no convergence result is provided regarding the optimization procedure proposed by the authors. Here we also consider the alternate strategy for which the following proposition can be stated:

Proposition 2 *Let $n \geq q$. Criterion (7) is jointly convex in $(\Omega_{xy}, \Omega_{yy})$. Moreover, the alternate optimization*

$$\hat{\Omega}_{yy}^{(k+1)} = \arg \min_{\Omega_{yy} > 0} J_{\lambda_1 \lambda_2}(\hat{\Omega}_{xy}^{(k)}, \Omega_{yy}), \tag{9a}$$

$$\hat{\Omega}_{xy}^{(k+1)} = \arg \min_{\Omega_{xy}} J_{\lambda_1 \lambda_2}(\Omega_{xy}, \hat{\Omega}_{yy}^{(k+1)}). \tag{9b}$$

leads to the optimal solution.

The proof is given in ‘‘Appendix’’, and relies on the fact that efficient procedures exist to solve the two convex sub-problems (9a) and (9b).

Because our procedure relies on alternating optimization, it is difficult to give either a global rate of convergence or a

complexity bound. Nevertheless, the complexity of each iteration is easy to derive, since it amounts to two well-known problems: the main computational cost in (9a) is due to the SVD of a $q \times q$ matrix, which costs $\mathcal{O}(q^3)$. Concerning (9b), it amounts to the resolution of an Elastic-Net problem with $p \times q$ variables and $n \times q$ samples. If the final number of nonzero entries in $\hat{\Omega}_{xy}$ is k , a good implementation with Cholesky update/downdate is roughly in $\mathcal{O}(npq^2k)$ (see, e.g. Bach et al. 2012). Since we typically assumed that $p \geq n \geq q$, the global cost of a single iteration of the alternating scheme is thus $\mathcal{O}(npq^2k)$, and we theoretically can treat problems with large p when k remains moderate.

Finally, we typically want to compute a series of solutions along the regularization path of Problem (7), i.e. for various values of (λ_1, λ_2) . To this end, we simply choose a grid of penalties $\Lambda_1 \times \Lambda_2 = \{\lambda_1^{\min}, \dots, \lambda_1^{\max}\} \times \{\lambda_2^{\min}, \dots, \lambda_2^{\max}\}$. The process is easily distributed on different computer cores, each core corresponding to a value picked in Λ_2 . Then on each core—i.e. for a fix $\lambda_2 \in \Lambda_2$ —we cover all the values of $\lambda_1 \in \Lambda_1$ relying on the warm start strategy frequently used to go through the regularization path of ℓ_1 -penalized problems (see Osborne et al. 2000).

3.2 Model selection and parameter tuning

Model selection amounts here to choosing a couple of tuning parameters (λ_1, λ_2) among the grid $\Lambda_1 \times \Lambda_2$, where λ_1 tunes the sparsity while λ_2 tunes the amount of structural regularization. We aim to pick either the model with minimum prediction error, or the one closest to the true model, assuming Eq. (3) holds. These two perspectives generally do not lead to the same model choices: when looking for the model minimizing the prediction error, K -fold cross-validation is the recommended option (Hesterberg et al. 2008) despite its additional computational cost. Letting $\kappa : \{1, \dots, n\} \rightarrow \{1, \dots, K\}$ the function indexing the fold to which observation i is allocated, the CV-choices for $(\lambda_1^{cv}, \lambda_2^{cv})$ are the ones that minimize

$$\frac{1}{n} \sum_{i=1}^n \left\| \mathbf{x}_i^T \hat{\mathbf{B}}_{-\kappa(i)}^{\lambda_1, \lambda_2} - \mathbf{y}_i \right\|_2^2,$$

where $\hat{\mathbf{B}}_{-\kappa(i)}^{\lambda_1, \lambda_2}$ minimizes the objective (7), once data $\kappa(i)$ has been removed. In place of the prediction error, other quantities can be cross-validated in our context, e.g. the negative log-likelihood (5). In the remainder of the paper, however, we only consider cross-validating the prediction error, though.

As an alternative to cross-validation, penalized criteria provide a fast way to perform model selection and are sometimes more suited to the selection of the true underlying model. Although relying on asymptotic derivations, they often offer good practical performance when the sample size

n is not too small compared to the problem dimension. For penalized methods, a general form for various information criteria is expressed as a function of the likelihood L (defined by (5) here) and the effective degrees of freedom:

$$-2 \log L \left(\hat{\boldsymbol{\Omega}}_{\mathbf{xy}}^{\lambda_1, \lambda_2}, \hat{\boldsymbol{\Omega}}_{\mathbf{yy}}^{\lambda_1, \lambda_2} \right) + \text{pen} \cdot \text{df}_{\lambda_1, \lambda_2}. \tag{10}$$

Setting $\text{pen} = 2$ or $\log(n)$ respectively leads to AIC or BIC. For the practical evaluation of (10), we must give some sense to the effective degrees of freedom $\text{df}_{\lambda_1, \lambda_2}$. We use the now classical definition of Efron (2004) and rely on the work of Tibshirani and Taylor (2012) to derive the following Proposition, the proof of which is postponed to ‘‘Appendix’’.

Proposition 3 *An unbiased estimator of $\text{df}_{\lambda_1, \lambda_2}$ for our fitting procedure is*

$$\hat{\text{df}}_{\lambda_1, \lambda_2} = \text{card}(\mathcal{A}) - \lambda_2 \text{tr} \left(\left(\hat{\mathbf{R}} \otimes \mathbf{L} \right)_{\mathcal{A}, \mathcal{A}} \left(\hat{\mathbf{R}} \otimes (\mathbf{S}_{\mathbf{xx}} + \lambda_2 \mathbf{L}) \right)_{\mathcal{A}, \mathcal{A}}^{-1} \right),$$

where $\mathcal{A} = \left\{ j : \text{vec} \left(\hat{\boldsymbol{\Omega}}_{\mathbf{xy}}^{\lambda_1, \lambda_2} \right) \neq 0 \right\}$ is the set of nonzero entries in $\hat{\boldsymbol{\Omega}}_{\mathbf{xy}}^{\lambda_1, \lambda_2}$.

Note that we can compute this expression at no additional cost, relying on computations already made during the optimization process.

4 Simulation studies

In this section, we would like to illustrate the new features of our proposal compared to several baselines in well controlled settings. To this end, we perform three simulation studies to evaluate (i) the gain brought by the estimation of the residual covariance \mathbf{R} , (ii) the gain brought by the inclusion of informative structure on the predictors via \mathbf{L} and (iii) the behavior of the method in presence of both residual covariance and structure.

4.1 Implementation details

In our experiments, performance are compared with well-established regularization methods, whose implementation is easily accessible: the LASSO (Tibshirani 1996), the fused-LASSO (Tibshirani et al. 2005), the multitask group-LASSO, MRCE (Rothman et al. 2010), the Elastic-Net (Zou and Hastie 2005) and the Structured Elastic-Net (Slawski et al. 2010). LASSO and group-LASSO are fitted with the R-package `glmnet` (Friedman et al. 2010), fused-LASSO with the R-package `FusedLasso` (Hoeffling 2010) and MRCE with Rothman et al. (2010)’s package. All other methods are fitted

using our own code. When applied to multiple outcomes, all univariate methods (LASSO, fused-LASSO, MRCE, Elastic-Net and Structured Elastic-Net) were run on each output separately. The tuning parameter was selected by a cross-validation step per output. Our own procedure is available as an R-package called `spring`, distributed on the R-forge.¹ As such, we sometimes refer to our method as ‘SPRING’ in the simulation part.

4.1.1 Data generation

Artificial datasets are generated according to the multivariate regression model (1). We assume that the decomposition $\mathbf{B} = \boldsymbol{\Omega}_{\mathbf{xy}} \boldsymbol{\Omega}_{\mathbf{yy}}^{-1} = \boldsymbol{\Omega}_{\mathbf{xy}} \mathbf{R}$ holds for the regression coefficients. We control the sparsity pattern of $\boldsymbol{\Omega}_{\mathbf{xy}}$ by arranging non null entries according to a possible structure of the predictors along the rows of $\boldsymbol{\Omega}_{\mathbf{xy}}$. We always use uncorrelated Gaussian predictors $\mathbf{x}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ in order not to promote excessively the methods that take this structure into account. Strength of the relationships between the outputs are tuned by the covariance matrix \mathbf{R} . We measure the performance of the learning procedures thanks to the prediction error (PE) estimated using a large test set of observations generated according to the true model. When relevant, mean squared error (MSE) of the regression coefficients \mathbf{B} is also presented. For conciseness, it is eluded when it shows results which are quantitatively identical to PE.

4.2 Influence of covariance between outcomes

The first simulation study aims to illustrate the advantage of splitting \mathbf{B} into $-\boldsymbol{\Omega}_{\mathbf{xy}} \mathbf{R}$ over working directly on \mathbf{B} . We set $p = 40$ predictors, $q = 5$ outcomes and randomly select 25 non null entries in $\{-1, 1\}$ to fill $\boldsymbol{\Omega}_{\mathbf{xy}}$. We do not put any structure along the predictors as this study intends to measure the gain of using the cGGM approach. The covariance follows a Toeplitz scheme:² one has $\mathbf{R}_{ij} = \tau^{|i-j|}$, for $i, j = 1, \dots, q$. We consider three scenarios tuned by $\tau \in \{.1, .5, .9\}$ corresponding to an increasing correlation between the outcomes that eventually makes the cGGM more relevant. These settings have been used to generate panel (a) of Fig. 1. For each covariance scenario, we generate a training set with size $n = 50$ and a test set with size $n = 1000$. We assess the performance of SPRING (taking $\mathbf{L} = \mathbf{I}$) by comparison with three baselines: (i) the LASSO, (ii) the ℓ_1/ℓ_2 multitask group-LASSO and (iii) SPRING with known covariance matrix \mathbf{R} . As it corresponds to the best fit we can obtain with our proposal, we call this variant the ‘oracle’ mode of SPRING. The

¹ <https://r-forge.r-project.org/projects/spring-pkg/>.

² We set \mathbf{R} a correlation matrix in order not to excessively penalize the LASSO or the group-LASSO, which both use the same tuning parameter λ_1 across the outcomes (and thus the same variance estimator).




Fig. 2 Illustration of the influence of correlations between outcomes. Scenarios {low, med, high} map to $\tau \in \{.1, .5, .9\}$ in the Toeplitz-shaped covariance matrix \mathbf{R}

final estimators are obtained by fivefold cross-validation on the training set. Figure 2 gives the boxplots of PE obtained for 100 replicates. As expected, the advantage of taking the covariance into account becomes more and more important for maintaining a low PE when τ increases. When correlation is low, the LASSO dominates the group-LASSO; this is the other way around in the high correlation setup, where the latter takes full advantage of its grouping effect along the outcomes. Still, our proposal remains significantly better as soon as τ is substantial enough. We also note that our iterative algorithm does a good job since SPRING remains close to its oracle variant.

4.3 Structure integration and robustness

The second simulation study is designed to measure the impact of introducing an informative prior along the predictors via \mathbf{L} . To remove any effect induced by the covariance between the outputs, we set $q = 1$. In this case the criterion is written as in Expression (8): \mathbf{R} boils down to a scalar σ^2 and $\mathbf{B}, \mathbf{\Omega}_{xy}$ turn to two p -size vectors sharing the same sparsity pattern such that $\boldsymbol{\beta} = -\boldsymbol{\omega}\sigma^2$. In this situation, SPRING is close to Slawski et al. (2010)’s structured Elastic-Net, except that we hope for a better estimation of the coefficients thanks to the estimation of σ . For comparison, we thus draw inspiration from the simulation settings originally used to illustrate the structured Elastic-Net: we set $\boldsymbol{\omega} = (\omega_j)_{j=1}^p$, with $p = 100$, so that we observe a sparse vector with two smooth bumps, one positive and the other one negative:

$$\omega_j = \begin{cases} -((30 - j)^2 - 100)/200 & j = 21, \dots, 39, \\ ((70 - j)^2 - 100)/200 & j = 61, \dots, 80, \\ 0 & \text{otherwise.} \end{cases}$$

Table 1 Illustrating gain brought by structure information toward statistical performances

| Method | Hamming dist. | MSE | PE |
|-------------------------------|---------------|-------------|--------------|
| SPRING ($\lambda_2 = 0.01$) | 14.31 (3.51) | 0.05 (0.01) | 30.46 (2.12) |
| SPRING ($\lambda_2 = 0$) | 20.66 (4.18) | 0.31 (0.07) | 55.73 (7.82) |
| LASSO | 20.47 (4.08) | 0.31 (0.07) | 55.66 (7.78) |

A natural choice for \mathbf{L} is to rely on the first forward difference operator as in the fused-Lasso (Tibshirani et al. 2005), or its smooth counterpart (Hebiri and Geer 2011). We thus set $\mathbf{L} = \mathbf{D}^T \mathbf{D}$ with

$$\mathbf{D}_{ij} = \begin{cases} 1 & \text{if } i = j, \\ -1 & \text{if } j = i + 1, \\ 0 & \text{otherwise} \end{cases}, \quad \begin{matrix} i = 1, \dots, p - 1, \\ j = 1, \dots, p. \end{matrix} \quad (11)$$

Hence, \mathbf{L} is the combinatorial Laplacian of a chain graph between the successive predictors.

Table 1 shows the typical gain brought by prior structure knowledge. We generate 100 learning samples of size $n = 100$ with $\sigma^2 = 5$ and reported the estimated PE, MSE and Hamming distance with the true coefficient) for the LASSO and for two versions of SPRING, with ($\lambda_2 = .01$) and without ($\lambda_2 = 0$) informative prior. We remind that the Hamming distance is $\sum_j s_j(1 - \hat{s}_j) + \hat{s}_j(1 - s_j)$, where s_j (resp. \hat{s}_j) is 1 if variable j has a non-zero coefficient in the true (resp. estimated) model, and 0 otherwise. It quantifies the distance between the true and estimated supports. Incorporating relevant structural information leads to a dramatic improvement in prediction but also in estimating the vector of parameters. As expected, univariate SPRING with no prior performs like the LASSO.

Now, the question is: what if we introduce a wrong prior, i.e. a matrix \mathbf{L} completely disconnected from the true structure of the coefficients? To answer this question, we use the same settings as above and randomly swap the entries in $\boldsymbol{\omega}$, using exactly the same \mathbf{x}_i to generate the \mathbf{y}_i .³ We then apply SPRING using respectively a non informative prior equal to the identity matrix (thus mimicking the Elastic-Net) and a ‘wrong’ prior \mathbf{L} whose rows and columns remain unswapped. We also try the LASSO, the Elastic-Net and the structured Elastic-Net. All methods are tuned with five-fold cross-validation on the learning set. Table 2 presents the results averaged over 100 runs both in terms of PE (using a size-1000 test set) and MSE.

As expected, the methods that do not integrate any structural information (LASSO, Elastic-Net and SPRING with $\mathbf{L} = \mathbf{I}$) are not affected by the permutation, and we avoid these redundancies in Table 2 to save space. Overall, these

³ We also used the same seed and CV-folds for both scenarios.

Table 2 Structure integration: performance and robustness

| Method | Scenario | MSE | PE |
|--|-----------|-------------|--------------|
| LASSO | – | .336 (.096) | 58.6 (10.2) |
| E-Net ($\mathbf{L} = \mathbf{I}$) | – | .340 (.095) | 59 (10.3) |
| SPRING ($\mathbf{L} = \mathbf{I}$) | – | .358 (.094) | 60.7 (10) |
| S. E-net | Unswapped | .163 (.036) | 41.3 (4.08) |
| ($\mathbf{L} = \mathbf{D}^T \mathbf{D}$) | Swapped | .352 (.107) | 60.3 (11.42) |
| SPRING | Unswapped | .062 (.022) | 31.4 (2.99) |
| ($\mathbf{L} = \mathbf{D}^T \mathbf{D}$) | Swapped | .378 (.123) | 62.9 (13.15) |

Swapped and unswapped results are the same for LASSO, Elastic-Net and SPRING for $\mathbf{L} = \mathbf{I}$




Fig. 3 Illustration of SPRING in a complete setting. Scenarios {low, med, high} sample size corresponds to $n \in \{50, 100, 200\}$

three methods share similar performance both in terms of PE and MSE. When the prior structure is relevant, SPRING, and to a lesser extent the structured Elastic-Net, clearly outperform the other competitors. Surprisingly, this is particularly true in terms of MSE, where SPRING also dominates the structured Elastic-Net that works with the same information. This means that the estimation of the variance also helped in the inference process. Finally, these results essentially support the robustness of the structured methods which are not much altered when using a wrong prior specification.

4.4 Assessing both covariance and structure estimation

This last simulation setting is made under a complete setting, including a structure between the predictors and correlations between the outcomes just like in Fig. 1, panel (b), with $p = 40$ and $q = 5$. The \mathbf{R} matrix has Toeplitz shape with parameter $\tau = 0.75$. The \mathbf{L} matrix is as described in Sect. 4.3. The results are displayed in Fig. 3. As expected, being the only method accounting for both the structure and the covariance, SPRING with structure performs well. Not accounting for the structure in SPRING (but keeping the covariance between outcomes) degrades the performances, which still remain comparable to these of the structured Elastic-net (that do account for the structure).

5 Application studies

In this section the flexibility of our proposal is illustrated by investigating three multivariate data problems from various contexts, namely spectroscopy, genetics and genomics, where we insist on the construction of the structuring matrix \mathbf{L} .

5.1 Near-infrared spectroscopy of cookie dough pieces

5.1.1 Context

In Near-Infrared (NIR) spectroscopy, one aims to predict one or several quantitative variables from the NIR spectrum of a given sample. Each sampled spectrum is a curve that represents the level of reflectance along the NIR region, that is, wavelengths from 800 to 2500 nanometers (nm). The quantitative variables are typically related to the chemical composition of the sample. The problem is then to select the most predictive region of the spectrum, i.e. some peaks that show good capabilities for predicting the response variable(s). This is known as a “calibration problem” in Statistics. NIR technique is used in fields as diverse as agronomy, astronomy or pharmacology. In such experiments, it is likely to encounter very strong correlations and structure along the predictors. In this perspective, Hans (2011) proposes to apply the Elastic-Net which is known to select simultaneously groups of correlated predictors. However it is not adapted to the prediction of several responses simultaneously. In Brown et al. (2001), an interesting wavelet regression model with Bayesian inference is introduced that enters the multivariate regression model, as does our proposal.

5.1.2 Description of the dataset

We consider the cookie dough data from Osborne et al. (1984). The data with the corresponding test and training sets are available in the `fds` R package. After data pretreatments as in Brown et al. (2001), we have $n = 39$ dough pieces in the training set: each sample consists in an NIR spectrum with $p = 256$ points measured from 1380 to 2400 nm (spaced by 4 nm), and in four quantitative variables that describe the percentages of fat, sugar flour and water of in the piece of dough.

5.1.3 Structure specification

We would like to account for the neighborhood structure between the predictors which is obvious in the context of NIR spectroscopy: since spectra are continuous curves, a smooth neighborhood prior will encourage predictors to be selected by “wave”, which seems more satisfactory than isolated peaks. Thus, we naturally define \mathbf{L} by means of the

Table 3 Prediction error for the cookie dough data

| Method | Fat | Sucrose | Flour | Water |
|--------------------|-------------|-------------|-------------|-------------|
| Stepwise MLR | .044 | 1.188 | .722 | .221 |
| Decision theory | .076 | .566 | .265 | .176 |
| PLS | .151 | .583 | .375 | .105 |
| PCR | .160 | .614 | .388 | .106 |
| Wavelet regression | .058 | .819 | .457 | .080 |
| LASSO | .044 | .853 | .370 | .088 |
| Fused-LASSO | .106 | 1.93 | .378 | .231 |
| Group-LASSO | .127 | .918 | .467 | .102 |
| MRCE | .151 | .821 | .321 | .081 |
| Structured E-net | .044 | .596 | .363 | .082 |
| SPRING (CV) | .065 | .397 | .237 | .083 |
| SPRING (BIC) | .048 | .389 | .243 | .066 |

Results of the best competitor appears in bold

first forward difference operator (11). We also tested higher orders of the operator to induce a stronger smoothing effect, but they do not lead to dramatic changes in terms of PE and we omit them. Order k is simply obtained by powering the first order matrix. Such techniques have been studied in a structured ℓ_1 -penalty framework in Kim et al. (2009), Tibshirani and Taylor (2011) and is known as *trend filtering*. Our approach however is different, though: it enters a multivariate framework and is based on a smooth ℓ_2 -penalty coupled with the ℓ_1 -penalty for selection.

5.1.4 Results

The predictive performance of a series of regression techniques is compared in Table 3, evaluated on the same test set.

The first four rows (namely stepwise multivariate regression, Bayesian decision theory approach, Partial Least Square and Principal Component Regression) correspond to calibration methods originally presented in Brown et al.

(2001); row five (Bayesian Wavelet Regression) is the original proposal from Brown et al. (2001); the remainder of the table is due to our own analysis. Among all the considered methods, only the group-Lasso imposes the same covariates to be active for all responses. We observe that, on this example, SPRING achieves extremely competitive results. The BIC performs notably well as a model selection criterion.

Figure 4 shows the regression coefficients adjusted with the penalized regression techniques: apart from LASSO which selects very isolated (and unstable) wavelengths, non-zero regression coefficients have quite a wide spread, and therefore hard to interpret. As expected, the multitask group-Lasso activates the same predictors across the responses, which is not a good idea when looking at the predictive performance in Table 3. The Fused-Lasso detects the most significant regions in the spectrum. However, the piecewise constant shape of the estimated signal leads to poor predictive performances. On the other hand, in Fig. 5, the direct effects selected by SPRING with BIC define predictive regions specific to each response which are well suited for interpretability purposes. Parameters fitted by our approach with BIC are represented on Fig. 5 with, from left to right, the estimators of the regression coefficients \mathbf{B} , of the direct effects $\mathbf{\Omega}_{xy}$ and of the residual covariance \mathbf{R} . As expected, the regression coefficients show no sparsity pattern since sparsity was imposed on direct links in the estimation step. One can observe a strong spatial structure along the spectrum, characterized by waves induced by a smooth first order difference prior. Concerning a potential structure between the outputs, we identify interesting regions where a strong correlation between the responses induces correlations between the regression coefficients. Consider for instance positions of the wavelength between 1750 and 2000 nm: the regression parameters related to “dry-flour” are clearly anti-correlated with those related to “sucrose”. Still, we cannot distinguish in $\hat{\mathbf{B}}$ for a direct effect of this region on either the flour or sucrose composition. Such a distinction is achieved on the middle panel where direct effects $\hat{\mathbf{\Omega}}_{xy}$ selected are plotted:




Fig. 4 Parameters estimated by the penalized regression methods for the cookie data




Fig. 5 Parameters estimated by our proposal for the cookie dough data (better seen in color). Since $\hat{\mathbf{B}}$ and $\hat{\mathbf{Q}}_{xy}$ are opposite in sign in our model, we represent $-\hat{\mathbf{Q}}_{xy}$ to ease the comparison between direct and indirect effects. (Color figure online)

it defines sparse predictive regions specific to each response which are well suited for interpretability purposes; in fact, it is now obvious that region 1750 to 2000 is rather linked to the sucrose.

5.2 Multi-trait genomic selection in *Brassica napus*

5.2.1 Context

In plant and animal genetics, deriving an accurate prediction of complex traits is of major importance for the early detection and selection of individuals of high genetic value. Due to the continued advancement of high-throughput genotyping and sequencing technologies, genetic information is now available at low cost, whereas acquiring trait information remains expensive. A typical experiment will only contain a few hundreds of phenotyped individuals genotyped for thousands of markers. Regularization methods such as ridge or Lasso regression or their Bayesian counterparts have been proposed (de los Campos et al. 2012) to predict one or several phenotypes based on the information of genetic markers. Still, in most studies only single trait genomic selection is performed, neglecting correlations between phenotypes and leading to poor prediction performance. Moreover, little attention has been devoted to the development of regularization methods including prior genetic knowledge.

5.2.2 Description of the dataset

We consider the *Brassica napus* dataset described in Ferreira et al. (1995) and Kole et al. (2002). Data consists in $n = 103$ double-haploid lines derived from 2 parent cultivars, ‘Stellar’ and ‘Major’, on which $p = 300$ genetic markers and $q = 8$ traits (responses) were recorded. Each marker is a

0/1 covariate with $x_i^j = 0$ if line i has the ‘Stellar’ allele at marker j , and $x_i^j = 1$ otherwise. Traits included are percent winter survival for 1992, 1993, 1994, 1997 and 1999 (surv92, surv93, surv94, surv97, surv99, respectively), and days to flowering after no vernalization (flower0), 4 weeks vernalization (flower4) or 8 weeks vernalization (flower8).

5.2.3 Structure specification

In a biparental line population, correlation between 2 markers depends on their genetic distance defined in terms of recombination fraction. As a consequence, one expects adjacent markers on the sequence to be correlated, yielding similar direct relationships with the phenotypic traits. Noting d_{12} the genetic distance between markers M_1 and M_2 , one has $\text{cor}(M_1, M_2) = \rho^{d_{12}}$, where $\rho = .98$.⁴ The covariance matrix \mathbf{L}^{-1} can hence be defined as $\mathbf{L}_{ij}^{-1} = \rho^{d_{ij}}$. Moreover, assuming recombination events are independent between M_1 and M_2 on the one hand, and M_2 and M_3 on the other hand, one has $d_{13} = d_{12} + d_{23}$ and matrix \mathbf{L}^{-1} exhibits an inhomogeneous AR(1) profile. As a consequence, \mathbf{L} is tridiagonal with general elements

$$w_{i,i} = \frac{1 - \rho^{2d_{i-1,i} + 2d_{i,i+1}}}{(1 - \rho^{2d_{i-1,i}})(1 - \rho^{2d_{i,i+1}})},$$

$$w_{i,i+1} = \frac{-\rho^{d_{i,i+1}}}{1 - \rho^{2d_{i,i+1}}}$$


and $w_{i,j} = 0$ if $|i - j| > 1$. For the first (resp. last) marker, the distance $d_{i-1,i}$ (resp. $d_{i,i+1}$) is infinite.

⁴ This value directly arises from the definition of the genetic distance itself.

Table 4 Estimated prediction error for the *Brassica napus* data for all traits analyzed jointly (standard error in parentheses)

| Method | surv92 | surv93 | surv94 | surv97 | surv99 | flower0 | flower4 | flower8 |
|-------------|--------------------|--------------------|--------------------|--------------------|--------------------|--------------------|--------------------|--------------------|
| LASSO | .730 (.011) | .977 (.009) | .943 (.010) | .947 (.009) | .916 (.010) | .609 (.011) | .501 (.011) | .744 (.011) |
| S. Enet | .697 (.011) | .987 (.009) | .941 (.011) | .945 (.009) | .911 (.010) | .577 (.011) | .478 (.010) | .727 (.012) |
| MRCE | .759 (.010) | .919 (.003) | .917 (.006) | .924 (.004) | .926 (.006) | .591 (.011) | .479 (.011) | .736 (.011) |
| group-lasso | .748 (.013) | .995 (.012) | .892 (.013) | .939 (.013) | .906 (.017) | .603 (.017) | .504 (.015) | .717 (.013) |
| SPRING | .724 (.010) | .948 (.008) | .848 (.010) | .940 (.006) | .907 (.009) | .489 (.010) | .419 (.009) | .616 (.012) |

Results of the best competitor appears in bold

**Fig. 6** Brassica study: residual covariance estimation

5.2.4 Results

To compare SPRING and its competitors in terms of predictive performance, PE is estimated by randomly splitting the 103 samples into training and test sets with sizes 93 and 10. Before adjusting the models, we first scale the outcomes on the training and test sets to facilitate interpretability. Fivefold cross-validation is used on the training set to choose the tuning parameters. Two hundred random samplings of the test and training sets were conducted to estimate the PE given in Table 4.

All methods provide similar results although SPRING provides the smallest error for half of the traits. A picture of the between-response covariance matrix estimated with SPRING is given in Fig. 6. It reflects the correlation between the traits, which are either explained by an unexplored part of the genotype, by the environment or by some interaction between the two. The residuals of the flowering times exhibit strong correlations, whereas correlations between the survival rates are weak. It also shows that the survival traits have a larger residual variability than do the flowering traits, suggesting a higher sensitivity to environmental conditions.

As suggested by a reviewer, we ran the analysis on survival and flowering responses separately. Indeed, merging

the two subsets of traits may hamper the performances of the group Lasso, as it forces the active covariates to be same for all responses. The result may also change for SPRING as the correlation between the two subsets of traits are not accounted for when analyzed separately. The other (univariate) methods are not affected by this change. The results are summarized in Table 5. No dramatic improvement or degradation is observed for any of the two methods and SPRING still displays better results.

We then turn to the effects of each marker on the different traits. The left panels of Fig. 7 give both the regression coefficients (top) and the direct effects (bottom). The gray zones correspond to chromosomes 2, 8 and 10, respectively. The exact location of the markers within these chromosomes is displayed in the right panels, where the size of the dots reflects the absolute value of the regression coefficients (top) and of the direct effects (bottom). The interest of considering direct effects rather than regression coefficients appears clearly here, if one looks for example at chromosome 2. Three large overlapping regions are observed in the coefficient plot, for each flowering trait. A straightforward interpretation would suggest that the corresponding region controls the general flowering process. The direct effect plot allows one to go deeper and shows that these three responses are actually controlled by three separate sub-regions within this chromosome. The confusion in the coefficient plot only results from the strong correlations observed among the three flowering traits.

5.3 Selecting regulatory motifs from multiple microarrays

5.3.1 Context

In genomics, the expression of genes is initiated by transcription factors that bind to the DNA upstream from the coding regions, called *regulatory regions*. This binding occurs when a given factor recognizes a certain (small) sequence called a *regulatory motif*. Genes hosting the same regulatory motif will be jointly expressed under certain conditions. As the binding relies on chemical affinity, some degeneracy can be tolerated in the motif definition, and motifs similar but for

Table 5 Estimated prediction error for the *Brassica napus* data for the flowering traits analyzed separately (standard error in parentheses)

| Method | surv92 | surv93 | surv94 | surv97 | surv99 | flower0 | flower4 | flower8 |
|-------------|--------------------|--------------------|--------------------|--------------------|--------------------|--------------------|--------------------|--------------------|
| group-lasso | .783 (.012) | .959 (.011) | .907 (.011) | .950 (.010) | .911 (.011) | .597 (.019) | .536 (.016) | .757 (.022) |
| SPRING | .732 (.010) | .986 (.010) | .906 (.011) | .927 (.009) | .897 (.009) | .484 (.018) | .443 (.015) | .639 (.021) |

Results of the best competitor appears in bold

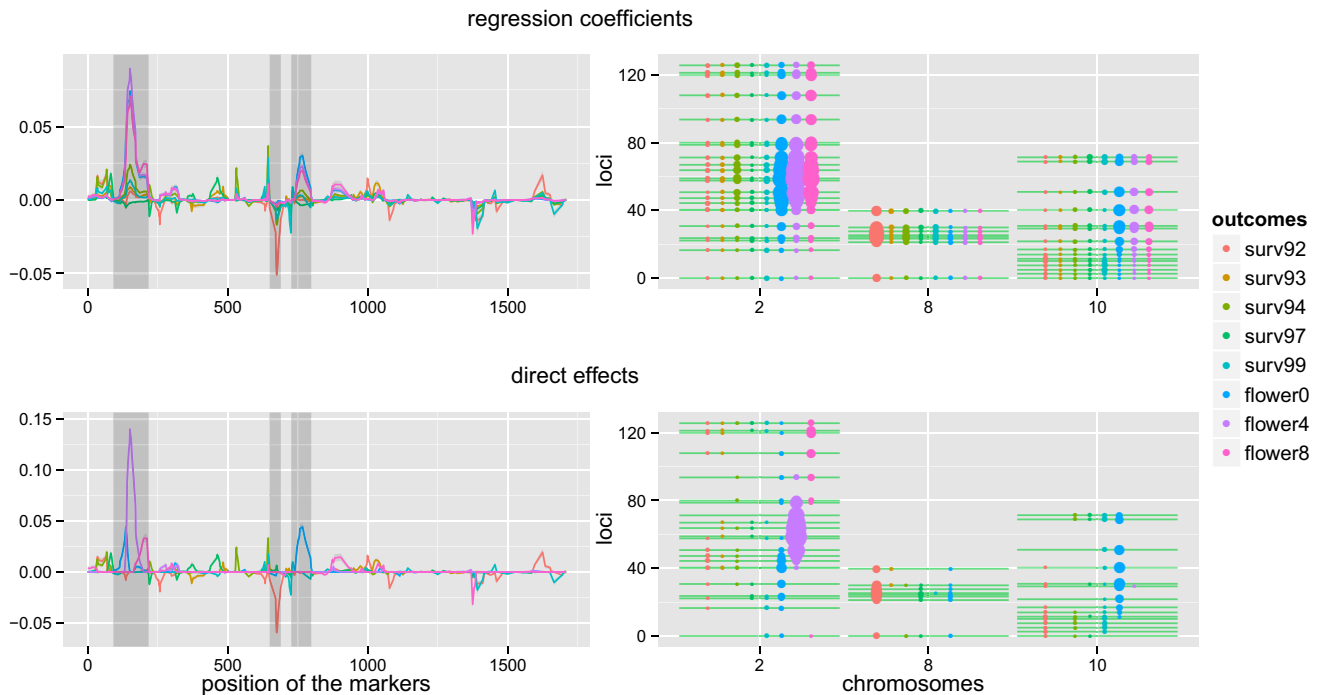


Fig. 7 Brassica study: direct and indirect genetic effects of the markers on the traits estimated by SPRING (better seen in color). (Color figure online)

small variations may share the same functional properties (see, e.g. Lajoie et al. 2012).

We are interested in the detection of such regulatory motifs, the presence of which controls the gene expression profile. To this aim we try to establish a relationship between the expression level of all genes across a series of conditions with the content of their respective regulatory regions in terms of motifs. In this context, we expect (i) the set of influential motifs to be small for each condition, (ii) the influential motifs for a given condition to be degenerate versions of each other, and (iii) the expression under similar conditions to be controlled by the same motifs.

5.3.2 Description of the dataset

In Gasch et al. (2000), a series of microarray experiments are conducted on yeast cells (*Saccharomyces cerevisiae*). Among these assays, we consider 12 time-course experiments profiling $n = 5883$ genes under various environmental changes as listed in Table 6. These expression sets form 12 potential response matrices \mathbf{Y} , the column number of which corresponds to the number of time points. Concerning the pre-

dictors, we consider the set of all motifs with length k formed with the four nucleotides, that is $\mathcal{M}_k = \{A, C, G, T\}^k$. There are $p = |\mathcal{M}_k| = 4^k$ such motifs. Unless otherwise stated, the motifs in \mathcal{M} are lined up in lexicographical order e.g., when $k = 2$, $AA, AC, AG, AT, CA, CC, \dots$ and so on. Then, the $n \times p$ matrix of predictors \mathbf{X} is filled such that X_{ij} equals the occurrence count of motif j in the regulatory region of gene i .

5.3.3 Structure specification

As we expect influential motifs for a given condition to be degenerate versions of each other, we first measure the similarity between any two motifs from \mathcal{M}_k with the Hamming defined as

$$\forall a, b \in \mathcal{M}_k, \quad \text{dist}(a, b) = \text{card}\{i : a_i \neq b_i\}.$$

For a fixed value of interest $0 \leq \ell \leq k$, we further define the ℓ -distance matrix $\mathbf{D}^{k,\ell} = (d_{ab}^{k,\ell})_{a,b \in \mathcal{M}_k}$ as

$$\forall a, b \in \mathcal{M}_k, \quad d_{ab}^{k,\ell} = \begin{cases} 1 & \text{if } \text{dist}(a, b) \leq \ell \\ 0 & \text{otherwise.} \end{cases}$$

Table 6 Time-course data from Gasch et al. (2000) considered for regulatory motif discovery

| Experiment | # time point | # motifs selected | | |
|---|--------------|-------------------|---------|---------|
| | | $k = 7$ | $k = 8$ | $k = 9$ |
| Heat shock | 8 | 30 | 68 | 43 |
| Shift from 37 to 25 °C | 5 | 3 | 11 | 33 |
| Mild heat shock | 4 | 24 | 13 | 23 |
| Response to H ₂ O ₂ | 10 | 15 | 10 | 21 |
| Menadione exposure | 9 | 16 | 1 | 7 |
| DDT exposure 1 | 8 | 15 | 10 | 30 |
| DDT exposure 2 | 7 | 11 | 33 | 21 |
| Diamide treatment | 8 | 45 | 25 | 35 |
| Hyperosmotic shock | 7 | 36 | 24 | 15 |
| Hypo-osmotic shock | 5 | 20 | 8 | 29 |
| Amino-acid starvation | 5 | 47 | 30 | 39 |
| Diauxic shift | 7 | 16 | 14 | 20 |
| Total number of unique motifs inferred | | 87 | 82 | 72 |

Table 7 Comparison of SPRING-selected motifs with MotifDB patterns. Each cell corresponds to a MotifDB pattern (top) compared to a set of aligned SPRING motifs with size 7 (down).

| | | | |
|---|---|---|---------------------------------|
| CTAAGCCAC TAGCCC GCGCCC | GCATGTGAA CCATATG TTGTGAG | CATGTAATT TGTAAAT TGTATAT | TGAAACA TTAGACC TAAAAAG |
| TGATCGGGCGCGCAGACGA GTATAAC GCGCCGT | TGCTGGTT GCTGGTT GCTGGTG | GATCGTATGATA ATCATAT TTGGTAT | ACGCGAAAA AACGAAA ACGAAAA |
| CCATACATCAC CATAGAC ATATCAC | ATTGACCTGGTC TCGACTT CGACTTG CCAGCTT | GACTAGATATATATATTCGAT ATATATT CATATAT ATATATG ATATATA | |

$D^{k,\ell}$ can be viewed as the adjacency matrix of a graph where the nodes are the motifs and where an edge is present between 2 nodes when the 2 motifs are at a Hamming distance less or equal to ℓ . We finally use the Laplacian of this graph as a structuring matrix $L^{k,\ell} = (\ell_{ab}^{k,\ell})_{a,b \in \mathcal{M}_k}$, that is

$$\ell_{ab}^{k,\ell} = \begin{cases} \sum_{c \in \mathcal{M}_k} d_{ac}^{k,\ell} & \text{if } a = b, \\ -1 & \text{if } d_{ab}^{k,\ell} = 1, \\ 0 & \text{otherwise.} \end{cases} \quad (12)$$

Note that we came across a similar proposal by Li et al. (2010) in the context of sparse, structured PLS. The derivation, however, is different, and the objective of this method is not the selection of motifs but of families of motifs via the compression performed by the PLS.

5.3.4 Results

We apply our methodology for candidate motifs from $\mathcal{M}_7, \mathcal{M}_8$ and \mathcal{M}_9 , which results in three lists of putative

motifs having a direct effect on gene expression. Due to the very large number of potential predictors that comes with a sparser matrix X when k increases, we first perform a screening step that keeps the 5000 motifs with the highest marginal correlations with Y . Second, SPRING is applied to each of the twelve time-course experiments described in Table 6. The selection of (λ_1, λ_2) is performed on a grid using the BIC (10). At the end of the day, the three lists corresponding to $k = 7, 8, 9$ include respectively 87, 82 and 72 motifs, for which at least one coefficient in the associated row $\hat{\Omega}_{xy}(j, \cdot)$ was found non-null for some of the twelve experiments, as detailed in Table 6.

To assess the relevance of the selected motifs, we compared them with the MotifDB patterns available in Bioconductor (Shannon 2013), where known transcription factor binding sites are recorded. There are 453 such reference motifs with size varying from 5 to 23 nucleotides. Consider the case of $k = 7$ for instance: among the 87 SPRING motifs, 62 match one MotifDB pattern each and 25 are clustered into 11 MotifDB patterns as depicted in Table 7. As seen in this table, the clusters of motifs selected by SPRING corre-

spond to sets of variants of the same pattern. These clusters consist of motifs that are close according to the similarity encoded in the structure matrix \mathbf{L} . In this example, the ability of SPRING to use domain-specific definitions of the structure between the predictors oriented the regression problem to account for motif degeneracy and helped in selecting motifs that are consistent known binding sites.

6 Discussion

We introduced a general framework for multivariate regression accounting for correlation between the outcomes, for a possible structure between the predictors, and inducing sparsity on the direct links between the predictors and the outcomes. The whole procedure is available through the R package **spring**. This provides a generic framework that can accommodate to a broad variety of structures between the covariates via the matrix \mathbf{L} . This comes to the price of the construction of this matrix that is specific to each application. A parallel can be made with other generic machine learning approaches such as SVM, where the kernel has to be tailored for a given application to ensure good performances. Although we recommend the \mathbf{L} matrix to be carefully designed to fully include the prior knowledge, some typical shapes have been made available in the **spring** package.

The proposed approach is valid for a limited number of outcomes, typically $q < n$. A natural extension of this work is to consider the case of large dimension outputs. This requires to also impose sparsity on the covariance matrix \mathbf{R} between the responses. The most natural way is to use an additional graphical Lasso-type penalty term. This can be achieved through a modification of step (9a) in the alternating optimization procedure. This is a work in progress.

Acknowledgements We would like to thank Mathieu Lajoie and Laurent Bréhélin for kindly sharing the dataset from [Gasch et al. \(2000\)](#). We also thank the reviewers for their questions and remarks, which helped us to improve our manuscript. This project was conducted in the framework of the project AMAIZING funded by the French ANR. This work has been partially supported by the GRANT Reg4Sel from the French INRA-SelGen metaprogram.

Appendix: Proofs

Most of the proofs rely on basic algebra and properties of the trace and the vec operators (see e.g. [Harville 1997](#)), in particular

$$\begin{aligned}\text{tr}(\mathbf{A}^T \mathbf{BCD}^T) &= \text{vec}(\mathbf{A})^T (\mathbf{D} \otimes \mathbf{B}) \text{vec}(\mathbf{C}), \\ \text{vec}(\mathbf{ABC}) &= (\mathbf{C}^T \otimes \mathbf{A}) \text{vec}(\mathbf{B}), \\ (\mathbf{A} \otimes \mathbf{B})(\mathbf{C} \otimes \mathbf{D}) &= (\mathbf{AC}) \otimes (\mathbf{BD}).\end{aligned}$$

Derivation of Proposition 1

Concerning the two regularization terms in (7), we have $\|\boldsymbol{\Omega}_{\mathbf{xy}}\|_1 = \|\boldsymbol{\omega}\|_1$ since the ℓ_1 -norm applies element-wise here and

$$\text{tr}(\boldsymbol{\Omega}_{\mathbf{yx}} \mathbf{L} \boldsymbol{\Omega}_{\mathbf{xy}} \boldsymbol{\Omega}_{\mathbf{yy}}^{-1}) = \boldsymbol{\omega}^T (\boldsymbol{\Omega}_{\mathbf{yy}}^{-1} \otimes \mathbf{L}) \boldsymbol{\omega}.$$

As for the log-likelihood (5), we work on the trace term:

$$\begin{aligned}\text{tr} \left(\left(\mathbf{Y} + \mathbf{X} \boldsymbol{\Omega}_{\mathbf{xy}} \boldsymbol{\Omega}_{\mathbf{yy}}^{-1} \right)^T \left(\mathbf{Y} + \mathbf{X} \boldsymbol{\Omega}_{\mathbf{xy}} \boldsymbol{\Omega}_{\mathbf{yy}}^{-1} \right) \boldsymbol{\Omega}_{\mathbf{yy}} \right) \\ = \text{vec} \left(\mathbf{Y} + \mathbf{X} \boldsymbol{\Omega}_{\mathbf{xy}} \boldsymbol{\Omega}_{\mathbf{yy}}^{-1} \right)^T (\boldsymbol{\Omega}_{\mathbf{yy}} \otimes \mathbf{I}_n) \text{vec} \left(\mathbf{Y} + \mathbf{X} \boldsymbol{\Omega}_{\mathbf{xy}} \boldsymbol{\Omega}_{\mathbf{yy}}^{-1} \right) \\ = \left\| (\boldsymbol{\Omega}_{\mathbf{yy}} \otimes \mathbf{I}_n)^{1/2} \left(\text{vec}(\mathbf{Y}) + (\boldsymbol{\Omega}_{\mathbf{yy}}^{-1} \otimes \mathbf{X}) \text{vec}(\boldsymbol{\Omega}_{\mathbf{xy}}) \right) \right\|_2^2 \\ = \left\| \text{vec}(\mathbf{Y} \boldsymbol{\Omega}_{\mathbf{yy}}^{1/2}) + (\boldsymbol{\Omega}_{\mathbf{yy}}^{-1/2} \otimes \mathbf{X}) \boldsymbol{\omega} \right\|_2^2.\end{aligned}$$

The rest of the proof is straightforward.

Convexity lemma

Lemma 1 *The function*

$$-\frac{1}{n} \log L(\boldsymbol{\Omega}_{\mathbf{xy}}, \boldsymbol{\Omega}_{\mathbf{yy}}) + \frac{\lambda_2}{2} \text{tr} \left(\boldsymbol{\Omega}_{\mathbf{yx}} \mathbf{L} \boldsymbol{\Omega}_{\mathbf{xy}} \boldsymbol{\Omega}_{\mathbf{yy}}^{-1} \right)$$

is jointly convex in $(\boldsymbol{\Omega}_{\mathbf{xy}}, \boldsymbol{\Omega}_{\mathbf{yy}})$ and admits at least one global minimum which is unique when $n \geq q$ and $(\lambda_2 \mathbf{L} + \mathbf{S}_{\mathbf{xx}})$ is positive definite.

The convexity of $-\frac{1}{n} \log L(\boldsymbol{\Omega}_{\mathbf{xy}}, \boldsymbol{\Omega}_{\mathbf{yy}})$ is proved in [Yuan and Zhang \(2014\)](#) (Proposition 1). Similar arguments can be straightforwardly applied in the case at hand. Existence of the global minimum is related to strict convexity in both $\boldsymbol{\Omega}_{\mathbf{xy}}$ and $\boldsymbol{\Omega}_{\mathbf{yy}}$, where direct differentiation leads to the corresponding conditions.

Proof of Proposition 2

The convexity of criterion $J(\boldsymbol{\Omega}_{\mathbf{xy}}, \boldsymbol{\Omega}_{\mathbf{yy}})$ in $(\boldsymbol{\Omega}_{\mathbf{xy}}, \boldsymbol{\Omega}_{\mathbf{yy}})$ is straightforward thanks to Lemma 1 and considering the fact that $\|\boldsymbol{\Omega}_{\mathbf{xy}}\|_1$ is also convex. One can then apply the results developed in [Tseng and Yun \(2009\)](#), [Tseng \(2001\)](#) on the convergence of block coordinate descent for the minimization of nonsmooth separable function. Since (7) is clearly separable in $(\boldsymbol{\Omega}_{\mathbf{xy}}, \boldsymbol{\Omega}_{\mathbf{yy}})$ for the nonsmooth part induced by the ℓ_1 -norm, the alternating scheme is guaranteed to converge to the unique global minimum under the assumption of Lemma 1. It remains to show that the two convex optimization subproblems (9a) and (9b) can be (efficiently) solved in practice.

Firstly, (9b) can be recast as an Elastic-Net problem, which in turn can be recast as a LASSO problem (see, e.g. [Zou and](#)

Hastie 2005; Slawski et al. 2010). This is straightforward thanks to Proposition 1: when $\hat{\Omega}_{yy}$ is fixed, solution to (9b) can be obtained via

$$\hat{\omega} \left(= \text{vec}(\hat{\Omega}_{xy}) \right) = \arg \min_{\omega \in \mathbb{R}^{pq}} \frac{1}{2} \|\mathbf{A}\omega - \mathbf{b}\|_2^2 + \frac{\lambda_2}{2} \omega^T \tilde{\mathbf{L}}\omega + \lambda_1 \|\omega\|_1, \quad (13)$$

where \mathbf{A} , \mathbf{b} and $\tilde{\mathbf{L}}$ are defined by

$$\mathbf{A} = \left(\hat{\Omega}_{yy}^{-1/2} \otimes \mathbf{X}/\sqrt{n} \right),$$

$$\mathbf{b} = -\text{vec} \left(\mathbf{Y}\hat{\Omega}_{yy}^{1/2} \right) / \sqrt{n} \text{ and } \tilde{\mathbf{L}} = \hat{\Omega}_{yy}^{-1} \otimes \mathbf{L}.$$

Secondly, we can solve analytically (9a) with simple matrix algebra. By differentiation of the objective (7) over Ω_{yy}^{-1} we obtain the quadratic form

$$\Omega_{yy} \mathbf{S}_{yy} \Omega_{yy} - \Omega_{yy} = \Omega_{yx} (\lambda_2 \mathbf{L} + \mathbf{S}_{xx}) \Omega_{xy}.$$

After right multiplying both sides by \mathbf{S}_{yy} , it becomes obvious that $\Omega_{yy} \mathbf{S}_{yy}$ and $\Omega_{yx} (\lambda_2 \mathbf{L} + \mathbf{S}_{xx}) \Omega_{xy} \mathbf{S}_{yy}$ commute and thus share the same eigenvectors \mathbf{U} . Besides, it induces the relationship $\eta_j^2 - \eta_j = \zeta_j$ between their respective eigenvalues η_j and ζ_j , and we are looking for the positive solution of η_j . To do so, first note that we may assume that $\Omega_{yx} (\lambda_2 \mathbf{L} + \mathbf{S}_{xx}) \Omega_{xy}$ and \mathbf{S}_{yy} are positive definite, when $\Omega_{xy} \neq \mathbf{0}$ and $n \geq q$; and second, recall that if a matrix is the product of two positive definite matrices then its eigenvalues are positive. Hence, $\zeta_j > 0$ and the positive solution of η_j is $\eta_j = (1 + \sqrt{1 + 4\zeta_j})/2$. We thus obtain

$$\hat{\Omega}_{yy} = \mathbf{U} \text{diag}(\eta) \mathbf{U}^{-1} \mathbf{S}_{yy}^{-1}. \quad (14)$$

Direct inversion yields

$$\hat{\Omega}_{yy} = \mathbf{U} \text{diag}(\eta/\zeta) \mathbf{U}^{-1} \hat{\Omega}_{yx} (\lambda_2 \mathbf{L} + \mathbf{S}_{xx}) \hat{\Omega}_{xy} (= \hat{\mathbf{R}}^{-1}),$$

To get an expression for $\hat{\Omega}_{yy}$ which does not require additional matrix inversion, just note that

$$\mathbf{S}_{yy}^{-1} = \left(\Omega_{yx} \hat{\Sigma}_{xx}^{\lambda_2} \Omega_{xy} \mathbf{S}_{yy} \right)^{-1} \Omega_{yx} \hat{\Sigma}_{xx}^{\lambda_2} \Omega_{xy}$$

$$= \mathbf{U} \text{diag}(\zeta^{-1}) \mathbf{U}^{-1} \Omega_{yx} \hat{\Sigma}_{xx}^{\lambda_2} \Omega_{xy}$$

where $\hat{\Sigma}_{xx}^{\lambda_2} = (\lambda_2 \mathbf{L} + \mathbf{S}_{xx})$. Combined with (14), this last equality leads to

$$\hat{\Omega}_{yy}^{-1} = \mathbf{S}_{yy} \mathbf{U} \text{diag}(\eta^{-1}) \mathbf{U}^{-1} (= \hat{\mathbf{R}}).$$

Finally, in the particular case where $\hat{\Omega}_{xy} = \mathbf{0}$, $\hat{\Omega}_{yy} = \mathbf{S}_{yy}^{-1}$.

Derivation of Proposition 3

To apply the results developed in Tibshirani and Taylor (2012) that rely on the well-known Stein’s Lemma (Stein 1981), we basically need to recast our problem as a classical LASSO applied on a Gaussian linear regression model such that the response vector follows a normal distribution of the form $\mathcal{N}(\boldsymbol{\mu}, \sigma^2 \mathbf{I})$. This is straightforward by means of Proposition 1: in the same way as we derive expression (13), we can go a little further and reach the following LASSO formulation

$$\arg \min_{\omega \in \mathbb{R}^{pq}} \frac{1}{2} \left\| \begin{pmatrix} \mathbf{A} \\ \sqrt{\lambda_2} \tilde{\mathbf{L}}^{1/2} \end{pmatrix} \omega - \begin{pmatrix} \mathbf{b} \\ \mathbf{0} \end{pmatrix} \right\|_2^2 + \lambda_1 \|\omega\|_1, \quad (15)$$

with \mathbf{A} , \mathbf{b} and $\tilde{\mathbf{L}}$ defined as in (13). From model (4), it is not difficult to see that \mathbf{b} corresponds to an uncorrelated vector form of \mathbf{Y} so as

$$\begin{pmatrix} \mathbf{b} \\ \mathbf{0} \end{pmatrix} = \begin{pmatrix} -\text{vec}(\mathbf{Y}\Omega_{yy}^{1/2}) \\ \mathbf{0} \end{pmatrix} \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{I}_{nq}).$$

The explicit form of $\boldsymbol{\mu}$ is of no interest here. The point is essentially to underline that the response vector is uncorrelated in (15), which allows us to apply Theorem 2 of Tibshirani and Taylor (2012) and results therein, notably for the Elastic-Net. By these means, an unbiased estimator of the degrees of freedom in (15) can be written as a function of the active set \mathcal{A} in $\hat{\omega}^{\lambda_1, \lambda_2}$:

$$\hat{df}_{\lambda_1, \lambda_2} = \text{tr} \left(\mathbf{A}_{\mathcal{A}} \left((\mathbf{A}^T \mathbf{A} + \lambda_2 \tilde{\mathbf{L}})_{\mathcal{A}, \mathcal{A}} \right)^{-1} \mathbf{A}_{\mathcal{A}}^T \right).$$

Routine simplifications lead to the desired result for the degrees of freedom.

References

Bach, F., Jenatton, R., Mairal, J., Obozinski, G.: Optimization with sparsity-inducing penalties. *Found. Trends Mach. Learn.* **4**(1), 1–106 (2012)

Brown, P., Vannucci, M., Fearn, T.: Multivariate bayesian variable selection and prediction. *J. R. Stat. Soc. B* **60**(3), 627–641 (1998)

Brown, P., Fearn, T., Vannucci, M.: Bayesian wavelet regression on curves with applications to a spectroscopic calibration problem. *J. Am. Stat. Assoc.* **96**, 398–408 (2001)

Chiquet, J., Grandvalet, Y., Ambroise, C.: Inferring multiple graphical structures. *Stat. Comput.* **21**(4), 537–553 (2011)

de los Campos, G., Hickey, J., Pong-Wong, R., Daetwyler, H., Calus, M.: Whole genome regression and prediction methods applied to plant and animal breeding. *Genetics* **193**(2), 327–345 (2012)

Efron, B.: The estimation of prediction error: covariance penalties and cross-validation (with discussion). *J. Am. Stat. Assoc.* **99**, 619–642 (2004)

- Ferreira, M., Satagopan, J., Yandell, B., Williams, P., Osborn, T.: Mapping loci controlling vernalization requirement and flowering time in *brassica napus*. *Theor. Appl. Genet.* **90**, 727–732 (1995)
- Friedman, J., Hastie, T., Tibshirani, R.: Regularization paths for generalized linear models via coordinate descent. *J. Stat. Softw.* **33**, 1–22 (2010)
- Gasch, A., Spellman, P., Kao, C., Carmel-Harel, O., Eisen, M.B., Storz, G., Botstein, D., Brown, P.: Genomic expression programs in the response of yeast cells to environmental changes. *Mol. Biol. Cell* **11**(12), 4241–4257 (2000)
- Hans, C.: Elastic net regression modeling with the orthant normal prior. *J. Am. Stat. Assoc.* **106**, 1383–1393 (2011)
- Harville, D.: *Matrix Algebra from a Statistician's Perspective*. Springer, New York (1997)
- Hebiri, M., van De Geer, S.: The smooth-lasso and other $l_1 + l_2$ penalized methods. *Electron. J. Stat.* **5**, 1184–1226 (2011)
- Hesterberg, T., Choi, N.M., Meier, L., Fraley, C.: Least angle and l_1 penalized regression: a review. *Stat. Surv.* **2**, 61–93 (2008)
- Hoefling, H.: A path algorithm for the fused lasso signal approximator. *J. Comput. Graph. Stat.* **19**(4), 984–1006 (2010)
- Kim, S., Xing, E.: Statistical estimation of correlated genome associations to a quantitative trait network. *PLoS Genet.* **5**(8), e1000587 (2009)
- Kim, S., Xing, E.: Tree-guided group lasso for multi-task regression with structured sparsity. In: *Proceedings of the 27th International Conference on Machine Learning*, pp. 543–550 (2010)
- Kim, S.J., Koh, K., Boyd, S., D. G.: l_1 trend filtering. *SIAM Rev.* **51**(2), 339–360 (2009)
- Kole, C., Thorman, C., Karlsson, B., Palta, J., Gaffney, P., Yandell, B., Osborn, T.: Comparative mapping of loci controlling winter survival and related traits in oilseed *brassica rapa* and *B. napus*. *Mol. Breed.* **1**, 201–210 (2002)
- Krishna, A., Bondell, H., Ghosh, S.: Bayesian variable selection using an adaptive powered correlation prior. *J. Stat. Plan. Inference* **139**(8), 2665–2674 (2009)
- Lajoie, M., Gascuel, O., Lefort, V., Brehelin, L.: Computational discovery of regulatory elements in a continuous expression space. *Genome Biol.* **13**(11), R109 (2012). doi:[10.1186/gb-2012-13-11-r109](https://doi.org/10.1186/gb-2012-13-11-r109)
- Li, C., Li, H.: Variable selection and regression analysis for graph-structured covariates with an application to genomics. *Ann. Appl. Stat.* **4**(3), 1498–1516 (2010)
- Li, X., Panea, C., Wiggins, C., Reinke, V., Leslie, C.: Learning “graph-mer” motifs that predict gene expression trajectories in development. *PLoS Comput. Biol.* **6**(4), e1000761 (2010)
- Lorbert, A., Eis, D., Kostina, V., Blei, D., Ramadge, P.: Exploiting covariate similarity in sparse regression via the pairwise elastic net. In: Teh, Y.W., Titterton, D.M. (eds.) *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics (AISTATS-10)*, vol. 9, pp. 477–484 (2010)
- Mardia, K., Kent, J., Bibby, J.: *Multivariate Analysis*. Academic Press, London (1979)
- Marin, J.M., Robert, C.P.: *Bayesian Core: A Practical Approach to Computational Bayesian Statistics*. Springer, New York (2007)
- Obozinski, G., Wainwright, M., Jordan, M.: Support union recovery in high-dimensional multivariate regression. *Ann. Stat.* **39**(1), 1–47 (2011)
- Osborne, B., Fearn, T., Miller, A., Douglas, S.: Application of near infrared reflectance spectroscopy to compositional analysis of biscuits and biscuit doughs. *J. Sci. Food Agric.* **35**, 99–105 (1984)
- Osborne, M.R., Presnell, B., Turlach, B.A.: On the lasso and its dual. *J. Comput. Graph. Stat.* **9**(2), 319–337 (2000)
- Park, T., Casella, G.: The Bayesian lasso. *J. Am. Stat. Assoc.* **103**, 681–686 (2008)
- Rapaport, F., Zinovyev, A., Dutreix, M., Barillot, E., Vert, J.P.: Classification of microarray data using gene networks. *BMC Bioinform.* **8**, 35 (2007)
- Rothman, A., Levina, E., Zhu, J.: Sparse multivariate regression with covariance estimation. *J. Comput. Graph. Stat.* **19**(4), 947–962 (2010)
- Shannon, P.: MotifDb: An Annotated Collection of Protein-DNA Binding Sequence Motifs. R package version 1.4.0 (2013)
- Slawski, M., W. Zu Castell, Tutz, G.: Feature selection guided by structural information. *Ann. Appl. Stat.* **4**, 1056–1080 (2010)
- Sohn, K., Kim, S.: Joint estimation of structured sparsity and output structure in multiple-output regression via inverse-covariance regularization. *JMLR W&CP*(22), 1081–1089 (2012)
- Städler, N., Bühlmann, P., Geer, S.: l_1 -penalization for mixture regression models. *Test* **19**(2), 209–256 (2010). doi:[10.1007/s11749-010-0197-z](https://doi.org/10.1007/s11749-010-0197-z)
- Stein, C.: Estimation of the mean of a multivariate normal distribution. *Ann. Stat.* **9**, 1135–1151 (1981)
- Tibshirani, R.: Regression shrinkage and selection via the lasso. *J. R. Stat. Soc. B* **58**, 267–288 (1996)
- Tibshirani, R., Taylor, J.: The solution path of the generalized lasso. *Ann. Stat.* **39**(3), 1335–1371 (2011). doi:[10.1214/11-AOS878](https://doi.org/10.1214/11-AOS878)
- Tibshirani, R., Taylor, J.: Degrees of freedom in lasso problems. *Ann. Stat.* **40**, 1198–1232 (2012)
- Tibshirani, R., Saunders, M., Rosset, S., Zhu, J., Knight, K.: Sparsity and smoothness via the fused lasso. *J. R. Stat. Soc. B* **67**, 91–108 (2005)
- Tseng, P.: Convergence of a block coordinate descent method for nondifferentiable minimization. *J. Optim. Theory Appl.* **109**(3), 475–494 (2001)
- Tseng, P., Yun, S.: A coordinate gradient descent method for nonsmooth separable minimization. *Math. Program.* **117**, 387–423 (2009)
- Yin, J., Li, H.: A sparse conditional Gaussian graphical model for analysis of genetical genomics data. *Ann. Appl. Stat.* **5**, 2630–2650 (2011)
- Yuan, X.T., Zhang, T.: Partial Gaussian graphical model estimation. *IEEE Trans. Inform. Theory* **60**(3), 1673–1687 (2014)
- Zou, H., Hastie, T.: Regularization and variable selection via the elastic net. *J. R. Stat. Soc. B* **67**, 301–320 (2005)