Méthodes régularisées pour l'analyse de données multivariées en grande dimension : théorie et applications.

Marie Perrot-Dockès Perrot-Dockès

▶ **To cite this version:**

Marie Perrot-Dockès Perrot-Dockès. Méthodes régularisées pour l'analyse de données multivariées en grande dimension : théorie et applications.. Applications [stat.AP]. Université Paris Saclay (COmUE), 2019. Français. NNT : 2019SACLS304 . tel-02384541

HAL Id: tel-02384541

https://theses.hal.science/tel-02384541v1

Submitted on 28 Nov 2019

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.

université PARIS-SACLAY

AgroParisTech
INSTITUT DES SCIENCES ET INDUSTRIES DU VIVANT ET DE L'ENVIRONNEMENT
PARIS INSTITUTE OF TECHNOLOGY FOR LIFE, FOOD AND ENVIRONMENTAL SCIENCES

UNIVERSITÉ PARIS SUD
Comprendre le monde, construire l'avenir

NNT : 2019SACLS304

Thèse de doctorat

Méthodes régularisées pour l'analyse de données multivariées en grande dimension: théorie et applications.

Thèse de doctorat de l'Université Paris-Saclay
préparée à l'Université de Paris-Sud et au sein d'AgroParisTech

École doctorale n°574 Mathématiques Hadamard (EDMH)
Spécialité de doctorat : Mathématiques aux interfaces

Thèse présentée et soutenue à Paris, le 8 octobre 2019, par

MARIE PERROT-DOCKÈS

Composition du Jury :

Liliane BEL Professeure, AgroParisTech (MIA 518)	Présidente
David CAUSEUR Professeur, Agrocampus Ouest (IRMAR UMR 6625)	Rapporteur
Jean-Professeur BARDET Professeur, Université Paris 1 (SAMM)	Examinateur
Pierre NEUVIAL Chargé de recherche, Institut de mathématiques de Toulouse (Statistique et Probabilités)	Examinateur
Loïc RAJJOU Professeur, AgroParisTech (UMR1318)	Examinateur
Vassili SOUMELIS Professeur des universités- praticien hospitalier, Hôpital Saint-Louis (Immunologie humaine et mécanismes inflammatoires)	Examinateur
Céline Lévy-Leduc Professeure, AgroParisTech (MIA 518)	Directrice de thèse
Julien Chiquet Chargé de recherche, AgroParisTech-INRA (MIA 518)	Co-directeur de thèse
Laure Sansonnet Maître de conférence, AgroParisTech (MIA 518)	Invité

Table des matières

Introduction

1.1 Contexte biologique

Comprendre comment certains traits phénotypiques sont transmis de génération en génération est un « mystère » qui a longtemps intrigué les scientifiques et les philosophes [1]. Les « spermatistes » pensaient que chaque spermatozoïde contenait un embryon qui n'avait plus qu'à grandir au sein de la mère, qui ne lui apportait alors que la nourriture essentielle à son développement. À l'inverse les « ovistes » pensaient que le bébé était contenu dans l'ovule et que le sperme permettait uniquement de déclencher le processus. Cependant, aucune de ces deux théories ne permet de répondre à la question récurrente *pourquoi l'enfant ressemble t-il à ses deux parents ?* Selon Darwin, les traits phénotypiques et les fonctions de chaque cellule sont expliqués par des particules très petites appelées « gemmules ». Les « gémmules » des deux parents se retrouvent ensuite dans les cellules de leur enfant qui a donc des caractéristiques moyennes entre ses deux parents. Cependant, selon cette théorie, la diversité ne ferait que diminuer d'année en année pour arriver à une espèce constituée de nombreux êtres identiques. Mendel émet alors une nouvelle théorie qui va révolutionner la génétique : les caractéristiques des individus ne sont pas gouvernées par un mais par deux facteurs, l'un venant de la mère et l'autre du père. Cette théorie a, par la suite, été vérifiée : chacune de nos caractéristiques étant bien gouvernée non par un mais par deux facteurs génotypiques. Il reste ensuite à comprendre comment ces facteurs génotypiques (maintenant connus comme gènes) influencent nos caractéristiques phénotypiques c'est pourquoi dans cette thèse nous nous intéresserons au passage du génotype au phénotype. Ces étapes sont nombreuses, complexes et la connaissance que nous en avons ne cesse d'évoluer depuis de nombreuses années. Une description détaillée de ces étapes et de l'évolution de nos connaissances est disponible dans le chapitre VI de Raven et al. (2017). Nous commençons par présenter le dogme central de la biologie moléculaire et nous donnons ensuite quelques exemples qui montrent que la vision séquentielle qu'il propose est en réalité simpliste.

Le dogme central de la biologie moléculaire décrit le passage du génotype au phénotype de manière séquentielle : l'ADN contenu dans chacun de nos chromosomes est **transcrit** en ARN qui est ensuite **traduit** en protéines qui vont directement influencer le phénotype. Plus précisément un chromosome contient deux brins d'ADN. Chacun de ces brins contient des nucléotides : l'adénine (A), la thymine

1. Un lecteur intéressé pourra se référer à « Éloge de la différence. La génétique et les hommes » de Jacquard (1978) qui propose une très belle introduction sur le sujet.

(T), la cytosine (C), la guanine (G). Un de ces brins est appelé le brin codant et l'autre est appelé le brin matrice. La séquence de nucléotides du brin codant est le complémentaire de la séquence de nucléotides du brin matrice, c'est-à-dire que la thymine sera remplacée par l'adénine, la cytosine par la guanine et réciproquement. Les gènes sont codés sur des fragments de ces brins d'ADN. Des ARN polymérases se déplacent le long du brin matrice de l'ADN pour le transcrire en ARN messager (ARNm). La séquence de nucléotides de l'ARNm est la même que celle du brin codant sauf que la thymine (T) est remplacée par l'uracile (U). Un ribosome va ensuite traduire cet ARNm en chaînes d'acides aminés. Chaque groupe de trois nucléotides parmi les 4 suivants : A, U, C, G, sera appelé un **codon**. Comme les chaînes d'ARNm sont composées avec quatre nucléotides différents, $4^3 = 64$ codons existent donc pour coder les acides aminés. Trois codons sont des codons stop qui arrêtent la traduction d'ARNm en protéines, les autres sont tous liés à un acide aminé, certains sont codés par plusieurs codons. Une fois formée une chaîne d'acides aminés se replie finissant ainsi de créer une protéine. Lorsqu'elles ne sont plus fonctionnelles, ou plus utilisées, les protéines sont ensuite dégradées par notre système et sont transformées en acides aminés et petites molécules qui font partie des métabolites.

Cette vision séquentielle du passage du génotype au phénotype est en réalité beaucoup trop simpliste. On a par exemple découvert des « petits » ARN qui seraient répresseurs de la traduction de certaines parties de l'ARNm. Ces petits ARN sont codés à divers emplacements du génome, y compris dans les zones que l'on pensait silencieuses pour la transcription. De plus les protéines agissent sur la transcription de l'ARNm. Le passage du génotype au phénotype est donc loin d'être séquentiel puisque les différentes étapes s'influencent mutuellement. De plus la connaissance que l'on en a ne cesse d'évoluer et est probablement encore simpliste par rapport à la réalité.

On étudie néanmoins les différences entre des individus d'une même espèce au cours de ces différentes étapes. Certains traits phénotypiques peuvent par exemple être expliqués par l'étude des séquences d'ADN (appelée la génomique). Plus précisément, on peut étudier des différences au sein des nucléotides qui varient fréquemment d'un individu à un autre (Yang et al., 2010). Ce polymorphisme d'un seul nucléotide est appelé SNP (pour *Single Nucleotide polymorphism*). La variation d'un seul nucléotide peut avoir des effets importants puisqu'elle peut modifier l'acide aminé, donc la protéine et les métabolites qui en découlent. De même, des locus (des régions d'ADN, pouvant coder un ou plusieurs gènes) variant d'un individu à l'autre peuvent avoir des effets importants sur le phénotype de l'individu. Ces études peuvent aussi être effectuées sur les brins d'ARN transcrits. On parlera alors de transcriptomique. La protéomique (respectivement la métabolomique) qui étudie l'abondance de protéines (resp. de métabolites) peut également permettre d'expliquer des différences entre individus d'une même espèce, qui ne seraient pas expliquées par l'étude des gènes ou des transcripts. En effet, il a été montré que

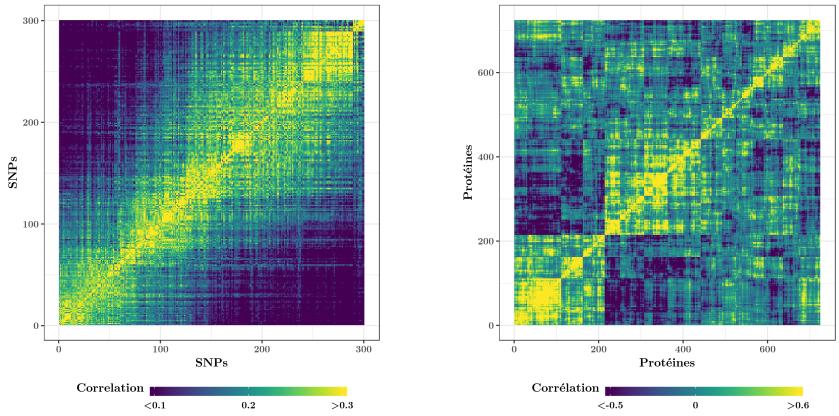

FIGURE 1.1 – Matrice de corrélation de SNPs le long du génome (Wittenburg et al., 2016) à gauche et de protéines à droite.

des variations dans l'expression des gènes n'entraînent pas nécessairement des variations proportionnelles dans l'abondance des métabolites (Riekeberg & Powers, 2017). Plus généralement, pour expliquer un phénomène d'intérêt, les différents types de jeu de données « -omiques » (génomique, transcriptomique, protéomique, métabolomique) ont chacun des avantages et inconvénients qui sont détaillés dans la table 1 de Karahalil (2016). C'est pourquoi dans la suite nous proposerons des méthodes générales permettant d'étudier les liens entre des marqueurs des données « -omiques » (les gènes, les protéines, les métabolites, les SNPs, les locus, ...) et certains traits phénotypiques.

Pour proposer de telles méthodes, il est important de bien comprendre les liens qui existent entre les marqueurs eux-mêmes et qui peuvent être dus au type de données considérées. Par exemple, durant la méiose, des chromosomes d'une même paire échangent un fragment de leur chromatine. On appelle cela l'entrecroisement chromosomique. Ainsi, plus des SNPs ou des gènes sont proches le long du brin d'ADN, plus ils vont être sur le même fragment de chromatine lors d'entrecroisements chromosomiques. La dépendance de deux SNPs ou gènes sera donc d'autant plus grande qu'ils sont physiquement proches (Gianola et al., 2003; Wittenburg et al., 2016; Chiquet et al., 2016). Cette dépendance le long du génome est représentée dans la figure 1.1 par la matrice de corrélation de données de SNPs, ordonnés le long du génome, de 106 vaches Holstein-Friesian disponible dans Wittenburg et al. (2016). De même les gènes ou les protéines impliqués dans une même voie métabolique sont souvent co-exprimés (Stuart et al., 2003). Une telle co-expression est représentée dans la figure 1.1 sur un cas de protéomique décrit plus en détails dans la section 1.3 et dans le chapitre 4.

L'objectif de cette thèse est de proposer des méthodes permettant de sélectionner un nombre restreint de marqueurs liés à un ou plusieurs traits d'intérêt tout en

prenant en compte la dépendance qui existe entre les marqueurs. Pour cela nous disposons de jeux de données sur lesquels on mesure q marqueurs sur n échantillons. Le nombre de marqueurs à mesurer peut être très grand (jusqu'à plusieurs dizaines de milliers) et les mesurer sur un échantillon peut être coûteux. Il est alors fréquent d'avoir un nombre de marqueurs à étudier beaucoup plus élevé que le nombre d'échantillons sur lesquels ils sont mesurés. Ainsi, nous proposerons dans cette thèse des méthodes adaptées à de tels cas aussi bien pour étudier la dépendance entre les marqueurs que pour étudier la présence ou l'absence de lien avec un phénomène d'intérêt.

1.2 Sélection de variables dans le modèle linéaire multivarié

Pour expliquer les liens entre différents marqueurs et différents traits phénotypiques on modélise les valeurs des marqueurs comme une combinaison linéaire du ou des traits phénotypiques. Supposons que nous cherchions les liens existant entre q marqueurs et p traits phénotypiques et que pour n échantillons nous ayons la valeur de ces q marqueurs et de ces p traits phénotypiques. Le mot « valeur » est ici utilisé pour généraliser une abondance, une quantité ou simplement une valeur nulle ou égale à 1 symbolisant l'appartenance ou non d'un échantillon à un groupe. Notons $\boldsymbol{X}_{i,k}$ la valeur du trait phénotypique k de l'échantillon i et $\boldsymbol{Y}_{i,j}$ la valeur du marqueur j de l'échantillon i. Si le trait phénotypique est catégoriel $\boldsymbol{X}_{i,k}$ vaut 1 si l'échantillon i est de la catégorie k et 0 sinon. On cherche alors à écrire $\boldsymbol{Y}_{i,j}$ par

$$\boldsymbol{Y}_{i,j} = \sum_{k=1}^{p} \boldsymbol{X}_{i,k}\boldsymbol{B}_{k,j} + \boldsymbol{E}_{i,j}, \forall i \in [\![1, n]\!], \forall j \in [\![1, q]\!], \tag{1.1}$$

où $\boldsymbol{B}_{k,j}$ est un coefficient associant le marqueur j au trait phénotypique k et $\boldsymbol{E}_{i,j}$ est un terme d'erreur aléatoire. En considérant les matrices $\boldsymbol{Y}$, $\boldsymbol{X}$, $\boldsymbol{B}$ et $\boldsymbol{E}$ qui sont telles que pour une matrice $\boldsymbol{A}$, $\boldsymbol{A}_{i,j}$ représente l'élément i de la colonne j de $\boldsymbol{A}$, le modèle (1.1) peut s'écrire comme suit :

$$\boldsymbol{Y} = \boldsymbol{X}\boldsymbol{B} + \boldsymbol{E}. \tag{1.2}$$

Dans ce modèle, $\boldsymbol{Y}$ est une matrice aléatoire de réponses de taille $n \times q$, $\boldsymbol{X}$ une matrice de design de taille $n \times p$ contenant les caractéristiques du phénomène d'intérêt, $\boldsymbol{B}$ une matrice de taille $p \times q$ contenant les valeurs des coefficients liant les réponses et le phénomène d'intérêt et $\boldsymbol{E}$ une matrice aléatoire d'erreur de taille $n \times q$. Au cours de cette thèse nous ferons l'hypothèse que les lignes de $\boldsymbol{E}$ sont indépendantes

et identiquement distribuées mais que ses colonnes ne le sont, en général, pas. Ainsi,

$$\forall i \in [\![1, n]\!],\ (E_{i,1}, \dots, E_{i,q}) \overset{iid}{\sim} \mathcal{N}(0, \boldsymbol{\Sigma}), \tag{1.3}$$

où $\mathcal{N}(0, \boldsymbol{\Sigma})$ désigne la loi d'un vecteur gaussien d'espérance nulle et de matrice de covariance $\boldsymbol{\Sigma}$.

Pour diminuer radicalement le nombre d'associations potentielles entre des marqueurs et des traits phénotypiques on s'intéresse à la mise en place de méthodes de sélection de variables dans le modèle linéaire général (1.2), voir par exemple Mardia et al. (1980). Faire de la sélection de variables dans ce modèle revient à proposer $\widehat{\boldsymbol{B}}$ un estimateur parcimonieux de $\boldsymbol{B}$. Un tel estimateur permet de mettre en avant des potentiels liens pertinents entre les réponses et les variables explicatives. En effet, un coefficient $\widehat{\boldsymbol{B}}_{i,j}$ positif (resp. négatif) indique un potentiel lien positif (resp. négatif) entre la caractéristique i du phénomène d'intérêt et le marqueur j. À l'inverse, un coefficient nul indique une potentielle absence de lien. On recherche donc un estimateur de $\boldsymbol{B}$ qui a des coefficients du même signe (positif, négatif ou nul) que $\boldsymbol{B}$. Pour cela, on définit la **consistance en signe** d'un estimateur $\widehat{C}$ de C par

$$\mathbb{P}\left(\operatorname{sign}(\widehat{C}) = \operatorname{sign}(C)\right) \to 1, \text{ lorsque } n \to \infty, \text{ où } \operatorname{sign}(x) = \begin{cases} 1 & \text{si } x > 0 \\ -1 & \text{si } x < 0 \\ 0 & \text{si } x = 0 \end{cases}.$$

Si $\widehat{\boldsymbol{B}}$ est un estimateur de $\boldsymbol{B}$ qui vérifie cette propriété, cela indique qu'avec une probabilité tendant vers un les valeurs positives (resp. négatives, resp. nulles) de $\widehat{\boldsymbol{B}}$ sont bien des valeurs positives (resp. négatives, resp. nulles) de $\boldsymbol{B}$. Une valeur non nulle de $\widehat{\boldsymbol{B}}$ indique donc une association réelle entre une variable réponse et une variable explicative. Une telle approche permet donc de diminuer drastiquement le nombre de validations expérimentales à effectuer et ce d'autant plus que $\widehat{\boldsymbol{B}}$ est parcimonieux.

1.2.1 État de l'art

Une approche possible pour sélectionner des variables dans le modèle linéaire multivarié (1.2) est de le faire indépendamment dans les q modèles linéaires univariés :

$$\boldsymbol{Y}_{\bullet,r} = \boldsymbol{X}\boldsymbol{B}_{\bullet,r} + \boldsymbol{E}_{\bullet,r},\ \forall r \in [\![1, q]\!], \tag{1.4}$$

où $\boldsymbol{Y}_{\bullet,r}$, $\boldsymbol{B}_{\bullet,r}$ et $\boldsymbol{E}_{\bullet,r}$ sont respectivement les r^{e} colonnes de $\boldsymbol{Y}$, $\boldsymbol{B}$ et $\boldsymbol{E}$ définies dans le modèle (1.2). Pour simplifier les notations on notera le modèle (1.4) :

$$y = \boldsymbol{X}b + e, \tag{1.5}$$

où $y \in \mathbb{R}^n$, $b \in \mathbb{R}^p$ et $e \in \mathbb{R}^n$. Faire de la sélection de variables dans ce modèle revient à proposer $\hat{b}$ un estimateur parcimonieux de b. Les variables ayant un coefficient nul ne sont pas sélectionnées alors que celles qui ont un coefficient non nul le sont. Il reste à savoir comment choisir les composantes de $\hat{b}$ qui seront mises à 0. Les coefficients du vecteur b peuvent être estimés par maximum de vraisemblance (voir par exemple Mardia et al., 1980; Draper & Smith, 1998), cependant cette méthode ne mène pas à une solution parcimonieuse. Différentes méthodes, pour effectuer de la sélection de variables dans le modèle (1.2) sont détaillées dans les travaux suivants : Thompson (1978), Zheng & Loh (1995), Chen et al. (2014), Heinze et al. (2018), Desboulets (2018). Certaines consistent à tester la nullité d'un ou plusieurs coefficients, ou à comparer deux modèles emboîtés à l'aide de tests de Student ou de Fisher (voir Mardia et al., 1980; Draper & Smith, 1998, par exemple). D'autres méthodes permettent de faire un compromis entre la diminution de l'erreur quadratique et le nombre de valeurs non nulles de $\hat{b}$. Les méthodes les plus fréquemment utilisées sont le critère AIC (Akaike, 1970), le C_p de Mallows (Mallows, 1973), le critère BIC (Schwarz et al., 1978) ou plus généralement le critère GIC (Nishii, 1984). Les méthodes de régression pénalisées consistent, quant à elles, à ajouter une pénalité sur les paramètres afin de sélectionner des variables. Une approche très fréquemment utilisée est la méthode Lasso introduite par Tibshirani (1996) qui propose l'estimateur suivant :

$$\widehat{b}(\lambda) = \text{Argmin}_{b\in\mathbb{R}^p} \left\{ \|y - \boldsymbol{X}b\|_2^2 + \lambda\|b\|_1 \right\}, \tag{1.6}$$

où $\|x\|_2^2$ est la norme ℓ_2 de $x = (x_1, \ldots, x_n)$, définie par $\|x\|_2^2 = \sum_{i=1}^n x_i^2$ et où $\|b\|_1$ est la norme ℓ_1 de $b = (b_1, \ldots, b_p)$, définie par $\|b\|_1 = \sum_{i=1}^p |b_i|$. Cette pénalité mène à une solution parcimonieuse. Zhao & Yu (2006) ont établi la consistance en signe de l'estimateur $\widehat{b}(\lambda)$:

Theorème 1.1. *Soit y un vecteur de $\mathbb{R}^n$ vérifiant (1.5). Supposons aussi qu'il existe des constantes positives M_1, M_2, M_3 et des nombres positifs c_1, c_2 tels que $0 < c_1 < c_2 < 1$ satisfaisant :*

(i) $\forall j \in [\![1, p]\!]$, $\frac{1}{n}(\boldsymbol{X}_{\bullet,j})^\intercal \boldsymbol{X}_{\bullet,j} \leq M_1$, *où $A^\intercal$ est la transposée de la matrice A et où $\boldsymbol{X}_{\bullet,j}$ est la colonne j de $\boldsymbol{X}$,*

(ii) $\alpha^\intercal \left(\frac{1}{n}(\boldsymbol{X}^\intercal\boldsymbol{X})_{J,J}\right)\alpha \geq M_2$, *pour tout $\alpha \in \mathbb{R}^{|J|}$ tel que $||\alpha||_2^2 = 1$, où $J = \{i, b_i \neq 0\}$, et $(\boldsymbol{X}^\intercal\boldsymbol{X})_{J,J}$ est la matrice $\boldsymbol{X}^\intercal\boldsymbol{X}$ restreinte aux lignes et colonnes dans J et $|J|$ est le cardinal de J,*

(iii) $|J| = O(n^{c_1})$,

(iv) $n^{\frac{1-c_2}{2}} \min_{j\in J} |b_j| \geq M_3$.

Supposons également qu'il existe une constante c_3 telle que $0 \leq c_3 < c_2 - c_1$ et $p = O(e^{n^{c_3}})$ et que la condition d'irreprésentabilité suivante soit satisfaite : il existe un vecteur positif et constant η tel que :

$$\left|(\boldsymbol{X}^\intercal\boldsymbol{X})_{J^c,J}((\boldsymbol{X}^\intercal\boldsymbol{X})_{J,J})^{-1}\,\text{sign}(b_J)\right| \leq \mathbf{1} - \eta, \tag{IC}$$

où l'inégalité est à comprendre composante par composante, $\mathbf{1}$ *est un vecteur de 1 de taille* $(p - |J|)$ *et* J^c *est le complémentaire de* J *dans* $[\![1, p]\!]$.

Alors on a, pour tout $\lambda = O(n^{\frac{1+c_4}{2}})$*, où* $c_3 < c_4 < c_2 - c_1$ *:*

$$\mathbb{P}\left(\text{sign}(\widehat{b}(\lambda)) = \text{sign}(b)\right) = 1 - o(e^{-n^{c_3}}) \to 1, \quad \textit{lorsque } n \to \infty,$$

où $\widehat{b}(\lambda)$ *est défini par* (1.6).

La condition (i) est satisfaite, par exemple, lorsque les colonnes de la matrice $\boldsymbol{X}$ sont centrées réduites. En forçant les valeurs singulières de la matrice de design restreinte aux variables pertinentes à ne pas être trop faibles devant n, la condition (ii) indique que la matrice $(\boldsymbol{X}^\intercal \boldsymbol{X})_{J,J}/n$ doit être inversible. Les conditions (iii) et (iv) indiquent respectivement que les valeurs non nulles de b ne doivent être ni trop nombreuses ni trop faibles par rapport au nombre n d'échantillons. Si ces différentes conditions sont satisfaites au vu des données on peut appliquer indépendamment le modèle (1.5) à chaque colonne de $\boldsymbol{Y}$. Cependant, cette méthode ne prend pas en compte la dépendance qui peut exister entre les différentes variables réponses. Nous nous intéresserons ici à des méthodes qui cherchent les positions non nulles de la matrice $\boldsymbol{B}$ définie dans le modèle (1.2), en prenant en compte cette dépendance. Pour cela, dans le cas gaussien, ces méthodes minimisent l'opposé de la log-vraisemblance du modèle et donc la fonction

$$\ell(\boldsymbol{B}, \boldsymbol{\Omega}) = \text{tr}\left(\frac{1}{\text{n}}(\boldsymbol{Y} - \boldsymbol{X}\boldsymbol{B})\boldsymbol{\Omega}(\boldsymbol{Y} - \boldsymbol{X}\boldsymbol{B})^\intercal\right) - \log(|\boldsymbol{\Omega}|), \tag{1.7}$$

où tr (A) désigne la trace de A, $\boldsymbol{\Omega} = \boldsymbol{\Sigma}^{-1}$ est la matrice de précision et $|\boldsymbol{\Omega}|$ est son déterminant. Mardia et al. (1980) montrent que l'estimateur de $\boldsymbol{B}$ qui minimise ℓ à $\boldsymbol{\Omega}$ fixée est indépendant de $\boldsymbol{\Sigma}$ et est donc le même que celui qui minimise l'erreur quadratique pour chaque colonne de $\boldsymbol{Y}$ indépendamment. Cependant cette méthode ne fait pas de sélection de variables c'est pourquoi on ajoute à la fonction (1.7) une contrainte induisant de la parcimonie sur les coefficients de $\boldsymbol{B}$. Les estimateurs ainsi obtenus ne sont plus indépendants de $\boldsymbol{\Sigma}$.

Rothman et al. (2010) proposent d'estimer à la fois $\boldsymbol{B}$ et $\boldsymbol{\Omega}$ de manière parcimonieuse. Pour cela ils minimisent une fonction de coût qui ajoute à ℓ deux pénalités : une sur la norme ℓ_1 des valeurs de $\boldsymbol{B}$ et une sur la norme ℓ_1 des valeurs extra-diagonales de $\boldsymbol{\Omega}$. Pour résoudre ce problème ils proposent une méthode itérative. Dans une première étape ils estiment d'abord $\boldsymbol{B}$ à $\boldsymbol{\Omega}$ fixé. Le problème devient alors un problème Lasso classique qui est résolu en utilisant un algorithme de descente par coordonnée. Dans une seconde étape ils estiment $\boldsymbol{\Omega}$ à $\boldsymbol{B}$ fixé en utilisant l'estimateur du *graphical-Lasso* proposé par Banerjee et al. (2008b) obtenu à l'aide de l'algorithme de Friedman et al. (2008).

Cette méthode a été étendue par Lee & Liu (2012) qui proposent deux autres

approches pour l'estimation de $\boldsymbol{B}$. La première consiste à estimer $\boldsymbol{\Omega}$ dans un premier temps puis à estimer $\boldsymbol{B}$ à $\boldsymbol{\Omega}$ fixé. Pour l'estimation de $\boldsymbol{\Omega}$ ils considèrent tout d'abord la matrice $\boldsymbol{X}$ comme aléatoire et donc les variables aléatoires

$$Z_i = (\boldsymbol{Y}_{i,\bullet}^\intercal, \boldsymbol{X}_{i,\bullet}^\intercal) \overset{iid}{\sim} \mathcal{N}(0, \boldsymbol{\Sigma_Z}), \forall i \in [\![1, n]\!],$$

où

$$\boldsymbol{\Sigma}_Z = \begin{pmatrix} \boldsymbol{\Sigma}_{y,y} & \boldsymbol{\Sigma}_{y,x} \\ \boldsymbol{\Sigma}_{x,y} & \boldsymbol{\Sigma}_{x,x} \end{pmatrix}.$$

En remarquant que la matrice de covariance $\boldsymbol{\Sigma}$ est la matrice de covariance des $\boldsymbol{Y}_{i,\bullet}$ conditionnellement aux $\boldsymbol{X}_{i,\bullet}$ on obtient

$$\boldsymbol{\Sigma} = \boldsymbol{\Sigma}_{y,y} - \boldsymbol{\Sigma}_{y,x}\boldsymbol{\Sigma}_{x,x}^{-1}\boldsymbol{\Sigma}_{y,x}.$$

Ainsi, estimer $\boldsymbol{\Sigma}_Z$ permet d'obtenir $\boldsymbol{\Sigma}$. Plus précisément, estimer $\boldsymbol{\Sigma}_Z^{-1}$ à l'aide du *graphical-lasso* permet d'obtenir $\boldsymbol{\Omega}$. La seconde méthode, proche de celle de Rothman et al. (2010), estime de manière itérative $\boldsymbol{\Omega}$ et $\boldsymbol{B}$. Cette méthode a une fonction de coût qui n'est pas convexe en $(\boldsymbol{B}, \boldsymbol{\Omega})$ et peut être instable lorsque $q \geq n$. Ils ont cependant mis en évidence à l'aide de simulations numériques que cette méthode peut être plus performante que celle où on estime $\boldsymbol{\Omega}$ dans un premier temps. Ils montrent ensuite que sous certaines conditions on retrouve la consistance en signe de ces estimateurs lorsque q est fixé.

Zhang et al. (2017) proposent une méthode qui tout comme la première méthode de Lee & Liu (2012) estime $\boldsymbol{\Omega}$ dans un premier temps avant d'estimer $\boldsymbol{B}$. Pour estimer $\boldsymbol{\Omega}$ ils utilisent la méthode CLIME (*constrained* l_1 *minimization for inverse matrix estimation*) proposée par Cai et al. (2011). Cette méthode définit

$$\widehat{\boldsymbol{\Omega}}_1 = \text{Argmin}_{\boldsymbol{\Omega}} ||\boldsymbol{\Omega}||_1, \text{avec } \boldsymbol{\Omega}\boldsymbol{S} - \boldsymbol{I}_q \leq \lambda, \tag{1.8}$$

où λ est un paramètre à fixer et $\boldsymbol{S}$ la matrice de covariance empirique. Pour assurer la symétrie, l'estimateur $\widehat{\boldsymbol{\Omega}} = (w_{i,j})_{1 \leq i,j \leq q}$ final est ensuite obtenu en prenant $w_{i,j} = w_{j,i} = \min(w_{i,j}, w_{j,i})$. Ils utilisent ensuite l'algorithme du *Weighted Square root Lasso* (Belloni et al., 2011) pour estimer $\boldsymbol{B}$ à $\boldsymbol{\Sigma}$ fixé. Ils établissent ensuite une inégalité de type oracle pour leur estimateur. Enfin, ils comparent par simulation les résultats obtenus par leur méthode et ceux obtenus à l'aide du *Square root Lasso* et du Lasso. Dans leurs différents scénarios de simulations le taux de faux positifs obtenu par leur méthode est plus faible que ceux obtenus avec les autres méthodes. Cependant, leur méthode ne dépend pas d'un unique paramètre mais de deux paramètres à calibrer en pratique.

Enfin Molstad et al. (2018) proposent une méthode pour estimer $\boldsymbol{B}$ sous l'hypothèse que plus deux réponses sont expliquées de la même façon par les prédicteurs plus elles vont être corrélées. Ainsi, ils proposent de remplacer $\boldsymbol{\Omega}$ dans ℓ par $[\boldsymbol{B}'\boldsymbol{B} +$

$\tau \boldsymbol{I}_q]^{-1}$, ou $\tau \geq 0$. Ils n'ont alors plus que la matrice $\boldsymbol{B}$ à estimer. Ils proposent ensuite un algorithme permettant de minimiser $\ell(\boldsymbol{B})$ sous différentes contraintes sur $\boldsymbol{B}$. Cet estimateur n'a pour l'instant pas été étudié théoriquement mais a été validé par des expériences numériques. Cependant cet estimateur ne prend pas en compte des dépendances dues à des facteurs qui ne sont pas dans $\boldsymbol{X}$.

1.2.2 Contributions des chapitres 2 et 3

Production scientifique

Cette section résume les publications suivantes :

- **M. Perrot-Dockès, C. Lévy-Leduc, L. Sansonnet, J. Chiquet.** Variable selection in multivariate linear models with high-dimensional covariance matrix estimation. *Journal of Multivariate Analysis, 166:78 – 97, 2018.*
- **M. Perrot-Dockès, C. Lévy-Leduc, J. Chiquet, L. Sansonnet, M. Brégère, M.-P. Étienne, S. Robin, G. Genta-Jouve.** A variable selection approach in the multivariate linear model: An application to LC-MS metabolomics data. *Statistical Applications in Genetics and Molecular Biology, 17:5, 2018.*

La méthode décrite est implémentée dans le paquet R :

- **M. Perrot-Dockès, C. Lévy-Leduc, J. Chiquet.** MultiVarSel. *R package version 1.1.2, 2018.*

Nous proposons d'estimer tout d'abord la matrice de covariance $\boldsymbol{\Sigma}$ puis de l'utiliser pour estimer $\boldsymbol{B}$. Notre méthode se rapproche donc de la première proposition de Lee & Liu (2012) avec cependant deux différences majeures. La première vient de notre cadre asymptotique. En effet, Lee & Liu (2012) étudient les propriétés de leur estimateur lorsque p peut tendre vers l'infini et que q est fixé. Nous autorisons au contraire q à tendre vers l'infini même potentiellement à une puissance de n, p étant fixé. En effet, dans notre cadre le nombre de réponses q peut être beaucoup plus grand que le nombre d'échantillons n. La seconde différence vient de notre estimation de $\boldsymbol{\Omega}$. Dans le chapitre 2 nous proposons des conditions que doit vérifier l'estimateur de la matrice de covariance pour conserver la consistance en signe de notre estimateur. Dans le chapitre 3 nous proposons une nouvelle méthode d'estimation de la matrice de covariance $\boldsymbol{\Sigma}$ lorsque celle-ci est supposée être une matrice de Toeplitz symétrique, ce qui revient à supposer que chaque ligne de la matrice $\boldsymbol{E}$ est une réalisation d'un processus stationnaire au second ordre (voir Brockwell & Davis, 1990). Ces deux chapitres sont résumés ci-après.

Contributions du chapitre 2

Afin de retirer toute dépendance au sein des colonnes de $\boldsymbol{E}$, notre méthode consiste tout d'abord à « blanchir », c'est-à-dire à « décorréler », les données en appliquant la transformation suivante :

$$\boldsymbol{Y}\,\boldsymbol{\Sigma}^{-1/2} = \boldsymbol{X}\boldsymbol{B}\,\boldsymbol{\Sigma}^{-1/2} + \boldsymbol{E}\,\boldsymbol{\Sigma}^{-1/2}. \tag{1.9}$$

En utilisant l'opérateur de vectorisation (vec) ce problème peut se réécrire comme suit :

$$\mathcal{Y} = \mathcal{X}\mathcal{B} + \mathcal{E}, \tag{1.10}$$

avec $\mathcal{Y} = \mathrm{vec}(\boldsymbol{Y}\boldsymbol{\Sigma}^{-1/2})$, $\mathcal{X} = (\boldsymbol{\Sigma}^{-1/2})^\intercal \otimes \boldsymbol{X}$, $\mathcal{B} = \mathrm{vec}(\boldsymbol{B})$ et $\mathcal{E} = \mathrm{vec}(\boldsymbol{E}\boldsymbol{\Sigma}^{-1/2})$ où $\otimes$ désigne le produit de Kronecker. Puisque $\mathcal{B} = \mathrm{vec}(\boldsymbol{B})$, estimer $\mathcal{B}$ dans le modèle (1.10) revient à estimer $\boldsymbol{B}$ dans le modèle (1.2). En utilisant le critère Lasso pour estimer $\mathcal{B}$ on obtient l'estimateur :

$$\widehat{\mathcal{B}}(\lambda) = \mathrm{Argmin}_{\mathcal{B}} \left\{ \|\mathcal{Y} - \mathcal{X}\mathcal{B}\|_2^2 + \lambda \|\mathcal{B}\|_1 \right\}. \tag{1.11}$$

En s'inspirant de Zhao & Yu (2006) nous avons établi la consistance en signe de l'estimateur $\widehat{\mathcal{B}}$ (voir théorème 2.1 du chapitre 2).

En pratique $\boldsymbol{\Sigma}$ est inconnue. Nous proposons donc de l'estimer dans une étape préalable et appelons cet estimateur $\widehat{\boldsymbol{\Sigma}}$. En remplaçant $\boldsymbol{\Sigma}^{-1/2}$ par $\widehat{\boldsymbol{\Sigma}}^{-1/2}$ dans le modèle (1.9) on obtient :

$$\boldsymbol{Y}\,\widehat{\boldsymbol{\Sigma}}^{-1/2} = \boldsymbol{X}\boldsymbol{B}\,\widehat{\boldsymbol{\Sigma}}^{-1/2} + \boldsymbol{E}\,\widehat{\boldsymbol{\Sigma}}^{-1/2}. \tag{1.12}$$

Une fois encore nous réécrivons le modèle (1.12) comme

$$\widetilde{\mathcal{Y}} = \widetilde{\mathcal{X}}\mathcal{B} + \widetilde{\mathcal{E}}, \tag{1.13}$$

où $\widetilde{\mathcal{Y}} = \mathrm{vec}(\boldsymbol{Y}\widehat{\boldsymbol{\Sigma}}^{-1/2})$, $\mathcal{X} = (\widehat{\boldsymbol{\Sigma}}^{-1/2})^\intercal \otimes \boldsymbol{X}$, $\mathcal{B} = \mathrm{vec}(\boldsymbol{B})$ et $\widetilde{\mathcal{E}} = \mathrm{vec}(\boldsymbol{E}\widehat{\boldsymbol{\Sigma}}^{-1/2})$. En utilisant le critère Lasso pour estimer $\mathcal{B}$ on obtient l'estimateur :

$$\widetilde{\mathcal{B}}(\lambda) = \mathrm{Argmin}_{\mathcal{B}} \left\{ \|\widetilde{\mathcal{Y}} - \widetilde{\mathcal{X}}\mathcal{B}\|_2^2 + \lambda \|\mathcal{B}\|_1 \right\}. \tag{1.14}$$

Le théorème ci-dessous (théorème 2.5 du chapitre 2) établit la consistance en signe de $\widetilde{\mathcal{B}}(\lambda)$ sous certaines conditions sur $\boldsymbol{X}$, $\boldsymbol{B}$, $\boldsymbol{\Sigma}$, $\boldsymbol{\Omega}$ et $\widehat{\boldsymbol{\Sigma}}$.

Theorème 1.2. *Soit $\boldsymbol{Y}$ vérifiant le modèle (1.2) sous l'hypothèse (1.3). Supposons que la condition d'irreprésentabilité (IC) soit vérifiée. Supposons de plus qu'il existe des constantes positives M_4, M_5, M_6 et M_7 et c_1, c_2 avec $0 < c_1 + c_2 < 1/2$, telles que*

(i) $\|(\boldsymbol{X}^\intercal\boldsymbol{X})/n\|_\infty \leq M_4$,

(ii) $\lambda_{\min}((\boldsymbol{X}^\intercal \boldsymbol{X})/n) \geq M_5$,

(iii) $|J| = O(q^{c_1})$, *où* $J = \{i,\ tel\ que\ \mathcal{B}_i \neq 0\}$ *et* $|J|$ *est le cardinal de* J,

(iv) $q^{c_2} \min_{j\in J} |\mathcal{B}_j| \geq M_3$.

(v) $\lambda_{\max}(\Sigma^{-1}) \leq M_6$,

(vi) $\lambda_{\min}(\Sigma^{-1}) \geq M_7$.

Supposons aussi que lorsque n *tend vers l'infini, on ait :*

(vii) $\|\Sigma^{-1} - \widehat{\Sigma}^{-1}\|_\infty = O_P((nq)^{-1/2})$,

(viii) $\rho(\Sigma - \widehat{\Sigma}) = O_P((nq)^{-1/2})$.

Ainsi, pour tout λ *tel que*

$$q = q_n = o\left(n^{\frac{1}{2(c_1+c_2)}}\right), \quad \frac{\lambda}{\sqrt{n}} \to \infty \ \text{ et } \ \frac{\lambda}{n} = o\left(q^{-(c_1+c_2)}\right), \text{ lorsque } n \to \infty,$$

on a

$$\mathbb{P}\left(\operatorname{sign}(\widetilde{\mathcal{B}}(\lambda)) = \operatorname{sign}(\mathcal{B})\right) \to 1, \ \text{ lorsque } n \to \infty,$$

où $\widetilde{\mathcal{B}}(\lambda)$ *est défini dans* (1.14). *Ici,* $\lambda_{\max}(A)$, $\lambda_{\min}(A)$, $\rho(A)$ *et* $\|A\|_\infty$ *sont respectivement la plus grande, la plus petite valeur propre, le rayon spectral et la norme infinie de* A.

Les conditions (i) à (iv) du Théorème 1.2 sont similaires aux conditions (i) à (iv) du Théorème 1.1. Les conditions (v) et (vi) du Théorème 1.2 indiquent que les valeurs propres de $\boldsymbol{\Sigma}$ et $\boldsymbol{\Omega}$ sont minorées par une constante strictement positive. Enfin, les conditions (vii) et (viii) indiquent que ni la norme infinie de l'erreur de prédiction de la matrice de précision ni le rayon spectral de l'erreur de prédiction de la matrice de covariance ne peuvent être trop grands.

Contributions du chapitre 3

Pour l'estimation de la matrice de covariance et le calcul de notre estimateur de $\boldsymbol{B}$ lorsque la matrice $\boldsymbol{\Sigma}$ est supposée Toeplitz symétrique, nous proposons une méthode en quatre étapes.

Première étape : Estimation de $\boldsymbol{E}$
On estime $\boldsymbol{E}$ par

$$\begin{aligned} \widehat{\boldsymbol{E}} &= \boldsymbol{Y} - \boldsymbol{X}(\boldsymbol{X}^\intercal \boldsymbol{X})^{-1}\boldsymbol{X}^\intercal \boldsymbol{Y} \\ &= (Id - P_{\boldsymbol{X}})\boldsymbol{Y} \\ &= P_{\boldsymbol{X}^\perp}\boldsymbol{Y}, \end{aligned} \tag{1.15}$$

où $P_{\boldsymbol{X}}$ (respectivement $P_{\boldsymbol{X}^\perp}$) est la matrice de projection orthogonale sur $\mathrm{Vect}(\boldsymbol{X})$ qui désigne le sous espace engendré par les colonnes de $\boldsymbol{X}$ (respectivement l'orthogonal de $\mathrm{Vect}(\boldsymbol{X})$). $\widehat{\boldsymbol{E}}$ est donc le projeté orthogonal de $\boldsymbol{Y}$ sur l'orthogonal de $\mathrm{Vect}(\boldsymbol{X})$.

Deuxième étape : Estimation de $\boldsymbol{\Sigma}$
Nous proposons plusieurs estimateurs pour $\boldsymbol{\Sigma}$ correspondant à différents modèles de processus stationnaires pour les lignes de $\boldsymbol{E}$.

Le modèle le plus simple est tel que chaque ligne de $\boldsymbol{E}$ est modélisée comme un **processus autorégressif d'ordre 1 (AR(1))**. Cela signifie que pour tout i de $[\![1, n]\!]$ et pour tout t de $[\![2, q]\!]$ $(E_{i,t})$ est tel que

$$E_{i,t} - \phi_1 E_{i,t-1} = W_{i,t}, \text{ avec } W_{i,t} \sim BB(0, \sigma^2),$$

où $|\phi_1| < 1$ et les $W_{i,t}$ sont des bruits blancs de variance σ^2 que l'on note $BB(0, \sigma^2)$. Lorsque $\sigma^2 = 1$ la matrice $\boldsymbol{\Omega}^{1/2}$ a la forme explicite suivante :

$$\boldsymbol{\Omega}^{1/2} = \begin{pmatrix} \sqrt{1-\phi_1^2} & -\phi_1 & 0 & \cdots & 0 \\ 0 & 1 & -\phi_1 & \cdots & 0 \\ 0 & 0 & \ddots & \ddots & \vdots \\ \vdots & \vdots & \ddots & \ddots & -\phi_1 \\ 0 & 0 & \cdots & 0 & 1 \end{pmatrix}. \tag{1.16}$$

Lorsque σ^2 n'est pas égale à 1, la matrice de covariance de chaque ligne de $\boldsymbol{E}\boldsymbol{\Omega}^{1/2}$ est égale à $\sigma^2 I_d$ et la matrice de corrélation est égale à la matrice identité. On obtient alors $\widehat{\boldsymbol{\Sigma}}^{-1/2}$ un estimateur de $\boldsymbol{\Sigma}^{-1/2}$ en remplaçant ϕ_1 par $\widehat{\phi_1}$ dans (1.16), où l'estimateur $\widehat{\phi_1}$ est obtenu à l'aide des équations de Yule Walker (décrites dans Walker, 1964) et de $\widehat{\boldsymbol{E}}$:

$$\widehat{\phi}_1 = \frac{\sum_{i=1}^{n}\sum_{\ell=2}^{q} \widehat{E}_{i,\ell}\widehat{E}_{i,\ell-1}}{\sum_{i=1}^{n}\sum_{\ell=1}^{q-1} \widehat{E}_{i,\ell}^2}.$$

Nous avons montré dans le chapitre 2 que la matrice $\widehat{\boldsymbol{\Sigma}}^{-1/2}$ ainsi obtenue vérifie les hypothèses du Théorème 1.2.

Des modèles un peu plus généraux sont tels que chaque ligne de $\boldsymbol{E}$ est modélisée comme un processus **autorégressif à moyenne mobile ARMA(p ,q)**. Dans ce cas pour tout i de $[\![1, n]\!]$ et pour tout t, $E_{i,t}$ est solution de :

$$E_{i,t} - \phi_1 E_{i,t-1} - \cdots - \phi_{\mathtt{p}} E_{i,t-\mathtt{p}} = W_{i,t} + \theta_1 W_{i,t-1} + \cdots + \theta_{\mathtt{q}} W_{i,t-\mathtt{q}},$$

avec $W_{i,t} \sim BB(0, \sigma^2)$ et où les ϕ_j et les θ_j sont des paramètres réels.

Dans le cas où la modélisation par un processus ARMA n'est pas appropriée, on peut modéliser chaque ligne de $\boldsymbol{E}$ comme un processus faiblement stationnaire général et estimer $\boldsymbol{\Sigma}$ comme une matrice Toeplitz symétrique c'est à dire

comme

$$\widehat{\mathbf{\Sigma}} = \begin{pmatrix} \widehat{\gamma}(0) & \widehat{\gamma}(1) & \cdots & \widehat{\gamma}(q-1) \\ \widehat{\gamma}(1) & \widehat{\gamma}(0) & \cdots & \widehat{\gamma}(q-2) \\ \vdots & & & \\ \widehat{\gamma}(q-1) & \widehat{\gamma}(q-2) & \cdots & \widehat{\gamma}(0) \end{pmatrix}, \tag{1.17}$$

où

$$\widehat{\gamma}(h) = \frac{1}{n}\sum_{i=1}^{n} \widehat{\gamma}_i(h),$$

et $\widehat{\gamma}_i(h)$ est l'estimateur de l'autocovariance $\gamma_i(h) = \mathbb{E}(E_{i,t}E_{i,t+h})$, pour tout t et tout h dans $\mathbb{Z}$.

Pour sélectionner l'estimateur le plus adapté, nous proposons dans le chapitre 3 un test statistique qui est une adaptation du test de Portmanteau lui même fondé sur le Théorème de Bartlett (Brockwell & Davis, 1991). Nous appelons $\widehat{\mathbf{\Sigma}}$ l'estimateur final de $\mathbf{\Sigma}$ ainsi obtenu.

Troisième étape : Blanchiment des données

Afin de décorréler au mieux les colonnes de $\boldsymbol{E}$ on multiplie l'équation (1.2) à droite par l'estimateur $\widehat{\mathbf{\Sigma}}^{-1/2}$ (voir l'équation (1.12)).

Quatrième étape : Sélection de variables

Le critère Lasso, décrit dans (1.14), est appliqué aux données ainsi transformées. Cependant cet estimateur dépend d'un paramètre λ qui permet de régler le niveau de parcimonie de $\widetilde{\mathcal{B}}(\lambda)$ et qui est inconnu. Pour sélectionner les coefficients de $\boldsymbol{B}$ qui sont non nuls on propose une adaptation de la *stability selection* (proposée par Meinshausen & Bühlmann, 2010) qui utilise la méthode de validation croisée.

Plus précisément nous choisissons un sous-échantillonnage aléatoire de taille $nq/2$ auquel nous appliquons les trois premières étapes décrites ci-dessus pour obtenir des données « blanchies ». Nous calculons sur ces données λ_{CV} qui minimise l'erreur de validation croisée dans le modèle (1.14). Nous obtenons ainsi $\widetilde{\mathcal{B}}(\lambda_{CV})$ qui est un estimateur parcimonieux de $\boldsymbol{B}$ et enregistrons les positions des valeurs non nulles de cet estimateur. Nous répétons ces étapes de sous-échantillonnage N fois. A l'issue de ces différentes étapes, nous avons accès au nombre de fois N_i où chaque composante $\widetilde{\mathcal{B}}_i$ de $\widetilde{\mathcal{B}}$ a été estimée comme étant non nulle. Nous conservons les composantes i dont la fréquence N_i/N est plus grande qu'un seuil donné. L'influence du choix de N et du seuil est étudié dans dans la section 3.3.3 du chapitre 3.

Notre méthode a été de plus comparée numériquement à des méthodes existantes et semble être la plus adaptée dans les scénarios étudiés. Le détail de ces comparaisons est disponible dans la section 3.3 du chapitre 3. Enfin, elle a été appliquée à une étude de métabolomique non ciblée dans la section 3.4 du chapitre 3.

1.3 Estimation de matrice de corrélation par blocs

1.3.1 État de l'art

La section 1.2.2 présente une méthode permettant de modéliser des marqueurs « -omiques » par des variables explicatives (traits phénotypiques par exemple) en prenant en compte la covariance qui existe entre ces marqueurs. La méthode proposée pour l'estimation de la matrice de covariance fait l'hypothèse qu'il s'agit d'une matrice de Toeplitz symétrique. Cependant, les marqueurs « -omiques » (gènes, protéines, métabolites, ...) sont souvent co-exprimés / co-accumulés par groupe (voir Section 1.1 et Stuart et al., 2003) ce qui se traduit par une structure de matrice de covariance par blocs.

Cette structure par blocs peut venir de facteurs sous-jacents tels que l'appartenance à une même voie métabolique ou la stimulation des marqueurs par les mêmes molécules. Plus généralement, les liens entre des variables peuvent être dus à des facteurs peu nombreux qui sont observables ou latents. Pour étudier de tels liens Fan et al. (2016) présentent des méthodes de modèle à facteurs qui décrivent chaque ligne i de $\boldsymbol{E}$, matrice de taille $n \times q$, comme :

$$\boldsymbol{E}_i = f_i \boldsymbol{B}_f + \boldsymbol{U}_i, \tag{1.18}$$

où $\boldsymbol{B}_f$ est une matrice de coefficients de taille $k \times q$ et pour tout i, f_i est un vecteur aléatoire de facteurs (observés ou latents) de taille k avec $k \leq \min(q, n)$ (les f_i sont supposés indépendants et identiquement distribués), les $\boldsymbol{U}_i$ sont des vecteurs aléatoires d'erreur de taille q, indépendants et identiquement distribués et sont supposés indépendant de $\boldsymbol{B}_f$. Sous cette hypothèse d'indépendance on a :

$$\boldsymbol{\Sigma} = \boldsymbol{B}_f^\intercal \mathrm{Cov}(f) \boldsymbol{B}_f + \boldsymbol{\Sigma}_u, \tag{1.19}$$

où $\boldsymbol{\Sigma}$ est la matrice de covariance des lignes de $\boldsymbol{E}$, $\boldsymbol{\Sigma}_u$ est la matrice de covariance des U_i et $\mathrm{Cov}(f)$ est la matrice de covariance des f_i. Le nombre de facteurs k étant plus petit que n, $\mathrm{Cov}(f)$ peut être estimée par la matrice de covariance empirique de f. Il ne reste plus qu'à estimer $\boldsymbol{\Sigma}_u$ par exemple de manière parcimonieuse, voire diagonale.

Lorsque les facteurs f sont des vecteurs gaussiens aléatoires d'espérance nulle et de variance l'identité Blum et al. (2016b) proposent une méthode qui estime $\boldsymbol{B}_f$ de manière parcimonieuse, amenant la matrice $\boldsymbol{\Sigma}$ à l'être aussi. Ils estiment $\boldsymbol{B}_f$ et $\boldsymbol{\Sigma}$ à l'aide d'un algorithme EM en ajoutant à l'étape de maximisation une pénalité sur la norme ℓ_1 des valeurs de $\boldsymbol{B}$ afin d'assurer la parcimonie. Cette étape peut aussi leur permettre d'estimer le nombre k de facteurs. En effet, ils appliquent leur algorithme avec un certain nombre k_{max} de facteurs, puis ne gardent ensuite que ceux qui ont au moins une composante non nulle.

Hosseini & Lee (2016) proposent aussi une méthode alliant parcimonie et modèle à facteurs permettant l'estimation de la matrice de précision, et non pas la matrice de covariance, comme une matrice parcimonieuse par blocs avec des blocs pouvant se chevaucher. Une telle matrice est très utile par exemple pour faire de la sélection de variables dans le modèle linéaire général comme on l'a vu dans la section 1.2.2.

Enfin, Perthame et al. (2016) allient sélection de variables et modèles à facteurs afin de prendre en compte la dépendance au sein des variables réponses lorsque $\boldsymbol{X}$ est une matrice d'ANOVA. Pour cela ils proposent le modèle :

$$\boldsymbol{Y} = \boldsymbol{X}\boldsymbol{B} + \boldsymbol{Z}\boldsymbol{B}_Z^\intercal + \boldsymbol{\mathcal{E}}, \tag{1.20}$$

où $\boldsymbol{B}$ et $\boldsymbol{B}_Z$ sont des matrices de coefficients de taille $p \times q$ et $k \times q$, $\boldsymbol{Z}$ est une matrice de variables latentes de taille $n \times k$ avec $k \leq q$ et où

$$\forall i \in [\![1, n]\!],\ (\boldsymbol{\mathcal{E}}_{i,1}, \ldots, \boldsymbol{\mathcal{E}}_{i,q}) \overset{iid}{\sim} \mathcal{N}(0, \boldsymbol{\Sigma}_u), \tag{1.21}$$

avec $\boldsymbol{\Sigma}_u$ une matrice diagonale. Chaque ligne de $\boldsymbol{Z}$ est supposée distribuée selon une loi normale d'espérance nulle et de matrice de covariance l'identité. Ils proposent une méthode itérative qui alterne l'estimation de $\boldsymbol{B}$, $\boldsymbol{B}_Z$ et $\boldsymbol{\Sigma}_u$ et l'inférence de $\boldsymbol{Z}$. Avec la transformation $\boldsymbol{Y} - \boldsymbol{Z}\boldsymbol{B}_Z^\intercal$ ils retirent la dépendance existant entre les différentes variables réponses et appliquent des méthodes de sélection de variables comme celles décrites dans la section 1.2.1.

1.3.2 Contributions du chapitre 4

Production scientifique

Cette section résume la publication :

- **M. Perrot-Dockès, C. Lévy-Leduc, L. Rajjou** Estimation of large block structured covariance matrices: Application to "multi-omic" approaches to study seed quality
  Soumise

La méthode décrite est implémentée dans le paquet du logiciel R :

- **M. Perrot-Dockès, C. Lévy-Leduc** BlockCov. *R package version 0.1.1, 2019.*

Nous nous intéressons ici à l'estimation de matrice de corrélation par blocs parcimonieuse en utilisant un modèle à facteur et des méthodes d'estimation de matrices parcimonieuses.

Nous supposons que $\boldsymbol{\Sigma}$ s'écrit comme suit :

$$\boldsymbol{\Sigma} = \boldsymbol{Z}\boldsymbol{Z}^{\intercal} + \boldsymbol{D}, \tag{1.22}$$

où $\boldsymbol{Z}$ est une matrice de taille $q \times k$ avec $k \ll q$ et $\boldsymbol{D}$ est une matrice diagonale de telle sorte que tous les termes diagonaux de $\boldsymbol{\Sigma}$ soient égaux à 1. Des exemples, dans le cas de blocs diagonaux et le cas de blocs extra-diagonaux, sont donnés dans la figure 1.2.

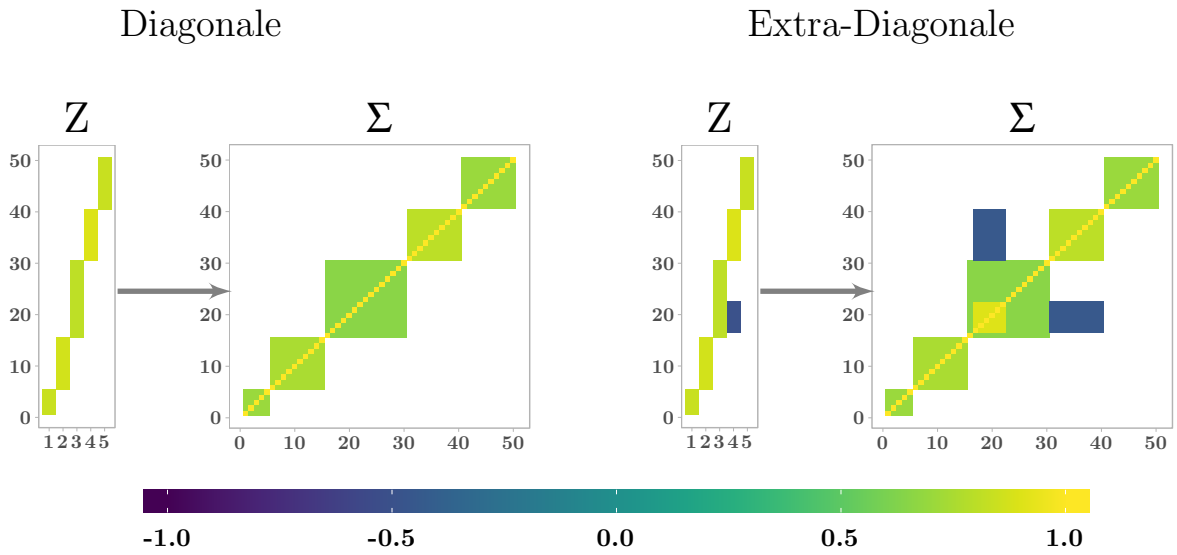

FIGURE 1.2 – Exemples de matrices $\boldsymbol{\Sigma}$ générées à partir de différentes matrices $\boldsymbol{Z}$.

Notre méthode comporte trois étapes :

Première étape : Approximation de rang faible d'une matrice contenant les coefficients extra-diagonaux de $\boldsymbol{\Sigma}$.
On définit la matrice $\boldsymbol{\Gamma}$ de taille $(q-1) \times (q-1)$ (voir figure 1.3) à partir de $\boldsymbol{\Sigma}$ telle que :

$$\boldsymbol{\Gamma}_{i,j} = \begin{cases} \boldsymbol{\Sigma}_{i,j+1}, & \text{si } 1 \leq i \leq j \leq q-1 \\ \boldsymbol{\Gamma}_{j,i}, & \text{si } \leq j < i \leq q-1 \end{cases} \tag{1.23}$$

Contrairement à $\boldsymbol{\Sigma}$ qui est de plein rang $\boldsymbol{\Gamma}$ a un rang $k \ll q$. De plus $\boldsymbol{\Gamma}$ contient tous les coefficients extra-diagonaux de $\boldsymbol{\Sigma}$. Comme tous les éléments diagonaux de $\boldsymbol{\Sigma}$ valent 1 estimer $\boldsymbol{\Gamma}$ suffit pour estimer $\boldsymbol{\Sigma}$. Le passage de $\boldsymbol{\Sigma}$ à $\boldsymbol{\Gamma}$ est présenté dans la figure 1.3. En pratique $\boldsymbol{\Sigma}$ est inconnue, on propose donc $\widetilde{\boldsymbol{\Gamma}}$ l'estimateur de $\boldsymbol{\Gamma}$ en remplaçant $\boldsymbol{\Sigma}$ par $\boldsymbol{R}$ la matrice de corrélation empirique dans l'équation (1.23). Nous approchons $\widetilde{\boldsymbol{\Gamma}}$ par $\widetilde{\boldsymbol{\Gamma}}^{(k)}$, son approximation de rang k, obtenue en utilisant sa décomposition en valeurs singulières. Le rang k est la plupart du temps inconnu. Deux méthodes pour l'estimation de k sont proposées et comparées dans le chapitre 4. La première utilise le critère de Cattell

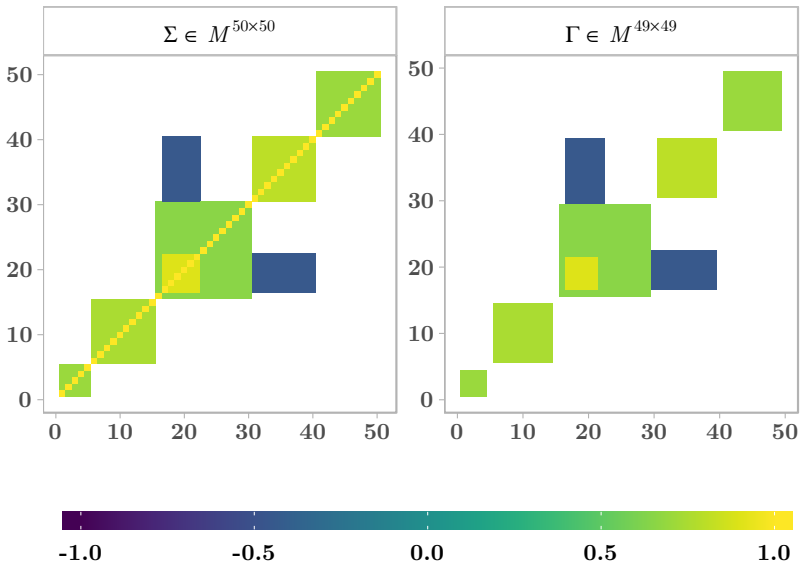

FIGURE 1.3 – Passage de $\mathbf{\Sigma}$ à $\mathbf{\Gamma}$

décrit dans Cattell (1966) et la seconde utilise la méthode de permutation PA proposée par Horn (1965) et étudiée théoriquement par Dobriban (2018). Ces deux méthodes sont comparées dans divers scénarios. L'estimateur ainsi obtenu est noté r.

Deuxième étape : Détecter les positions de valeurs non nulles de $\mathbf{\Gamma}$
Afin d'obtenir une matrice parcimonieuse on applique un critère lasso sur les valeurs de $\widetilde{\mathbf{\Gamma}}^{(r)}$, les coefficient gardés comme non nuls sont ensuite ré-estimés avec la méthode des moindres carrés. La matrice obtenue en appliquant le critère lasso avec un paramètre λ sera ci-après notée $\widetilde{\mathbf{\Gamma}}^{(r)}(\lambda)$. Deux méthodes pour l'estimation du seuil sont proposées et comparées dans le chapitre 4. La première est celle proposée par Bickel & Levina (2008) qui est fondée sur une méthode de validation croisée. La seconde méthode cherche une rupture dans la pente liant l'erreur $\|\widetilde{\mathbf{\Gamma}} - \widetilde{\mathbf{\Gamma}}^{(r)}(\lambda)\|_F$ au nombre de valeur non nulles dans $\widetilde{\mathbf{\Gamma}}^{(r)}(\lambda)$, où $\|\boldsymbol{M}\|_F$ désigne la norme de Frobenius de la matrice $\boldsymbol{M}$.
L'estimateur ainsi obtenu est noté : $\widehat{\mathbf{\Gamma}}^{(r)}$. On obtient alors $\widetilde{\mathbf{\Sigma}}$ un premier estimateur de $\mathbf{\Sigma}$ comme

$$\widetilde{\mathbf{\Sigma}}_{i,j} = \begin{cases} \widehat{\mathbf{\Gamma}}^{(r)}_{i,j-1} & \text{si } 1 \leq i < j \leq q \\ 1 & \text{si } 1 \leq i = j \leq q \\ \widetilde{\mathbf{\Sigma}}_{j,i} & \text{si } 1 \leq j < i \leq q \end{cases}$$

Troisième étape : Assurer la positivité de l'estimateur $\widehat{\mathbf{\Sigma}}$
On obtient $\widehat{\mathbf{\Sigma}}$ en appliquant à $\widetilde{\mathbf{\Sigma}}$ la méthode proposée par Higham (2002). Cette

méthode calcule $\widehat{\mathbf{\Sigma}}$ la matrice définie positive avec des coefficients diagonaux égaux à 1 la plus proche de $\widetilde{\mathbf{\Sigma}}$.

Nous proposons ensuite une méthode permettant d'obtenir $\widehat{\mathbf{\Omega}}^{-1/2}$, un estimateur de $\mathbf{\Omega}^{-1/2}$, à partir de $\widehat{\mathbf{\Sigma}}$.

Cette stratégie a été étudiée à l'aide d'expériences numériques. Différents scénarios de simulations faisant varier la forme de $\mathbf{\Sigma}$, le nombre de variables et le nombre d'échantillons sont étudiés dans la section 4.3.3 du chapitre 4. Dans ces différents scénarios l'estimateur $\widehat{\mathbf{\Sigma}}$ obtenu par notre méthode et implémenté dans le package `BlockCov` est, selon nos critères, plus performant que ceux obtenus à l'aide de diverses méthodes de clustering (kmeans, spectral clustering, clustering hiérarchique) ainsi que ceux proposés par Blum et al. (2016b) et Rothman (2012). De même, dans les différents scénarios, les performances de notre estimateur de $\widehat{\mathbf{\Omega}}^{-1/2}$ ont été comparées à celles des estimateurs obtenus à l'aide des méthodes de clustering et aux performances de l'estimateur proposé par Hosseini & Lee (2016). Dans la plupart des scénarios envisagés notre estimateur est plus performant que les autres. L'estimateur de $\widehat{\mathbf{\Omega}}^{-1/2}$ peut aussi être utilisé dans la méthode décrite section 1.2.2 pour faire de la sélection de variables. Des courbes ROC et des courbes précision-rappel, comparant les variables sélectionnées grâce à notre estimateur et celles sélectionnées par la méthode décrite dans Perthame et al. (2016) sont proposées à la fin du chapitre 4.

Cette méthode a ensuite été appliquée à une étude de métabolomique et de protéomique visant à étudier l'impact de la température de production sur la capacité germinative de graines d'Arabidopsis thaliana.

1.4 Une autre application : étude du dialogue entre les cellules dendritiques et les lymphocytes Th

Ces méthodes ont ensuite été appliquées sur un jeu de données immunologiques afin de mieux comprendre le dialogue existant entre les cellules dendritiques et les lymphocytes Th, deux types de cellules essentielles à notre système immunitaire.

1.4.1 Introduction à l'immunologie

Les cellules dendritiques (DC) sont des cellules sentinelles. Elles sont présentes dans toutes les portes d'entrée empruntées par les agents infectieux (antigènes). Lorsque l'un d'entre eux s'invite dans notre organisme il est ingéré par une DC qui émet alors des signaux biochimiques afin de stimuler l'ensemble du système immunitaire et notamment les lymphocytes T-helper (Th) pour combattre cet agent infectieux. Avant de recevoir ces signaux biochimiques le lymphocyte Th est dit naïf. En recevant ces signaux il va se différencier et renvoyer des signaux spécifiques pour

combattre ce type d'antigène. En fonction de ces signaux les lymphocytes Th ont été catégorisés en différent profils. Les deux premiers profils, caractérisés dans la littérature, sont les profils Th1 et Th2 (Mosmann et al., 1986; Mosmann & Coffman, 1989). En présence d'un pathogène intracellulaire, donc d'une maladie auto-immune, la DC va sécréter de l'interleukine 12 (IL12), une fois capté par le lymphocyte T naïf celui-ci va se différencier en lymphocytes Th1 et sécréter notamment de l'interféron gamma (IFNg) pour combattre cette maladie auto-immune. De même, en présence d'un parasite extra-cellulaire, les DC vont sécréter des signaux qui vont amener le lymphocyte T naïf à se différencier en un lymphocyte Th2 qui va sécréter des interleukines 4, 5, 13 (IL4, IL5, IL13). De nombreuses études ont mis en avant de nouveaux profils tel que le profil Th17 induit par la présence d'IL6, TNFa, IL23 TGFb qui sécrètent IL17A et IL17F pour répondre à la présence de bactéries et de champignons extérieurs (voir Park et al., 2005). La figure 1.4 qui est une version simplifiée de la figure 1 de Leung et al. (2010) montre différents profils Th décrits.

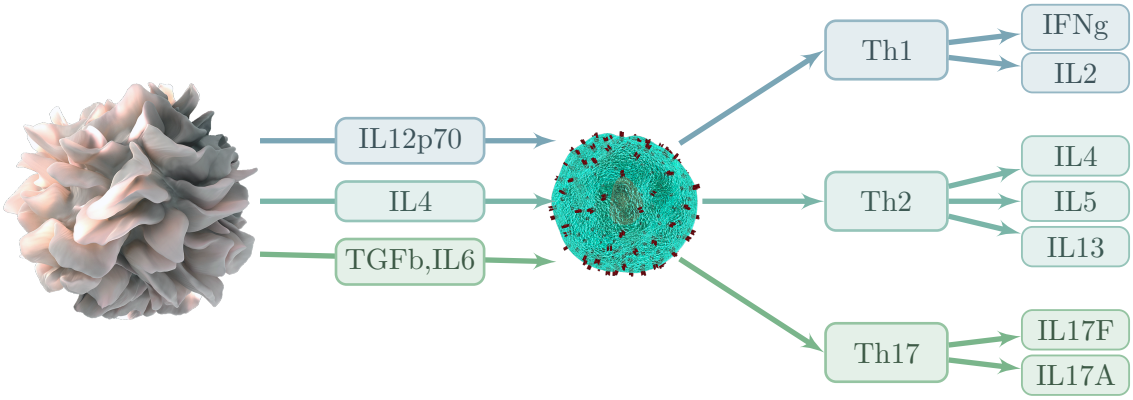

FIGURE 1.4 – Les différents profils Th

En pratique, en présence d'un antigène, les DC vont sécréter de nombreux signaux qui vont être captés par des lymphocytes Th qui vont eux-mêmes sécréter de nombreux signaux, et ne seront pas caractérisés par un simple profil. Afin de mieux comprendre la réponse immunitaire il est important de mettre en évidence quels signaux émis par les DC induisent un signal spécifique des lymphocytes Th. Volpe et al. (2009) ont étudié 5 signaux des cellules dendritiques visant à expliquer la différenciation des lymphocytes T et ont mis en évidence la synergie de 4 de ces paramètres dans la génération de Th17. Cependant, un grand nombre de paramètres (plus de 40) sont connus pour induire la réponse Th et n'ont jamais été étudiés tous ensemble jusqu'à présent.

1.4.2 Contributions du chapitre 5

Production scientifique

Cette section résume la publication suivante :

- **M. Grandclaudon*, M. Perrot-Dockes*, C. Trichot, O. Mostafa-Abouzid, W. Abou-Jaoudé, F. Berger, P. Hupé, D. Thieffry, L. Sansonnet, J. Chiquet, C. Lévy-Leduc, V. Soumelis** A quantitative multivariate model of human dendritic cell-T helper cell communication. *Acceptée au journal Cell et en annexe de la thèse.*

** : ces auteurs ont contribué de manière égale à cette publication*

Dans cette partie nous étudions des données mesurant la réponse des lymphocytes Th aux signaux des DC sous diverses conditions de perturbations visant à reproduire in vitro, l'environnement in situ et in vivo dans lequel baignent les cellules dendritiques et les lymphocytes Th. Ce jeu de données contient pour 428 couples groupe de DC / groupe de lymphocytes Th les valeurs de 36 signaux sécrétés par les DC et de 18 signaux provenant des lymphocytes Th.

De ces expériences découle une matrice $\boldsymbol{X}$ (resp. une matrice $\boldsymbol{Y}$) de taille 428×36 contenant les valeurs des 36 signaux des DC (resp. des 18 signaux de Th) pour les 428 échantillons. Pour expliquer la matrice $\boldsymbol{Y}$ en fonction de la matrice $\boldsymbol{X}$ nous avons appliqué la méthodologie décrite dans le chapitre 2 et dans la section 1.2.2. En observant les valeurs non nulles de $\widehat{\boldsymbol{B}}$ il est alors possible de mettre en évidence des associations potentielles entre les signaux des DC et les signaux des Th. Certaines sont connues de la littérature, d'autres ne le sont pas. Par exemple, le modèle suggère une association entre IL12p70 et IL17F. Or IL12p70 est connu comme induisant les profils Th1 alors qu'IL17F est caractéristique des Th17. Des expériences ont été effectuées et il s'est avéré qu'en effet IL12p70 peut influencer IL17F dans certains contextes. Cette étude a donné lieu à un article qui est disponible en annexe, un résumé détaillé est proposé dans le chapitre 5 de cette thèse.

Chapter 2

Variable selection in multivariate linear models with high-dimensional covariance matrix estimation

Scientific production

The content of this chapter is contained in the article: M. Perrot-Dockès, C. Lévy-Leduc, L. Sansonnet, J. Chiquet, *"Variable selection in multivariate linear models with high-dimensional covariance matrix estimation"* Journal of Multivariate Analysis, 166:78 – 97, 2018. The method which is presented is implemented in the MultiVarSel R package available from the CRAN.

Abstract

In this paper, we propose a novel variable selection approach in the framework of multivariate linear models taking into account the dependence that may exist between the responses. It consists in estimating beforehand the covariance matrix Σ of the responses and to plug this estimator in a Lasso criterion, in order to obtain a sparse estimator of the coefficient matrix. The properties of our approach are investigated both from a theoretical and a numerical point of view. More precisely, we give general conditions that the estimators of the covariance matrix and its inverse have to satisfy in order to recover the positions of the null and non null entries of the coefficient matrix when the size of Σ is not fixed and can tend to infinity. We prove that these conditions are satisfied in the particular case of some Toeplitz matrices. Our approach is implemented in the R package `MultiVarSel` available from the Comprehensive R Archive Network (CRAN) and is very attractive since it benefits from a low computational load. We also assess the performance of our methodology using synthetic data and compare it with alternative approaches. Our numerical experiments show that including the estimation of the covariance matrix in the Lasso criterion dramatically improves the variable selection performance in many cases.

2.1 Introduction

The multivariate linear model consists in generalizing the classical linear model, in which a single response is explained by p variables, to the case where the number q of responses is larger than 1. Such a general modeling can be used in a wide variety of applications ranging from econometrics Lütkepohl (2005) to bioinformatics Meng et al. (2014). In the latter field, for instance, multivariate models have been used to gain insight into complex biological mechanisms like metabolism or gene regulation. This has been made possible thanks to recently developed sequencing technologies. For further details, we refer the reader to Mehmood et al. (2012). However, the downside of such a technological expansion is to include irrelevant variables in the statistical models. To circumvent this, devising efficient variable selection approaches in the multivariate setting has become a growing concern.

A first naive approach to deal with the variable selection issue in the multivariate setting consists in applying classical univariate variable selection strategies to each response separately. Some well-known variable selection methods include the least absolute shrinkage and selection operator (LASSO) proposed by Tibshirani Tibshirani (1996) and the smoothly clipped absolute deviation (SCAD) approach devised by Fan and Li Fan & Li (2001). However, such a strategy does not take into account the dependence that may exist between the different responses.

In this paper, we shall consider the following multivariate linear model:

$$Y = XB + E, \tag{2.1}$$

where $Y = (Y_{i,j})_{1\leq i\leq n, 1\leq j\leq q}$ denotes the $n \times q$ random response matrix, X denotes the $n \times p$ design matrix, B denotes a $p \times q$ coefficient matrix and $E = (E_{i,j})_{1\leq i\leq n, 1\leq j\leq q}$ denotes the $n \times q$ random error matrix, where n is the sample size. In order to model the potential dependence that may exist between the columns of E, we shall assume that for each i in $\{1, \ldots, n\}$,

$$(E_{i,1}, \ldots, E_{i,q}) \sim \mathcal{N}(0, \Sigma), \tag{2.2}$$

where Σ denotes the covariance matrix of the ith row of the error matrix E. We shall moreover assume that the different rows of E are independent. With such assumptions, there is some dependence between the columns of E but not between the rows. Our goal is here to design a variable selection approach which is able to identify the positions of the null and non null entries in the sparse matrix B by taking into account the dependence between the columns of E.

This issue has recently been considered by Lee and Liu Lee & Liu (2012) who extended the approach of Rothman et al. Rothman et al. (2010). More precisely, Lee and Liu Lee & Liu (2012) proposed three approaches for dealing with this issue based on penalized maximum likelihood with a weighted ℓ_1 regularization. In their

first approach B is estimated by using a plug-in estimator of Σ^{-1}, in the second one, Σ^{-1} is estimated by using a plug-in estimator of B and in the third one, Σ^{-1} and B are estimated simultaneously. Lee and Liu Lee & Liu (2012) also investigate the asymptotic properties of their methods when the sample size n tends to infinity and the number of rows and columns q of Σ is fixed.

In this paper, we propose to estimate Σ beforehand and to plug this estimator in a Lasso criterion, in order to obtain a sparse estimator of B. Hence, our methodology is close to the first approach of Lee and Liu Lee & Liu (2012). However, there are two main differences.

The first one is the asymptotic framework in which our theoretical results are established: q is assumed to depend on n and to tend to infinity at a rate which can be larger than n as n tends to infinity. Moreover, p is assumed to be fixed. In this framework, we give general conditions that the estimators of Σ and Σ^{-1} have to satisfy in order to be able to recover the support of B that is to find the positions of the null and non null entries of the matrix B. Such a framework in which q is much larger than n and p is fixed may for instance occur in metabolomics which aims to provide a global snapshot of the metabolism. In a typical metabolomic experiment, we have access to the responses of q metabolites (features) for n samples belonging to different groups. This information can be summarized in a $n \times q$ data matrix where $q \approx 5000$ and $n \approx 30$. The goal is then to identify the most important features distinguishing the different groups. Hence, this problem can be modeled by (3.2) where X is the design matrix of a one-way ANOVA where p corresponds to the number of groups.

The second main difference between Lee & Liu (2012) and our approach is the strategy that we use for estimating Σ. In Lee & Liu (2012), Σ^{-1} is estimated by using an adaptation of the Graphical Lasso (GLASSO) proposed by Friedman et al. (2008). This technique has also been considered in Yuan & Lin (2007); Banerjee et al. (2008a); Rothman et al. (2008). Here, we propose to estimate Σ beforehand from the empirical covariance matrix of the residuals assuming that Σ has a particular structure, Toeplitz for instance. We prove its efficiency in the particular case of an AR(1) process in Section 2.2.3. Such a process is used, for instance, in population genetics for modeling the phenomenon of recombination as shown in Chiquet et al. (2016). More generally, we give general conditions that the estimators of Σ and Σ^{-1} have to satisfy in order to be able to recover the support of B. Hence, any approach providing estimators satisfying these conditions can be used.

Let us now describe more precisely our methodology. We start by "whitening" the observations Y by applying the following transformation to Model (3.2):

$$Y\,\Sigma^{-1/2} = XB\,\Sigma^{-1/2} + E\,\Sigma^{-1/2}. \tag{2.3}$$

The goal of such a transformation is to remove the dependence between the columns

of Y. Then, for estimating B, we proceed as follows. Let us observe that (3.5) can be rewritten as:

$$\mathcal{Y} = \mathcal{X}\mathcal{B} + \mathcal{E}, \tag{2.4}$$

with

$$\mathcal{Y} = \text{vec}(Y\,\Sigma^{-1/2}), \quad \mathcal{X} = (\Sigma^{-1/2})^\top \otimes X, \quad \mathcal{B} = \text{vec}(B) \quad \text{and} \quad \mathcal{E} = \text{vec}(E\,\Sigma^{-1/2}), \tag{2.5}$$

where vec denotes the vectorization operator and $\otimes$ the Kronecker product.

With Model (4.4), estimating B is equivalent to estimate $\mathcal{B}$ since $\mathcal{B} = \text{vec}(B)$. Then, for estimating $\mathcal{B}$, we use the classical LASSO criterion defined as follows for a nonnegative λ:

$$\widehat{\mathcal{B}}(\lambda) = \text{Argmin}_{\mathcal{B}}\{\|\mathcal{Y} - \mathcal{X}\mathcal{B}\|_2^2 + \lambda\|\mathcal{B}\|_1\}, \tag{2.6}$$

where $\|\cdot\|_1$ and $\|\cdot\|_2$ denote the classical ℓ_1-norm and ℓ_2-norm, respectively. Inspired by Zhao & Yu (2006), Theorem 2.1 established some conditions under which the positions of the null and non null entries of $\mathcal{B}$ can be recovered by using $\widehat{\mathcal{B}}$.

In practical situations, the covariance matrix Σ is generally unknown and has thus to be estimated. Let $\widehat{\Sigma}$ denote an estimator of Σ. For a description of the methodology that we propose for estimating Σ, we refer the reader to the end of Section 2.2.2. Then, the estimator $\widehat{\Sigma}^{-1/2}$ of $\Sigma^{-1/2}$ is such that

$$\widehat{\Sigma}^{-1} = \widehat{\Sigma}^{-1/2}(\widehat{\Sigma}^{-1/2})^\top.$$

When $\Sigma^{-1/2}$ is replaced by $\widehat{\Sigma}^{-1/2}$, (3.5) becomes

$$Y\,\widehat{\Sigma}^{-1/2} = X B\,\widehat{\Sigma}^{-1/2} + E\,\widehat{\Sigma}^{-1/2}, \tag{2.7}$$

which can be rewritten as

$$\widetilde{\mathcal{Y}} = \widetilde{\mathcal{X}}\mathcal{B} + \widetilde{\mathcal{E}}, \tag{2.8}$$

where

$$\widetilde{\mathcal{Y}} = \text{vec}(Y\,\widehat{\Sigma}^{-1/2}), \quad \widetilde{\mathcal{X}} = (\widehat{\Sigma}^{-1/2})^\top \otimes X, \quad \mathcal{B} = \text{vec}(B) \quad \text{and} \quad \widetilde{\mathcal{E}} = \text{vec}(E\,\widehat{\Sigma}^{-1/2}).$$

In Model (2.8), $\mathcal{B}$ is estimated by

$$\widetilde{\mathcal{B}}(\lambda) = \text{Argmin}_{\mathcal{B}}\{\|\widetilde{\mathcal{Y}} - \widetilde{\mathcal{X}}\mathcal{B}\|_2^2 + \lambda\|\mathcal{B}\|_1\}. \tag{2.9}$$

By extending Theorem 2.1, Theorem 2.5 gives some conditions on the eigenvalues of Σ^{-1} and on the convergence rate of $\widehat{\Sigma}$ and its inverse to Σ and Σ^{-1}, respectively, under which the positions of the null and non null entries of $\mathcal{B}$ can be recovered by using $\widetilde{\mathcal{B}}$.

We prove in Section 2.2.3 that when Σ is a particular Toeplitz matrix, namely

the covariance matrix of an AR(1) process, the assumptions of Theorem 2.5 are satisfied. This strategy has been implemented in the R package `MultiVarSel`, which is available on the Comprehensive R Archive Network (CRAN), for more general Toeplitz matrices Σ such as the covariance matrix of ARMA processes or general stationary processes. For a successful application of this methodology to particular "-omic" data, namely metabolomic data, we refer the reader to Perrot-Dockès et al. (2018). For a review of the most recent methods for estimating high-dimensional covariance matrices, we refer the reader to Pourahmadi (2013).

The paper is organized as follows. Section 2.2 is devoted to the theoretical results of the paper. The assumptions under which the positions of the non null and null entries of $\mathcal{B}$ can be recovered are established in Theorem 2.1 when Σ is known and in Theorem 2.5 when Σ is unknown. Section 2.2.3 studies the specific case of the AR(1) model. We present in Section 4.3 some numerical experiments in order to support our theoretical results. The proofs of our main theoretical results are given in Section 2.5.

2.2 Theoretical results

2.2.1 Case where Σ is known

Let us first introduce some notations. Let

$$C = \frac{1}{nq}\mathcal{X}^\top \mathcal{X} \quad \text{and} \quad J = \{j : 1 \leq j \leq pq, \mathcal{B}_j \neq 0\}, \tag{2.10}$$

where $\mathcal{X}$ is defined in (2.5) and where $\mathcal{B}_j$ denotes the jth component of the vector $\mathcal{B}$ defined in (2.5).

Let also define

$$C_{J,J} = \frac{1}{nq}(\mathcal{X}_{\bullet,J})^\top \mathcal{X}_{\bullet,J} \quad \text{and} \quad C_{J^c,J} = \frac{1}{nq}(\mathcal{X}_{\bullet,J^c})^\top \mathcal{X}_{\bullet,J}, \tag{2.11}$$

where $\mathcal{X}_{\bullet,J}$ and $\mathcal{X}_{\bullet,J^c}$ denote the columns of $\mathcal{X}$ belonging to the set J defined in (2.10) and to its complement J^c, respectively.

More generally, for any matrix A, $A_{I,J}$ denotes the partitioned matrix extracted from A by considering the rows of A belonging to the set I and the columns of A belonging to the set J, with $\bullet$ indicating all the rows or all the columns.

The following theorem gives some conditions under which the estimator $\widehat{\mathcal{B}}$ defined in (4.5) is sign-consistent as defined by Zhao & Yu (2006), namely,

$$\lim_{n\to\infty} \mathbb{P}\{\text{sign}(\widehat{\mathcal{B}}) = \text{sign}(\mathcal{B})\} = 1,$$

where the sign function maps positive entries to 1, negative entries to -1 and zero to 0.

Theorème 2.1. *Assume that* $\mathcal{Y} = (\mathcal{Y}_1, \ldots, \mathcal{Y}_{nq})^\top$ *satisfies Model* (4.4). *Assume also that there exist some positive constants* M_1, M_2, M_3 *and positive numbers* c_1, c_2 *such that* $c_1 + c_2 \in (0, 1/2)$ *satisfying*

(A1) *For all* $n \in \mathbb{N}$ *and* $j \in \{1, \ldots, pq\}$, $n^{-1}(\mathcal{X}_{\bullet,j})^\top \mathcal{X}_{\bullet,j} \leq M_1$, *where* $\mathcal{X}_{\bullet,j}$ *is the* j*th column of* $\mathcal{X}$ *defined in* (2.5).

(A2) *For all* $n \in \mathbb{N}$, $n^{-1}\lambda_{\min}\{(\mathcal{X}^\top \mathcal{X})_{J,J}\} \geq M_2$, *where* $\lambda_{\min}(A)$ *denotes the smallest eigenvalue of* A.

(A3) $|J| = O(q^{c_1})$, *where* J *is defined in* (2.10) *and* $|J|$ *is the cardinality of the set* J,

(A4) $q^{c_2} \min_{j \in J} |\mathcal{B}_j| \geq M_3$.

Assume also that the following strong Irrepresentable Condition holds:

(IC) *There exists a positive constant vector* η *such that*

$$|(\mathcal{X}^\top \mathcal{X})_{J^c,J}\{(\mathcal{X}^\top \mathcal{X})_{J,J}\}^{-1} \operatorname{sign}(\mathcal{B}_J)| \leq \mathbf{1} - \eta,$$

where $\mathbf{1}$ *is a* $(pq - |J|) \times 1$ *vector of 1s and the inequality holds element-wise.*

Then, for all λ *that satisfies*

(L) $q = q_n = o\{n^{\frac{1}{2(c_1+c_2)}}\}$, $\lambda/\sqrt{n} \to \infty$ *and* $\lambda/n = o\{q^{-(c_1+c_2)}\}$, *as* $n \to \infty$,

we have

$$\lim_{n \to \infty} \mathbb{P}[\operatorname{sign}\{\widehat{\mathcal{B}}(\lambda)\} = \operatorname{sign}(\mathcal{B})] = 1,$$

where $\widehat{\mathcal{B}}(\lambda)$ *is defined by* (4.5).

Remark 2.1. *Observe that if* $c_1 + c_2 < 1/(2k)$, *for some positive* k, *then the first condition of* (L) *becomes* $q = o(n^k)$. *Hence for large values of* k, *the size* q *of* Σ *is much larger than* n.

Theorem 2.1 is established under similar assumptions as those required by Theorem 4 of Zhao & Yu (2006). However, Assumptions (5), (6), (7) and (8) of Zhao & Yu (2006) had to be adapted to our multivariate framework where q tends to infinity and were replaced by (A1), (A2), (A3) and (A4) in order to allow q and n to grow to infinity with different rates. Moreover, the assumptions on λ had also to be adapted to deal with our specific framework.

The proof of Theorem 2.1 is given in Section 2.5. It is based on the following Proposition 2.2 which is an adaptation to the multivariate case of Proposition 1 in Zhao & Yu (2006). More precisely, in order to prove Theorem 2.1, we show that $\mathbb{P}(A_n^c)$ and $\mathbb{P}(B_n^c)$ tend to 0 as $n \to \infty$.

Proposition 2.2. *Let $\widehat{\mathcal{B}}(\lambda)$ be defined by (4.5). Then*

$$\mathbb{P}[\operatorname{sign}\{\widehat{\mathcal{B}}(\lambda)\} = \operatorname{sign}(\mathcal{B})] \geq \mathbb{P}(A_n \cap B_n),$$

where

$$A_n = \left\{ |(C_{J,J})^{-1} W_J| < \sqrt{nq} \left(|\mathcal{B}_J| - \frac{\lambda}{2nq} |(C_{J,J})^{-1} \operatorname{sign}(\mathcal{B}_J)| \right) \right\} \tag{2.12}$$

and

$$B_n = \left\{ |C_{J^c,J}(C_{J,J})^{-1} W_J - W_{J^c}| \leq \frac{\lambda}{2\sqrt{nq}} \left(\mathbf{1} - |C_{J^c,J}(C_{J,J})^{-1} \operatorname{sign}(\mathcal{B}_J)| \right) \right\}, \tag{2.13}$$

with $W = \mathcal{X}'\mathcal{E}/\sqrt{nq}$. In (2.12) and (2.13), $C_{J,J}$ and $C_{J^c,J}$ are defined in (2.11) and W_J and W_{J^c} denote the components of W being in J and J^c, respectively. Note that the previous inequalities hold element-wise.

The proof of Proposition 2.2 is given in Section 2.5.

In the following proposition, which is also proved in Section 2.5, we give some conditions on X and Σ under which Assumptions (A1) and (A2) of Theorem 2.1 hold.

Proposition 2.3. *If there exist some positive constants M_1', M_2', m_1, m_2 such that, for all $n \in \mathbb{N}$,*

(C1) *For all $j \in \{1, \ldots, p\}$, $n^{-1}(X_{\bullet,j})^\top X_{\bullet,j} \leq M_1'$,*

(C2) $n^{-1}\lambda_{\min}(X^\top X) \geq M_2'$,

(C3) $\lambda_{\max}(\Sigma^{-1}) \leq m_1$,

(C4) $\lambda_{\min}(\Sigma^{-1}) \geq m_2$,

then Assumptions (A1) and (A2) of Theorem 2.1 are satisfied.

Remark 2.2. *Observe that (C1) and (C2) hold in the case where the columns of the matrix X are orthogonal.*

We give in Proposition 2.6 in Section 2.2.3 some conditions under which Condition (IC) holds in the specific case where Σ is the covariance matrix of an AR(1) process.

2.2.2 Case where Σ is unknown

Similarly as in (2.10) and (2.11), we introduce the notations

$$\widetilde{C} = \frac{1}{nq} \widetilde{\mathcal{X}}^\top \widetilde{\mathcal{X}} \tag{2.14}$$

and

$$\widetilde{C}_{J,J} = \frac{1}{nq}(\widetilde{\mathcal{X}}_{\bullet,J})^\top \widetilde{\mathcal{X}}_{\bullet,J} \quad \text{and} \quad \widetilde{C}_{J^c,J} = \frac{1}{nq}(\widetilde{\mathcal{X}}_{\bullet,J^c})^\top \widetilde{\mathcal{X}}_{\bullet,J}, \tag{2.15}$$

where $\widetilde{\mathcal{X}}_{\bullet,J}$ and $\widetilde{\mathcal{X}}_{\bullet,J^c}$ denote the columns of $\widetilde{\mathcal{X}}$ belonging to the set J defined in (2.10) and to its complement J^c, respectively.

A straightforward extension of Proposition 2.2 leads to the following proposition for Model (2.8).

Proposition 2.4. *Let $\widetilde{\mathcal{B}}(\lambda)$ be defined by (2.9). Then*

$$\mathbb{P}[\operatorname{sign}\{\widetilde{\mathcal{B}}(\lambda)\} = \operatorname{sign}(\mathcal{B})] \geq \mathbb{P}(\widetilde{A_n} \cap \widetilde{B_n}),$$

where

$$\widetilde{A_n} = \left\{ \left| (\widetilde{C}_{J,J})^{-1} \widetilde{W}_J \right| < \sqrt{nq} \left(|\mathcal{B}_J| - \frac{\lambda}{2nq} |(\widetilde{C}_{J,J})^{-1} \operatorname{sign}(\mathcal{B}_J)| \right) \right\} \tag{2.16}$$

and

$$\widetilde{B_n} = \left\{ \left| \widetilde{C}_{J^c,J} (\widetilde{C}_{J,J})^{-1} \widetilde{W}_J - \widetilde{W}_{J^c} \right| \leq \frac{\lambda}{2\sqrt{nq}} \left(\mathbf{1} - \left| \widetilde{C}_{J^c,J} (\widetilde{C}_{J,J})^{-1} \operatorname{sign}(\mathcal{B}_J) \right| \right) \right\}, \tag{2.17}$$

with $\widetilde{W} = \widetilde{\mathcal{X}}'\widetilde{\mathcal{E}}/\sqrt{nq}$. In (2.16) and (2.17), $\widetilde{C}_{J,J}$ and $\widetilde{C}_{J^c,J}$ are defined in (2.15) and $\widetilde{W}_J$ and $\widetilde{W}_{J^c}$ denote the components of $\widetilde{W}$ being in J and J^c, respectively. Note that the previous inequalities hold element-wise.

The following theorem extends Theorem 2.1 to the case where Σ is unknown and gives some conditions under which the estimator $\widetilde{\mathcal{B}}$ defined in (2.9) is sign-consistent.

Theorème 2.5. *Assume that Assumptions (A3), (A4), (IC) and (L) of Theorem 2.1 hold. Assume also that, there exist some positive constants M_4, M_5, M_6 and M_7, such that for all $n \in \mathbb{N}$,*

(A5) $\|(X^\top X)/n\|_\infty \leq M_4$,

(A6) $\lambda_{\min}\{(X^\top X)/n\} \geq M_5$,

(A7) $\lambda_{\max}(\Sigma^{-1}) \leq M_6$,

(A8) $\lambda_{\min}(\Sigma^{-1}) \geq M_7$.

Suppose also that

(A9) $\|\Sigma^{-1} - \widehat{\Sigma}^{-1}\|_\infty = O_P\{(nq)^{-1/2}\}$, *as* $n \to \infty$,

(A10) $\rho(\Sigma - \widehat{\Sigma}) = O_P\{(nq)^{-1/2}\}$, *as* $n \to \infty$.

Let $\widetilde{\mathcal{B}}(\lambda)$ be defined by (2.9), then

$$\lim_{n\to\infty} \mathbb{P}[\operatorname{sign}\{\widetilde{\mathcal{B}}(\lambda)\} = \operatorname{sign}(\mathcal{B})] = 1.$$

In the previous assumptions, $\lambda_{\max}(A)$, $\lambda_{\min}(A)$, $\rho(A)$ and $\|A\|_\infty$ denote the largest eigenvalue, the smallest eigenvalue, the spectral radius and the infinite norm (induced by the associated vector norm) of the matrix A.

Remark 2.3. *In order to distinguish the assumptions that are required for the design matrix X and for the estimator $\widehat{\Sigma}$ of Σ, the assumptions of Theorem 2.5 only involve X, Σ and $\Sigma - \widehat{\Sigma}$ but not $\widetilde{\mathcal{X}}$.*

Remark 2.4. *Observe that Assumptions (A5) and (A6) hold in the case where the columns of the matrix X are orthogonal. Note also that (A7) and (A8) are the same as (C3) and (C4) in Proposition 2.3.*

The proof of Theorem 2.5 is given in Section 2.5 and is based on Proposition 2.4. In order to prove Theorem 2.5, it is enough to show that $\mathbb{P}(\widetilde{A_n}^c)$ and $\mathbb{P}(\widetilde{B_n}^c)$ tend to 0 as $n \to \infty$. The idea of the proof consists in rewriting $\mathbb{P}(\widetilde{A_n}^c)$ (resp. $\mathbb{P}(\widetilde{B_n}^c)$) by adding terms depending on $\Sigma - \widehat{\Sigma}$ to $\mathbb{P}(A_n^c)$ (resp. $\mathbb{P}(B_n^c)$) and to prove that these additional terms tend to zero as $n \to \infty$.

In order to estimate Σ, we propose the following strategy:

(a) Fitting a classical linear model to each column of the matrix Y in order to have access to an estimation $\widehat{E}$ of the random error matrix E. It is possible since p is assumed to be fixed and smaller than n.

(b) Estimating Σ from $\widehat{E}$ by assuming that Σ has a particular structure, Toeplitz for instance.

More precisely, $\widehat{E}$ defined in the first step is such that:

$$\widehat{E} = \{\mathrm{Id}_{\mathbb{R}^n} - X(X^\top X)^{-1}X^\top\}E \equiv \Pi E, \tag{2.18}$$

which implies that

$$\widehat{\mathcal{E}} = \mathrm{vec}(\widehat{E}) = (\mathrm{Id}_{\mathbb{R}^q} \otimes \Pi)\mathcal{E}, \tag{2.19}$$

where $\mathcal{E}$ is defined in (2.5).

We prove in Proposition 2.7 below that our strategy for estimating Σ provides an estimator satisfying the assumptions of Theorem 2.5 in the case where $(E_{1,t})_t, \ldots, (E_{n,t})_t$ are assumed to be independent AR(1) processes.

2.2.3 The AR(1) case

Sufficient conditions for Assumption (IC) of Theorem 2.1

The following proposition gives some conditions under which the strong Irrepresentable Condition (IC) of Theorem 2.1 holds.

Proposition 2.6. *Assume that $(E_{1,t})_t, \dots, (E_{n,t})_t$ in Model (3.2) are independent AR(1) processes such that, for all $i \in \{1, \dots, n\}$, $E_{i,t} - \phi_1 E_{i,t-1} = Z_{i,t}$, where the $Z_{i,t}$s are zero-mean iid Gaussian random variables with variance $\sigma^2 = 1$ and $|\phi_1| < 1$. Assume also that X defined in (3.2) is such that $X^\top X = \nu Id_{\mathbb{R}^p}$, where ν is a positive constant. Moreover, suppose that if $j \in J$, then $j > p$ and $j < pq - p$. Suppose also that for all j, $j-p$ or $j+p$ is not in J. Then, the strong Irrepresentable Condition (IC) of Theorem 2.1 holds.*

The proof of Proposition 2.6 is given in Section 2.5.

Sufficient conditions for Assumptions (A7), (A8), (A9) and (A10) of Theorem 2.5

The following proposition establishes that in the particular case where the $(E_{1,t})_t, \dots, (E_{n,t})_t$ are independent AR(1) processes, our strategy for estimating Σ provides an estimator satisfying the assumptions of Theorem 2.5.

Proposition 2.7. *Assume that $(E_{1,t})_t, \dots, (E_{n,t})_t$ in Model (3.2) are independent AR(1) processes such that, for all $i \in \{1, \dots, n\}$, $E_{i,t} - \phi_1 E_{i,t-1} = Z_{i,t}$, where the $Z_{i,t}$s are zero-mean iid Gaussian random variables with variance $\sigma^2 = 1$ and $|\phi_1| < 1$. Let*

$$\widehat{\Sigma} = \frac{1}{1-\widehat{\phi}_1^2} \begin{pmatrix} 1 & \widehat{\phi}_1 & \widehat{\phi}_1^2 & \dots & \widehat{\phi}_1^{q-1} \\ \widehat{\phi}_1 & 1 & \widehat{\phi}_1 & \dots & \widehat{\phi}_1^{q-2} \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ \widehat{\phi}_1^{q-1} & \dots & \dots & \dots & 1 \end{pmatrix},$$

where

$$\widehat{\phi}_1 = \frac{\sum_{i=1}^n \sum_{\ell=2}^q \widehat{E}_{i,\ell} \widehat{E}_{i,\ell-1}}{\sum_{i=1}^n \sum_{\ell=1}^{q-1} \widehat{E}_{i,\ell}^2}, \tag{2.20}$$

where $\widehat{E} = (\widehat{E}_{i,\ell})_{1 \le i \le n, 1 \le \ell \le q}$ is defined in (2.18). Then, Assumptions (A7), (A8), (A9) and (A10) of Theorem 2.5 are valid.

The proof of Proposition 2.7 is given in Section 2.5. It is based on the following lemma.

Lemma 2.8. *Assume that $(E_{1,t})_t, \dots, (E_{n,t})_t$ in Model (3.2) are independent AR(1) processes such that, for all $i \in \{1, \dots, n\}$, $E_{i,t} - \phi_1 E_{i,t-1} = Z_{i,t}$, where the $Z_{i,t}$s are zero-mean iid Gaussian random variables with variance σ^2 and $|\phi_1| < 1$. Let*

$$\widehat{\phi}_1 = \frac{\sum_{i=1}^n \sum_{\ell=2}^q \widehat{E}_{i,\ell} \widehat{E}_{i,\ell-1}}{\sum_{i=1}^n \sum_{\ell=1}^{q-1} \widehat{E}_{i,\ell}^2},$$

where $\widehat{E} = (\widehat{E}_{i,\ell})_{1\leq i\leq n, 1\leq \ell\leq q}$ *is defined in* (2.18). *Then, as* $n \to \infty$, $\sqrt{nq_n}\,(\widehat{\phi}_1 - \phi_1) = O_p(1)$.

Lemma 2.8 is proved in Section 2.5. Its proof is based on Lemma 2.10 in Section 2.6.

2.3 Numerical experiments

The goal of this section is twofold: (i) to provide sanity checks for our theoretical results in a well-controlled framework; and (ii) to investigate the robustness of our estimator to some violations of the assumptions of our theoretical results. The latter may reveal a broader scope of applicability for our method than the one guaranteed by the theoretical results.

We investigate (i) in the AR(1) framework presented in Section 2.2.3. Indeed, all assumptions made in Theorems 2.1 and 2.5 can be specified with well-controllable simulation parameters in the AR(1) case with balanced design matrix X.

Point (ii) aims to explore the limitations of our theoretical framework and assess its robustness. To this end, we propose two numerical studies relaxing some of the assumptions of our theorems: first, we study the effect of an unbalanced design — which violates the sufficient condition of the irrepresentability condition (IC) given in Proposition 2.6 — on the sign-consistency; and second, we study the effect of other types of dependence than an AR(1).

In all experiments, the performance are assessed in terms of sign-consistency. In other words, we evaluate the probability for the sign of various estimators to be equal to $\text{sign}(\mathcal{B})$. More precisely, we investigate for each estimator if their exists at least one λ such that $\text{sign}\{\widehat{\mathcal{B}}(\lambda)\} = \text{sign}(\mathcal{B})$. We compare the performance of three different estimators:

(i) $\widehat{\mathcal{B}}$ defined in (4.5), which corresponds to the LASSO criterion applied to the data whitened with the true covariance matrix Σ; we call this estimator `oracle`. Its theoretical properties are established in Theorem 2.1.

(ii) $\widetilde{\mathcal{B}}$ defined in (2.9), which corresponds to the LASSO criterion applied to the data whitened with an estimator of the covariance matrix $\widehat{\Sigma}$; we refer to this estimator as `whitened-lasso`. Its theoretical properties are established in Theorem 2.5.

(iii) the LASSO criterion applied to the raw data, which we call `raw-lasso` hereafter. Its theoretical properties are established only in the univariate case in Alquier & Doukhan (2011).

2.3.1 AR(1) dependence structure with balanced one-way ANOVA

In this section, we consider Model (3.2) where X is the design matrix of a one-way ANOVA with two balanced groups. Each row of the random error matrix E is distributed as a centered Gaussian random vector as in Eq. (3.3) where the matrix Σ is the covariance matrix of an AR(1) process defined in Section 2.2.3.

In this setting, Assumptions (A1), (A2) and Condition (IC) of Theorem 2.1 are satisfied, see Propositions 2.3 and 2.6. The three remaining assumptions (A3), (A4) and (L) are related to more practical quantities: (A3) controls the sparsity level of the problem, involving c_1; (A4) basically controls the signal-to-noise ratio, involving c_2 and (L) links the sample size n, q and the two constants c_1, c_2, so that an appropriate range of penalty λ exists for having a large probability of support recovery. This latter assumption is used in our experiments to tune the difficulty of the support recovery as follows: we consider different values of n, q, c_1, c_2. For each 4-tuple (n, q, c_1, c_2), we generated the $2q$-vector $\mathcal{B}$ as follows: the absolute values of its $|J| = q^{c_1}$ non-null components are sampled from a uniform distribution on the interval $[1/q^{c_2}, 2/q^{c_2}]$. Thus, Assumptions (A3) and (A4) are fulfilled. Hence, the problem difficulty is essentially driven by the validity of Assumption (L) where $q = o(n^k)$ with $c_1 + c_2 = 1/2k$, and so by the relationship between n, q and k.

We consider a large range of sample sizes n varying from 10 to 1000 and three different values for $q \in \{10, 50, 1000\}$. The constants c_1, c_2 are chosen such that $c_1 + c_2 = 1/2k$ with $c_1 = c_2$ and $k \in \{1, 2, 4\}$. Additional values of c_1 and c_2 have also been considered and the corresponding results are available upon request. Finally, we consider two values for the parameter ϕ_1 appearing in the definition of the AR(1) process: $\phi_1 \in \{0.5, 0.95\}$.

Note that in this AR(1) setting with the estimator $\widehat{\phi}_1$ of ϕ_1 defined in (2.20), all the assumptions of Theorem 2.5 are fulfilled, see Proposition 2.7.

The frequencies of support recovery for the three estimators averaged over 1000 replications is displayed in Figure 2.1.

We observe from Figure 2.1 that `whitened-lasso` and `oracle` have similar performance since ϕ_1 is well estimated. These two approaches always exhibit better performance than `raw-lasso`, especially when $\phi_1 = 0.95$. In this case, the sample size n required to reach the same performance is indeed ten time larger for `raw-lasso` than for `oracle` and `whitened-lasso`.

Finally, the performance of all estimators are altered when n is too small, especially in situations where the signal to noise ratio (SNR) is small and the signal is not sparse enough, these two characteristics corresponding to small values of k.

2.3.2 Robustness to unbalanced designs and correlated features

The goal of this section is to study some particular design matrices X in Model (3.2) that may lead to violation of the Irrepresentability Condition (IC).

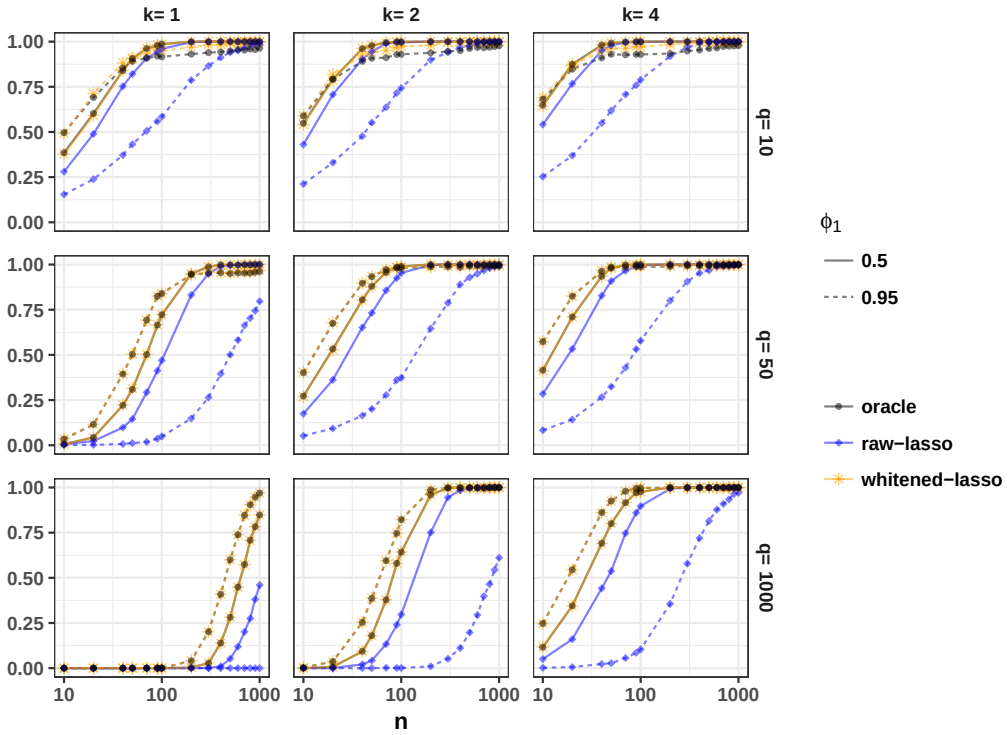

Figure 2.1 – Frequencies of support recovery in a multivariate one-way ANOVA model with two balanced groups and an AR(1) dependence.

To this end, we consider the multivariate linear model (3.2) with the same AR(1) dependence as the one considered in Section 2.3.1. Then, two different matrices X are considered: First, an one-way ANOVA model with two unbalanced groups with respective sizes n_1 and n_2 such that $n_1 + n_2 = n$; and second, a multiple regression model with p correlated Gaussian predictors such that the rows of X are iid $\mathcal{N}(0, \Sigma^X)$.

For the one-way ANOVA, violation of (IC) may occur when $r = n_1/n$ is too different from $1/2$, as stated in Proposition 2.6. For the regression model, we choose for Σ^X a 9×9 matrix ($p = 9$) such that $\Sigma^X_{i,i} = 1$, $\Sigma^X_{i,j} = \rho$, when $i \neq j$.

The other simulation parameters are fixed as in Section 2.3.1.

We report in Figure 2.2 the results for the case where $q = 1000$ and $k = 2$ both for unbalanced one-way ANOVA (top panels) and regression with correlated predictors (bottom panels). For the one-way ANOVA, r varies in $\{0.4, 0.2, 0.1\}$. For the regression case, ρ varies in $\{0.2, 0.6, 0.8\}$. In both cases, the gray lines correspond to the ideal situation that is, either balanced ($r = 0.5$) or uncorrelated (ρ=0) in the legend of Figure 2.2. The probability of support recovery is estimated over 1000 runs.

From this figure, we note that correlated features or unbalanced designs deteri-

orate the support recovery of all estimators. This was expected for these LASSO-based methods which all suffer from the violation of the irrepresentability condition (IC). However, we also note that `whitened-lasso` and `oracle` have similar performance, which means that the estimation of Σ is not altered, and that whitening always improves the support recovery.

2.3.3 Robustness to more general autoregressive processes

In this section, we consider the case where X is the design matrix of a one-way ANOVA with two balanced groups and where Σ is the covariance matrix of an AR(m) process with $m \in \{5, 10\}$. Figure 2.3 displays the performance of the different estimators when $q = 500$ and $k = 2$. Here, for computing $\widehat{\Sigma}$ in `whitened-lasso`, the parameters $\phi_1, \ldots, \phi_m$ of the AR(m) process are estimated as follows. They are obtained by averaging over the n rows of $\widehat{E}$ defined in (2.18) the estimations $\widehat{\phi}_1^{(i)}, \ldots, \widehat{\phi}_m^{(i)}$ obtained for the ith row of $\widehat{E}$ by using standard estimation approaches for AR processes described in Brockwell & Davis (1990). As previously, we observe from this figure that `whitened-lasso` and `oracle` have better performance than `raw-lasso`.

2.4 Discussion

In this paper, we proposed a variable selection approach for multivariate linear models taking into account the dependence that may exist between the responses

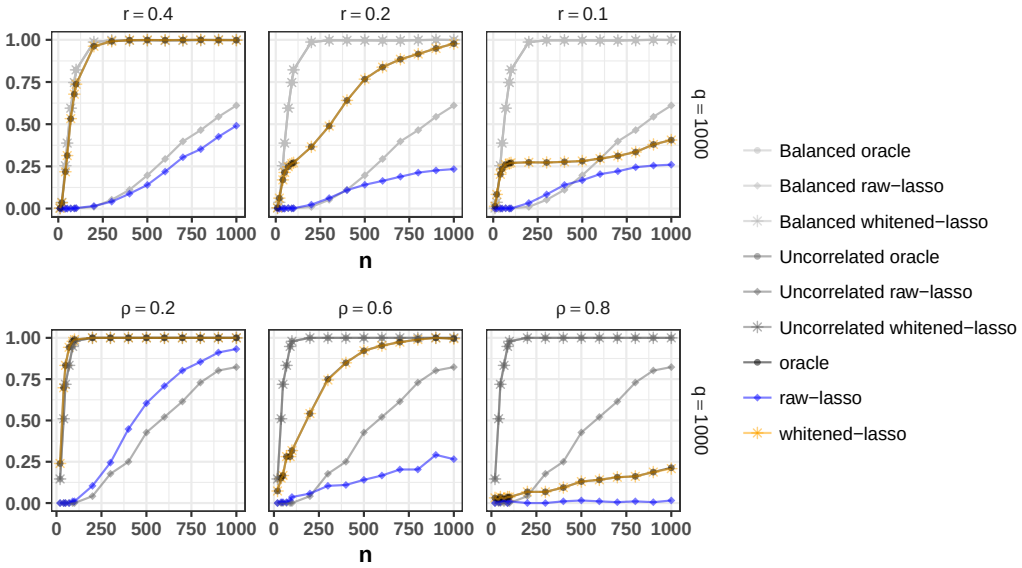

Figure 2.2 – Frequencies of support recovery in general linear models with unbalanced designs: one-way ANOVA and regression.

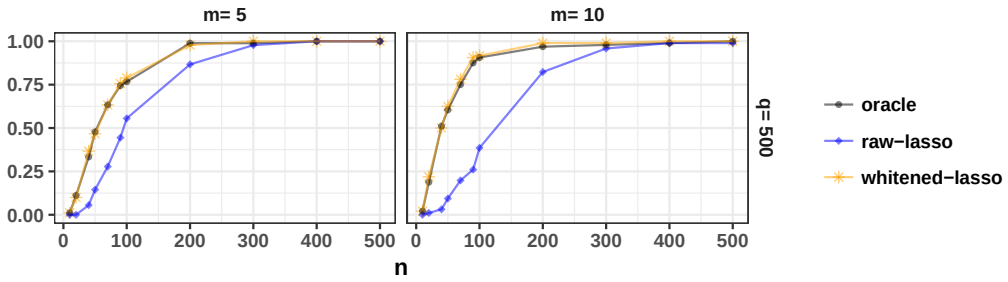

Figure 2.3 – Frequencies of support recovery in one-way ANOVA with AR(m) covariance matrix.

and establish its theoretical properties. More precisely, our method consists in estimating the covariance matrix Σ which models the dependence between the responses and to plug this estimator in a Lasso criterion, in order to obtain a sparse estimator of the coefficient matrix. Then, we give general conditions that the estimators of the covariance matrix and its inverse have to satisfy in order to recover the positions of the null and non null entries of the coefficient matrix when the size of Σ is not fixed and can tend to infinity. In particular, we prove that these general conditions are satisfied for some specific Toeplitz matrices such as the covariance matrix of an AR(1) process. Note that our approach has been successfully applied to a data set coming from a metabolomic experiment. For further details, we refer the reader to Perrot-Dockès et al. (2018). Since, in this paper, we used general Toeplitz covariance matrices such as those of ARMA(p,q) processes or of weakly dependent stationary processes, it would be interesting to prove that the strategy that we used provides estimators of Σ satisfying the assumptions of Theorem 2.5. Moreover, it would be interesting to see if other types of structured covariance matrices would satisfy the assumptions of our theorems.

2.5 Proofs

Proof of Proposition 2.2. For a fixed nonnegative λ, by (4.5),

$$\widehat{\mathcal{B}} = \widehat{\mathcal{B}}(\lambda) = \text{Argmin}_{\mathcal{B}}\{\|\mathcal{Y} - \mathcal{X}\mathcal{B}\|_2^2 + \lambda\|\mathcal{B}\|_1\}.$$

Denoting $\widehat{u} = \widehat{\mathcal{B}} - \mathcal{B}$, we get

$$\begin{aligned}\|\mathcal{Y} - \mathcal{X}\widehat{\mathcal{B}}\|_2^2 + \lambda\|\widehat{\mathcal{B}}\|_1 &= \|\mathcal{X}\mathcal{B} + \mathcal{E} - \mathcal{X}\widehat{\mathcal{B}}\|_2^2 + \lambda\|\widehat{u} + \mathcal{B}\|_1 = \|\mathcal{E} - \mathcal{X}\widehat{u}\|_2^2 + \lambda\|\widehat{u} + \mathcal{B}\|_1 \\ &= \|\mathcal{E}\|_2^2 - 2\widehat{u}'\mathcal{X}^\top\mathcal{E} + \widehat{u}'\mathcal{X}^\top\mathcal{X}\widehat{u} + \lambda\|\widehat{u} + \mathcal{B}\|_1.\end{aligned}$$

Thus,

$$\widehat{u} = \text{Argmin}_u\, V(u),$$

where

$$V(u) = -2(\sqrt{nq}u)^\top W + (\sqrt{nq}u)^\top C(\sqrt{nq}u) + \lambda\|u + \mathcal{B}\|_1.$$

Since the first derivative of V with respect to u is equal to

$$2\sqrt{nq}\{C(\sqrt{nq}u) - W\} + \lambda \operatorname{sign}(u + \mathcal{B}),$$

$\widehat{u}$ satisfies

$$C_{J,J}(\sqrt{nq}\,\widehat{u}_J) - W_J = -\frac{\lambda}{2\sqrt{nq}} \operatorname{sign}(\widehat{u}_J + \mathcal{B}_J) = -\frac{\lambda}{2\sqrt{nq}} \operatorname{sign}(\widehat{\mathcal{B}}_J)$$

if $\widehat{u}_J + \mathcal{B}_J = \widehat{\mathcal{B}}_J \neq 0$, and

$$|C_{J^c,J}(\sqrt{nq}\,\widehat{u}_J) - W_{J^c}| \leq \frac{\lambda}{2\sqrt{nq}}.$$

Note that, if $|\widehat{u}_J| < |\mathcal{B}_J|$, then $\widehat{\mathcal{B}}_J \neq 0$ and $\operatorname{sign}(\widehat{\mathcal{B}}_J) = \operatorname{sign}(\mathcal{B}_J)$.

Let us now prove that when A_n and B_n, defined in (2.12) and (2.13), are satisfied then there exists $\widehat{u}$ satisfying

$$C_{J,J}(\sqrt{nq}\,\widehat{u}_J) - W_J = -\frac{\lambda}{2\sqrt{nq}} \operatorname{sign}(\mathcal{B}_J), \tag{2.21}$$

$$|\widehat{u}_J| < |\mathcal{B}_J|, \tag{2.22}$$

$$\left|C_{J^c,J}(\sqrt{nq}\,\widehat{u}_J) - W_{J^c}\right| \leq \frac{\lambda}{2\sqrt{nq}}. \tag{2.23}$$

Note that A_n implies

$$\begin{aligned}\sqrt{nq}\left\{-|\mathcal{B}_J| + \frac{\lambda}{2nq}(C_{J,J})^{-1}\operatorname{sign}(\mathcal{B}_J)\right\} &< (C_{J,J})^{-1}W_J \\ &< \sqrt{nq}\left\{|\mathcal{B}_J| + \frac{\lambda}{2nq}(C_{J,J})^{-1}\operatorname{sign}(\mathcal{B}_J)\right\}.\end{aligned} \tag{2.24}$$

By denoting

$$\widehat{u}_J = \frac{1}{\sqrt{nq}}(C_{J,J})^{-1}W_J - \frac{\lambda}{2nq}(C_{J,J})^{-1}\operatorname{sign}(\mathcal{B}_J), \tag{2.25}$$

we obtain from (2.24) that (2.21) and (2.22) hold. Note that B_n implies

$$\begin{aligned}-\frac{\lambda}{2\sqrt{nq}}\{\mathbf{1} - C_{J^c,J}(C_{J,J})^{-1}\operatorname{sign}(\mathcal{B}_J)\} &\leq C_{J^c,J}(C_{J,J})^{-1}W_J - W_{J^c} \\ &\leq \frac{\lambda}{2\sqrt{nq}}\{\mathbf{1} + C_{J^c,J}(C_{J,J})^{-1}\operatorname{sign}(\mathcal{B}_J)\}.\end{aligned}$$

Hence,

$$\left| C_{J^c,J} \left\{ (C_{J,J})^{-1} W_J - \frac{\lambda}{2\sqrt{nq}} (C_{J,J})^{-1} \mathrm{sign}(\mathcal{B}_J) \right\} - W_{J^c} \right| \leq \frac{\lambda}{2\sqrt{nq}},$$

which is (2.23) by (2.25). This concludes the proof. □

Proof of Theorem 2.1. By Proposition 2.2,

$$\mathbb{P}[\mathrm{sign}\{\widehat{\mathcal{B}}(\lambda)\} = \mathrm{sign}(\mathcal{B})] \geq \mathbb{P}(A_n \cap B_n) = 1 - \mathbb{P}(A_n^c \cup B_n^c) \geq 1 - \mathbb{P}(A_n^c) - \mathbb{P}(B_n^c),$$

where A_n and B_n are defined in (2.12) and (2.13). It is thus enough to prove that $\mathbb{P}(A_n^c)$ and $\mathbb{P}(B_n^c)$ tend to zero as $n \to \infty$. By definition of A_n,

$$\begin{aligned} \mathbb{P}(A_n^c) &= \mathbb{P}\left[\left| (C_{J,J})^{-1} W_J \right| \geq \sqrt{nq} \left\{ |\mathcal{B}_J| - \frac{\lambda}{2nq} |(C_{J,J})^{-1} \mathrm{sign}(\mathcal{B}_J)| \right\} \right] \\ &\leq \sup_{j \in J} \mathbb{P}\left\{ |\xi_j| \geq \sqrt{nq} \left(|\mathcal{B}_j| - \frac{\lambda}{2nq} |b_j| \right) \right\}, \end{aligned}$$

where

$$\xi = (\xi_j)_{j \in J} = (C_{J,J})^{-1} W_J = \frac{1}{\sqrt{nq}} (C_{J,J})^{-1} (\mathcal{X}_{\bullet,J})^\top \mathcal{E} \equiv H_A \, \mathcal{E},$$

and $b = (b_j)_{j \in J} = (C_{J,J})^{-1} \mathrm{sign}(\mathcal{B}_J)$. By definition of B_n and (IC),

$$\begin{aligned} \mathbb{P}(B_n^c) &= \mathbb{P}\left[|C_{J^c,J}(C_{J,J})^{-1} W_J - W_{J^c}| > \frac{\lambda}{2\sqrt{nq}} \left\{ \mathbf{1} - |C_{J^c,J}(C_{J,J})^{-1} \mathrm{sign}(\mathcal{B}_J)| \right\} \right] \\ &\leq \mathbb{P}\left\{ |C_{J^c,J}(C_{J,J})^{-1} W_J - W_{J^c}| > \frac{\lambda}{2\sqrt{nq}} \eta \right\} \\ &\leq \sup_{j \in J^c} \mathbb{P}\left(|\zeta_j| > \frac{\lambda}{2\sqrt{nq}} \eta \right), \end{aligned}$$

where

$$\begin{aligned} \zeta = (\zeta_j)_{j \in J^c} &= C_{J^c,J}(C_{J,J})^{-1} W_J - W_{J^c} \\ &= \frac{1}{\sqrt{nq}} \{ C_{J^c,J}(C_{J,J})^{-1} (\mathcal{X}_{\bullet,J})^\top - (\mathcal{X}_{\bullet,J^c})^\top \} \mathcal{E} \\ &\equiv H_B \, \mathcal{E}. \end{aligned}$$

Note that, for all $j \in J$,

$$|b_j| \leq \sum_{j \in J} |b_j| \leq \sqrt{|J|} \left(\sum_{j \in J} b_j^2 \right)^{1/2} = \sqrt{|J|} \, \|b\|_2.$$

Moreover,

$$\|b\|_2 = \|(C_{J,J})^{-1}\text{sign}(\mathcal{B}_J)\|_2 \leq \|(C_{J,J})^{-1}\|_2\sqrt{|J|} \equiv \lambda_{\max}\{(C_{J,J})^{-1}\}\sqrt{|J|},$$

where $\lambda_{\max}(A)$ denotes the largest eigenvalue of the matrix A. Observe that

$$\lambda_{\max}\{(C_{J,J})^{-1}\} = \frac{1}{\lambda_{\min}(C_{J,J})} = \frac{q}{\lambda_{\min}\{(\mathcal{X}^\top\mathcal{X})_{J,J}\}/n} \leq \frac{q}{M_2}, \tag{2.26}$$

by Assumption (A2) of Theorem 2.1. Thus, for all $j \in J$, $|b_j| \leq q|J|/M_2$. By Assumption (A4) of Theorem 2.1, we get thus that for all $j \in J$,

$$\sqrt{nq}\left[|\mathcal{B}_j| - \frac{\lambda}{2nq}\left|\{(C_{J,J})^{-1}\text{sign}(\mathcal{B}_J)\}_j\right|\right] = \sqrt{nq}\left(|\mathcal{B}_j| - \frac{\lambda}{2nq}|b_j|\right) \tag{2.27}$$

$$\geq \sqrt{nq}\left(M_3 q^{-c_2} - \frac{\lambda q|J|}{2nqM_2}\right). \tag{2.28}$$

Thus,

$$\mathbb{P}(A_n^c) \leq \sup_{j\in J}\mathbb{P}\left\{|\xi_j| \geq \sqrt{nq}\left(M_3 q^{-c_2} - \frac{\lambda q|J|}{2nqM_2}\right)\right\}.$$

Since $\mathcal{E}$ is a centered Gaussian random vector having a covariance matrix equal to identity, $\xi = H_A\,\mathcal{E}$ is a centered Gaussian random vector with a covariance matrix equal to:

$$H_A H_A^\top = \frac{1}{nq}(C_{J,J})^{-1}(\mathcal{X}_{\bullet,J})^\top\mathcal{X}_{\bullet,J}(C_{J,J})^{-1} = (C_{J,J})^{-1}.$$

Hence, by (2.26), we get that for all $j \in J$, $\text{Var}(\xi_j) = ((C_{J,J})^{-1})_{jj} \leq \lambda_{\max}(C_{J,J}^{-1}) \leq q/M_2$. Thus,

$$\mathbb{P}\left\{|\xi_j| \geq \sqrt{nq}\left(M_3 q^{-c_2} - \frac{\lambda q|J|}{2nqM_2}\right)\right\} \leq \mathbb{P}\left\{|Z| \geq \frac{\sqrt{M_2}}{\sqrt{q}}\left(M_3 q^{-c_2}\sqrt{nq} - \frac{\lambda q|J|}{2\sqrt{nq}M_2}\right)\right\},$$

where Z is a standard Gaussian random variable. By Chernoff's inequality, we thus obtain that for all $j \in J$,

$$\mathbb{P}\left\{|\xi_j| \geq \sqrt{nq}\left(M_3 q^{-c_2} - \frac{\lambda q|J|}{2nqM_2}\right)\right\} \leq 2\exp\left\{-\frac{M_2}{2q}\left(M_3 q^{-c_2}\sqrt{nq} - \frac{\lambda q|J|}{2\sqrt{nq}M_2}\right)^2\right\}.$$

By Assumption (A3) of Theorem 2.1, we get that under the last condition of (L), as $n \to \infty$,

$$\lambda q|J|/\sqrt{nq} = o\left(q^{-c_2}\sqrt{nq}\right), \tag{2.29}$$

Thus,

$$\lim_{n\to\infty}\mathbb{P}(A_n^c) = 0. \tag{2.30}$$

Let us now bound $\mathbb{P}(B_n^c)$. Observe that $\zeta = H_B\,\mathcal{E}$ is a centered Gaussian random vector with a covariance matrix equal to

$$\begin{aligned} H_B H_B^\top &= \frac{1}{nq}\{C_{J^c,J}(C_{J,J})^{-1}(\mathcal{X}_{\bullet,J})^\top - \mathcal{X}_{\bullet,J^c}^\top\}\{\mathcal{X}_{\bullet,J}(C_{J,J})^{-1}C_{J,J^c} - \mathcal{X}_{\bullet,J^c}\} \\ &= C_{J^c,J^c} - C_{J^c,J}(C_{J,J})^{-1}C_{J,J^c} \\ &= \frac{1}{nq}(\mathcal{X}_{\bullet,J^c})^\top[\mathrm{Id}_{\mathbb{R}^{nq}} - \mathcal{X}_{\bullet,J}\{(\mathcal{X}_{\bullet,J})^\top\mathcal{X}_{\bullet,J}\}^{-1}(\mathcal{X}_{\bullet,J})^\top]\mathcal{X}_{\bullet,J^c} \\ &= \frac{1}{nq}(\mathcal{X}_{\bullet,J^c})^\top\left(\mathrm{Id}_{\mathbb{R}^{nq}} - \Pi_{\mathrm{Im}(\mathcal{X}_{\bullet,J})}\right)\mathcal{X}_{\bullet,J^c}, \end{aligned}$$

where $\Pi_{\mathrm{Im}(\mathcal{X}_{\bullet,J})}$ denotes the orthogonal projection onto the column space of $\mathcal{X}_{\bullet,J}$. Note that, for all $j \in J^c$,

$$\begin{aligned} \mathrm{Var}(\zeta_j) &= \frac{1}{nq}((\mathcal{X}_{\bullet,J^c})^\top\left(\mathrm{Id}_{\mathbb{R}^{nq}} - \Pi_{\mathrm{Im}(\mathcal{X}_{\bullet,J})}\right)\mathcal{X}_{\bullet,J^c})_{jj} \\ &= \frac{1}{nq}((\mathcal{X}_{\bullet,J^c})^\top\mathcal{X}_{\bullet,J^c})_{jj} - \frac{1}{nq}((\mathcal{X}_{\bullet,J^c})^\top\,\Pi_{\mathrm{Im}(\mathcal{X}_{\bullet,J})}\mathcal{X}_{\bullet,J^c})_{jj} \\ &\le \frac{1}{nq}((\mathcal{X}_{\bullet,J^c})^\top\mathcal{X}_{\bullet,J^c})_{jj} \le \frac{M_1}{q}, \end{aligned}$$

where the inequalities come from Lemma 2.9 and Assumption (A1) of Theorem 2.1. Thus, for all $j \in J^c$,

$$\mathbb{P}\left(|\zeta_j| > \frac{\lambda}{2\sqrt{nq}}\eta\right) \le \mathbb{P}\left(|Z| > \frac{\lambda\sqrt{q}}{2\sqrt{M_1}\sqrt{nq}}\eta\right),$$

where Z is a standard Gaussian random variable. By Chernoff's inequality, for all $j \in J^c$,

$$\mathbb{P}\left(|\zeta_j| > \frac{\lambda}{2\sqrt{nq}}\eta\right) \le 2\exp\left\{-\frac{1}{2}\left(\frac{\lambda}{2\sqrt{M_1}\sqrt{n}}\eta\right)^2\right\}.$$

Hence, under the assumption that $\lambda/\sqrt{n} \to \infty$, which is the second condition of (L),

$$\lim_{n\to\infty}\mathbb{P}(B_n^c) = 0. \tag{2.31}$$

This completes the argument. □

Proof of Proposition 2.3. Let us first prove that (C1) and (C3) imply (A1). For $j \in \{1,\dots,pq\}$, by considering the Euclidian division of $j-1$ by p given by $j-1 =$

$pk_j + r_j$, we observe that

$$\begin{aligned}
(\mathcal{X}_{\bullet,j})^\top \mathcal{X}_{\bullet,j} &= \{((\Sigma^{-1/2})^\top \otimes X)_{\bullet,j}\}^\top ((\Sigma^{-1/2})^\top \otimes X)_{\bullet,j} \\
&= \{(\Sigma^{-1/2}) \otimes X^\top)_{j,\bullet}\}((\Sigma^{-1/2})^\top \otimes X)_{\bullet,j} \\
&= \{(\Sigma^{-1/2})_{k_j+1,\bullet} \otimes (X_{\bullet,r_j+1})^\top\}[\{(\Sigma^{-1/2})_{\bullet,k_j+1}\}^\top \otimes X_{\bullet,r_j+1}\}] \\
&= (\Sigma^{-1/2})_{k_j+1,\bullet}\{(\Sigma^{-1/2})_{\bullet,k_j+1}\}^\top \otimes (X_{\bullet,r_j+1})^\top X_{\bullet,r_j+1} \\
&= (\Sigma^{-1})_{k_j+1,k_j+1} \otimes (X_{\bullet,r_j+1})^\top X_{\bullet,r_j+1} \\
&= (\Sigma^{-1})_{k_j+1,k_j+1}(X_{\bullet,r_j+1})^\top X_{\bullet,r_j+1}.
\end{aligned}$$

Hence, using (C1), we get that for all $j \in \{1, \ldots, pq\}$,

$$\begin{aligned}
\frac{1}{n}(\mathcal{X}_{\bullet,j})^\top \mathcal{X}_{\bullet,j} &\le M_1'(\Sigma^{-1})_{k_j+1,k_j+1} \\
&\le M_1' \sup_{k\in\{0,\ldots,q-1\}} \{(\Sigma^{-1})_{k+1,k+1}\} \le M_1'\lambda_{\max}(\Sigma^{-1}) \le M_1' m_1,
\end{aligned}$$

where the last inequality comes from (C3), which gives (A1).

Let us now prove that (C2) and (C4) imply (A2). Note that

$$\begin{aligned}
(\mathcal{X}^\top \mathcal{X})_{J,J} &= \{((\Sigma^{-1/2})^\top \otimes X\}^\top ((\Sigma^{-1/2})^\top \otimes X))_{J,J} \\
&= (\Sigma^{-1/2}(\Sigma^{-1/2})^\top \otimes X^\top X)_{J,J} \\
&= (\Sigma^{-1} \otimes X^\top X)_{J,J}.
\end{aligned}$$

Then, by Theorem 4.3.15 of Horn & Johnson (1986),

$$\begin{aligned}
\lambda_{\min}\{(\mathcal{X}^\top \mathcal{X})_{J,J}\} &= \lambda_{\min}\{(\Sigma^{-1} \otimes X^\top X)_{J,J}\} \\
&\ge \lambda_{\min}(\Sigma^{-1} \otimes X^\top X) \\
&= \lambda_{\min}(X^\top X)\lambda_{\min}(\Sigma^{-1}).
\end{aligned}$$

Finally, by using Conditions (C2) and (C4), we obtain

$$\frac{1}{n}\lambda_{\min}(\mathcal{X}^\top \mathcal{X})_{J,J} \ge \frac{1}{n}\lambda_{\min}(X^\top X)\lambda_{\min}(\Sigma^{-1}) \ge M_2' m_2,$$

which gives (A2). □

Proof of Theorem 2.5. By Proposition 2.4,

$$\mathbb{P}[\operatorname{sign}\{\widetilde{\mathcal{B}}(\lambda)\} = \operatorname{sign}(\mathcal{B})] \ge \mathbb{P}(\widetilde{A}_n \cap \widetilde{B}_n) = 1 - \mathbb{P}(\widetilde{A}_n^c \cup \widetilde{B}_n^c) \ge 1 - \mathbb{P}(\widetilde{A}_n^c) - \mathbb{P}(\widetilde{B}_n^c),$$

where $\widetilde{A}_n$ and $\widetilde{B}_n$ are defined in (2.16) and (2.17). By definition of $\widetilde{A}_n$, we get

$$\mathbb{P}(\widetilde{A}_n^c) = \mathbb{P}\left[\left\{\left|(\widetilde{C}_{J,J})^{-1}\widetilde{W}_J\right| \ge \sqrt{nq}\left(|\mathcal{B}_J| - \frac{\lambda}{2nq}|(\widetilde{C}_{J,J})^{-1}\operatorname{sign}(\mathcal{B}_J)|\right)\right\}\right].$$

Observing that

$$\begin{aligned}(\widetilde{C}_{J,J})^{-1}\widetilde{W}_J = (C_{J,J})^{-1}W_J &+ (C_{J,J})^{-1}(\widetilde{W}_J - W_J) \\ &+ \{(\widetilde{C}_{J,J})^{-1} - (C_{J,J})^{-1}\}W_J \\ &+ \{(\widetilde{C}_{J,J})^{-1} - (C_{J,J})^{-1}\}(\widetilde{W}_J - W_J),\end{aligned}$$

and

$$(\widetilde{C}_{J,J})^{-1}\mathrm{sign}(\mathcal{B}_J) = (C_{J,J})^{-1}\mathrm{sign}(\mathcal{B}_J) + \{(\widetilde{C}_{J,J})^{-1} - (C_{J,J})^{-1}\}\mathrm{sign}(\mathcal{B}_J),$$

and using the triangle inequality, we obtain

$$\begin{aligned}\mathbb{P}(\widetilde{A}_n^c) \leq\ & \mathbb{P}\left[\left\{\left|(C_{J,J})^{-1}W_J\right| \geq \frac{\sqrt{nq}}{5}\left(|\mathcal{B}_J| - \frac{\lambda}{2nq}\left|(C_{J,J})^{-1}\mathrm{sign}(\mathcal{B}_J)\right|\right)\right\}\right] \\ &+ \mathbb{P}\left[\left\{\left|(C_{J,J})^{-1}\left(\widetilde{W}_J - W_J\right]\right| \geq \frac{\sqrt{nq}}{5}\left(|\mathcal{B}_J| - \frac{\lambda}{2nq}\left|(C_{J,J})^{-1}\mathrm{sign}(\mathcal{B}_J)\right|\right)\right\}\right) \\ &+ \mathbb{P}\left[\left\{\left|\left((\widetilde{C}_{J,J})^{-1} - (C_{J,J})^{-1}\right)W_J\right| \geq \frac{\sqrt{nq}}{5}\left(|\mathcal{B}_J| - \frac{\lambda}{2nq}\left|(C_{J,J})^{-1}\mathrm{sign}(\mathcal{B}_J)\right|\right)\right\}\right] \\ &+ \mathbb{P}\left[\left\{\left|\left((\widetilde{C}_{J,J})^{-1} - (C_{J,J})^{-1}\right)\left(\widetilde{W}_J - W_J\right)\right| \geq \frac{\sqrt{nq}}{5}\left(|\mathcal{B}_J| - \frac{\lambda}{2nq}\left|(C_{J,J})^{-1}\mathrm{sign}(\mathcal{B}_J)\right|\right)\right\}\right] \\ &+ \mathbb{P}\left[\left\{\frac{\lambda}{2\sqrt{nq}}\left|\left((\widetilde{C}_{J,J})^{-1} - (C_{J,J})^{-1}\right)\mathrm{sign}(\mathcal{B}_J)\right| \geq \frac{\sqrt{nq}}{5}\left(|\mathcal{B}_J| - \frac{\lambda}{2nq}\left|(C_{J,J})^{-1}\mathrm{sign}(\mathcal{B}_J)\right|\right)\right\}\right]\end{aligned} \tag{2.32}$$

The first term in the right-hand side of (2.32) tends to 0 by the definition of A_n^c and (2.30). By (2.27), the last term of (2.32) satisfies, for all $j \in J$,

$$\begin{aligned}&\mathbb{P}\left[\left|\{(\widetilde{C}_{J,J})^{-1} - (C_{J,J})^{-1}\}\mathrm{sign}(\mathcal{B}_J)\right| \geq \frac{2nq}{5\lambda}\left\{|\mathcal{B}_J| - \frac{\lambda}{2nq}|(C_{J,J})^{-1}\mathrm{sign}(\mathcal{B}_J)|\right\}\right] \\ &\quad\leq \mathbb{P}\left[\left|\left(\left\{(\widetilde{C}_{J,J})^{-1} - (C_{J,J})^{-1}\right\}\mathrm{sign}(\mathcal{B}_J)\right)_j\right| \geq \frac{2nq}{5\lambda}\left\{M_3 q^{-c_2} - \frac{\lambda q|J|}{2nqM_2}\right)\right].\end{aligned}$$

Let $U = (\widetilde{C}_{J,J})^{-1} - (C_{J,J})^{-1}$ and $s = \mathrm{sign}(\mathcal{B}_J)$ then for all $j \in J$,

$$|(Us)_j| = \left|\sum_{k\in J} U_{jk}s_k\right| \leq \sqrt{|J|}\|U\|_2. \tag{2.33}$$

We focus on

$$\begin{aligned}\|(\widetilde{C}_{J,J})^{-1} - (C_{J,J})^{-1}\|_2 &= \|(\widetilde{C}_{J,J})^{-1}(C_{J,J} - \widetilde{C}_{J,J})(C_{J,J})^{-1}\|_2 \\ &\leq \|(\widetilde{C}_{J,J})^{-1}\|_2 \, \|C_{J,J} - \widetilde{C}_{J,J}\|_2 \, \|(C_{J,J})^{-1}\|_2 \\ &\leq \frac{\rho(C_{J,J} - \widetilde{C}_{J,J})}{\lambda_{\min}(\widetilde{C}_{J,J})\lambda_{\min}(C_{JJ})} \\ &\leq \frac{\rho(C_{J,J} - \widetilde{C}_{J,J})}{\lambda_{\min}(\widetilde{C}_{J,J})(M_2/q)},\end{aligned}$$

where the last inequality comes from Assumption (A2) of Theorem 2.1, which gives

$$\|(C_{J,J})^{-1}\|_2 \leq q/M_2. \tag{2.34}$$

Using Theorem 4.3.15 of Horn & Johnson (1986), we get

$$\|(\widetilde{C}_{J,J})^{-1} - (C_{J,J})^{-1}\|_2 \leq \frac{q\rho(C - \widetilde{C})}{\lambda_{\min}(\widetilde{C})M_2}.$$

By definition of C and $\widetilde{C}$ given in (2.10) and (2.14), respectively, we get

$$C = \frac{\Sigma^{-1} \otimes (X^\top X)}{nq} \quad \text{and} \quad \widetilde{C} = \frac{\widehat{\Sigma}^{-1} \otimes (X^\top X)}{nq}. \tag{2.35}$$

By using that the eigenvalues of the Kronecker product of two matrices is equal to the product of the eigenvalues of the two matrices, we obtain

$$\begin{aligned}\|(\widetilde{C}_{J,J})^{-1} - (C_{J,J})^{-1}\|_2 &\leq \frac{\rho(\Sigma^{-1} - \widehat{\Sigma}^{-1})\lambda_{\max}\{(X^\top X)/n\}q}{\lambda_{\min}(\widehat{\Sigma}^{-1})\lambda_{\min}\{(X^\top X)/n\}M_2} \\ &\leq \frac{\rho(\Sigma^{-1} - \widehat{\Sigma}^{-1})\lambda_{\max}(\widehat{\Sigma})\lambda_{\max}\{(X^\top X)/n\}q}{\lambda_{\min}\{(X^\top X)/n\}M_2} \\ &\leq \frac{\rho(\Sigma^{-1} - \widehat{\Sigma}^{-1})\{\rho(\widehat{\Sigma} - \Sigma) + \lambda_{\max}(\Sigma)\}\lambda_{\max}\{(X^\top X)/n\}q}{\lambda_{\min}\{(X^\top X)/n\}M_2},\end{aligned}$$

where the last inequality follows from Theorem 4.3.1 of Horn & Johnson (1986). Thus, by Assumptions (A5), (A6), (A8), (A9) and (A10), we get that, as $n \to \infty$,

$$\|(\widetilde{C}_{J,J})^{-1} - (C_{J,J})^{-1}\|_2 = O_P\{q(nq)^{-1/2}\}. \tag{2.36}$$

Hence, by (2.33), we get that, for all $j \in J$,

$$\mathbb{P}\left[\left|\left(\{(\widetilde{C}_{J,J})^{-1} - (C_{J,J})^{-1}\}\mathrm{sign}(\mathcal{B}_J)\right)_j\right| \geq \frac{2nq}{5\lambda}\left(M_3 q^{-c_2} - \frac{\lambda q|J|}{2nqM_2}\right)\right]$$
$$\leq \mathbb{P}\left\{\sqrt{|J|}\,\|(\widetilde{C}_{J,J})^{-1} - (C_{J,J})^{-1}\|_2 \geq \frac{2\sqrt{nq}}{5\lambda}\left(M_3 q^{-c_2}\sqrt{nq} - \frac{\lambda q|J|}{2\sqrt{nq}M_2}\right)\right\}.$$

By (2.29), (2.36) and (A3), it is enough to prove that

$$\lim_{n\to\infty} \mathbb{P}\left\{q^{c_1/2} q (nq)^{-1/2} \geq nqq^{-c_2}\lambda\right\} = 0.$$

By the last condition of (L), $(nqq^{-c_2}/\lambda)/q^{1+c_1} \to \infty$ as $n \to \infty$ and the result follows since $n \to \infty$. Hence, the last term of (2.32) tends to zero as $n \to \infty$.

Let us now study the second term in the right-hand side of (2.32). We have

$$\widetilde{W}_J - W_J = \frac{1}{\sqrt{nq}}\left\{\left(\widetilde{\mathcal{X}}^\top\widetilde{\mathcal{E}}\right)_J - \left(\mathcal{X}^\top\mathcal{E}\right)_J\right\} = \frac{1}{\sqrt{nq}}\left(\widetilde{\mathcal{X}}^\top\widetilde{\mathcal{E}} - \mathcal{X}^\top\mathcal{E}\right)_J$$
$$= \frac{1}{\sqrt{nq}}\left((\widehat{\Sigma}^{-1/2}\otimes X^\top)\left\{(\widehat{\Sigma}^{-1/2})^\top \otimes \mathrm{Id}_{\mathbb{R}^n}\right\}\mathrm{vec}(E) - (\Sigma^{-1/2}\otimes X^\top)\left\{(\Sigma^{-1/2})^\top \otimes \mathrm{Id}_{\mathbb{R}^n}\right\}\mathrm{vec}(E)\right)_J$$
$$= \frac{1}{\sqrt{nq}}\left(\left\{(\widehat{\Sigma}^{-1} - \Sigma^{-1})\otimes X^\top\right\}\mathrm{vec}(E)\right)_J \stackrel{d}{=} AZ, \tag{2.37}$$

where Z is a centered Gaussian random vector having a covariance matrix equal to the identity and

$$A = \frac{1}{\sqrt{nq}}\left[\left\{(\widehat{\Sigma}^{-1} - \Sigma^{-1})\otimes X^\top\right\}\left\{(\Sigma^{1/2})^\top \otimes \mathrm{Id}_{\mathbb{R}^n}\right\}\right]_{J,\bullet}. \tag{2.38}$$

By the Cauchy–Schwarz inequality, we get for every $K \times nq$ matrix B, every $nq \times 1$ vector U, and every $k \in \{1, \ldots, K\}$,

$$|(BU)_k| = \left|\sum_{\ell=1}^{nq} B_{k,\ell}U_\ell\right| \leq \|B\|_2\,\|U\|_2. \tag{2.39}$$

Thus, for all $j \in J$, for all γ in $\mathbb{R}$ and every $|J| \times |J|$ matrix D,

$$\mathbb{P}\left\{\left|\left(D\left(\widetilde{W}_J - W_J\right)\right)_j\right| \geq \gamma\right\} = \mathbb{P}\left\{\left|(DAZ)_j\right| \geq \gamma\right\} \leq \mathbb{P}\left(\|D\|_2\,\|A\|_2\|Z\|_2 \geq \gamma\right), \tag{2.40}$$

where A is defined in (2.38) and Z is a centered Gaussian random vector having a

covariance matrix equal to the identity. Hence, for all $j \in J$,

$$\mathbb{P}\left[\left|\left((C_{J,J})^{-1}(\widetilde{W}_J - W_J)\right)_j\right| \geq \frac{\sqrt{nq}}{5}\left\{|\mathcal{B}_j| - \frac{\lambda}{2nq}\left|\left((C_{J,J})^{-1}\mathrm{sign}(\mathcal{B}_J)\right)_j\right|\right\}\right]$$
$$\leq \mathbb{P}\left\{\|(C_{J,J})^{-1}\|_2\|A\|_2\|Z\|_2 \geq \frac{\sqrt{nq}}{5}\left(M_3 q^{-c_2} - \frac{\lambda q|J|}{2nqM_2}\right)\right\}.$$

Let us bound $\|A\|_2$. Observe that

$$\left\|\left[\left\{(\widehat{\Sigma}^{-1} - \Sigma^{-1}) \otimes X^\top\right\}\left\{(\Sigma^{1/2})^\top \otimes \mathrm{Id}\right\}\right]_{J,\bullet}\right\|_2 \tag{2.41}$$
$$= \rho\left\{\left((\widehat{\Sigma}^{-1} - \Sigma^{-1})\Sigma(\widehat{\Sigma}^{-1} - \Sigma^{-1}) \otimes (X^\top X)\right)_{J,J}\right\}^{1/2}$$
$$\leq \rho\left\{(\widehat{\Sigma}^{-1} - \Sigma^{-1})\Sigma(\widehat{\Sigma}^{-1} - \Sigma^{-1})\right\}^{1/2}\lambda_{\max}\left(X^\top X\right)^{1/2}$$
$$\leq \rho(\widehat{\Sigma}^{-1} - \Sigma^{-1})\lambda_{\max}(\Sigma)^{1/2}\lambda_{\max}(X^\top X)^{1/2}, \tag{2.42}$$

where the first inequality comes from Theorem 4.3.15 of Horn & Johnson (1986). Hence, by (A5), (A8) and (A9)

$$\|A\|_2 = \frac{1}{\sqrt{nq}}\left\|\left(\left\{(\widehat{\Sigma}^{-1} - \Sigma^{-1}) \otimes X^\top\right\}\left\{(\Sigma^{1/2})^\top \otimes \mathrm{Id}\right\}\right)_{J,\bullet}\right\|_2 = O_P\{q^{-1/2}(nq)^{-1/2}\}. \tag{2.43}$$

By (2.29), (2.34) and (2.43), it is enough to prove that

$$\lim_{n\to\infty}\mathbb{P}\left(\sum_{k=1}^{nq} Z_k^2 \geq nq\, n\, q^{-2c_2}\right) = 0.$$

The result follows from Markov's inequality and the first condition of (L).

Let us now study the 3rd term in the right-hand side of (2.32). Observe that

$$W_J = \frac{1}{\sqrt{nq}}\left((\Sigma^{-1/2} \otimes X^\top)\left\{(\Sigma^{-1/2})^\top \otimes \mathrm{Id}_{\mathbb{R}^n}\right\}\mathrm{vec}(E)\right)_J$$
$$\stackrel{d}{=} \frac{1}{\sqrt{nq}}\left[(\Sigma^{-1} \otimes X^\top)\left\{(\Sigma^{1/2})^\top \otimes \mathrm{Id}_{\mathbb{R}^n}\right\}\right]_{J,\bullet} Z \equiv A_1 Z,$$

where Z is a centered Gaussian random vector having a covariance matrix equal to identity and

$$A_1 = \frac{1}{\sqrt{nq}}\left[(\Sigma^{-1} \otimes X^\top)\left\{(\Sigma^{1/2})^\top \otimes \mathrm{Id}_{\mathbb{R}^n}\right\}\right]_{J,\bullet}. \tag{2.44}$$

Using (2.39), we get for every $j \in J$, every $\gamma \in \mathbb{R}$, and every $|J| \times |J|$ matrix D,

$$\mathbb{P}\{|\,(D\,W_J)_j\,| \geq \gamma\} = \mathbb{P}\{|\,(DA_1Z)_j\,| \geq \gamma\} \leq \mathbb{P}\left(\|D\|_2\,\|A_1\|_2\,\|Z\|_2 \geq \gamma\right), \tag{2.45}$$

where A_1 is defined in (2.44) and Z is a centered Gaussian random vector having a covariance matrix equal to identity. Hence, for all $j \in J$,

$$\mathbb{P}\left[\left|\left(\{(\widetilde{C}_{J,J})^{-1} - (C_{J,J})^{-1}\}W_J\right)_j\right| \geq \frac{\sqrt{nq}}{5}\left\{|\mathcal{B}_j| - \frac{\lambda}{2nq}\left|\left((C_{J,J})^{-1}\mathrm{sign}(\mathcal{B}_J)\right)_j\right|\right\}\right]$$
$$\leq \mathbb{P}\left\{\left\|(\widetilde{C}_{J,J})^{-1} - (C_{J,J})^{-1}\right\|_2 \|A_1\|_2 \|Z\|_2 \geq \frac{\sqrt{nq}}{5}\left(M_3 q^{-c_2} - \frac{\lambda q|J|}{2nqM_2}\right)\right\}.$$

Let us now bound $\|A_1\|_2$. Note that

$$\left\|\left[(\Sigma^{-1}\otimes X^\top)\{(\Sigma^{1/2})^\top \otimes \mathrm{Id}_{\mathbb{R}^n}\}\right]_{J,\bullet}\right\|_2 = \left\|\left[(\Sigma^{-1/2}\otimes X^\top)\right]_{J,\bullet}\right\|_2$$
$$= \rho\left\{\left(\Sigma^{-1}\otimes(X^\top X)\right)_{J,J}\right\}^{1/2} \leq \rho[\{\Sigma^{-1}\otimes(X^\top X)\}]^{1/2} \leq \lambda_{\max}(\Sigma^{-1})^{1/2}\lambda_{\max}(X^\top X)^{1/2},$$

where the first inequality comes from Theorem 4.3.15 of Horn & Johnson (1986). Hence, by (A5) and (A7),

$$\|A_1\|_2 \leq \frac{1}{nq}\lambda_{\max}(\Sigma^{-1})^{1/2}\lambda_{\max}(X^\top X)^{1/2} = O_P(q^{-1/2}). \tag{2.46}$$

By (2.29), (2.36) and (2.46) it is thus enough to prove that

$$\lim_{n\to\infty}\mathbb{P}\left(\sum_{k=1}^{nq} Z_k^2 \geq nq\, n\, q^{-2c_2}\right) = 0.$$

The result follows from Markov's inequality and the first condition of (L).

Let us now study the 4th term in the right-hand side of (2.32). By (2.40), for all $j \in J$,

$$\mathbb{P}\left[\left|\left(\left\{(\widetilde{C}_{J,J})^{-1} - C_{J,J})^{-1}\right\}\left(\widetilde{W}_J - W_J\right)\right)_j\right| \geq \frac{\sqrt{nq}}{5}\left\{|\mathcal{B}_j| - \frac{\lambda}{2nq}\left|\left((C_{J,J})^{-1}\mathrm{sign}(\mathcal{B}_J)\right)_j\right|\right\}\right]$$
$$\leq \mathbb{P}\left\{\left\|(\widetilde{C}_{J,J})^{-1} - C_{J,J})^{-1}\right\|_2 \|A\|_2 \|Z\|_2 \geq \frac{\sqrt{nq}}{5}\left(M_3 q^{-c_2} - \frac{\lambda q|J|}{2nqM_2}\right)\right\},$$

where A is defined in (2.38). By (2.29), (2.36) and (2.43), it is thus enough to prove that

$$\lim_{n\to\infty}\mathbb{P}\left\{\sum_{k=1}^{nq} Z_k^2 \geq (nq)\, n^2\, q^{1-2c_2}\right\} = 0.$$

The result follows from the Markov inequality and the fact that $c_2 < 1/2$.

Let us now study $\mathbb{P}(\widetilde{B}_n)$. By definition of $\widetilde{B}_n$, we get that

$$\mathbb{P}(\widetilde{B}_n^c) = \mathbb{P}\left[\left\{|\widetilde{C}_{J^c,J}(\widetilde{C}_{J,J})^{-1}\widetilde{W}_J - \widetilde{W}_{J^c}| \geq \frac{\lambda}{2\sqrt{nq}}\left(1 - |\widetilde{C}_{J^c,J}(\widetilde{C}_{J,J})^{-1}\mathrm{sign}(\mathcal{B}_J)|\right)\right\}\right].$$

Observe that

$$\begin{aligned}\widetilde{C}_{J^c,J}(\widetilde{C}_{J,J})^{-1}\widetilde{W}_J - \widetilde{W}_{J^c} &= C_{J^c,J}(C_{J,J})^{-1}W_J - W_{J^c} + C_{J^c,J}(C_{J,J})^{-1}(\widetilde{W}_J - W_J)\\ &+ C_{J^c,J}\{(\widetilde{C}_{J,J})^{-1} - (C_{J,J})^{-1}\}W_J + C_{J^c,J}\{(\widetilde{C}_{J,J})^{-1} - (C_{J,J})^{-1}\}(\widetilde{W}_J - W_J)\\ &+ (\widetilde{C}_{J^c,J} - C_{J^c,J})(C_{J,J})^{-1}W_J + (\widetilde{C}_{J^c,J} - C_{J^c,J})(C_{J,J})^{-1}(\widetilde{W}_J - W_J)\\ &+ (\widetilde{C}_{J^c,J} - C_{J^c,J})\{(\widetilde{C}_{J,J})^{-1} - (C_{J,J})^{-1}\}W_J\\ &+ (\widetilde{C}_{J^c,J} - C_{J^c,J})\{(\widetilde{C}_{J,J})^{-1} - (C_{J,J})^{-1}\}(\widetilde{W}_J - W_J) + W_{J^c} - \widetilde{W}_{J^c}.\end{aligned}$$

Moreover,

$$\begin{aligned}&\widetilde{C}_{J^c,J}(\widetilde{C}_{J,J})^{-1}\mathrm{sign}(\mathcal{B}_J) = C_{J^c,J}(C_{J,J})^{-1}\mathrm{sign}(\mathcal{B}_J) + C_{J^c,J}\{(\widetilde{C}_{J,J})^{-1} - (C_{J,J})^{-1}\}\mathrm{sign}(\mathcal{B}_J)\\ &+(\widetilde{C}_{J^c,J} - C_{J^c,J})(C_{J,J})^{-1}\mathrm{sign}(\mathcal{B}_J) + (\widetilde{C}_{J^c,J} - C_{J^c,J})\{(\widetilde{C}_{J,J})^{-1} - (C_{J,J})^{-1}\}\mathrm{sign}(\mathcal{B}_J).\end{aligned}$$

By (IC) and the triangle inequality, we obtain that

$$\begin{aligned}\mathbb{P}(\widetilde{B}_n^c) \le\ &\mathbb{P}\left(\left|C_{J^c,J}(C_{J,J})^{-1}W_J - W_{J^c}\right| \ge \frac{\lambda}{24\sqrt{nq}}\eta\right)\\ &+ \mathbb{P}\left\{|C_{J^c,J}(C_{J,J})^{-1}(\widetilde{W}_J - W_J)| \ge \frac{\lambda}{24\sqrt{nq}}\eta\right\}\\ &+ \mathbb{P}\left[|C_{J^c,J}\{(\widetilde{C}_{J,J})^{-1} - (C_{J,J})^{-1}\}W_J| \ge \frac{\lambda}{24\sqrt{nq}}\eta\right]\\ &+ \mathbb{P}\left[|C_{J^c,J}\{(\widetilde{C}_{J,J})^{-1} - (C_{J,J})^{-1}\}(\widetilde{W}_J - W_J)| \ge \frac{\lambda}{24\sqrt{nq}}\eta\right]\\ &+ \mathbb{P}\left\{|(\widetilde{C}_{J^c,J} - C_{J^c,J})(C_{J,J})^{-1}W_J| \ge \frac{\lambda}{24\sqrt{nq}}\eta\right\}\\ &+ \mathbb{P}\left\{|(\widetilde{C}_{J^c,J} - C_{J^c,J})(C_{J,J})^{-1}(\widetilde{W}_J - W_J)| \ge \frac{\lambda}{24\sqrt{nq}}\eta\right\}\\ &+ \mathbb{P}\left\{|(\widetilde{C}_{J^c,J} - C_{J^c,J})\{(\widetilde{C}_{J,J})^{-1} - (C_{J,J})^{-1}\}W_J| \ge \frac{\lambda}{24\sqrt{nq}}\eta\right\}\\ &+ \mathbb{P}\left[|(\widetilde{C}_{J^c,J} - C_{J^c,J})\{(\widetilde{C}_{J,J})^{-1} - (C_{J,J})^{-1}\}(\widetilde{W}_J - W_J)| \ge \frac{\lambda}{24\sqrt{nq}}\eta\right]\\ &+ \mathbb{P}\left\{|W_{J^c} - \widetilde{W}_{J^c}| \ge \frac{\lambda}{24\sqrt{nq}}\eta\right\}\\ &+ \mathbb{P}\left[|C_{J^c,J}\{(\widetilde{C}_{J,J})^{-1} - (C_{J,J})^{-1}\}\mathrm{sign}(\mathcal{B}_J)| \ge \frac{\eta}{12}\right]\\ &+ \mathbb{P}\left\{|(\widetilde{C}_{J^c,J} - C_{J^c,J})(C_{J,J})^{-1}\mathrm{sign}(\mathcal{B}_J)| \ge \frac{\eta}{12}\right\}\\ &+ \mathbb{P}\left[|(\widetilde{C}_{J^c,J} - C_{J^c,J})\{(\widetilde{C}_{J,J})^{-1} - (C_{J,J})^{-1}\}\mathrm{sign}(\mathcal{B}_J)| \ge \frac{\eta}{12}\right]. \qquad (2.47)\end{aligned}$$

The first term in the right-hand side of (2.47) tends to 0 by (2.31). Let us now

study the 2nd term of (2.47). By (2.40), we get that for all $j \in J^c$,

$$\mathbb{P}\left\{\left|\left(C_{J^c,J}(C_{J,J})^{-1}(\widetilde{W}_J - W_J)\right)_j\right| \geq \frac{\lambda}{24\sqrt{nq}}\eta\right\} \leq \mathbb{P}\left\{\|C_{J^c,J}\|_2 \, \|(C_{J,J})^{-1}\|_2 \, \|A\|_2 \, \|Z\|_2 \geq \frac{\lambda}{24\sqrt{nq}}\eta\right\}.$$

Observe that

$$\begin{aligned}\|C_{J^c,J}\|_2 &= \rho\left\{\frac{(\mathcal{X}_{\bullet,J^c})^\top \mathcal{X}_{\bullet,J}}{nq}\frac{(\mathcal{X}_{\bullet,J})^\top \mathcal{X}_{\bullet,J^c}}{nq}\right\}^{1/2} = \frac{1}{nq}\left\|(\mathcal{X}_{\bullet,J^c})^\top \mathcal{X}_{\bullet,J}\right\|_2 \leq \frac{\left\|(\mathcal{X}_{\bullet,J^c})^\top\right\|_2}{\sqrt{nq}}\frac{\|\mathcal{X}_{\bullet,J}\|_2}{\sqrt{nq}} \\ &\leq \rho\left\{\frac{(\mathcal{X}_{\bullet,J^c})^\top \mathcal{X}_{\bullet,J^c}}{nq}\right\}^{1/2}\rho\left\{\frac{(\mathcal{X}_{\bullet,J})^\top \mathcal{X}_{\bullet,J}}{nq}\right\}^{1/2} \\ &= \rho(C_{J^cJ^c})^{1/2}\rho(C_{J,J})^{1/2} \leq \rho(C) = \frac{\lambda_{\max}(\Sigma^{-1})}{q}\lambda_{\max}(X^\top X/n) = O_P(q^{-1}).\end{aligned} \tag{2.48}$$

In (2.48) the last inequality and the fourth equality come from Theorem 4.3.15 of Horn & Johnson (1986) and (2.35), respectively. The last equality comes from (A5) and (A7). By (2.34), (2.43) and (2.48), it is thus enough to prove that

$$\lim_{n\to\infty}\mathbb{P}\left[\sum_{k=1}^{nq} Z_k^2 \geq \left\{(nq)^{1/2}\sqrt{q}\frac{\lambda}{\sqrt{nq}}\right\}^2\right] = \lim_{n\to\infty}\mathbb{P}\left\{\sum_{k=1}^{nq} Z_k^2 \geq (nq)\left(\frac{\lambda}{\sqrt{n}}\right)^2\right\} = 0$$

which holds true by the second condition of (L) and Markov's inequality. Hence, the second term of (2.47) tends to zero as $n \to \infty$.

Let us now study the 3rd term of (2.47). By (2.45), we get that for all $j \in J^c$,

$$\mathbb{P}\left\{\left|\left(C_{J^c,J}\{(\widetilde{C}_{J,J})^{-1} - (C_{J,J})^{-1}\}W_J\right)_j\right| \geq \frac{\lambda}{24\sqrt{nq}}\eta\right\} \leq \mathbb{P}\left\{\|C_{J^c,J}\|_2 \, \|(\widetilde{C}_{J,J})^{-1} - (C_{J,J})^{-1}\|_2 \, \|A_1\|_2 \, \|Z\|_2 \geq \frac{\lambda}{24\sqrt{nq}}\eta\right\}.$$

By (2.36), (2.46) and (2.48), it is thus enough to prove that

$$\lim_{n\to\infty}\mathbb{P}\left[\sum_{k=1}^{nq} Z_k^2 \geq \left\{(nq)^{1/2}\sqrt{q}\frac{\lambda}{\sqrt{nq}}\right\}^2\right] = \lim_{n\to\infty}\mathbb{P}\left\{\sum_{k=1}^{nq} Z_k^2 \geq (nq)\left(\frac{\lambda}{\sqrt{n}}\right)^2\right\} = 0,$$

which holds true by the second condition of (L) and Markov's inequality. Hence, the 3rd term of (2.47) tends to zero as $n \to \infty$.

Let us now study the 4th term of (2.47). By (2.40), it amounts to prove that

$$\lim_{n\to\infty} \mathbb{P}\left\{ \|C_{J^c,J}\|_2 \, \|(\widetilde{C}_{J,J})^{-1} - (C_{J,J})^{-1}\|_2 \, \|A\|_2 \, \|Z\|_2 \geq \frac{\lambda}{24\sqrt{nq}}\eta \right\} = 0.$$

By (2.48), (2.36) and (2.43) it is enough to prove that

$$\lim_{n\to\infty} \mathbb{P}\left\{ \sum_{k=1}^{nq} Z_k^2 \geq (nq)\,(nq) \left(\frac{\lambda}{\sqrt{n}}\right)^2 \right\} = 0,$$

which holds true by the second condition of (L). Hence, the 4th term of (2.47) tends to zero as $n \to \infty$.

Let us now study the 5th term of (2.47). By (2.45), proving that the 5th term of (2.47) tends to 0 amounts to proving that

$$\lim_{n\to\infty} \mathbb{P}\left\{ \|C_{J^c,J} - \widetilde{C}_{J^c,J}\|_2 \|(C_{J,J})^{-1}\|_2 \, \|A_1\|_2 \, \|Z\|_2 \geq \frac{\lambda}{24\sqrt{nq}}\eta \right\} = 0.$$

Let us now bound $\|C_{J^c,J} - \widetilde{C}_{J^c,J}\|_2$ as follows:

$$\begin{aligned}
\|C_{J^c,J} - \widetilde{C}_{J^c,J}\|_2 &= \|(C - \widetilde{C})_{J^c,J}\|_2 = \rho\{(C - \widetilde{C})_{J^c,J}(C - \widetilde{C})_{J^c,J}\}^{1/2} \\
&\leq \|(C - \widetilde{C})_{J^c,J}(C - \widetilde{C})_{J^c,J}\|_\infty^{1/2} \leq \|(C - \widetilde{C})(C - \widetilde{C})\|_\infty^{1/2} \leq \|C - \widetilde{C}\|_\infty \\
&= \frac{1}{q}\|\Sigma^{-1} - \widehat{\Sigma}^{-1}\|_\infty \, \|X^\top X/n\|_\infty = O_P\{q^{-1}(nq)^{-1/2}\}, \qquad (2.49)
\end{aligned}$$

as $n \to \infty$, where the last equality comes from (A5) and (A9).

By (2.34), (2.46) and (2.49), to prove that the 5th term of (2.47) tends to zero as $n \to \infty$, it is enough to prove that

$$\lim_{n\to\infty} \mathbb{P}\left\{ \sum_{k=1}^{nq} Z_k^2 \geq nq \left(\frac{\lambda}{\sqrt{n}}\right)^2 \right\} = 0,$$

which can be deduced from Markov's inequality and the second condition of (L).

Using similar arguments as those used for proving that the second, third and fourth terms of (2.47) tend to zero, we get that the 6th,7th and 8th terms of (2.47) tend to zero, as $n \to \infty$, by replacing (2.48) by (2.49).

Let us now study the 9th term of (2.47). Replacing J by J^c in (2.37), (2.38), (2.40), (2.41) and (2.43) in order to prove that the ninth term of (2.47) tends to 0 it is enough to prove that

$$\lim_{n\to\infty} \mathbb{P}\left\{ \sum_{k=1}^{nq} Z_k^2 \geq nq \left(\frac{\lambda}{\sqrt{n}}\right)^2 \right\} = 0.$$

which holds using Markov's inequality and the second condition of (L).

Let us now study the 10th term of (2.47). Using the same idea as the one used for proving (2.33), we get that

$$\mathbb{P}\left[|C_{J^c,J}\{(\widetilde{C}_{J,J})^{-1}-(C_{J,J})^{-1}\}\mathrm{sign}(\mathcal{B}_J)|\geq\frac{\eta}{12}\right] \\ \leq\mathbb{P}\left\{\sqrt{|J|}\|C_{J^c,J}\|_2\,\|(\widetilde{C}_{J,J})^{-1}-(C_{J,J})^{-1}\|_2\geq\frac{\eta}{12}\right\},$$

which tends to zero as $n\to\infty$ by (A3), (2.36), (2.48) and the fact that $c_1<1/2$.

Let us now study the 11th term of (2.47). Using the same idea as the one used for proving (2.33), we get that

$$\mathbb{P}\left\{|(\widetilde{C}_{J^c,J}-C_{J^c,J})(C_{J,J})^{-1}\mathrm{sign}(\mathcal{B}_J)|\geq\frac{\eta}{12}\right\}\leq\mathbb{P}\left\{\sqrt{|J|}\,\|\widetilde{C}_{J^c,J}-C_{J^c,J}\|_2\,\|(C_{J,J})^{-1}\|_2\geq\frac{\eta}{12}\right\},$$

which tends to zero as $n\to\infty$ by (A3), (2.34) and (2.49) and the fact that $c_1<1/2$.

Finally, the 12th term of (2.47) can be bounded as follows:

$$\mathbb{P}\left[|(\widetilde{C}_{J^c,J}-C_{J^c,J})\{(\widetilde{C}_{J,J})^{-1}-(C_{J,J})^{-1}\}\mathrm{sign}(\mathcal{B}_J)|\geq\frac{\eta}{12}\right] \\ \leq\mathbb{P}\left\{\sqrt{|J|}\,\|\widetilde{C}_{J^c,J}-C_{J^c,J}\|_2\,\|(\widetilde{C}_{J,J})^{-1}-(C_{J,J})^{-1}\|_2\geq\frac{\eta}{12}\right\},$$

which tends to zero as $n\to\infty$ by (A3), (2.36) and (2.49) and the fact that $c_1<1/2$. □

Proof of Proposition 2.6. Observe that

$$\Sigma^{-1}=\begin{pmatrix}1&-\phi_1&0&\cdots&0\\-\phi_1&1+\phi_1^2&-\phi_1&\cdots&0\\0&-\phi_1&\ddots&\ddots&\vdots\\\vdots&\vdots&\ddots&1+\phi_1^2&-\phi_1\\0&0&\cdots&-\phi_1&1\end{pmatrix}. \tag{2.50}$$

Let $S=\mathcal{X}^\top\mathcal{X}=\Sigma^{-1}\otimes X^\top X$. Then,

$$S_{i,j}=\begin{cases}n_{r_i+1} & \text{if } j=i \text{ and } k_i\in\{0,q-1\},\\(1+\phi_1^2)n_{r_i+1} & \text{if } j=i \text{ and } k_i\notin\{0,q-1\},\\-\phi_1 n_{r_i+1} & \text{if } j=i+p \text{ or if } j=i-p,\\0 & \text{otherwise,}\end{cases}$$

where $i-1=(p-1)k_i+r_i$ corresponds to the Euclidean division of $i-1$ by $p-1$.

In order to prove (IC), it is enough to prove that $\|S_{J^c,J}(S_{J,J})^{-1}\|_\infty\leq 1-\eta$, where $\eta\in(0,1)$. Since for all j, we have $(j-p)\in J^c$ or $(j+p)\in J^c$, it follows that

$\|S_{J^c,J}\|_\infty = \nu|\phi_1|$. Let $A = S_{J,J}$. Since $A = (a_{i,j})$ is a diagonally dominant matrix, then, by Theorem 1 of Varah (1975),

$$\|A^{-1}\|_\infty \leq 1/\min_k \left(a_{k,k} - \sum_{\substack{1 \leq j \leq |J| \\ j \neq k}} a_{k,j} \right).$$

Using that for all j, $(j-p) \in J^c$ or $(j+p) \in J^c$,

$$\sum_{\substack{1 \leq j \leq |J| \\ j \neq k}} a_{k,j} \leq \nu|\phi_1|.$$

If $k \in J$ then $k > p$ and $k < pq - p$. Thus, $a_{k,k} \geq \nu(1+\phi_1^2)$. Hence, $\|A^{-1}\|_\infty \leq 1/\{\nu(1+\phi_1^2 - |\phi_1|)\}$ and

$$\|S_{J^c,J}(S_{J,J})^{-1}\|_\infty \leq \|S_{J^c,J}\|_\infty \|(S_{J,J})^{-1}\|_\infty \leq \frac{|\phi_1|}{1+\phi_1^2 - |\phi_1|}.$$

Since $|\phi_1| < 1$, the strong Irrepresentability Condition holds when $|\phi_1| \leq (1-\eta)(1+|\phi_1|^2 - |\phi_1|)$, which is true for a small enough η. □

Proof of Proposition 2.7. Since $|\phi_1| < 1$, $\|\Sigma^{-1}\|_\infty \leq |\phi_1| + |1+\phi_1^2| \leq 3$, which gives (A7) by Theorem 5.6.9 of Horn & Johnson (1986). Observe that

$$\|\Sigma\|_\infty \leq \frac{1}{1-\phi_1^2}\left(1 + 2\sum_{h=1}^{q-1}|\phi_1|^h\right) \leq \frac{1}{1-\phi_1^2}\left(1 + \frac{2}{1-|\phi_1|}\right) = \frac{3-|\phi_1|}{1-\phi_1^2} \leq \frac{3}{1-\phi_1^2},$$

which gives (A8) by Theorem 5.6.9 of Horn & Johnson (1986).

Since $\widehat{\Sigma}^{-1}$ has the same expression as Σ^{-1} defined in (2.50) except that ϕ_1 is replaced by $\widehat{\phi_1}$ defined in (2.20), we get that

$$\|\Sigma^{-1} - \widehat{\Sigma}^{-1}\|_\infty \leq 2|\phi_1 - \widehat{\phi}_1| + (\phi_1 - \widehat{\phi}_1)^2,$$

which implies Assumption (A9) of Theorem 2.5 by Lemma 2.8.

Let us now check Assumption (A10) of Theorem 2.5. Since, by Theorem 5.6.9 of Horn & Johnson (1986), $\rho(\Sigma - \widehat{\Sigma}) \leq \|\Sigma - \widehat{\Sigma}\|_\infty$, it is enough to prove that, as

$n \to \infty$, $\|\Sigma - \widehat{\Sigma}\|_\infty = O_P\{(nq)^{-1/2}\}$. Observe that

$$\begin{aligned}
\|\Sigma - \widehat{\Sigma}\|_\infty &\leq \left|\frac{1}{1-\phi_1^2} - \frac{1}{1-\widehat{\phi}_1^2}\right| + 2\sum_{h=1}^{q-1}\left|\frac{\phi_1^h}{1-\phi_1^2} - \frac{\widehat{\phi}_1^h}{1-\widehat{\phi}_1^2}\right| \\
&\leq \left|\frac{\phi_1^2 - \widehat{\phi}_1^2}{(1-\phi_1^2)(1-\widehat{\phi}_1^2)}\right| + 2\sum_{h=1}^{q-1}\left|\frac{\phi_1^h - \widehat{\phi}_1^h}{1-\phi_1^2}\right| + 2\sum_{h=1}^{q-1}\left|\widehat{\phi}_1^h\left(\frac{1}{1-\phi_1^2} - \frac{1}{1-\widehat{\phi}_1^2}\right)\right| \\
&\leq \left|\frac{(\phi_1 - \widehat{\phi}_1)(\phi_1 + \widehat{\phi}_1)}{(1-\phi_1^2)(1-\widehat{\phi}_1^2)}\right| + 2\sum_{h=1}^{q-1}\left|\frac{\phi_1^h - \widehat{\phi}_1^h}{1-\phi_1^2}\right| + 2\sum_{h=1}^{q-1}\left|\left(\widehat{\phi}_1^h - \phi_1^h\right)\left(\frac{1}{1-\phi_1^2} - \frac{1}{1-\widehat{\phi}_1^2}\right)\right| \\
&\qquad + 2\sum_{h=1}^{q-1}\left|\phi_1^h\left(\frac{1}{1-\phi_1^2} - \frac{1}{1-\widehat{\phi}_1^2}\right)\right| \\
&\leq \left|\frac{(\phi_1 - \widehat{\phi}_1)(\phi_1 + \widehat{\phi}_1)}{(1-\phi_1^2)(1-\widehat{\phi}_1^2)}\right|\left(1 + \frac{2}{1-|\phi_1|}\right) \\
&\qquad + 2\left(\frac{1}{|1-\phi_1^2|} + \left|\frac{(\phi_1 - \widehat{\phi}_1)(\phi_1 + \widehat{\phi}_1)}{(1-\phi_1^2)(1-\widehat{\phi}_1^2)}\right|\right)\sum_{h=1}^{q-1}\left|\widehat{\phi}_1^h - \phi_1^h\right|.
\end{aligned}$$

Moreover,

$$\begin{aligned}
\sum_{h=1}^{q-1}\left|\widehat{\phi}_1^h - \phi_1^h\right| \leq \left|\widehat{\phi}_1 - \phi_1\right|\sum_{h=1}^{q-1}\sum_{k=0}^{h-1}|\phi_1|^k|\widehat{\phi}_1|^{h-k-1} &\leq \left|\widehat{\phi}_1 - \phi_1\right|\left(\frac{1-|\widehat{\phi}_1|^{q-1}}{1-|\widehat{\phi}_1|}\right)\left(\frac{1-|\phi_1|^{q-1}}{1-|\phi_1|}\right) \\
&\leq \left|\widehat{\phi}_1 - \phi_1\right|\left(\frac{1}{1-|\widehat{\phi}_1|}\right)\left(\frac{1}{1-|\phi_1|}\right).
\end{aligned}$$

Thus, by Lemma 2.8, $\|\Sigma - \widehat{\Sigma}\|_\infty = O_P\{(nq)^{-1/2}\}$, which implies Assumption (A10) of Theorem 2.5. □

Proof of Lemma 2.8. In the following, for notational simplicity, $q = q_n$. Observe that

$$\sqrt{nq}\widehat{\phi}_1 = \frac{\frac{1}{\sqrt{nq}}\sum_{i=1}^n\sum_{\ell=2}^q \widehat{E}_{i,\ell}\widehat{E}_{i,\ell-1}}{\frac{1}{nq}\sum_{i=1}^n\sum_{\ell=1}^{q-1}\widehat{E}_{i,\ell}^2}.$$

By (2.18),

$$\begin{aligned}
\sum_{i=1}^n\sum_{\ell=2}^q \widehat{E}_{i,\ell}\widehat{E}_{i,\ell-1} &= \sum_{\ell=2}^q(\widehat{E}_{\bullet,\ell})^\top\widehat{E}_{\bullet,\ell-1} = \sum_{\ell=2}^q(\Pi E_{\bullet,\ell})^\top(\Pi E_{\bullet,\ell-1}) \\
&= \sum_{\ell=2}^q(\phi_1\Pi E_{\bullet,\ell-1} + \Pi Z_{\bullet,\ell})^\top(\Pi E_{\bullet,\ell-1}) \qquad (2.51) \\
&= \phi_1\sum_{\ell=1}^{q-1}(\Pi E_{\bullet,\ell})^\top(\Pi E_{\bullet,\ell}) + \sum_{\ell=2}^q(\Pi Z_{\bullet,\ell})^\top(\Pi E_{\bullet,\ell-1}),
\end{aligned}$$

where (2.51) comes from the definition of $(E_{i,t})$. Hence,

$$\sqrt{nq}\,(\widehat{\phi}_1 - \phi_1) = \frac{\frac{1}{\sqrt{nq}}\sum_{\ell=2}^{q}(\Pi Z_{\bullet,\ell})^\top(\Pi E_{\bullet,\ell-1})}{\frac{1}{nq}\sum_{i=1}^{n}\sum_{\ell=1}^{q-1}\widehat{E}_{i,\ell}^2}.$$

In order to prove that $\sqrt{nq}\,(\widehat{\phi}_1 - \phi_1) = O_P(1)$, it is enough to prove that

$$\frac{1}{nq}\sum_{i=1}^{n}\sum_{\ell=1}^{q-1}E_{i,\ell}^2 - \frac{1}{nq}\sum_{i=1}^{n}\sum_{\ell=1}^{q-1}\widehat{E}_{i,\ell}^2 = o_P(1), \text{ as } n\to\infty, \tag{2.52}$$

by Lemma 2.10 and, as $n\to\infty$,

$$\frac{1}{\sqrt{nq}}\sum_{\ell=2}^{q}(\Pi Z_{\bullet,\ell})^\top(\Pi E_{\bullet,\ell-1}) = O_P(1). \tag{2.53}$$

Let us first prove (2.52). By (2.19), $\widehat{\mathcal{E}} = (\mathrm{Id}_{\mathbb{R}^q}\otimes\Pi)\mathcal{E} \equiv A\mathcal{E}$. Note that $\mathrm{Cov}(\widehat{\mathcal{E}}) = A(\Sigma\otimes\mathrm{Id}_{\mathbb{R}^n})A^\top = \Sigma\otimes\Pi$. Hence, for all $i\in\{1,\dots,n\}$,

$$\mathrm{Var}(\widehat{\mathcal{E}}_i) \le \lambda_{\max}(\Sigma).$$

Since the covariance matrix of $\mathcal{E}$ is equal to $\Sigma\otimes\mathrm{Id}_{\mathbb{R}^n}$, it follows thay, for all i, $\mathrm{Var}(\mathcal{E}_i)\le\lambda_{\max}(\Sigma)$. By Markov's inequality,

$$\begin{aligned}\frac{1}{nq}\sum_{i=1}^{n}\sum_{\ell=1}^{q-1}E_{i,\ell}^2 - \frac{1}{nq}\sum_{i=1}^{n}\sum_{\ell=1}^{q-1}\widehat{E}_{i,\ell}^2 &= \frac{1}{nq}\sum_{i=1}^{n}\sum_{\ell=1}^{q}E_{i,\ell}^2 - \frac{1}{nq}\sum_{i=1}^{n}\sum_{\ell=1}^{q}\widehat{E}_{i,\ell}^2 + o_P(1)\\ &= \frac{1}{nq}(\|\mathcal{E}\|_2^2 - \|\widehat{\mathcal{E}}\|_2^2) + o_P(1).\end{aligned}$$

Observe that

$$\begin{aligned}\|\mathcal{E}\|_2^2 - \|\widehat{\mathcal{E}}\|_2^2 = \|\mathcal{E}\|_2^2 - \|A\mathcal{E}\|_2^2 &= \mathcal{E}^\top\mathcal{E} - \mathcal{E}^\top A^\top A\mathcal{E} = \mathcal{E}^\top\left(\mathrm{Id}_{\mathbb{R}^{nq}} - \mathrm{Id}_{\mathbb{R}^q}\otimes\Pi\right)\mathcal{E}\\ &= \mathcal{E}^\top\{\mathrm{Id}_{\mathbb{R}^q}\otimes(\mathrm{Id}_{\mathbb{R}^n} - \Pi)\}\mathcal{E} = \sum_{i=1}^{pq}\widetilde{\mathcal{E}}_i^2,\end{aligned}$$

where $\widetilde{\mathcal{E}} = O\mathcal{E}$, where O is an orthogonal matrix. Using the fact that

$$\mathbb{E}(\widetilde{\mathcal{E}}_i^2) = \mathrm{Cov}(\widetilde{\mathcal{E}})_{i,i} \le \lambda_{\max}(\Sigma),$$

and Markov's inequality, we get (2.52).

Let us now prove (2.53). By definition of $(E_{i,t})$ and since $|\phi_1|<1$,

$$\mathbb{E}\{(\Pi Z_{\bullet,\ell})^\top(\Pi E_{\bullet,\ell-1})\} = 0.$$

Moreover,

$$\mathbb{E}\left[\left\{\sum_{\ell=2}^{q}(\Pi Z_{\bullet,\ell})^\top(\Pi E_{\bullet,\ell-1})\right\}^2\right] = \mathbb{E}\left[\left\{\sum_{\ell=2}^{q}\sum_{i=1}^{n}\left(\sum_{k=1}^{n}\Pi_{i,k}Z_{k,\ell}\right)\left(\sum_{j=1}^{n}\Pi_{i,j}E_{j,\ell-1}\right)\right\}^2\right]$$
$$= \sum_{2\le\ell,\ell'\le q}\ \sum_{1\le i,j,k,i',j',k'\le n}\Pi_{i,k}\Pi_{i',k'}\Pi_{i,j}\Pi_{i',j'}\mathbb{E}(Z_{k,\ell}Z_{k',\ell'}E_{j,\ell-1}E_{j',\ell'-1})$$
$$= \sum_{2\le\ell,\ell'\le q}\ \sum_{1\le i,j,k,i',j',k'\le n}\Pi_{i,k}\Pi_{i',k'}\Pi_{i,j}\Pi_{i',j'}\sum_{r,s\ge 0}\phi_1^r\phi_1^s\mathbb{E}(Z_{k,\ell}Z_{k',\ell'}Z_{j,\ell-1-r}Z_{j',\ell'-1-s}),$$

since the $(E_{i,t})$ are AR(1) processes with $|\phi_1|<1$. Note that

$$\mathbb{E}(Z_{k,\ell}Z_{k',\ell'}Z_{j,\ell-1-r}Z_{j',\ell'-1-s}) = 0$$

except when $\ell=\ell'$, $k=k'$, $j=j'$ and $r=s$. Thus,

$$\mathbb{E}\left[\left(\sum_{\ell=2}^{q}(\Pi Z_{\bullet,\ell})^\top(\Pi E_{\bullet,\ell-1})\right)^2\right] = \sigma^4\left(\sum_{r\ge 0}\phi_1^{2r}\right)\sum_{\ell=2}^{q}\sum_{1\le i,j,k,i'\le n}\Pi_{i,k}\Pi_{i',k}\Pi_{i,j}\Pi_{i',j}$$
$$= \frac{q\sigma^4}{1-\phi_1^2}\mathrm{tr}(\Pi) \le \frac{nq\sigma^4}{1-\phi_1^2},$$

where $\mathrm{tr}(\Pi)$ denotes the trace of Π, which concludes the proof of (2.53) by Markov inequality. □

2.6 Technical lemmas

Lemma 2.9. *Let $A\in\mathcal{M}_n(\mathbb{R})$ and Π an orthogonal projection matrix. For any $j\in\{1,\dots,n\}$, $(A^\top\Pi A)_{jj}\ge 0$.*

Proof. Observe that $(A^\top\Pi A) = A^\top\Pi^\top\Pi A = (\Pi A)^\top(\Pi A)$, since Π is an orthogonal projection matrix. Moreover, $(A^\top\Pi A)_{jj} = e_j^\top(\Pi A)^\top(\Pi A)e_j \ge 0$, since $(\Pi A)^\top(\Pi A)$ is a positive semidefinite symmetric matrix, where e_j is a vector containing null entries except the jth entry which is equal to 1. □

Lemma 2.10. *Assume that $(E_{1,t})_t,\dots,(E_{n,t})_t$ are independent AR(1) processes such that, for all $i\in\{1,\dots,n\}$, $E_{i,t}-\phi_1E_{i,t-1}=Z_{i,t}$, where the $Z_{i,t}$s are zero-mean iid Gaussian random variables with variance σ^2 and $|\phi_1|<1$. Then, as $n\to\infty$,*

$$\frac{1}{nq_n}\sum_{i=1}^{n}\sum_{\ell=1}^{q_n-1}E_{i,\ell}^2 \xrightarrow{P} \frac{\sigma^2}{1-\phi_1^2}.$$

Proof. In the following, for notational simplicity, $q=q_n$. Since $\mathbb{E}(E_{i,\ell}^2)=\sigma^2/(1-\phi_1^2)$,

it is enough to prove that, as $n \to \infty$,

$$\frac{1}{nq} \sum_{i=1}^{n} \sum_{\ell=1}^{q-1} \{E_{i,\ell}^2 - \mathbb{E}(E_{i,\ell}^2)\} \xrightarrow{P} 0.$$

Since

$$E_{i,\ell}^2 = \left(\sum_{j \geq 0} \phi_1^j Z_{i,\ell-j} \right)^2 = \sum_{j,j' \geq 0} \phi_1^j \phi_1^{j'} Z_{i,\ell-j} Z_{i,\ell-j'},$$

$$\begin{aligned} &\operatorname{Var}\left[\frac{1}{nq} \sum_{i=1}^{n} \sum_{\ell=1}^{q-1} \left\{ E_{i,\ell}^2 - \mathbb{E}(E_{i,\ell}^2) \right\} \right] = \frac{1}{(nq)^2} \sum_{i=1}^{n} \sum_{1 \leq \ell, \ell' \leq q-1} \operatorname{Cov}(E_{i,\ell}^2; E_{i,\ell'}^2) \\ &= \frac{1}{(nq)^2} \sum_{i=1}^{n} \sum_{1 \leq \ell, \ell' \leq q-1} \sum_{j,j' \geq 0} \sum_{k,k' \geq 0} \phi_1^j \phi_1^{j'} \phi_1^k \phi_1^{k'} \operatorname{Cov}(Z_{i,\ell-j} Z_{i,\ell-j'}; Z_{i,\ell'-k} Z_{i,\ell'-k'}). \end{aligned} \tag{2.54}$$

By the Cauchy–Schwarz inequality $|\operatorname{Cov}(Z_{i,\ell-j} Z_{i,\ell-j'}; Z_{i,\ell'-k} Z_{i,\ell'-k'})|$ is bounded by a positive constant. Moreover $\sum_{j \geq 0} |\phi_1|^j < \infty$, hence (2.54) tends to zero as $n \to \infty$, which concludes the proof of the lemma. □

Chapter 3

A variable selection approach in the multivariate linear model: An application to LC-MS metabolomics data

Scientific production

The content of this chapter is contained in the article: M. Perrot-Dockès, C. Lévy-Leduc, J. Chiquet, L. Sansonnet, M. Brégère, M.-P. Étienne, S. Robin, G. Genta-Jouve *"A variable selection approach in the multivariate linear model: An application to LC-MS metabolomics data"* Statistical Applications in Genetics and Molecular Biology, 17(5), 2018. The method which is presented is implemented in the MultiVarSel R package available from the CRAN.

Abstract

Omic data are characterized by the presence of strong dependence structures that result either from data acquisition or from some underlying biological processes. Applying statistical procedures that do not adjust the variable selection step to the dependence pattern may result in a loss of power and the selection of spurious variables. The goal of this paper is to propose a variable selection procedure within the multivariate linear model framework that accounts for the dependence between the multiple responses. We shall focus on a specific type of dependence which consists in assuming that the responses of a given individual can be modelled as a time series. We propose a novel Lasso-based approach within the framework of the multivariate linear model taking into account the dependence structure by using different types of stationary processes covariance structures for the random error matrix. Our numerical experiments show that including the estimation of the covariance matrix of the random error matrix in the Lasso criterion dramatically improves the variable selection performance. Our approach is successfully applied to an untargeted LC-MS (Liquid Chromatography-Mass Spectrometry) data set made of African copals samples. Our methodology is implemented in the R package `MultiVarSel` which is available from the Comprehensive R Archive Network (CRAN).

3.1 Introduction

Metabolomics aims to provide a global snapshot (quantitative or qualitative) of the metabolism at a given time and by extension phenotypic information (see Nicholson et al., 1999). It studies the concentration of small molecules called metabolites that are the end products of the enzymatic machinery of the cell. Indeed, minor variations in gene or protein expression levels that are not observable via high throughput experiments may have an influence on the metabolites and hence on the phenotype of interest. Thus, metabolomics is a promising approach that can advantageously complement usual transcriptomic and proteomic analyses. For further details on metabolomics, we refer the reader to Smith et al. (2014). The analysis of the metabolomic biological samples is often performed using High Resolution Mass Spectrometry (HRMS), Nuclear Magnetic Resonance (NMR) or Liquid Chromatography-Mass Spectrometry (LC-MS) and produces a large number of features (hundreds or thousands) that can explain a difference between two or more populations (see Zhang et al., 2012). It is well-known in the untargeted LC-MS data analysis that the identification of metabolites discriminating these populations remains a major bottleneck and therefore the selection of relevant features (metabolites) is a crucial step, as explained in Verdegem et al. (2016). Our goal is to tackle the task of feature selection by taking advantage of the specificities of the LC-MS spectra.

We consider a typical untargeted metabolomic experiment where LC-MS spectra (intensity vs m/z) are obtained from n samples, resulting in an $n \times q$ data matrix where the q columns are ordered according to their m/z ratio. Note that the abbreviation m/z represents the quantity formed by dividing the ratio of the mass of an ion to the unified atomic mass unit, by its charge number (regardless of sign). Figure 3.1 displays an example of such a spectrum. It has to be noticed that the data were first pre-processed using the methodology described in Section 3.4.1. We further assume that the n samples are collected under C conditions and denote n_c the number of samples from Condition c, hence $\sum_c n_c = n$. Multivariate ANOVA (MANOVA, see e.g. Mardia et al., 1980; Muller & Stewart, 2006) provides a natural framework to analyze such a data set. Denoting $\boldsymbol{Y}_{c,r}$ the q-dimensional vector corresponding to the spectrum from the rth replicate in Condition c, the MANOVA model assumes that

$$\boldsymbol{Y}_{c,r} = \boldsymbol{\mu}_c + \boldsymbol{E}_{c,r}, \tag{3.1}$$

where $\boldsymbol{\mu}_c$ is the theoretical mean spectrum in Condition c and $\boldsymbol{E}_{c,r}$ is a random q-dimensional error vector. Each metabolite corresponds to a given component of these three vectors. A "relevant" feature is then defined as the jth m/z value, the theoretical concentration $\mu_c^{(j)}$ of which significantly varies between conditions. Stacking the row vectors $\boldsymbol{Y}_{c,r}$ and $\boldsymbol{E}_{c,r}$, the MANOVA model can be rephrased as

follows:

$$\boldsymbol{Y} = \boldsymbol{X}\boldsymbol{B} + \boldsymbol{E}, \tag{3.2}$$

where $\boldsymbol{Y} = (Y_{i,j})_{1\leq i\leq n,\ 1\leq j\leq q}$ is the $n \times q$ observation matrix, $\boldsymbol{X}$ is the $n \times C$ design matrix of a one-way ANOVA model, $\boldsymbol{B} = (\mu_c^{(j)})_{1\leq c\leq C,\ 1\leq j\leq q}$ is the $C \times q$ coefficient matrix and $\boldsymbol{E} = (E_{i,j})_{1\leq i\leq n,\ 1\leq j\leq q}$ is the $n \times q$ random error matrix. Observe that C corresponds to the number of covariates. For notational simplicity, the samples indexed with (c, r) are now identified with a single index $i \in \{1, \dots n\}$, starting with the n_1 samples from Condition $c = 1$, then the n_2 samples from Condition $c = 2$, etc. In this framework, assuming that the mean spectrum $\overline{\boldsymbol{\mu}} = n^{-1} \sum_n n_c \boldsymbol{\mu}_c$ is set to zero, the problem of determining which metabolites are relevant boils down to finding the non null coefficients in the matrix $\boldsymbol{B}$ and hence can be seen as a variable selection problem in the multivariate linear model. Several approaches can be considered for solving this task: either *a posteriori* methods such as classical statistical tests in ANOVA models (see Mardia et al., 1980; Faraway, 2004) or methods embedding the variable selection such as Lasso-type methodologies (Tibshirani, 1996). However, a naive application of such approaches does not take into account the potential dependence between the different columns of $\boldsymbol{Y}$, which may affect the identification of the relevant features. This drawback will be illustrated in Section 4.3.

Different supervised machine learning approaches have been used to analyze "omics" data during the last few years (see Saccenti et al., 2013; Ren et al., 2015; Boccard & Rudaz, 2016; Zhang et al., 2017). Among them, in metabolomics, the most popular is the partial least squares-discriminant analysis (PLS-DA) which has recently been extended to sPLS-DA (sparse partial least squares-discriminant analysis) by Lê Cao et al. (2011) to include a variable selection step.

The originality of our approach lies in the modeling of the dependence that exists among the columns of $\boldsymbol{Y}$ which comes from the fact that usually biomarkers share biosynthetic pathways (Audoin et al. (2014)). To account for this dependence, we assume that the samples are all independent, namely, all the rows of $\boldsymbol{E}$ are independent and for each sample i, the noise vector $\boldsymbol{E}_i$ has a multivariate Gaussian distribution:

$$\boldsymbol{E}_i = (E_{i,1}, \dots, E_{i,q}) \sim \mathcal{N}(0, \boldsymbol{\Sigma}_q), \tag{3.3}$$

where $\boldsymbol{\Sigma}_q$ denotes the covariance matrix. The simplest assumption regarding the covariance matrix is $\boldsymbol{\Sigma}_q = \sigma^2 \boldsymbol{I}_q$, where $\boldsymbol{I}_q$ denotes the $q \times q$ identity matrix. In this case the different columns of $\boldsymbol{Y}$ are assumed to be independent. The other extreme assumption consists in letting $\boldsymbol{\Sigma}_q$ free, assuming no specific form for this dependence. However, in such a situation, $q(q+1)/2$ parameters should be estimated which is not possible when $n < q$, which is the most standard case. Our approach lies in between, assuming that some dependence exists but that it has a specific structure. The form we consider is motivated by the nature of LC-MS spectra, which can be seen as random functions of the m/z ratio. This suggests to consider

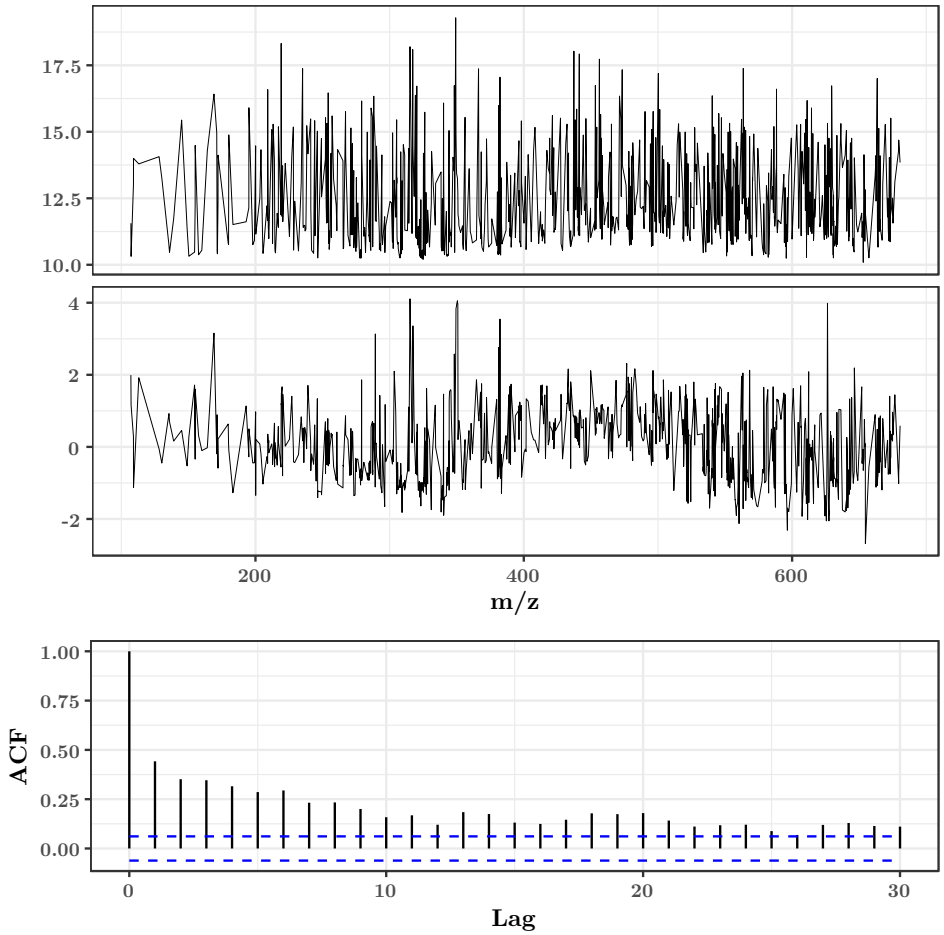

Figure 3.1 – An example of a LC-MS spectrum (an instance of $Y_{c,r}$) (top), the same spectrum centered and normalized (middle) and its empirical autocorrelation function (bottom).

each random vector $\boldsymbol{E}_i$ as a time-series and to borrow classical dependence structure from time-series analysis to model $\boldsymbol{\Sigma}_q$. This approach is consistent with the fact that the empirical autocorrelation function of LC-MS spectra (see Figure 3.1 for an example) displays the typical characteristics of most time-series such as vanishing autocorrelation when the lag increases.

On top of accounting for the dependence between the columns of $\boldsymbol{Y}$, our methodology can deal with a potentially high number of features (columns of $\boldsymbol{Y}$) thanks to the underlying Lasso-based feature selection and the modeling of the dependence which produces sparse estimates of $\boldsymbol{\Sigma}_q^{-1}$. We also couple the whole procedure to a stability selection step to ensure robustness of the selected features. This methodology is implemented in the R package `MultiVarSel` which is available from the Comprehensive R Archive Network (CRAN).

The rest of the paper is organized as follows. Our method is described in Section

4.2. Some numerical experiments on synthetic data are provided in Section 4.3. Finally, an application to a metabolomic data set made of African copals samples is given in Section 3.4.

3.2 Statistical inference

The strategy that we propose can be summarized as follows.

— First step: Fitting a one-way ANOVA to each column of the matrix $\boldsymbol{Y}$ in order to have access to an estimation $\widehat{\boldsymbol{E}}$ of the error matrix $\boldsymbol{E}$.

— Second step: Estimating the matrix $\boldsymbol{\Sigma}_q$ by using the methods described in Sections 3.2.1 and 3.2.1. Then, choosing the most convenient estimator $\widehat{\boldsymbol{\Sigma}}_q$ thanks to a statistical test described in Section 3.2.1.

— Third step: Thanks to $\widehat{\boldsymbol{\Sigma}}_q$, transforming the data in order to remove the dependence between the columns of $\boldsymbol{Y}$. Such a transformation will be called "whitening" hereafter.

— Fourth step: Applying to the transformed observations the Lasso approach described in Section 3.2.2.

The first step provides a first estimate $\widetilde{\boldsymbol{B}}$ of $\boldsymbol{B}$. An estimate of $\boldsymbol{E}$ is then defined as

$$\widehat{\boldsymbol{E}} = \boldsymbol{Y} - \boldsymbol{X}\widetilde{\boldsymbol{B}}. \tag{3.4}$$

In the following, we shall focus on the three other steps.

3.2.1 Estimation of the dependence structure of $\boldsymbol{E}$

We propose hereafter to model each row of $\boldsymbol{E}$ as a realization of a stationary process and hence we shall use time-series models in order to describe the dependence structure of $\boldsymbol{E}$. We refer the reader to Brockwell & Davis (1991) for further details on time series modeling.

We shall consider a large variety of models ranging from the simplest parametric to the most general nonparametric dependence modeling. In each case we focus on the estimation of $\boldsymbol{\Sigma}_q^{-1/2}$ since the use of the following transformation:

$$\boldsymbol{Y}\,\boldsymbol{\Sigma}_q^{-1/2} = \boldsymbol{X}\boldsymbol{B}\,\boldsymbol{\Sigma}_q^{-1/2} + \boldsymbol{E}\,\boldsymbol{\Sigma}_q^{-1/2} \tag{3.5}$$

removes the dependence between the columns of $\boldsymbol{Y}$. Indeed the covariance matrix of each row of $\boldsymbol{E}\boldsymbol{\Sigma}_q^{-1/2}$ is now equal to the identity matrix. Such a procedure will be called "whitening" hereafter.

Parametric dependence

The simplest model among the parametric models is the autoregressive process of order 1 denoted AR(1). More precisely, for each i in $\{1,\dots,n\}$, $E_{i,t}$ satisfies the following equation:

$$E_{i,t} - \phi_1 E_{i,t-1} = W_{i,t}, \text{ with } W_{i,t} \sim WN(0,\sigma^2), \tag{3.6}$$

where $|\phi_1| < 1$ and $WN(0,\sigma^2)$ denotes a zero-mean white noise process of variance σ^2, defined as follows,

$$Z_t \sim WN(0,\sigma^2) \text{ if } \begin{cases} \mathbb{E}(Z_t) = 0, \\ \mathbb{E}(Z_t Z_{t'}) = 0 \text{ if } t \neq t', \\ \mathbb{E}(Z_t^2) = \sigma^2. \end{cases} \tag{3.7}$$

Note that the closer to one the parameter ϕ_1 the stronger the dependence between the $E_{i,t}$'s.

In this case, the inverse of the square root of the covariance matrix $\mathbf{\Sigma}_q$ of $(E_{i,1},\dots,E_{i,q})$ has a simple closed-form expression given by

$$\mathbf{\Sigma}_q^{-1/2} = \begin{pmatrix} \sqrt{1-\phi_1^2} & -\phi_1 & 0 & \cdots & 0 \\ 0 & 1 & -\phi_1 & \cdots & 0 \\ 0 & 0 & \ddots & \ddots & \vdots \\ \vdots & \vdots & \ddots & \ddots & -\phi_1 \\ 0 & 0 & \cdots & 0 & 1 \end{pmatrix}. \tag{3.8}$$

Hence, to obtain the expression of $\widehat{\mathbf{\Sigma}}_q^{-1/2}$, it is enough to have an estimation of the parameter ϕ_1 and to replace it in (3.8). For this, we use the estimator $\widehat{\boldsymbol{E}}$ defined in (3.4) and obtained by fitting a standard ANOVA model to the observations, which corresponds to the first step of our method. Then ϕ_1 is estimated by $\widehat{\phi}_1$ defined by

$$\widehat{\phi}_1 = \frac{1}{n}\sum_{i=1}^{n} \widehat{\phi}_{1,i},$$

where $\widehat{\phi}_{1,i}$ denotes the estimator of ϕ_1 obtained by the classical Yule-Walker equations from $(\widehat{E}_{i,1},\dots,\widehat{E}_{i,q})$, see Brockwell & Davis (1991) for more details.

More generally, it is also possible to have access to $\mathbf{\Sigma}_q^{-1/2}$ for more complex processes such as the ARMA(p,q) process defined as follows: For each i in $\{1,\dots,n\}$,

$$E_{i,t} - \phi_1 E_{i,t-1} - \cdots - \phi_p E_{i,t-p} = W_{i,t} + \theta_1 W_{i,t-1} + \dots \theta_q W_{i,t-q}, \tag{3.9}$$

where $W_{i,t} \sim WN(0,\sigma^2)$, the ϕ_i's and the θ_i's are real parameters.

Nonparametric dependence case

In the situation where a parametric modeling is not relevant for $\mathbf{\Sigma}_q$, it can be estimated by

$$\widehat{\mathbf{\Sigma}}_q = \begin{pmatrix} \widehat{\gamma}(0) & \widehat{\gamma}(1) & \cdots & \widehat{\gamma}(q-1) \\ \widehat{\gamma}(1) & \widehat{\gamma}(0) & \cdots & \widehat{\gamma}(q-2) \\ \vdots & & & \\ \widehat{\gamma}(q-1) & \widehat{\gamma}(q-2) & \cdots & \widehat{\gamma}(0) \end{pmatrix}, \tag{3.10}$$

with

$$\widehat{\gamma}(h) = \frac{1}{n}\sum_{i=1}^{n} \widehat{\gamma}_i(h),$$

where $\widehat{\gamma}_i(h)$ is the standard autocovariance estimator of $\gamma_i(h) = \mathbb{E}(E_{i,t}E_{i,t+h})$, for all t. Usually, $\widehat{\gamma}_i(h)$ is referred to as the empirical autocovariance of the $\widehat{E}_{i,t}$'s at lag h (*i.e.* the empirical covariance between $(\widehat{E}_{i,1}, \ldots, \widehat{E}_{i,n-h})$ and $(\widehat{E}_{i,h+1}, \ldots, \widehat{E}_{i,n})$). For a definition of the standard autocovariance estimator we refer the reader to Chapter 7 of Brockwell & Davis (1991). The matrix $\widehat{\mathbf{\Sigma}}_q^{-1/2}$ is then obtained by inverting the Cholesky factor of $\widehat{\mathbf{\Sigma}}_q$.

Choice of the whitening modeling

In order to decide which dependence modeling better fits the data at hand we propose hereafter a statistical test. If the whitening modeling used is well chosen then each row of $\widetilde{\boldsymbol{E}} = \widehat{\boldsymbol{E}}\widehat{\mathbf{\Sigma}}_q^{-1/2}$ should be a white noise as defined in (3.7), where $\widehat{\boldsymbol{E}}$ is defined in (3.4).

One of the most popular approaches for testing whether a random process is a white noise or not, is the Portmanteau test which is based on the Bartlett theorem (for further details we refer the reader to Brockwell & Davis, 1991, Theorem 7.2.2). By this theorem, we get that under the null hypothesis (H_0): "For each i in $\{1, \ldots, n\}$, $(\widetilde{E}_{i,1}, \ldots, \widetilde{E}_{i,q})$ is a white noise",

$$q\sum_{h=1}^{H} \widehat{\rho}_i(h)^2 \approx \chi^2(H), \text{ as } q \to \infty, \tag{3.11}$$

for each i in $\{1, \ldots, n\}$, where $\widehat{\rho}_i(h)$ denotes the empirical autocorrelation of $(\widetilde{E}_{i,1}, \ldots, \widetilde{E}_{i,q})$ at lag h and $\chi^2(H)$ denotes the chi-squared distribution with H degrees of freedom. Thus, by (3.11), we have at our disposal a p-value for each i in $\{1, \ldots, n\}$ that we denote by Pval_i. In order to have a single p-value instead of n, we shall consider

$$q\sum_{i=1}^{n}\sum_{h=1}^{H} \widehat{\rho}_i(h)^2 \approx \chi^2(nH), \text{ as } q \to \infty, \tag{3.12}$$

where the approximation comes from the fact that the rows of $\widetilde{\boldsymbol{E}}$ are assumed to be independent. Equation (3.12) thus provides a p-value: Pval. Hence, if Pval $< \alpha$, the null hypothesis (H_0) is rejected at the level α, where α is usually equal to 5% and a large value of Pval indicates that the modeling for the dependence structure of $\boldsymbol{E}$ is well chosen.

3.2.2 Estimation of B

Lasso based approach

Let us first explain briefly the usual framework in which the Lasso approach is used. We consider a high-dimensional linear model of the following form

$$\mathcal{Y} = \mathcal{X}\mathcal{B} + \mathcal{E}, \tag{3.13}$$

where $\mathcal{Y}$, $\mathcal{B}$ and $\mathcal{E}$ are vectors. Note that, in high-dimensional linear models, the matrix $\mathcal{X}$ has usually more columns than rows which means that the number of variables is larger than the number of observations but $\mathcal{B}$ is usually a sparse vector, namely it contains a lot of null components.

In such models a very popular approach initially proposed by Tibshirani (1996) is the Least Absolute Shrinkage and Selection Operator (LASSO), which is defined as follows for a positive λ:

$$\widehat{\mathcal{B}}(\lambda) = \text{argmin}_{\mathcal{B}} \left\{ \|\mathcal{Y} - \mathcal{X}\mathcal{B}\|_2^2 + \lambda \|\mathcal{B}\|_1 \right\}, \tag{3.14}$$

where, for $u = (u_1, \ldots, u_n)$, $\|u\|_2^2 = \sum_{i=1}^n u_i^2$ and $\|u\|_1 = \sum_{i=1}^n |u_i|$, i.e. the ℓ_1-norm of the vector u. Observe that the first term of (3.14) is the classical least-squares criterion and that $\lambda\|\mathcal{B}\|_1$ can be seen as a penalty term. The interest of such a criterion is the sparsity enforcing property of the ℓ_1-norm ensuring that the number of non-zero components of the estimator $\widehat{\mathcal{B}}$ of $\mathcal{B}$ is small for large enough values of λ.

This methodology cannot be directly applied to our model since we have to deal with matrices and not with vectors. Nevertheless, as explained in Appendix A, Model (3.2) can be rewritten as in (3.13) where $\mathcal{Y}$, $\mathcal{B}$ and $\mathcal{E}$ are vectors of size nq, pq and nq, respectively. Hence, retrieving the positions of the non null components in $\mathcal{B}$ is a first approach for finding relevant variables. However, this approach does not take into account the dependence between the columns of $\boldsymbol{Y}$. Hence, we propose hereafter a modified version of the standard Lasso criterion (3.14) taking into account this potential dependence.

As explained previously, our contribution consists first in "whitening" the observations, namely removing the dependence that may exist within the observations matrix, by multiplying (3.2) on the right by $\widehat{\boldsymbol{\Sigma}}_q^{-1/2}$, see (3.5) where $\boldsymbol{\Sigma}_q^{-1/2}$ is re-

placed by $\widehat{\mathbf{\Sigma}}_q^{-1/2}$. By using the same vectorization trick that allows us to transform Model (3.2) into Model (3.13), the Lasso criterion can be applied to the vectorized version of Model (3.5) where $\mathbf{\Sigma}_q^{-1/2}$ is replaced by $\widehat{\mathbf{\Sigma}}_q^{-1/2}$. The specific expressions of $\mathcal{Y}$, $\mathcal{X}$, $\mathcal{B}$ and $\mathcal{E}$ are given in Appendix B.

Note that this idea of "whitening" the observations has also been proposed by Rothman et al. (2010) where the estimation of $\mathbf{\Sigma}_q$ and $\boldsymbol{B}$ is performed simultaneously. An implementation is available in the R package `MRCE`. In our approach, $\mathbf{\Sigma}_{\boldsymbol{q}}$ is estimated first and then its estimator is used in (3.5) instead of $\mathbf{\Sigma}_q$ before applying the Lasso criterion. Hence, our method can be seen as a variant of the MRCE method in which $\mathbf{\Sigma}_q$ is estimated beforehand. Moreover, after some numerical experiments, we observed that for the values of n and q that we aim at using, the computational burden of the approach designed by Rothman et al. (2010) is too high for addressing our datasets for fixed regularization parameters, contrary to ours. In addition, in practical situations, the regularization parameters of the MRCE approach have to be tuned. As a consequence, we have not been able to use the MRCE approach for the purpose we consider here.

Model selection issue

Estimator (3.14) depends on a parameter λ which tunes the sparsity level in $\widehat{\mathcal{B}}$. We propose to mix two standard approaches to estimate the positions of the non null components in $\mathcal{B}$: the 10-fold cross-validation method and the stability selection approach of Meinshausen & Buhlmann (2010) which guarantees the robustness of the selected variables.

We first divide our samples into ten groups and remove one group from the dataset thus creating 10 training sets: $Y^{\mathcal{D}_1}, \ldots, Y^{\mathcal{D}_{10}}$. For each training set $Y^{\mathcal{D}_k}$, we apply the first three steps of our approach and the Lasso criterion with a 10-fold cross-validation procedure to get $\lambda_{CV}^{(k)}$. Then, we randomly select a subsample of size $qn_k/2$, where n_k denotes the size of the vector of observations $\mathcal{Y}^{\mathcal{D}_k} = \mathrm{Vec}(Y^{\mathcal{D}_k})$. We then apply the Lasso criterion to this subsample with $\lambda = \lambda_{CV}^{(k)}$ and store the indices i of the non null $\widehat{B}_i$. This random selection of a subsample of the training set and the application of the Lasso criterion are repeated N times. At the end, we have access to the number of times $N_i^{(k)}$ where each component $\widehat{B}_i$ is non null among the N replications for each group k. We only keep in the final set of selected variables the indices i such that $(\sum_{k=1}^{10}(N_i^{(k)}/N))/10$ is larger than a given threshold.

The influence of N and the choice of the threshold will be investigated in Section 4.3.

For some theoretical results supporting our approach we refer the reader to Perrot-Dockès et al. (2018).

3.3 Simulation study

The goal of this section is to assess the statistical performance of our methodology implemented in the R package `MultiVarSel`. In order to emphasize the benefits of using a whitening approach from the variable selection point of view, we shall first compare our approach to standard methodologies. Then, we shall analyze the performance of our statistical test for choosing the best dependence modeling. Finally, we shall investigate the performance of our model selection criterion.

To assess the performance of the different methodologies, we generate observations $\boldsymbol{Y}$ according to Model (3.2) with $q = 1000$, $p = 3$, $n = 30$ ($n_1 = 9$, $n_2 = 8$ and $n_3 = 13$) and different dependence modelings, namely different matrices $\boldsymbol{\Sigma}_q$ corresponding to the AR(1) model described in (3.6) with $\sigma = 1$ and $\phi_1 = 0.7$ or 0.9. Note that the values of the parameters p, q and n are chosen in order to match the metabolomic data analyzed in Section 3.4.

We shall also investigate the effect of the sparsity and of the signal to noise ratio (SNR). The sparsity level s corresponds to the proportion of non null elements in $\mathcal{B}$. Different signal to noise ratios are obtained by multiplying $\boldsymbol{B}$ in (3.2) by a coefficient κ.

3.3.1 Variable selection performance

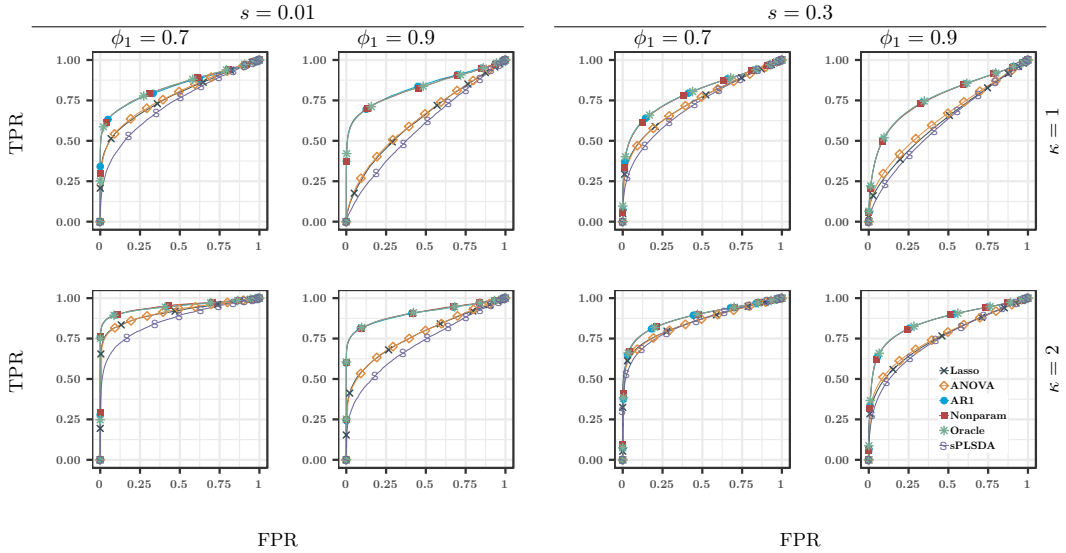

Figure 3.2 – Means of the ROC curves obtained from 200 replications for the different methodologies in the AR(1) dependence modeling; κ is linked to the signal to noise ratio (first row: $\kappa = 1$, second row $\kappa = 2$); ϕ_1 is the correlation level in the AR(1) and s the sparsity level (*i.e.* the fraction of nonzero elements in $\boldsymbol{B}$).

The goal of this section is to compare the performance of our different whitening strategies to standard existing methodologies. More precisely, we shall compare

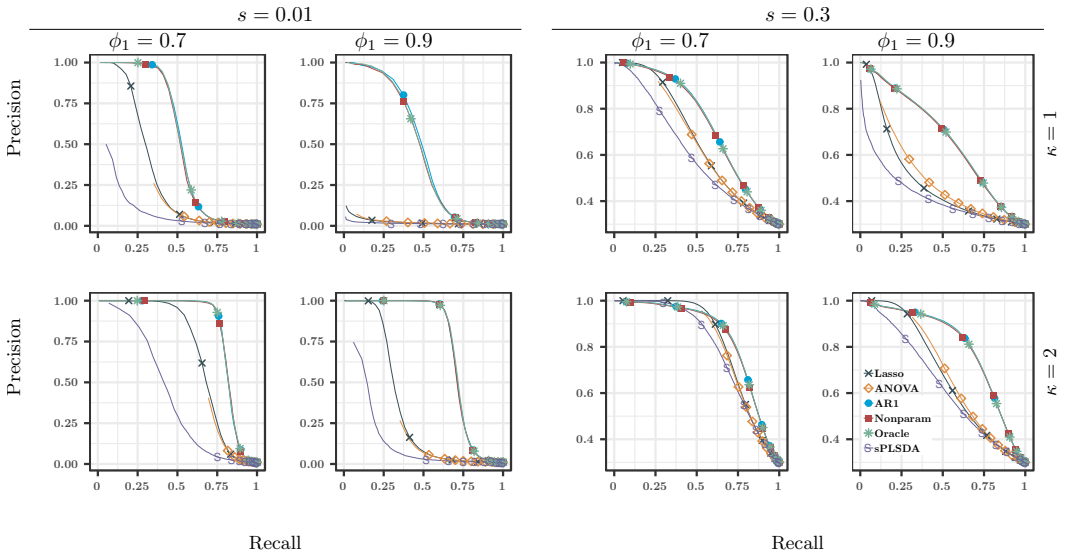

Figure 3.3 – Means of the precision-recall curves obtained from 200 replications for the different methodologies in the AR(1) dependence modeling; κ is linked to the signal to noise ratio (first row: $\kappa = 1$, second row $\kappa = 2$); ϕ_1 is the correlation level in the AR(1) and s the sparsity level (*i.e.* the fraction of nonzero elements in $\boldsymbol{B}$).

our approaches to the classical ANOVA method (denoted ANOVA), the standard Lasso (denoted Lasso), namely the Lasso approach without the whitening step and to sPLSDA (Lê Cao et al., 2011), implemented in the `mixOmics` R package and also in MetaboAnalyst, which is widely used in the metabolomics field. By ANOVA, we mean the classical one-way ANOVA applied to each column of the observations matrix $\boldsymbol{Y}$ without taking the dependence into account. Our different whitening approaches (described in Sections 3.2.1 and 3.2.1) are denoted by AR1 and Nonparam. These methods are also compared to the Oracle approach where the matrix $\boldsymbol{\Sigma}_q$ is known, which is never the case in practical situations.

We shall use three classical criteria for comparison: ROC curves, AUC (Area Under the ROC Curve) and Precision-Recall (PR) curves. ROC curves display the true positive rates (TPR) as a function of the false positive rates (FPR) and the closer to one the AUC the better the methodology. PR curves display the Precision as a function of the Recall. Since the features selected by sPLSDA are not assigned to a given condition c, we shall consider that as soon as a feature is selected it is a true positive, which gives a great advantage to sPLSDA.

We can see from Figures 3.2, 3.3 and Table 3.1 that in the case of an AR(1) dependence, taking into account this dependence provides better results than sPLSDA and than approaches that consider the columns of the matrix $\boldsymbol{E}$ as independent. Moreover, we observe that the performance of the non parametric modeling are on a par with those of the parametric and the oracle ones. We also note that the larger the sparsity level s the smaller the difference of performance between the approaches.

SNR	ϕ_1	s	Lasso	ANOVA	AR1	Nonpar	Oracle	sPLSDA
1	0.7	0.01	0.78	0.78	0.83	0.84	0.84	0.73
1	0.7	0.3	0.74	0.74	0.80	0.80	0.80	0.72
1	0.9	0.01	0.63	0.64	0.83	0.83	0.83	0.58
1	0.9	0.3	0.63	0.64	0.77	0.77	0.77	0.61
2	0.7	0.01	0.91	0.91	0.95	0.95	0.95	0.86
2	0.7	0.3	0.85	0.85	0.88	0.88	0.88	0.84
2	0.9	0.01	0.77	0.77	0.91	0.91	0.91	0.72
2	0.9	0.3	0.75	0.76	0.86	0.86	0.86	0.74

Table 3.1 – AUC of the different methods corresponding to Figure 3.2

As expected, the larger the signal to noise ratio κ the better the performance of the different methodologies. We also conducted numerical experiments in a balanced one-way ANOVA framework. Since the conclusions are similar, we did not report the results here but they are available upon request.

3.3.2 Choice of the dependence modeling

The goal of this section is to assess the performance of the whitening test proposed in Section 3.2.1. We generated observations $\boldsymbol{Y}$ as described at the beginning of Section 4.3, with AR(1) dependence, a sparsity level $s = 0.01$ and SNR such that $\kappa = 1$. The corresponding results are displayed in Figure 3.4.

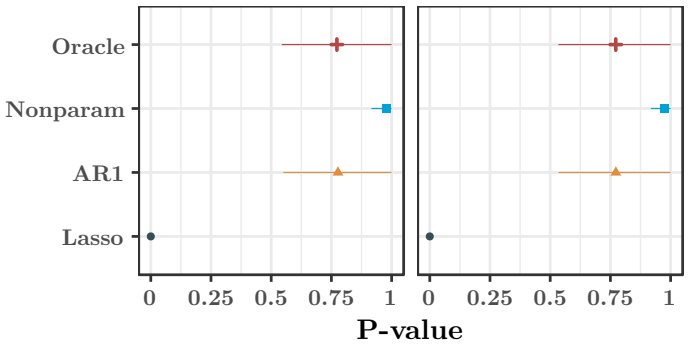

Figure 3.4 – Means and standard deviations of the p-values of the test described in Section 3.2.1 of the main paper for the different approaches in the AR(1) dependence modeling when $\phi_1 = 0.7$ (left) and $\phi_1 = 0.9$ (right).

We observe that our test behaves properly: it provides p-values close to zero in the case where no whitening strategy is used (Lasso) and that when one of the proposed whitening approaches is used the p-values are larger than 0.7.

3.3.3 Choice of the model selection criterion

We investigate here the performance of our model selection criterion described in Section 3.2.2. Figure 3.5 displays the TPR and the FPR for different values N of the sampling replicates and different thresholds. We can see from this figure than taking N larger than 1000 and a threshold of 0.999 ensures a small false positive rate and a large true positive rate.

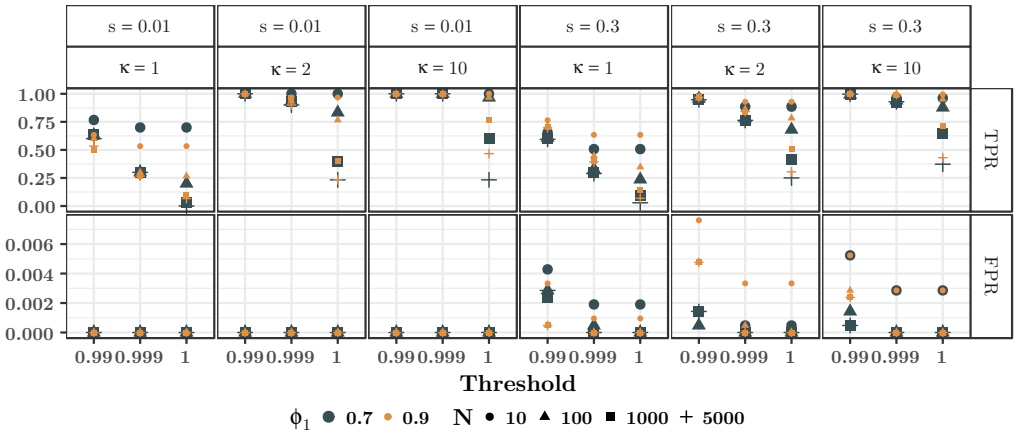

Figure 3.5 – Influence of the number of replications N and of the threshold.

Bullets ('•') in Figure 3.6 show the positions of the variables selected by our four-step approach for two possible thresholds (0.999 and 1) from $N = 1000$ replications. The positions of the non null coefficients in $\boldsymbol{B}$ are displayed with '+'. Here $\boldsymbol{Y}$ is generated with the parameters described at the beginning of Section 4.3 in the case of an AR(1) dependence with $\phi_1 = 0.9$ and $\kappa = 10$. We observe from this figure that the positions of the non null coefficients are recovered for both thresholds. However, the performance are slightly better when the threshold is equal to 0.999.

3.3.4 Numerical performance

In order to investigate the computational burden of our approach, we generated matrices $\boldsymbol{Y}$ satisfying Model (3.2) with $n = 30$ and $q \in \{100, 1000, 2000, \ldots, 5000\}$. Here, the rows of the matrix $\boldsymbol{E}$ are generated as realizations of an AR(1) process and the level of sparsity s of $\boldsymbol{B}$ is equal to 0.01. Figure 3.7 displays the computational times of `MultiVarSel`, including the model selection step described in Section 3.2.2, for different number of replications in the stability selection stage. Timings were obtained on a workstation with 16 GB of RAM and Intel Core i7 (3.66GHz) CPU, using 8 cores for parallel computing. Our implementation uses the R language (R Core Team, 2017) and relies on the `glmnet` and `Matrix` packages (Friedman et al., 2010a; Bates & Maechler, 2017). We can see from this figure that the computational

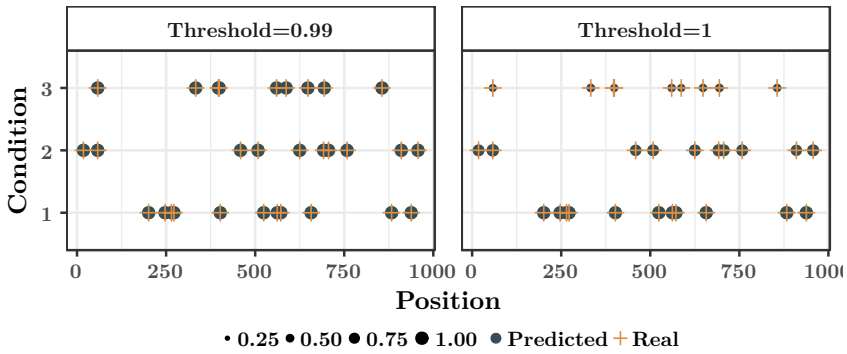

Figure 3.6 – Positions of the variables selected by our approach ('•') when $\kappa = 10$. Values on the y-axis correspond to the 3 conditions. The results obtained when the threshold is equal to 0.999 (resp. 1) are on the left (resp. on the right). The size of the bullets is all the more large that the selection frequency is high.

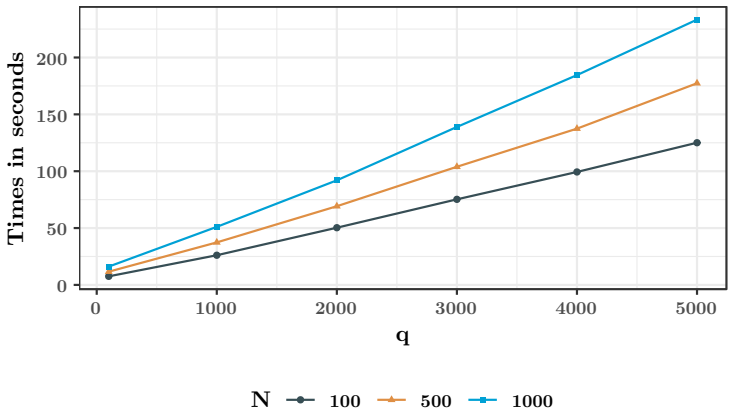

Figure 3.7 – Computational times (in seconds) of `MultiVarSel`. The number of replications corresponds to the number N of subsamplings in the stability selection step.

burden of `MultiVarSel` is very low and that it takes only a few minutes to analyze matrices having 5000 columns.

3.4 Application to the analysis of a LC-MS data set

In this section, `MultiVarSel` is applied to a LC-MS (Liquid Chromatography-Mass Spectrometry) data set made of African copals samples. The samples correspond to ethanolic extracts of copals produced by trees belonging to two genera *Copaifera* (C) and *Trachylobium* (T) with a second level of classification coming from the geographical provenance of the *Copaifera* samples (West (W) or East (E) Africa). Since all the *Trachylobium* samples come from East Africa, we can use the

modeling proposed in Equations (3.1) and (3.2) with $C = 3$ conditions: CE, CW and TE such that $n_{\text{CE}} = 9$, $n_{\text{CW}} = 8$ and $n_{\text{TE}} = 13$. Our goal is to identify the most important features (the m/z values) for distinguishing the different conditions. In this section, we also compare the performance of our method with those of other techniques which are widely used in metabolomics.

3.4.1 Data pre-processing

LC-MS chromatograms were aligned using the `R` package `XCMS` proposed by Smith et al. (2006) with the following parameters: a signal to noise ratio threshold of 10:1 for peak selection, a step size of 0.2 min and a minimum difference in m/z for peaks with overlapping retention times of 0.05 amu. Sample filtering was also performed: To be considered as informative, as suggested by Kirwan et al. (2013), a peak was required to be present in at least 80% of the samples. Missing values imputation was realized using the KNN algorithm described in Hrydziuszko & Viant (2012). Subsequently, the spectra were normalized to equalize signal intensities to the median profile in order to reduce any variance arising from differing dilutions of the biological extracts and probabilistic quotient normalization (PQN) was used, see Dieterle et al. (2006) for further details. In order to reduce the size of the data matrix which contains 6327 metabolites, selection of the adducts of interest $[\text{M+H}]^+$ was then performed using the CAMERA package of Kuhl et al. (2012). A $n \times q$ matrix $\boldsymbol{Y}$ was then obtained with $q = 1019$ and submitted to the statistical analyses.

3.4.2 Application of our four-step approach

The observations matrix $\boldsymbol{Y}$ is first centered and scaled.

— **First step:** A one-way ANOVA is fitted to each column of the observation matrix $\boldsymbol{Y}$ in order to have access to an estimation $\widehat{\boldsymbol{E}}$ of the matrix $\boldsymbol{E}$. Then, the test proposed in Section 3.2.1 is applied to $\widehat{\boldsymbol{E}}$ that is without "whitening" the observations. We found a p-value equal to zero which indicates that the columns of $\widehat{\boldsymbol{E}}$ cannot be considered as independent and hence that applying the whitening strategy should improve the results.

— **Second step:** The different whitening strategies described in Section 3.2.1 were applied and the highest p-value for the test described in Section 3.2.1 is obtained for the nonparametric whitening. More precisely, the p-values obtained for the AR(1) and the nonparametric dependence modeling are equal to 0 and 0.664, respectively. Hence, in the following we shall use the nonparametric modeling.

— **Third step:** Observations were whitened with $\widehat{\boldsymbol{\Sigma}}_q$ obtained by using the nonparametric modeling.

— **Fourth step:** The Lasso approach described in Section 3.2.2 was then applied to the whitened observations. The stability selection is used with $N = 1000$ replications and a threshold equal to 0.999.

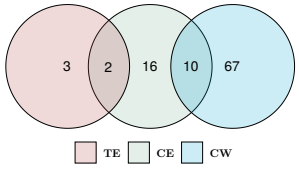

Figure 3.8 – Venn diagram of the features selected for each condition by `MultiVarSel`.

Figure 3.8 displays the Venn diagram of the features (m/z values) selected for each condition CE, TE and CW. Among the 1019 features, 98 features have been selected by MultiVarSel: 77 have been selected for Condition TE, 28 for Condition CW and 5 for Condition CE. Note that there were no features selected for all the conditions, 10 for both TE and CW and 2 for both CW and CE.

3.4.3 Comparison with existing methods

The goal of this section is to compare our approach with the sparse partial least square discriminant analysis (sPLS-DA) which is classically used in metabolomics.

Additional simulations

Since in the case of real data, the position of the relevant features is of course unknown, we propose the following additional simulations in order to further compare these two approaches. We start by applying the first step of our approach in order to get $\widehat{\boldsymbol{E}}$. Then, we perform M random samplings with replacement among the rows of $\widehat{\boldsymbol{E}}$. Let $\boldsymbol{E}^\star$ denote one of them, then we generate a new observation matrix $\boldsymbol{Y}^\star = \boldsymbol{X}^\star\boldsymbol{B} + \boldsymbol{E}^\star$, where $\boldsymbol{X}^\star$ is the same as $\boldsymbol{X}$ except that its rows are permuted in order to ensure a correspondence between the rows of $\boldsymbol{E}^\star$ and $\boldsymbol{X}^\star$. The matrix $\boldsymbol{B}$ is obtained as in Section 4.3 with $s = 0.01$ and $\kappa = 0.5$ and 1. ROC curves averaged over $M = 50$ random samplings are displayed in Figure 3.9. We can see from this figure that our approach outperforms the classical ones. Other values of s and κ have been tested. The corresponding results are not reported here but available upon request.

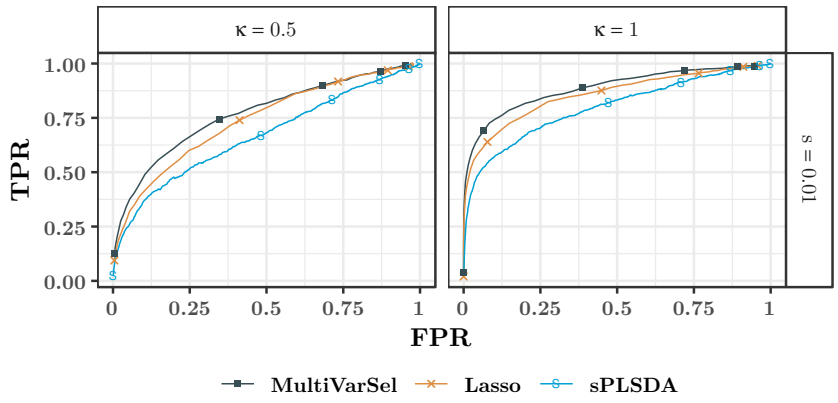

Figure 3.9 – Means of the ROC curves obtained by `MultiVarSel`, Lasso and sPLS-DA.

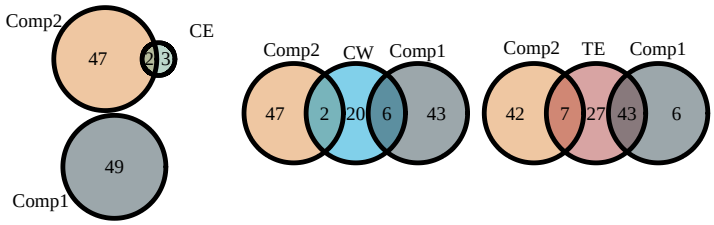

Figure 3.10 – Venn diagrams comparing the features selected by `MultiVarSel` in the three conditions with those selected by sPLS-DA in its two components.

Results on the LC-MS data set

As recommended by Lê Cao et al. (2011), we used two components for sPLS-DA. Moreover, in order to make sPLS-DA comparable with our approach, 49 variables are kept for each component. However, as explained in Section 4.3, the main difference between our approach and sPLSDA is that the features selected by sPLSDA are not assigned to a given condition c, and thus less interpretable.

Figure 3.11 displays the location of the features (m/z values) selected by our approach and sPLS-DA. We can see from this figure that the features selected for the condition TE are mainly located between 400 and 500 m/z whereas those selected for the condition CE are around 600 m/z. The features selected by the first component of the sPLS-DA are also mainly located between 400 and 500 m/z. However, as previously explained, the features selected by sPLSDA are assigned to a component built by the method and not to a condition of the experimental design. Venn diagrams comparing the features selected by both methods are available in Figure 3.10. We observe from these Venn diagrams that the features selected in each component of sPLS-DA do not characterize the conditions of the MANOVA model contrary to ours.

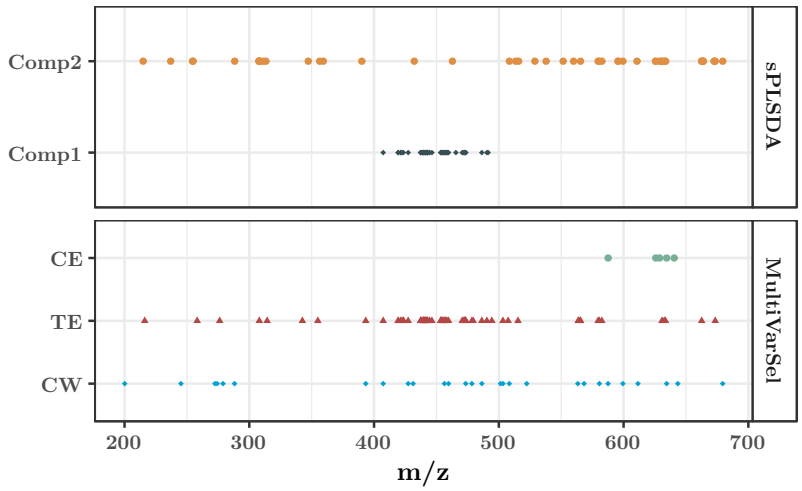

Figure 3.11 – Comparison of the features selected by `MultiVarSel` and `sPLS-DA`.

3.5 Conclusion

In this paper, we proposed a novel approach for feature selection taking into account the dependence that may exist between the columns of the observations matrix. Our approach is implemented in the R package `MultiVarSel` which is available from The Comprehensive R Archive Network (CRAN). We have shown that our method has two main features. Firstly, it is very efficient for selecting a restricted number of stable features characterizing each condition. Secondly, its very low computational burden makes its use possible on very large LC-MS metabolomics data.

Acknowledgment: *This project has been funded by La mission pour l'interdisciplinarité du CNRS in the frame of the DEFI ENVIROMICS (project AREA). The authors thank the Musée François Tillequin for providing the samples from the Guibourt Collection.*

Appendix A

Let $vec(\boldsymbol{A})$ denote the vectorization of the matrix $\boldsymbol{A}$ formed by stacking the columns of $\boldsymbol{A}$ into a single column vector. Let us apply the vec operator to Model (3.2), then

$$vec(\boldsymbol{Y}) = vec(\boldsymbol{XB} + \boldsymbol{E}) = vec(\boldsymbol{XB}) + vec(\boldsymbol{E}).$$

Let $\mathcal{Y} = vec(\boldsymbol{Y})$, $\mathcal{B} = vec(\boldsymbol{B})$ and $\mathcal{E} = vec(\boldsymbol{E})$. Hence,

$$\mathcal{Y} = vec(\boldsymbol{XB}) + \mathcal{E} = (\mathrm{I}_q \otimes \boldsymbol{X})\mathcal{B} + \mathcal{E},$$

where we used that

$$vec(AXB) = (B' \otimes A)vec(X),$$

see (Mardia et al., 1979, Appendix A.2.5). In this equation, B' denotes the transpose of the matrix B. Thus,

$$\mathcal{Y} = \mathcal{X}\mathcal{B} + \mathcal{E},$$

where $\mathcal{X} = \mathrm{I}_q \otimes \boldsymbol{X}$ and $\mathcal{Y}$, $\mathcal{B}$ and $\mathcal{E}$ are vectors of size nq, pq and nq, respectively.

Appendix B

Let us apply the *vec* operator to Model (3.5) where $\boldsymbol{\Sigma}_q^{-1/2}$ is replaced by $\widehat{\boldsymbol{\Sigma}}_q^{-1/2}$, then

$$\begin{aligned} vec(\boldsymbol{Y}\widehat{\boldsymbol{\Sigma}}_q^{-1/2}) &= vec(\boldsymbol{X}\boldsymbol{B}\widehat{\boldsymbol{\Sigma}}_q^{-1/2}) + vec(\boldsymbol{E}\widehat{\boldsymbol{\Sigma}}_q^{-1/2}) \\ &= ((\widehat{\boldsymbol{\Sigma}}_q^{-1/2})' \otimes \boldsymbol{X})vec(\boldsymbol{B}) + vec(\boldsymbol{E}\widehat{\boldsymbol{\Sigma}}_q^{-1/2}). \end{aligned}$$

Hence,

$$\mathcal{Y} = \mathcal{X}\mathcal{B} + \mathcal{E},$$

where $\mathcal{Y} = vec(\boldsymbol{Y}\widehat{\boldsymbol{\Sigma}}_q^{-1/2})$, $\mathcal{X} = (\widehat{\boldsymbol{\Sigma}}_q^{-1/2})' \otimes \boldsymbol{X}$ and $\mathcal{E} = vec(\boldsymbol{E}\widehat{\boldsymbol{\Sigma}}_q^{-1/2})$.

Chapter 4

Estimation of large block structured covariance matrices: Application to "multi-omic" approaches to study seed quality

Scientific production

The content of this chapter is contained in the article submitted for publication: M. Perrot-Dockès, C. Lévy-Leduc *"Estimation of large block structured covariance matrices: Application to "multi-omic" approaches to study seed quality"* https://arxiv.org/abs/1806.10093 The method which is presented is implemented in the BlockCov R package available from the CRAN.

Abstract

Motivated by an application in high-throughput genomics and metabolomics, we propose a novel, efficient and fully data-driven approach for estimating large block structured sparse covariance matrices in the case where the number of variables is much larger than the number of samples without limiting ourselves to block diagonal matrices. Our approach consists in approximating such a covariance matrix by the sum of a low-rank sparse matrix and a diagonal matrix. Our methodology also can deal with matrices for which the block structure appears only if the columns and rows are permuted according to an unknown permutation. Our technique is implemented in the R package `BlockCov` which is available from the Comprehensive R Archive Network (CRAN) and from GitHub. In order to illustrate the statistical and numerical performance of our package some numerical experiments are provided as well as a thorough comparison with alternative methods. Finally, our approach is applied to the use of "multi-omic" approaches for studying seed quality.

4.1 Introduction

Plant functional genomics refers to the description of the biological function of a single or a group of genes and both the dynamics and the plasticity of genome

expression to shape the phenotype. Combining multi-omics such as transcriptomic, proteomic or metabolomic approaches allows us to address in a new light the dimension and the complexity of the different levels of gene expression control and the delicacy of the metabolic regulation of plants under fluctuation environments. Thus, our era marks a real conceptual shift in plant biology where the individual is no longer considered as a simple sum of components but rather as a system with a set of interacting components to maximize its growth, its reproduction and its adaptation. Plant systems biology is therefore defined by multidisciplinary and multi-scale approaches based on the acquisition of a wide range of data as exhaustive as possible.

In this context, it is crucial to propose new methodologies for integrating heterogeneous data explaining the co-regulations/co-accumulations of products of gene expression (mRNA, proteins) and metabolites. In order to better understand these phenomena, our goal will thus be to propose a new approach for estimating block structured covariance matrix in a high-dimensional framework where the dimension of the covariance matrix is much larger than the sample size. In this setting, it is well known that the commonly used sample covariance matrix performs poorly. In recent years, researchers have proposed various regularization techniques to consistently estimate large covariance matrices or the inverse of such matrices, namely precision matrices. To estimate such matrices, one of the key assumptions made in the literature is that the matrix of interest is sparse, namely many entries are equal to zero. A number of regularization approaches including banding, tapering, thresholding and ℓ_1 minimization, have been developed to estimate large covariance matrices or their inverse such as, for instance, Ledoit & Wolf (2004), Bickel & Levina (2008), Banerjee et al. (2008b), Bien & Tibshirani (2011) and Rothman (2012) among many others. For further references, we refer the reader to Cai & Yuan (2012) and to the review of Fan et al. (2016).

In this paper, we shall consider the following framework. Let $\boldsymbol{E}_1, \boldsymbol{E}_2, \cdots, \boldsymbol{E}_n$, n zero-mean i.i.d. q-dimensional random vectors having a covariance matrix $\boldsymbol{\Sigma}$ such that the number q of its rows and columns is much larger than n. The goal of the paper is to propose a new estimator of $\boldsymbol{\Sigma}$ and of the square root of its inverse, $\boldsymbol{\Sigma}^{-1/2}$, in the particular case where $\boldsymbol{\Sigma}$ is assumed to have a block structure without limiting ourselves to diagonal blocks. An accurate estimator of $\boldsymbol{\Sigma}$ can indeed be very useful to better understand the links between the columns of the observation matrix and may highlight some biological processes. Moreover, an estimator of $\boldsymbol{\Sigma}^{-1/2}$ can be very useful in the general linear model in order to remove the dependence that may exist between the columns of the observation matrix. For further details on this point, we refer the reader to Perrot-Dockès et al. (2018), Perrot-Dockès et al. (2018) and to the R package `MultiVarSel` in which such an approach is proposed and implemented for performing variable selection in the multivariate linear model in the presence of dependence between the columns of the observation matrix.

More precisely, in this paper, we shall assume that

$$\boldsymbol{\Sigma} = \boldsymbol{Z}\boldsymbol{Z}' + \boldsymbol{D}, \tag{4.1}$$

where $\boldsymbol{Z}$ is a $q \times k$ sparse matrix with $k \ll q$, $\boldsymbol{Z}'$ denotes the transpose of the matrix $\boldsymbol{Z}$ and $\boldsymbol{D}$ is a diagonal matrix such that the diagonal terms of $\boldsymbol{\Sigma}$ are equal to one. Two examples of such matrices $\boldsymbol{Z}$ and $\boldsymbol{\Sigma}$ are given in Figure 4.1 in the case where $k = 5$ and $q = 50$ and in the case where the columns of $\boldsymbol{\Sigma}$ do not need to be permuted in order to see the block structure. Based on (4.1), our model could seem to be close to factor models described in Johnson & Wichern (1988) and Fan et al. (2016). However, in Johnson & Wichern (1988), the high-dimensional aspects are not considered and in Fan et al. (2016) the sparsity constraint is not studied. Blum et al. (2016b) proposed a methodology which is based on the factor model but with a sparsity constraint on the coefficients of $\boldsymbol{Z}$ which leads to a sparse covariance matrix. Note also that the block diagonal assumption has already been recently considered by Devijver & Gallopin (2018) for estimating the inverse of large covariance matrices in high-dimensional Gaussian Graphical Models (GGM).

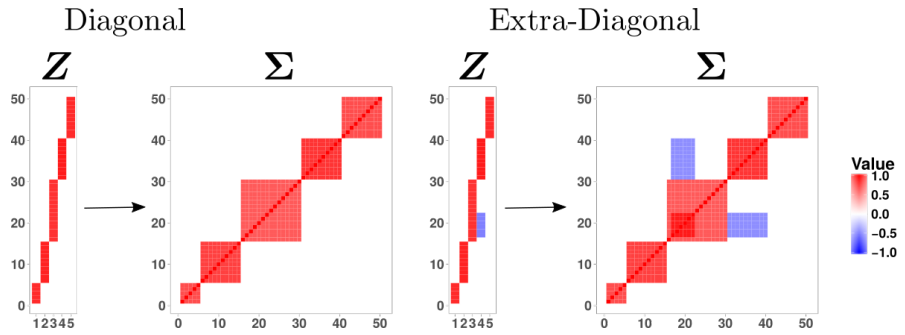

Figure 4.1 – Examples of matrices $\boldsymbol{\Sigma}$ generated from different matrices $\boldsymbol{Z}$ leading to a block diagonal or to a more general block structure (extra-diagonal blocks).

We also propose a methodology to estimate $\boldsymbol{\Sigma}$ in the case where the block structure is latent; that is, permuting the columns and rows of $\boldsymbol{\Sigma}$ renders visible its block structure. An example of such a matrix $\boldsymbol{\Sigma}$ is given in Figure 4.2 in the case where $k = 5$ and $q = 50$.

Our approach is fully data-driven and consists in providing a low rank matrix approximation of the $\boldsymbol{Z}\boldsymbol{Z}'$ part of $\boldsymbol{\Sigma}$ and then in using a ℓ_1 regularization to obtain a sparse estimator of $\boldsymbol{\Sigma}$. When the block structure is latent, a hierarchical clustering step must be applied first. With this estimator of $\boldsymbol{\Sigma}$, we explain how to obtain an estimator of $\boldsymbol{\Sigma}^{-1/2}$.

Our methodology is described in Section 4.2. Some numerical experiments on synthetic data are provided in Section 4.3. An application to the analysis of "-omic" data to study seed quality is performed in Section 4.4.

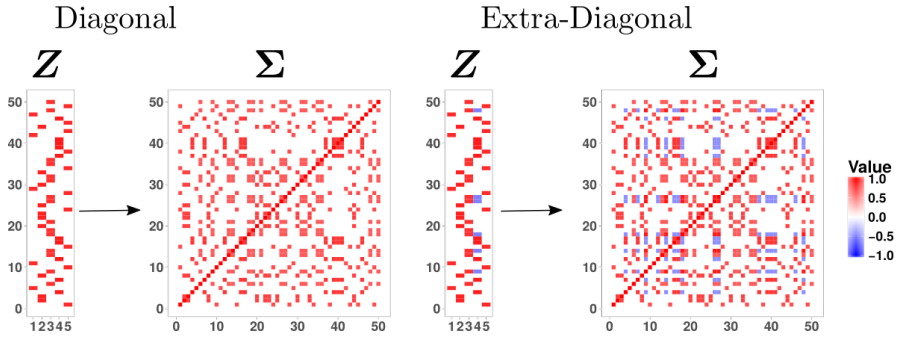

Figure 4.2 – Examples of matrices $\mathbf{\Sigma}$ of Figure 4.1 in which the columns and rows are randomly permuted.

4.2 Statistical inference

The strategy that we propose for estimating $\mathbf{\Sigma}$ and $\mathbf{\Sigma}^{-1/2}$ can be summarized as follows.

— First step: Low rank approximation. In this step, we propose to approximate the part $\boldsymbol{Z}\boldsymbol{Z}'$ of $\mathbf{\Sigma}$ by a low rank matrix using a Singular Value Decomposition (SVD).

— Second step: Detecting the position of the non null values. In this step, we use a Lasso criterion to yield a sparse estimator $\widetilde{\mathbf{\Sigma}}$ of $\mathbf{\Sigma}$.

— Third step: Positive definiteness. We apply the methodology of Higham (2002) to $\widetilde{\mathbf{\Sigma}}$ to ensure that the final estimator $\widehat{\mathbf{\Sigma}}$ of $\mathbf{\Sigma}$ is positive definite.

— Fourth step: Estimation of $\mathbf{\Sigma}^{-1/2}$. In this step, $\mathbf{\Sigma}^{-1/2}$ is estimated from the spectral decomposition of $\widehat{\mathbf{\Sigma}}$ obtained in the previous step.

4.2.1 Low rank approximation

By definition of $\boldsymbol{Z}$ in (4.1), $\boldsymbol{Z}\boldsymbol{Z}'$ is a $q \times q$ low rank matrix having its rank smaller or equal to $k \ll q$. In the first step, our goal is thus to propose a low rank approximation of an estimator of $\boldsymbol{Z}\boldsymbol{Z}'$.

Let $\boldsymbol{S}$ be the sample $q \times q$ covariance matrix defined by

$$\boldsymbol{S} = \frac{1}{n-1}\sum_{i=1}^{n}\left(\boldsymbol{E}_i - \overline{\boldsymbol{E}}\right)\left(\boldsymbol{E}_i - \overline{\boldsymbol{E}}\right)', \quad \text{with } \overline{\boldsymbol{E}} = \frac{1}{n}\sum_{i=1}^{n}\boldsymbol{E}_i,$$

where $\boldsymbol{E}_i = (E_{i,1}, \ldots, E_{i,q})'$. The corresponding $q \times q$ sample correlation matrix $\boldsymbol{R} = (R_{i,j})$ is defined by:

$$R_{i,j} = \frac{S_{i,j}}{\sigma_i \sigma_j}, \ \forall 1 \leq i, j \leq q, \tag{4.2}$$

where

$$\sigma_i^2 = \frac{1}{n-1}\sum_{\ell=1}^{n}(E_{\ell,i} - \overline{E}_i)^2, \quad \text{with } \overline{E}_i = \frac{1}{n}\sum_{\ell=1}^{n} E_{\ell,i},\ \forall 1 \leq i \leq q.$$

Let us also consider the $(q-1) \times (q-1)$ matrix $\mathbf{\Gamma}$ defined by:

$$\begin{aligned} \Gamma_{i,j} &= R_{i,j+1},\ \forall 1 \leq i \leq j \leq q-1, \\ \Gamma_{i,j} &= \Gamma_{j,i},\ \forall 1 \leq j < i \leq q-1. \end{aligned} \tag{4.3}$$

If $\boldsymbol{S}$ was the real matrix $\mathbf{\Sigma}$, the corresponding matrix Γ would have a rank less than or equal to k. Since $\boldsymbol{S}$ is an estimator of $\mathbf{\Sigma}$, we shall use a rank r approximation $\mathbf{\Gamma}_r$ of $\mathbf{\Gamma}$. This will be performed by considering in its singular value decomposition only the r largest singular values and by replacing the other ones by 0. By Eckart & Young (1936), this corresponds to the best rank r approximation of $\mathbf{\Gamma}$. The choice of r will be discussed in Section 4.2.5.

4.2.2 Detecting the position of the non null values

Let us first explain the usual framework in which the Lasso approach is used. We consider a linear model of the following form

$$\mathcal{Y} = \mathcal{X}\mathcal{B} + \mathcal{E}, \tag{4.4}$$

where $\mathcal{Y}$, $\mathcal{B}$ and $\mathcal{E}$ are vectors and $\mathcal{B}$ is sparse meaning that it has a lot of null components.

In such models a very popular approach initially proposed by Tibshirani (1996) is the Least Absolute Shrinkage eStimatOr (Lasso), which is defined as follows for a positive λ:

$$\widehat{\mathcal{B}}(\lambda) = \text{Argmin}_{\mathcal{B}} \left\{ \|\mathcal{Y} - \mathcal{X}\mathcal{B}\|_2^2 + \lambda \|\mathcal{B}\|_1 \right\}, \tag{4.5}$$

where, for $u = (u_1, \ldots, u_n)$, $\|u\|_2^2 = \sum_{i=1}^n u_i^2$ and $\|u\|_1 = \sum_{i=1}^n |u_i|$, *i.e.* the ℓ_1-norm of the vector u. Observe that the first term of (4.5) is the classical least-squares criterion and that $\lambda\|\mathcal{B}\|_1$ can be seen as a penalty term. The interest of such a criterion is the sparsity enforcing property of the ℓ_1-norm ensuring that the number of non-zero components of the estimator $\widehat{\mathcal{B}}$ of $\mathcal{B}$ is small for large enough values of λ. Let

$$\mathcal{Y} = vec_H(\mathbf{\Gamma}_r), \tag{4.6}$$

where vec_H defined in Section 16.4 of Harville (2001) is such that for a $n \times n$ matrix

A,

$$vec_H(A) = \begin{pmatrix} a_1* \\ a_2* \\ \vdots \\ a_n* \end{pmatrix},$$

where a_i* is the sub-vector of the column i of A obtained by striking out the $i-1$ first elements. In order to estimate the sparse matrix $\boldsymbol{Z}\boldsymbol{Z}'$, we need to propose a sparse estimator of $\mathbf{\Gamma}_r$. To do this we apply the Lasso criterion described in (4.5), where $\mathcal{X}$ is the identity matrix. In the case where $\mathcal{X}$ is an orthogonal matrix it has been shown in Giraud (2014) that the solution of (4.5) is:

$$\widehat{\mathcal{B}}(\lambda)_j = \begin{cases} \mathcal{X}_j'\mathcal{Y}(1 - \frac{\lambda}{2|\mathcal{X}_j'\mathcal{Y}|}), & \text{if } |\mathcal{X}_j'\mathcal{Y}| > \frac{\lambda}{2} \\ 0, & \text{otherwise,} \end{cases}$$

where $\mathcal{X}_j$ denotes the jth column of $\mathcal{X}$. Using the fact that $\mathcal{X}$ is the identity matrix we get

$$\widehat{\mathcal{B}}(\lambda)_j = \begin{cases} \mathcal{Y}_j(1 - \frac{\lambda}{2|\mathcal{Y}_j|}), & \text{if } |\mathcal{Y}_j| > \frac{\lambda}{2} \\ 0, & \text{otherwise.} \end{cases} \tag{4.7}$$

We then reestimate the non null coefficients using the least-squares criterion and get:

$$\widetilde{\mathcal{B}}(\lambda)_j = \begin{cases} \mathcal{Y}_j, & \text{if } |\mathcal{Y}_j| > \frac{\lambda}{2} \\ 0, & \text{otherwise,} \end{cases} \tag{4.8}$$

where $\mathcal{Y}$ is defined in (4.6).

It has to be noticed that $\widehat{\mathbf{\Gamma}}_r$ obtained in (4.7) satisfies the following criterion:

$$\widehat{\mathbf{\Gamma}}_r = \text{Argmin}_{\mathbf{\Theta}} \left\{ \|\mathbf{\Gamma}_r - \mathbf{\Theta}\|_F + \lambda |\mathbf{\Theta}|_1 \right\},$$

where $\|\cdot\|_F$ denotes the Frobenius norm defined for a matrix A by $\|A\|_F^2 = \text{Trace}(A'A)$, $|M|_1 = \|vec(M)\|_1$ denotes the ℓ_1-norm of the vector formed by stacking the columns of M. It is thus closely related to the generalized thresholding estimator defined in Wen et al. (2016) and to the one defined in Rothman (2012) with $\tau = 0$ except that in our case $|\mathbf{\Theta}^-|_1$ is replaced by $|\mathbf{\Theta}|_1$ where $\mathbf{\Theta}^-$ corresponds to the matrix $\mathbf{\Theta}$ in which the diagonal terms are replaced by 0. The diagonal terms of $\mathbf{\Sigma}$ were indeed already removed in $\mathbf{\Gamma}_r$. Hence, we get $\widehat{\mathbf{\Gamma}}_r$ by elementwise soft-thresholding that is by putting to zero the value of $\mathbf{\Gamma}_r$ that are under a given threshold and by multiplying the non null values by a coefficient containing this threshold.

Here, we choose to estimate $\mathbf{\Gamma}_r$ by $\widetilde{\mathbf{\Gamma}}_r(\lambda)$ defined through $\widetilde{\mathcal{B}}(\lambda)$ in (4.8) which corresponds to a hard-thresholding and we set the upper triangular part of the estimator $\widetilde{\mathbf{\Sigma}}(\lambda)$ of $\mathbf{\Sigma}$ to be equal to $\widetilde{\mathbf{\Gamma}}_r(\lambda)$. Since the diagonal terms of $\mathbf{\Sigma}$ are

assumed to be equal to 1, we take the diagonal terms of $\widetilde{\mathbf{\Sigma}}(\lambda)$ equal to 1. The lower triangular part of $\widetilde{\mathbf{\Sigma}}(\lambda)$ is then obtained by symmetry.

The choice of the best parameter λ denoted λ_{final} in the following will be discussed in Section 4.3.2.

4.2.3 Positive definiteness

To ensure the positive definiteness of our estimator $\widehat{\mathbf{\Sigma}}$ of $\mathbf{\Sigma}$, we consider the nearest correlation matrix to $\widetilde{\mathbf{\Sigma}}(\lambda_{\text{final}})$ which is computed by using the methodology proposed by Higham (2002) and which is implemented in the function `nearPD` of the R package `Matrix`, see Bates & Maechler (2017).

4.2.4 Estimation of $\mathbf{\Sigma}^{-1/2}$

Even if providing an estimator of a large covariance matrix can be very useful in practice, it may also be interesting to efficiently estimate $\mathbf{\Sigma}^{-1/2}$. Such an estimator can indeed be used in the general linear model in order to remove the dependence that may exist between the columns of the observations matrix. For further details on this point, we refer the reader to Perrot-Dockès et al. (2018), Perrot-Dockès et al. (2018) and to the R package `MultiVarSel` in which such an approach is proposed and implemented for performing variable selection in the multivariate linear model in the presence of dependence between the columns of the observation matrix.

Since $\widehat{\mathbf{\Sigma}}$ is a symmetric matrix, it can be rewritten as $\boldsymbol{U}\boldsymbol{D}\boldsymbol{U}'$, where $\boldsymbol{D}$ is a diagonal matrix and $\boldsymbol{U}$ is an orthogonal matrix. The matrix $\mathbf{\Sigma}^{-1/2}$ can thus be estimated by $\boldsymbol{U}\boldsymbol{D}^{-1/2}\boldsymbol{U}'$ where $\boldsymbol{D}^{-1/2}$ is a diagonal matrix having its diagonal terms equal to the square root of the inverse of the singular values of $\widehat{\mathbf{\Sigma}}$. However, inverting the square root of too small eigenvalues may lead to poor estimators of $\mathbf{\Sigma}^{-1/2}$. This is the reason why we propose to estimate $\mathbf{\Sigma}^{-1/2}$ by

$$\widehat{\mathbf{\Sigma}}_t^{-1/2} = \boldsymbol{U}\boldsymbol{D}_t^{-1/2}\boldsymbol{U}', \tag{4.9}$$

where $\boldsymbol{D}_t^{-1/2}$ is a diagonal matrix such that its diagonal entries are equal to the square root of the inverse of the diagonal entries of $\boldsymbol{D}$ except for those which are smaller than a given threshold t which are replaced by 0 in $\boldsymbol{D}_t^{-1/2}$. The choice of t will be further discussed in Section 4.3.7.

4.2.5 Choice of the parameters

Our methodology for estimating $\mathbf{\Sigma}$ depends on two parameters: The number r of singular values kept for defining $\mathbf{\Gamma}_r$ and the parameter λ which controls the sparsity level namely the number of zero values in $\widetilde{\mathcal{B}}(\lambda)$ defined in (4.8).

For choosing r, we shall compare two strategies in Section 4.3.1:

— The Cattell criterion based on the Cattell's scree plot described in Cattell (1966) and

— the PA permutation method proposed by Horn (1965) and recently studied from a theoretical point of view by Dobriban (2018).

To choose the parameter λ in (4.8), we shall compare two strategies in Section 4.3.2:

— The BL approach proposed in Bickel & Levina (2008) based on cross-validation and

— the Elbow method which consists in computing for different values of λ the Frobenius norm $\|\boldsymbol{R} - \widetilde{\boldsymbol{\Sigma}}(\lambda)\|_F$, where $\boldsymbol{R}$ and $\widetilde{\boldsymbol{\Sigma}}(\lambda)$ are defined in (4.2) and at the end of Section 4.2.2, respectively. Then, it fits two simple linear regressions and chooses the value of λ achieving the best fit.

4.3 Numerical experiments

Our methodology described in the previous section is implemented in the R package `BlockCov` and is available from the CRAN (Comprehensive R Achive Network) and from GitHub.

We propose hereafter to investigate the performance of our approach for different types of matrices $\boldsymbol{\Sigma}$ defined in (4.1) and for different values of n and q. The four following cases considered correspond to different types of matrices $\boldsymbol{Z}$, the matrices $\boldsymbol{D}$ being chosen accordingly to ensure that the matrix $\boldsymbol{\Sigma}$ has its diagonal terms equal to 1.

— **Diagonal-Equal** case. In this situation, $\boldsymbol{Z}$ has the structure displayed in the left part of Figure 4.1, namely it has 5 columns such that the numbers of the non values in the five columns are equal to $0.1 \times q$, $0.2 \times q$, $0.3 \times q$, $0.2 \times q$ and $0.2 \times q$, respectively and the non null values are equal to $\sqrt{0.7}$, $\sqrt{0.75}$, $\sqrt{0.65}$, $\sqrt{0.8}$ and $\sqrt{0.7}$, respectively.

— **Diagonal-Unequal** case. In this scenario, $\boldsymbol{Z}$ has the same structure as for the **Diagonal-Equal** case except that the non null values in the five columns are not fixed but randomly chosen in $[\sqrt{0.6}, \sqrt{0.8}]$ except for the third column for which its values are randomly chosen in $[\sqrt{0.3}, \sqrt{0.6}]$.

— **Extra-Diagonal-Equal** case. Here, $\boldsymbol{Z}$ has the structure displayed in the right part of Figure 4.1. The values of the columns of $\boldsymbol{Z}$ are the same as those of the **Diagonal-Equal** case except for the fourth column which is assumed to contain additional non values equal to -0.5 in the range $[0.35 \times q, 0.45 \times q]$.

— **Extra-Diagonal-Unequal** case. $\boldsymbol{Z}$ has the same structure as in the **Extra-Diagonal-Equal** case except that the values are randomly chosen as in the **Diagonal-Unequal** case except for the fourth column where the additional non values are still equal to -0.5 in the range $[0.35 \times q, 0.45 \times q]$.

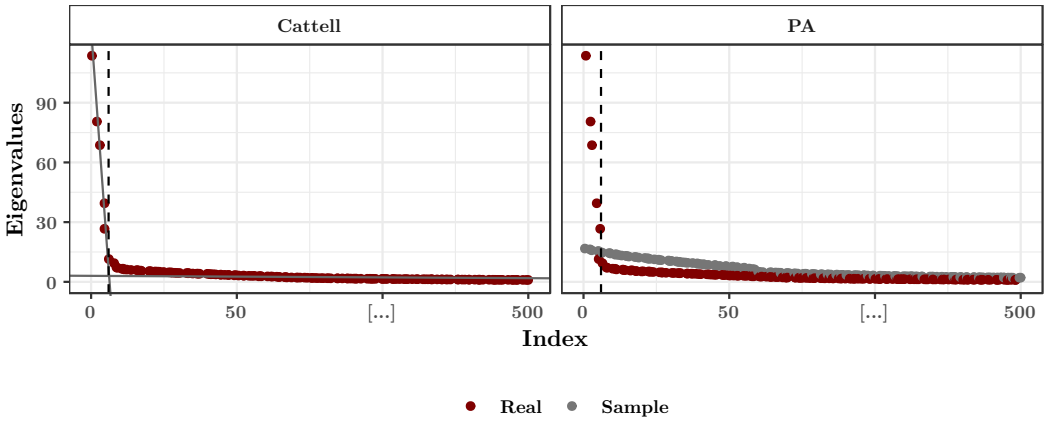

Figure 4.3 – Illustration of PA and Cattell criteria for choosing r when $q = 500$ and $n = 30$ in the Extra-Diagonal-Unequal case. The value of r found by both methodologies is displayed with a dotted line, the straight lines obtained for the Cattell criterion and the eigenvalues of the permuted matrices in the PA methodology are displayed in grey.

For $n \in \{10, 30, 50\}$ and $q \in \{100, 500\}$, 100 $n \times q$ matrices $\boldsymbol{E}$ were generated such that its rows $\boldsymbol{E}_1, \boldsymbol{E}_2, \cdots, \boldsymbol{E}_n$ are i.i.d. q-dimensional zero-mean Gaussian vectors having a covariance matrix $\boldsymbol{\Sigma}$ chosen according to the four previous cases: **Diagonal-Equal**, **Diagonal-Unequal**, **Extra-Diagonal-Equal** or **Extra-Diagonal-Unequal**.

4.3.1 Low rank approximation

The approaches for choosing r described in Section 4.2.5 are illustrated in Figure 4.3 in the Extra-Diagonal-Unequal case. We can see from this figure that both methodologies find the right value of r which is here equal to 5.

To go further, we investigate the behavior of our methodologies from 100 replications of the matrix $\boldsymbol{E}$ for the four different types of $\boldsymbol{\Sigma}$. Figure 4.4 displays the barplots associated to the estimation of r made in the different replications by the two approaches for the different scenarii. We can see from this figure that the PA criterion seems to be slightly more stable than the Cattell criterion when $n \geq 30$. However, in the case where $n = 10$, the PA criterion underestimates the value of r. Moreover, in terms of computational time, the performance of Cattell is much better, see Figure 4.5.

4.3.2 Positions of the non null values

For the four scenarios, the performance of the two approaches: BL and Elbow described in Section 4.2.5 for choosing λ and hence the number of non null values in

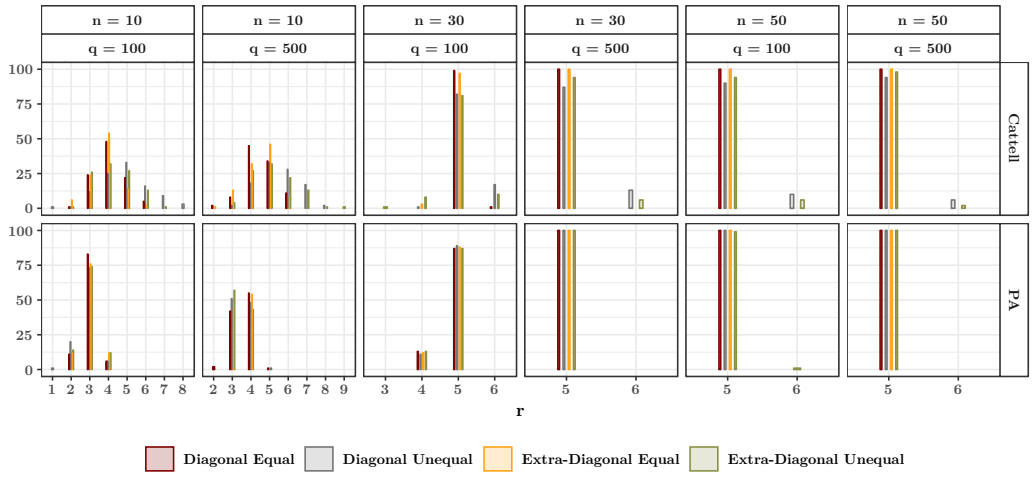

Figure 4.4 – Barplots corresponding to the number of times where each value of r is chosen in the low-rank approximation from 100 replications for the two methodologies in the different scenarii for the different values of n and q.

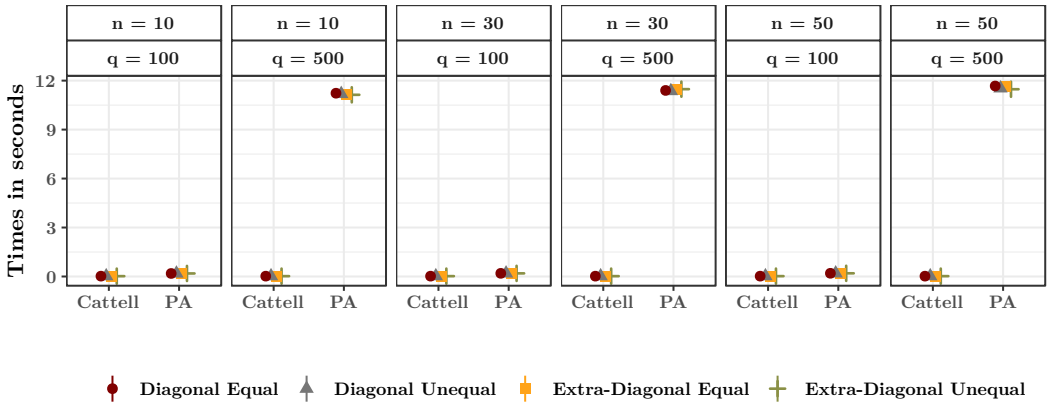

Figure 4.5 – Computational times of PA and Cattell criteria.

$\widetilde{\Sigma}(\lambda)$ is illustrated in Figure 4.6. This figure displays the True Positive Rate (TPR) and the False Positive Rate (FPR) of the methodologies from 100 replications of the matrix $\boldsymbol{E}$ for the four different types of $\boldsymbol{\Sigma}$ and for different values of n and q.

We can see from this figure that the performance of Elbow is on a par with the one of BL except for the case where $n = 10$ for which the performance of Elbow is slightly better in terms of True Positive Rate. Moreover, in terms of computational time, the performance of Elbow is much better, see Figure 4.7.

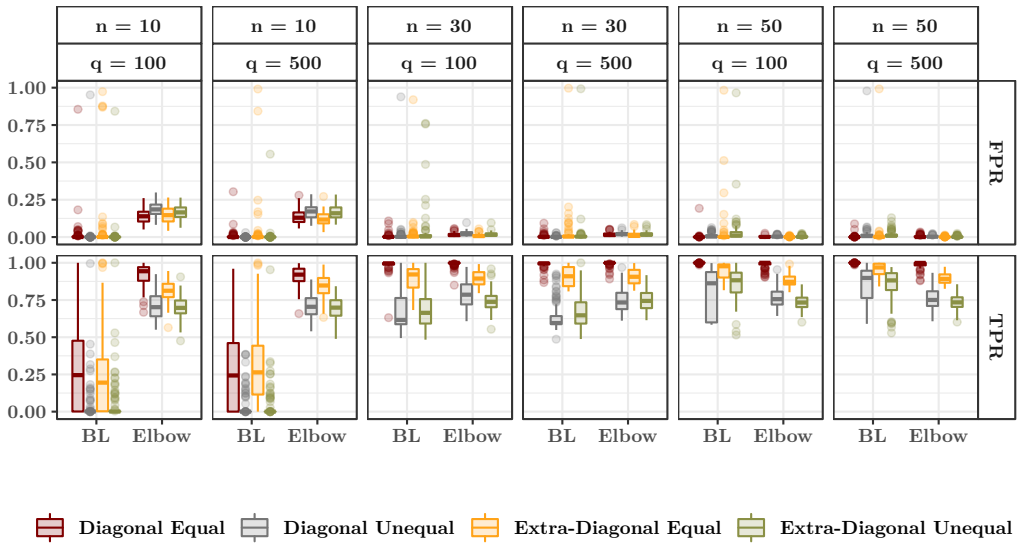

Figure 4.6 – Boxplots comparing the TPR (True Positive Rate) and the FPR (False positive Rate) of the two methodologies proposed to select the parameter λ from 100 replications in the different scenarii.

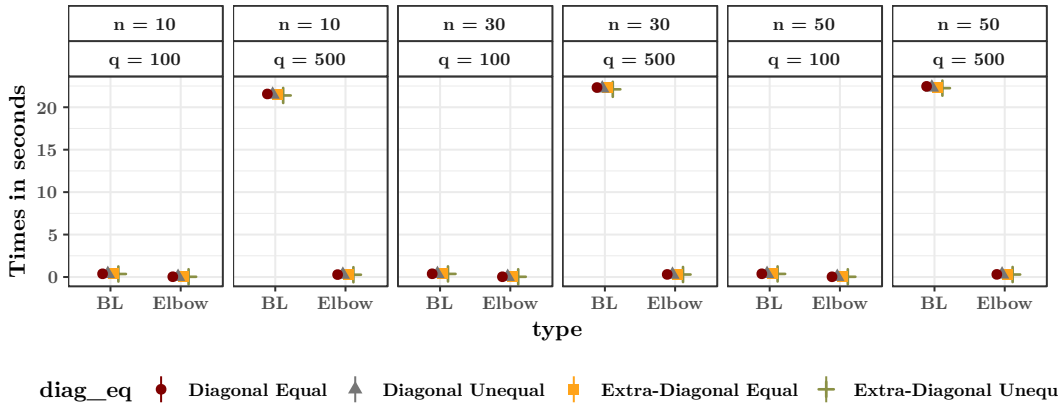

Figure 4.7 – Computational times of Elbow and BL criteria.

4.3.3 Comparison with other methodologies

The goal of this section is to compare the statistical performance of our approach with other methodologies.

Since our goal is to estimate a covariance matrix containing blocks, we shall compare our approach with clustering techniques. Once the groups or blocks have be obtained, $\mathbf{\Sigma}$ is estimated by assuming that the corresponding matrix estimator is block-wise constant except for the diagonal blocks for which the diagonal entries

are equal to 1 and the extra-diagonal terms are assumed to be equal. This gives a great advantage to these methodologies in the **Diagonal-Equal** and in the **Extra-Diagonal-Equal** scenarii. More precisely, let $\rho_{i,j}$ denote the value of the entries in the block having its rows corresponding to Group (or Cluster) i and its columns to Group (or Cluster) j. Then, for a given clustering C:

$$\rho_{i,j} = \begin{cases} \frac{1}{\#C(i)\#C(j)} \sum_{k \in C(i), \ell \in C(j)} R_{k,\ell}, & \text{if } C(i) \neq C(j) \\ \frac{1}{\#C(i)(\#C(i)-1)} \sum_{k \in C(i), \ell \in C(i), k \neq \ell} R_{k,\ell}, & \text{if } C(i) = C(j) \end{cases}, \tag{4.10}$$

where $C(i)$ denotes the cluster i, $\#C(i)$ denotes the number of elements in the cluster $C(i)$ and $R_{k,\ell}$ is the (k,ℓ) entry of the matrix $\boldsymbol{R}$ defined in Equation (4.2).

For the matrices $\boldsymbol{\Sigma}$ corresponding to the four scenarios previously described, we shall compare the statistical performance of the following methods:

— **empirical** which estimates $\boldsymbol{\Sigma}$ by $\boldsymbol{R}$ defined in (4.2),

— **blocks** which estimates $\boldsymbol{\Sigma}$ using the methodology described in this article with the criteria PA and BL for choosing r and λ, respectively,

— **blocks_fast** which estimates $\boldsymbol{\Sigma}$ using the methodology described in this article with the criteria Cattell and Elbow for choosing r and λ, respectively,

— **blocks_real** which estimates $\boldsymbol{\Sigma}$ using the methodology described in this article when r and the number of non null values are assumed to be known which gives access to the best value of λ,

— **hclust** which estimates $\boldsymbol{\Sigma}$ by determining clusters using a hierarchical clustering with the "complete" agglomeration method described in Hastie et al. (2001) and then uses Equation (4.10) to estimate $\boldsymbol{\Sigma}$,

— **Specc** which estimates $\boldsymbol{\Sigma}$ by determining clusters using spectral clustering described in von Luxburg (2007) and estimates $\boldsymbol{\Sigma}$ with Equation (4.10),

— **kmeans** which estimates $\boldsymbol{\Sigma}$ by determining clusters from a k-means clustering approach described in Hastie et al. (2001) and then uses Equation (4.10) to estimate $\boldsymbol{\Sigma}$.

In order to improve the performance of the clustering approaches: **hclust**, **Specc** and **kmeans**, the real number of clusters has been provided to these methods. The performance of the different approaches is assessed using the Frobenius norm of the difference between $\boldsymbol{\Sigma}$ and its estimator.

Figure 4.8 displays the mean and standard deviations of the Frobenius norm of the difference between $\boldsymbol{\Sigma}$ and its estimator for different values of n and q in the four different cases: **Diagonal-Equal**, **Diagonal-Unequal**, **Extra-Diagonal-Equal** and **Extra-Diagonal-Unequal**. We can see from this figure that in the

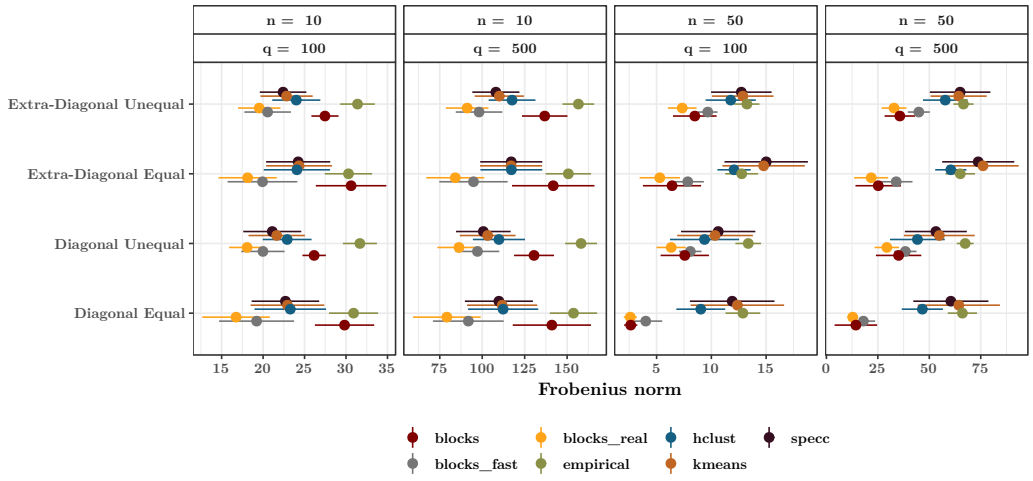

Figure 4.8 – Comparison of the Frobenius norm of $\boldsymbol{\Sigma} - \widehat{\boldsymbol{\Sigma}}$ for different estimators $\widehat{\boldsymbol{\Sigma}}$ of $\boldsymbol{\Sigma}$ and for different $\boldsymbol{\Sigma}$.

case where $n = 10$, the performance of **blocks_fast** is on a par with the one of **blocks_real** and is better than the one of **blocks**. In the case where $n = 50$, the performance of **blocks** is slightly better than the one of **blocks_fast** and is similar to the one of **blocks_real**. Moreover, in all cases, either **blocks_fast** or **blocks** outperforms the other approaches.

Then, the estimators of $\boldsymbol{\Sigma}$ derived from **blocks**, **blocks_fast** and **blocks_real** were compared to the **PDSCE** estimator proposed by Rothman (2012) and implemented in the R package **PDSCE** and to the estimator proposed by Blum et al. (2016b) and implemented in the **FANet** package Blum et al. (2016a). Since the computational burden of **PDSCE** is high for large values of q, we limit ourselves to the **Extra-Diagonal-Equal** case when $n = 30$ and $q = 100$ for the comparison. Figure 4.9 displays the results. We can see from this figure that **blocks**, **blocks_fast** and **blocks_real** provide better results than **PDSCE** and **FANet**. However, it has to be noticed that **PDSCE** is not designed for dealing with block structured covariance matrices but just for providing sparse estimators of large covariance matrices.

4.3.4 Columns permutation

In practice, it may occur that the columns of $\boldsymbol{E}$ consisting of the rows $\boldsymbol{E}_1, \boldsymbol{E}_2, \ldots, \boldsymbol{E}_n$ are not ordered in a way which makes blocks appear in the matrix $\boldsymbol{\Sigma}$. To address this issue, we propose to perform a hierarchical clustering on $\boldsymbol{E}$ beforehand and use the obtained permutation of the observations which guarantees that a cluster plot using this ordering will not have crossings of the branches. Let us denote $\boldsymbol{E}_{ord}$ the matrix $\boldsymbol{E}$ in which the columns have been permuted according to this ordering and $\boldsymbol{\Sigma}_{ord}$ the covariance matrix of each row of $\boldsymbol{E}_{ord}$. Then, we apply our methodology

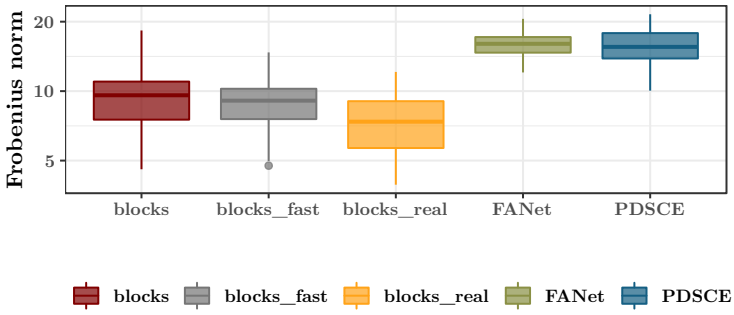

Figure 4.9 – Comparison of the Frobenius norm of $\widehat{\mathbf{\Sigma}} - \mathbf{\Sigma}$ in the **Extra-Diagonal-Equal** case for $n = 30$ and $q = 100$.

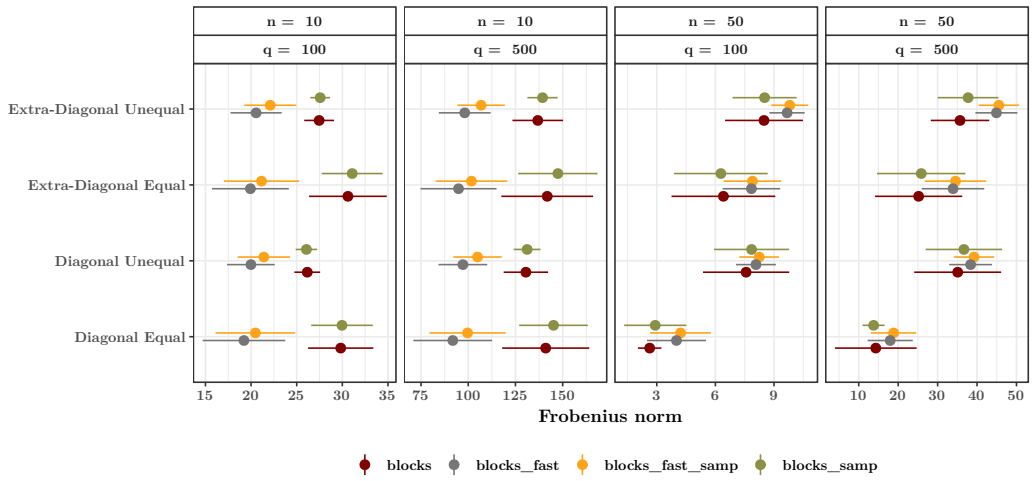

Figure 4.10 – Comparison of the Frobenius norm of $\mathbf{\Sigma} - \widehat{\mathbf{\Sigma}}$, and $\mathbf{\Sigma}_{perm} - \widehat{\mathbf{\Sigma}}_{perm}$.

to $\boldsymbol{E}_{ord}$ which should provide an efficient estimator of $\mathbf{\Sigma}_{ord}$. In order to get an estimator of $\mathbf{\Sigma}$ the columns and rows are permuted according to the ordering coming from the hierarchical clustering.

To assess the corresponding loss of performance, we generated for each matrix $\boldsymbol{E}$ used for making Figure 4.8 a matrix $\boldsymbol{E}_{perm}$ in which the columns of $\boldsymbol{E}$ were randomly permuted. The associated covariance matrix is denoted $\mathbf{\Sigma}_{perm}$. Then, we applied the methodology described in the previous paragraph denoted **blocks_samp** and **blocks_fast_samp** in Figure 4.10 thus providing $\widehat{\mathbf{\Sigma}}_{perm}$. The performance of this new methodology was compared to the methodology that we proposed in the previous sections (denoted **blocks** and **blocks_fast** in Figure 4.10) when the columns of $\boldsymbol{E}$ were not permuted. The results are displayed in Figure 4.10. We can see from this figure that the performance of our approach does not seem to be altered by the permutation of the columns.

4.3.5 Numerical performance

Figure 4.11 displays the computational times for estimating $\mathbf{\Sigma}$ with the methods **blocks** and **blocks_fast** for different values of q ranging from 100 to 3000 and $n = 30$. The timings were obtained on a workstation with 16 GB of RAM and Intel Core i7 (3.66GHz) CPU. Our methodology is implemented in the R package `BlockCov` which uses the R language (R Core Team, 2017) and relies on the R package `Matrix`. We can see from this figure that it takes around 3 minutes to estimate a 1000×1000 correlation matrix.

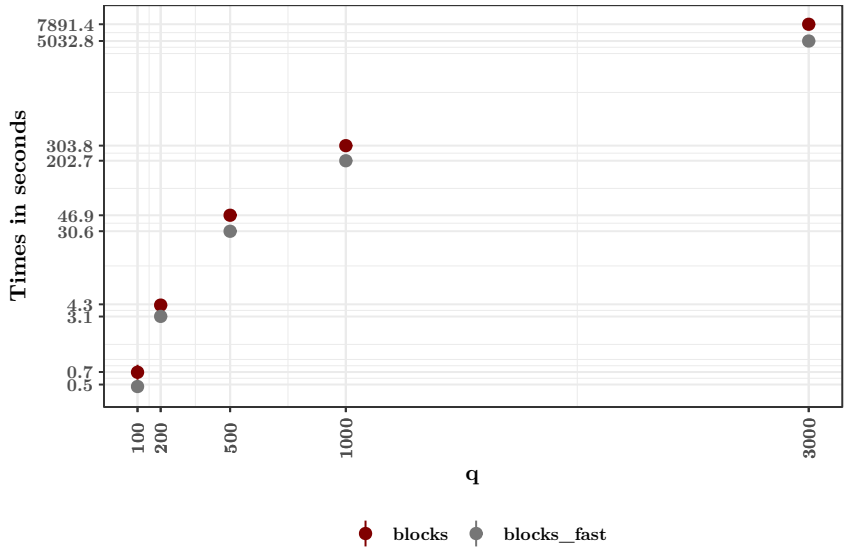

Figure 4.11 – Times in seconds to perform our methodology in the **Extra-Diagonal Unequal** case.

4.3.6 Choice of the threshold t for estimating $\mathbf{\Sigma}^{-1/2}$

Since we are interested in assessing the ability of $\widehat{\mathbf{\Sigma}}_t^{-1/2}$ defined in (4.9) to remove the dependence that may exist between the columns of $\boldsymbol{E}$, we shall consider the Frobenius norm of $\widehat{\mathbf{\Sigma}}_t^{-1/2}\mathbf{\Sigma}\widehat{\mathbf{\Sigma}}_t^{-1/2} - \mathrm{Id}_q$ which should be close to zero, where Id_q denotes the identity matrix of $\mathbb{R}^q$. Figure 4.12 displays the Frobenius norm of $\widehat{\mathbf{\Sigma}}_t^{-1/2}\mathbf{\Sigma}\widehat{\mathbf{\Sigma}}_t^{-1/2} - \mathrm{Id}_q$ for different threshold t. A threshold of 0.1 seems to provide a small error in terms of Frobenius norm. Hence, in the following, t will be equal to 0.1 and $\widehat{\mathbf{\Sigma}}_{0.1}^{-1/2}$ will be referred as $\widehat{\mathbf{\Sigma}}^{-1/2}$.

This technique was applied to all of the estimators of $\mathbf{\Sigma}$ discussed in Section 4.3.3 to get different estimators of $\mathbf{\Sigma}^{-1/2}$. The Frobenius norm of the error $\widehat{\mathbf{\Sigma}}^{-1/2}\mathbf{\Sigma}\widehat{\mathbf{\Sigma}}^{-1/2} - \mathrm{Id}_q$ is used to compare the different estimators obtained by considering the different estimators of $\mathbf{\Sigma}$. The results are displayed in Figure 4.13. We observe from

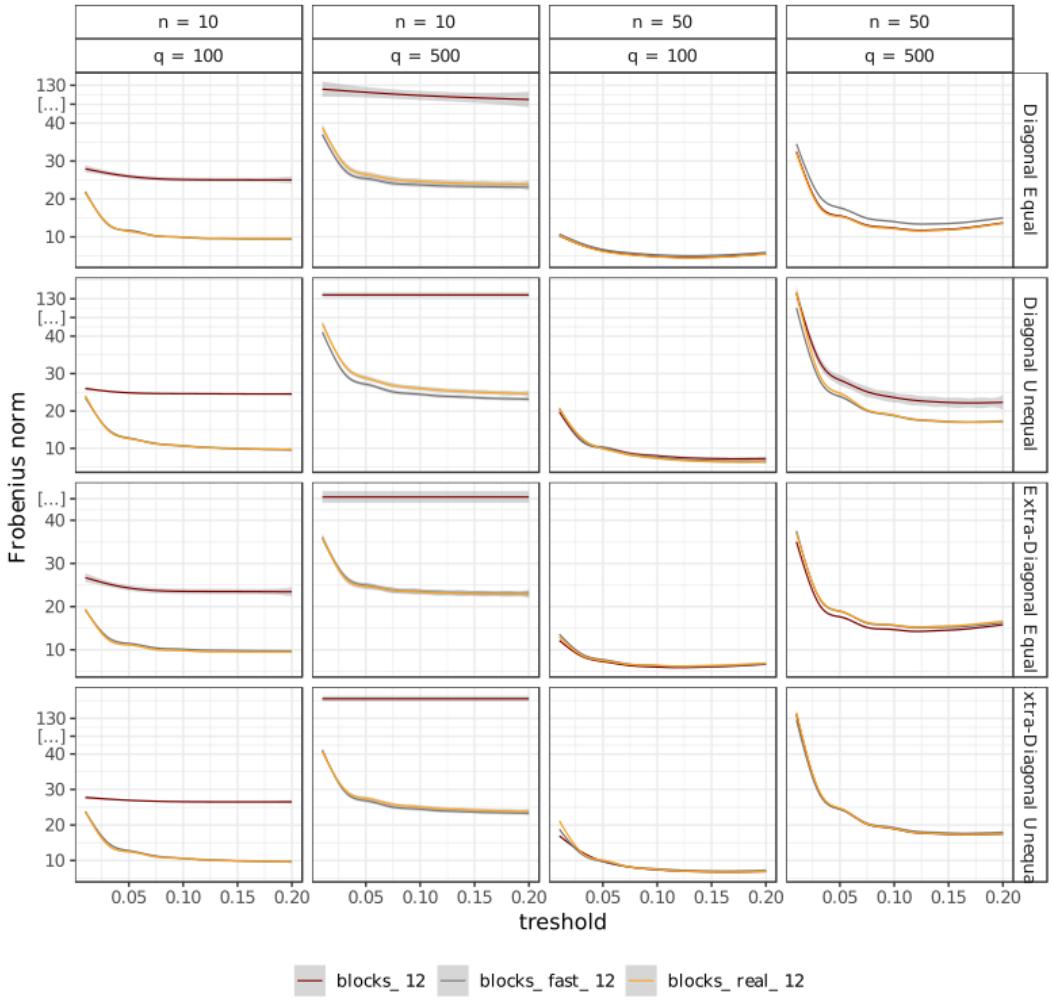

Figure 4.12 – Frobenius norm of $\widehat{\mathbf{\Sigma}}_t^{-1/2}\mathbf{\Sigma}\widehat{\mathbf{\Sigma}}_t^{-1/2} - \mathrm{Id}_q$, where $\widehat{\mathbf{\Sigma}}_t^{-1/2}$ is computed for different thresholds t.

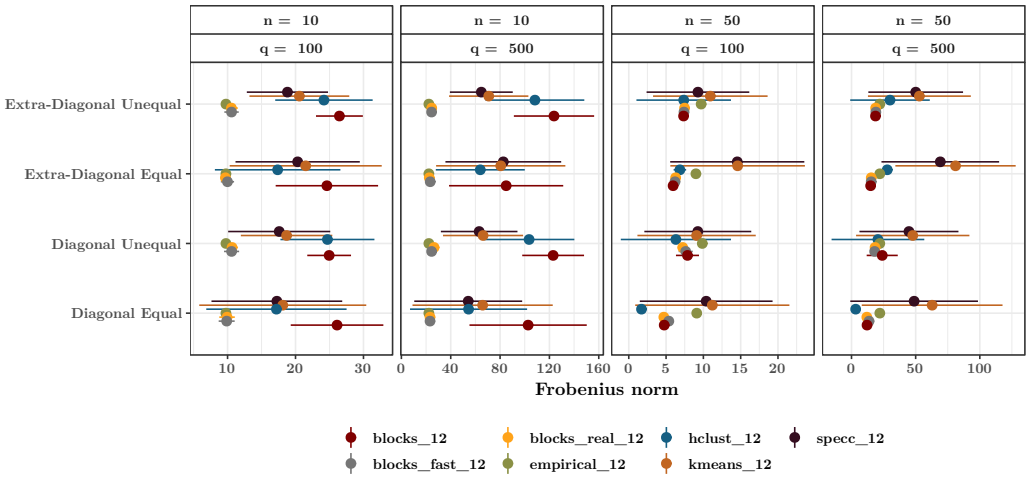

Figure 4.13 – Comparison of the Frobenius norm of the error $\widehat{\boldsymbol{\Sigma}}^{-1/2}\boldsymbol{\Sigma}\widehat{\boldsymbol{\Sigma}}^{-1/2} - \mathrm{Id}_q$, for different estimators $\widehat{\boldsymbol{\Sigma}}$ of $\boldsymbol{\Sigma}$.

this figure that in the case where $n = 10$ the estimators of $\boldsymbol{\Sigma}^{-1/2}$ derived from the **empirical**, the **blocks_fast** and the **blocks_real** estimators of $\boldsymbol{\Sigma}$ perform similarly and seem to be more adapted than the others to remove the dependence among the columns of $\boldsymbol{E}$. However, when $n = 50$, the behavior is completely different. Firstly, in the **Diagonal-Equal** case, the estimator of $\boldsymbol{\Sigma}^{-1/2}$ derived from the **hclust** estimator of $\boldsymbol{\Sigma}$ seems to perform better than the others. Secondly, in the **Diagonal-Unequal** case, the estimator derived from **blocks**, **blocks_fast** and **blocks_real** perform similarly than the one obtained from **hclust**. Thirdly, in the **Extra-Diagonal** case, the estimators derived from **blocks**, **blocks_fast** and **blocks_real** methodology perform better than the other estimators.

Then, the estimators of $\boldsymbol{\Sigma}^{-1/2}$ derived from **blocks**, **blocks_fast** and **blocks_real** were compared to the **GRAB** estimator proposed by Hosseini & Lee (2016). Since the computational burden of **GRAB** is high for large values of q, we limit ourselves to the **Extra-Diagonal-Equal** case when $n = 30$ and $q = 100$ for the comparison. Figure 4.14 displays the results. We can see that **blocks** and **blocks_real** provide better results than **GRAB**. However, it has to be noticed that the latter approach depends on a lot of parameters that were difficult to choose, thus we used the default ones.

4.3.7 Use of $\Sigma^{-1/2}$ to remove the dependence in multivariate linear models

Eventually, we assess the performance of the BlockCov methodology to remove the dependence in the columns of an observation matrix in order to be used for variable selection in the multivariate linear model as it is performed in the `MultiVarSel`

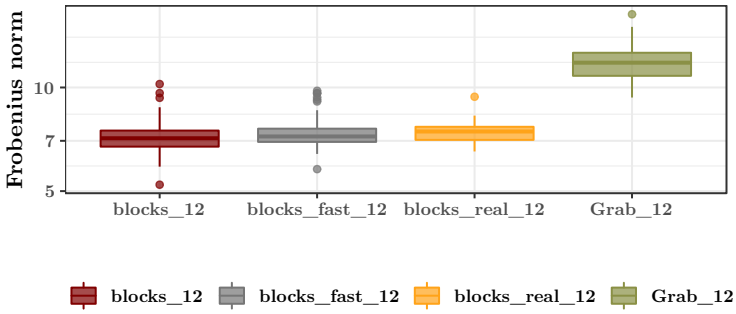

Figure 4.14 – Comparison of the Frobenius norm of $\widehat{\boldsymbol{\Sigma}}^{-1/2}\boldsymbol{\Sigma}\widehat{\boldsymbol{\Sigma}}^{-1/2} - \mathrm{Id}_q$ in the **Extra-Diagonal-Equal** case.

R package:

$$\boldsymbol{Y} = \boldsymbol{X}\boldsymbol{B} + \boldsymbol{E}, \tag{4.11}$$

where $\boldsymbol{Y}$ is a $n \times q$ response matrix, $\boldsymbol{X}$ is a $n \times p$ design matrix, $\boldsymbol{B}$ is a coefficients matrix and $\boldsymbol{E}$ is an error matrix. Here, $\boldsymbol{E}_1, \boldsymbol{E}_2, \cdots, \boldsymbol{E}_n$ are n zero-mean i.i.d. q-dimensional Gaussian random vectors having a covariance matrix $\boldsymbol{\Sigma}$. To achieve this goal, we generate observations $\boldsymbol{Y}$ according to this multivariate linear model. We choose $q = 100$, $p = 3$, $n = 30$ and bX is the design matrix of a one-way ANOVA model.We compared our methodology with the one proposed by Perthame et al. (2016) and implemented in the `FADA` R package Perthame et al. (2019). We shall investigate the effect of the sparsity of $\boldsymbol{B}$ and of the signal to noise ratio (SNR) for the four scenarii defining $\boldsymbol{\Sigma}$ on the selection of the non null values of $\boldsymbol{B}$ in (4.11). Different signal to noise ratios are obtained by multiplying $\boldsymbol{B}$ in (4.11) by a coefficient κ.

Since the results are barely influenced by the scenario chosen for $\boldsymbol{\Sigma}$, only the **Extradiagonal-Equal** case is displayed in Figure 4.15, the other scenarii are available in Annexe 4.6.1. We can see from this figure that when the signal to noise ratio is low and the value of s is high, meaning that there is a lot of non-zero values, the FADA methodology performs better than the BlockCov methodology. Nevertheless, in the three other cases the performance of BlockCov is either better or on a par with the one of FADA methodology.

4.4 Application to "multi-omic" approaches to study seed quality

Climate change could lead to major crop failures in world. In the present study, we addressed the impact of mother plant environment on seed composition. Indeed, seed quality is of paramount ecological and agronomical importance. They

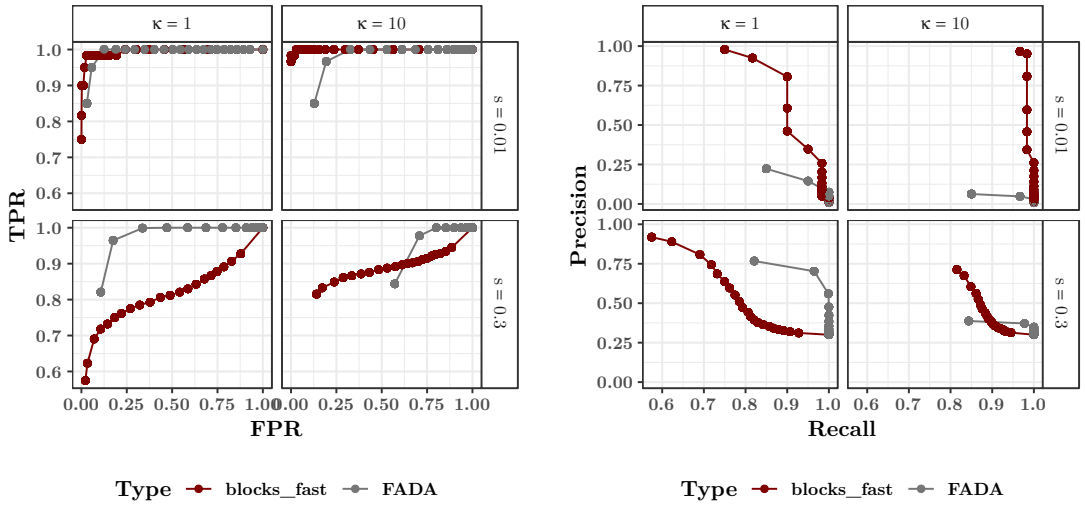

Figure 4.15 – Means of the ROC curves (left) and Precision Recall curves (right) obtained from 100 replications comparing the variables selected by the MultiVarSel strategy using either $\boldsymbol{\Sigma}^{-1/2}$ obtained by BlockCov to remove the dependence or the methodology proposed by FADA methodology. κ is linked to the signal to noise ratio and s denotes the sparsity levels *i.e* the fraction of non-zero elements in $\boldsymbol{B}$.

are the most efficient form of dispersal of flowering plants in the environment. Seeds are remarkably adapted to harsh environmental conditions as long as they are in a quiescent state. Dry mature seeds (so called "orthodox seeds") are an appropriate resource for preservation of plant genetic diversity in seedbanks. It has been reported that the temperature regime during seed production affects agronomical traits such as seed germination potential, see Huang et al. (2014),MacGregor et al. (2015) and Kerdaffrec & Nordborg (2017). In order to highlight biomarkers of seed quality according to thermal environment of the mother plant, Arabidopsis seeds were produced under three temperature regimes (14-16 ^{o}C, 18-22 ^{o}C or 25-28 ^{o}C under a long-day photoperiod). Dry mature seeds were analysed by shotgun proteomic and GC/MS-based metabolomics Durand et al. (2019). The choice to use the model plant, Arabidopsis, was motivated by the colossal effort of the international scientific community for its genome annotation. This plant remains at the forefront of modern genetics, genomics, plant modelling and system biology, see Provart et al. (2016). Arabidopsis provides a very useful basis to study gene regulatory networks, and develop modelling and systems biology approaches for translational research towards agricultural applications.

In this section, we apply our R packages `BlockCov` and `MultiVarSel` Perrot-Dockès et al. (2019) to metabolomic and proteomic data to better understand the impact of the temperature on the seed quality. More precisely, we use the following

modeling for our observations:

$$\boldsymbol{Y} = \boldsymbol{X}\boldsymbol{B} + \boldsymbol{E}, \tag{4.12}$$

where $\boldsymbol{Y}$ is a $n \times q$ matrix containing the responses of the q metabolites (resp. the q proteins) for the n samples with $n = 9$, $q = 199$ (resp. $q = 724$) for the metabolomic (resp. proteomic) dataset, $\boldsymbol{X}$ is a $n \times 3$ design matrix of a one-way ANOVA model, such that its first (resp. second, resp. third) column is a vector which is equal to 1 if the corresponding sample grows under low (resp medium, resp. elevated) temperatures and 0 otherwise. $\boldsymbol{B}$ is a coefficient matrix and $\boldsymbol{E}$ is such that its n rows $\boldsymbol{E}_1, \boldsymbol{E}_2, \cdots, \boldsymbol{E}_n$ are n zero-mean i.i.d. q-dimensional random vectors having a covariance matrix $\boldsymbol{\Sigma}$. We used our R package `BlockCov` to estimate $\boldsymbol{\Sigma}$ and $\boldsymbol{\Sigma}^{-1/2}$ assuming that there exists a latent block structure in the covariance matrix of the rows of $\boldsymbol{E}$. More precisely, we assume that there exists some groups of metabolites (resp. proteins) having the same behavior since they belong to the same biological process. Then, we plugged this estimator into our R package `MultiVarSel` to obtain a sparse estimation of $\boldsymbol{B}$. Thanks to this estimator of $\boldsymbol{B}$, we could identify the metabolites (resp. proteins) having a higher (resp. lower) concentration when the temperature is high or low.

4.4.1 Results obtained for the metabolomic data

We first estimated the matrices $\boldsymbol{\Sigma}$ and $\boldsymbol{\Sigma}^{-1/2}$ associated to $\boldsymbol{E}$ defined in Equation (4.12) by using the methodology developed in this paper, namely the `BlockCov` package. By the results of Section 4.3, we know that the PA and BL approaches performed poorly when $n = 10$. Since here $n = 9$, we used the Cattell and Elbow criteria to choose r and λ, respectively. The results are displayed in Figure 4.16. The Cattell criterion chooses $r = 7$ and the Elbow criterion chooses $\lambda = 0.472$, which implies that among the 19701 coefficients of the correlation matrix only 6696 values are considered as non null values.

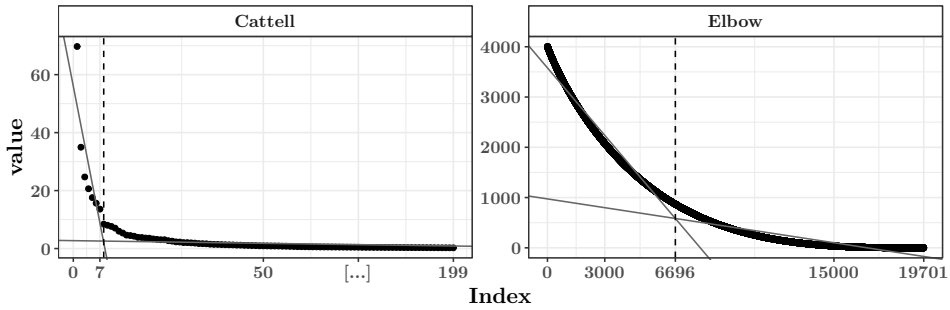

Figure 4.16 – Illustration of the Cattell and Elbow criteria.

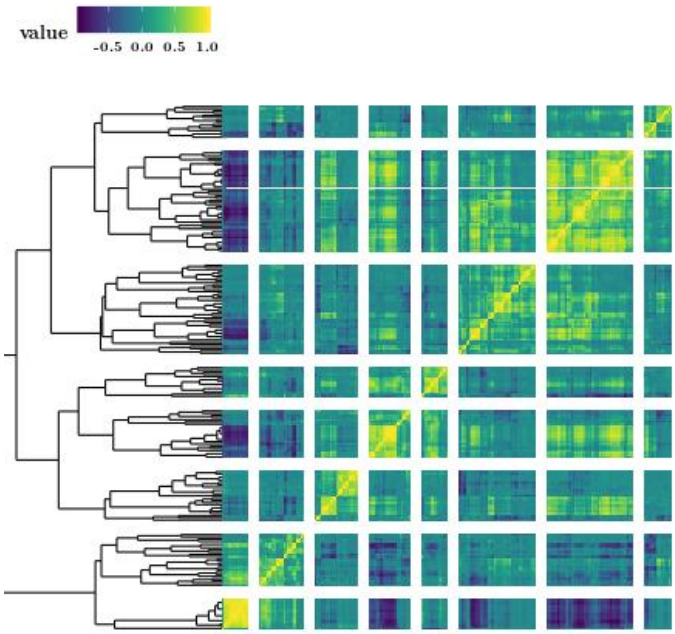

Figure 4.17 – Estimator of the correlation matrix $\boldsymbol{\Sigma}$ of the rows of $\boldsymbol{E}$ once the rows and the columns have been permuted according to the ordering provided by the hierarchical clustering.

The estimation of $\boldsymbol{\Sigma}$ obtained with our methodology is displayed in Figure 4.17 once the rows and the columns have been permuted according to the ordering provided by the hierarchical clustering to make visible the latent block structure.

Using the estimator of $\boldsymbol{\Sigma}^{-1/2}$ provided by the `BlockCov` package in the R package `MultiVarSel` provides the sparse estimator of the matrix $\boldsymbol{B}$ defined in Model 4.12 and displayed in Figure 4.18. We can see from this figure that for the metabolite X5MTP the coefficient of the matrix $\widehat{\boldsymbol{B}}$ is positive when the temperature is high which means that the production of the metabolite X5MTP is larger in high temperature conditions than in low temperature conditions.

In order to go further in the biological interpretation, we wanted to better understand the underlying block structure of the estimator of the correlation matrix of the residuals based on metabolite abundances $\widehat{\boldsymbol{\Sigma}}$. Thus, we applied a hierarchical clustering with 8 groups to this matrix in order to split it into blocks. The corresponding dendogram is on the left part of Figure 4.17. The matrix containing the correlation means within and between the blocks or groups of metabolites is displayed in Figure 4.19. The composition of the metabolites groups is available in Appendix 4.6.2.

Interestingly, we could observe that X5MTP belongs to Group 6 which displays

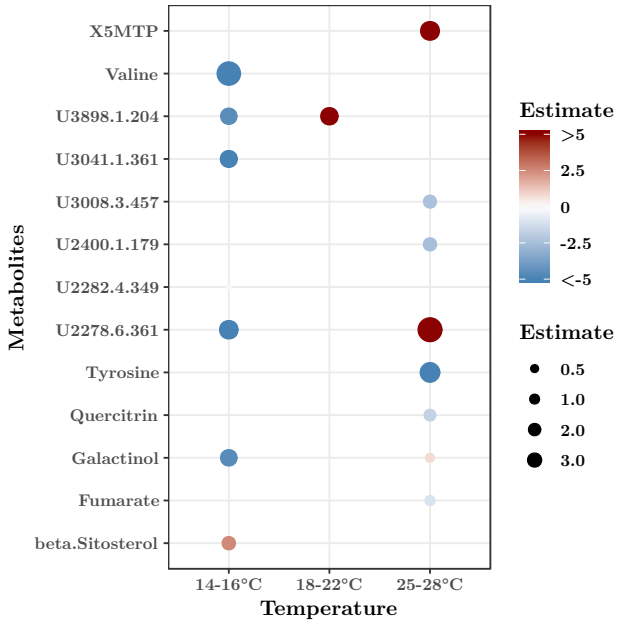

Figure 4.18 – Sparse estimator of the coefficients matrix $\boldsymbol{B}$ obtained thanks to the package `MultiVarSel` with a threshold of 0.95.

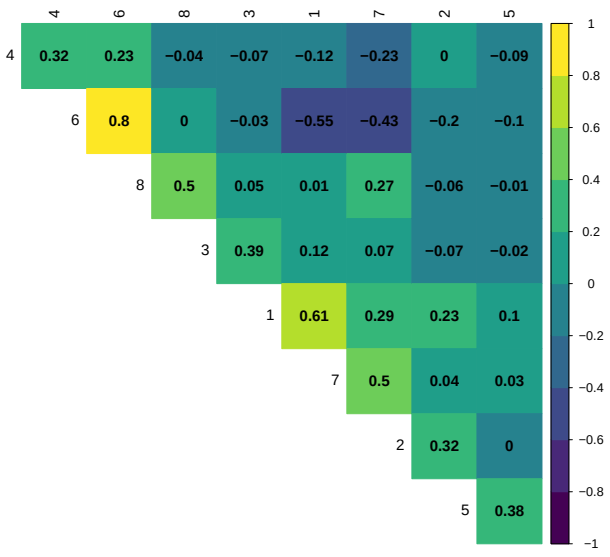

Figure 4.19 – Means of the correlations between the groups of metabolites.

an high correlations mean equal to 0.8 between the 14 metabolites that make it up. At least 6 metabolites of this group belong to the same family, namely glucosinolates (i.e. X4MTB, 4-methylthiobutyl glucosinolate; X5MTP, 5-methylthiopentyl glucosinolate; X6MTH, 6-methylthiohexyl glucosinolate; X7MTH, 7-methylthiohexyl glucosinolate; X8MTO, 8-methylthiooctyl glucosinolate; UGlucosinolate140.1, unidentified glucosinolate). Glucosinolates (GLS) are specialized metabolites found in Brassicaceae and related families (e.g. Capparaceae), containing a β-thioglucose moiety, a sulfonated oxime moiety, and a variable aglycone side chain derived from a α-amino acid. These compounds contribute to the plant's overall defense mechanism, see Wittstock & Halkier (2002). Methylthio-GLS are derivated from methionine. Methionine is elongated through condensation with acetyl CoA and then, are converted to aldoximes through the action of individual members of the cytochrome P450 enzymes belonging to CYP79 family, see Field et al. (2004). The aldoxime undergoes condensation with a sulfur donor, and stepwise converted to GLS, followed by the side chain modification. The present results suggest that the accumulation of methionine-derived glucosinolate family is strongly coordinated in Arabidopsis seed. Moreover, we can see that they are influenced by the effect of the mother plant thermal environment.

4.4.2 Results obtained for the proteomic data

The same study was conducted on the proteomic data. The estimator of the correlation matrix of the residuals based on proteine abundances $\widehat{\mathbf{\Sigma}}$ obtained with our methodology is displayed in Figure 4.20 once the rows and the columns have been permuted according to the ordering provided by the hierarchical clustering to make visible the latent block structure. To better understand the underlying block structure of $\widehat{\mathbf{\Sigma}}$, we applied a hierarchical clustering with 9 groups to this matrix in order to split it into blocks. The corresponding dendogram is on the left part of Figure 4.20.

The matrix containing the correlation means within and between the blocks or groups of proteins is displayed in Figure 4.21. We can see from this figure that Group 8 has the highest correlation mean equal to 0.47. It consists of 34 proteins which are given in Appendix 4.6.3.

A basic gene ontology analysis (http://geneontology.org/) showed that proteins involved in response to stress (biotic and abiotic), in nitrogen and phosphorus metabolic processes, in photosynthesis and carbohydrate metabolic process and in oxidation-reduction process are overrepresented in this group, see Figure 4.22. Thus, the correlation estimated within Group 8 seems to reflect a functional coherence of the proteins of this group.

The variable selection in the multivariate linear model using the R package `MultiVarSel` provided 31 proteins differentially accumulated in seeds produced un-

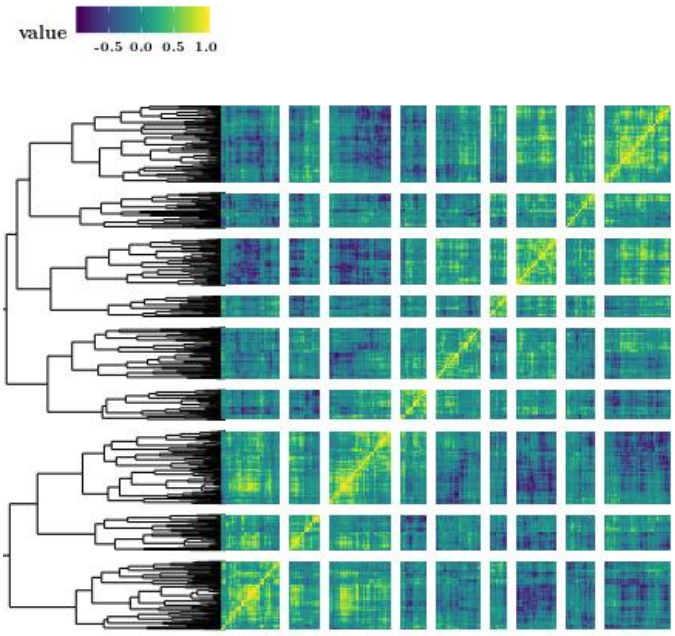

Figure 4.20 – Estimator of the correlation matrix of the residuals of the protein accumulation measures once the rows and the columns of the residual matrix have been permuted according to the ordering provided by the hierarchical clustering.

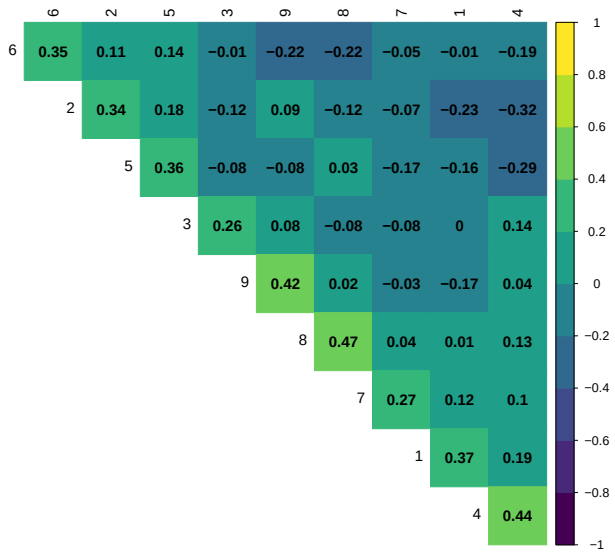

Figure 4.21 – Means of the correlations between the groups of proteins.

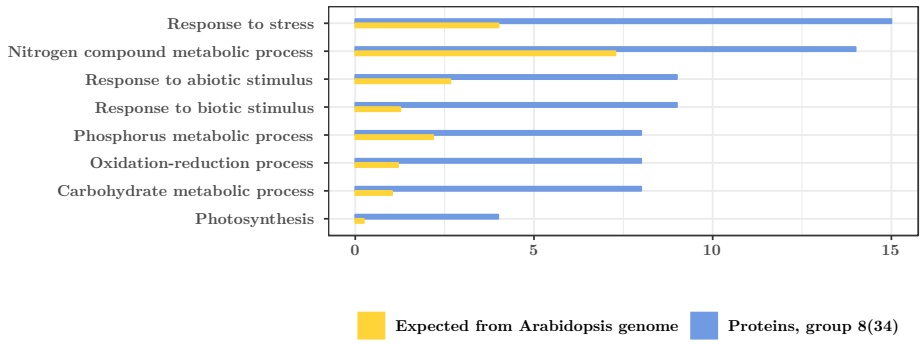

Figure 4.22 – Gene ontology (GO) term enrichment analysis of the 34 proteins belonging to Group 8. Data from PANTHER overrepresentation test (http://www.geneontology.org); One uploaded id (i.e. AT5G50370) mapped to two genes. Thus, GO term enrichment was performed on 35 elements. Blue bars: observed proteins in Group 8; Orange bars: expected result from the reference Arabidopsis genome.

der low, medium or elevated temperature. An aspartyl protease (AT3G54400), belongs to both, the Group 8 and to the proteins selected by `MultiVarSel`. This cell wall associated protein was up-acccumulated in dry seeds produced under low temperature. The gene encoding for this protease was described as a cold responsive gene assigned to the C-repeat binding factor (CBF) regulatory pathway, see Vogel et al. (2006). This pathway is requested for regulation of dormancy induced by low temperatures, see Kendall et al. (2011). Consistently, in Figure 4.23, two other proteins related to cell wall organization, a beta-glucosidase (BGLC1, AT5G20950) and a translation elongation factor (eEF-1Bβ1, AT1G30230) were differentially accumulated in seeds produced under contrasted temperature. eEF-1Bβ1 is associated to plant development and is involved in cell wall formation, see Hossain et al. (2012). These results suggest that cell wall rearrangements occur under temperature effect during seed maturation.

As displayed in Figure 4.23, 6 other proteins involved in mRNA translation: AT1G02780, AT1G04170, AT1G18070, AT1G72370, AT2G04390 and AT3G04840 were selected. The absolute failure of seed germination in the presence of protein synthesis inhibitors underlines the essential role of translation for achieving this developmental process, see Rajjou et al. (2004). Previous studies highlighted the importance of selective and sequential mRNA translation during seed germination and seed dormancy, see Galland et al. (2014), Bai et al. (2017) and Bai et al. (2018). Thus, exploring translational regulation during seed maturation and germination through the dynamic of mRNA recruitment on polysomes or either neosynthesized proteome are emerging fields in seed research.

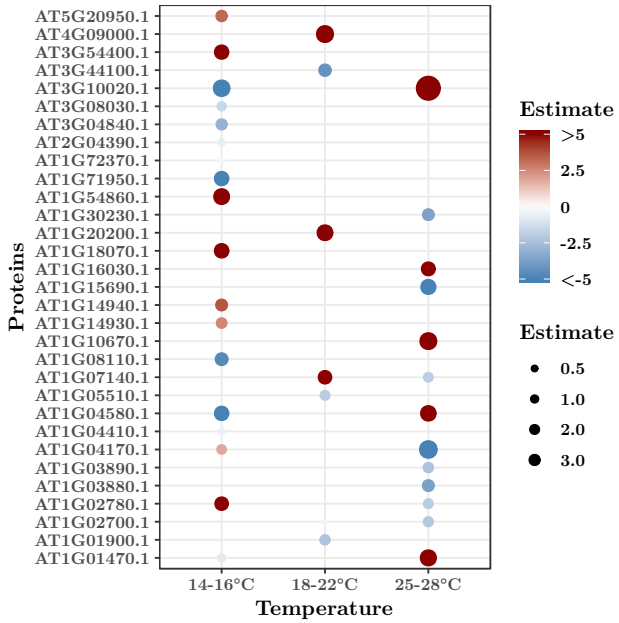

Figure 4.23 – Values of the coefficients obtained using the package `MultiVarSel` with a threshold of 0.95 on the proteomic dataset.

4.5 Conclusion

In this paper, we propose a fully data-driven methodology for estimating large block structured sparse covariance matrices in the case where the number of variables is much larger than the number of samples without limiting ourselves to block diagonal matrices. Our methodology can also deal with matrices for which the block structure only appears if the columns and rows are permuted according to an unknown permutation. Our technique is implemented in the R package `BlockCov` which is available from the Comprehensive R Archive Network and from GitHub. In the course of this study, we have shown that `BlockCov` is a very efficient approach both from the statistical and numerical point of view. Moreover, its very low computational load makes its use possible even for very large covariance matrices having several thousands of rows and columns.

Acknowledgments

We thank the members of the EcoSeed European project (FP7 Environment, Grant/Award Number: 311840 EcoSeed, Coord. I. Kranner). IJPB was supported by the Saclay Plant Sciences LABEX (ANR-10-LABX-0040-SPS). We also thank the people who produced the biological material and the proteomic and metabolomic analysis. In particular, we would like to thank the Warwick University (UWAR,

Finch-Savage WE and Awan S) for the production of seeds, the Plant Observatory-Biochemistry platform (IJPB, Versailles; Bailly M, Cueff G) for having prepared the samples for the proteomics and metabolomics, the PAPPSO Proteomic Plateform (GQE-Moulon; Balliau T, Zivy M) for mass spectrometry-based proteome analysis and the Plant Observatory-Chemistry/Metabolism platform (IJPB, Versailles; Clement G) for the analysis of GC/MS-based metabolome analyses.

4.6 Appendix

4.6.1 Variable selection performance

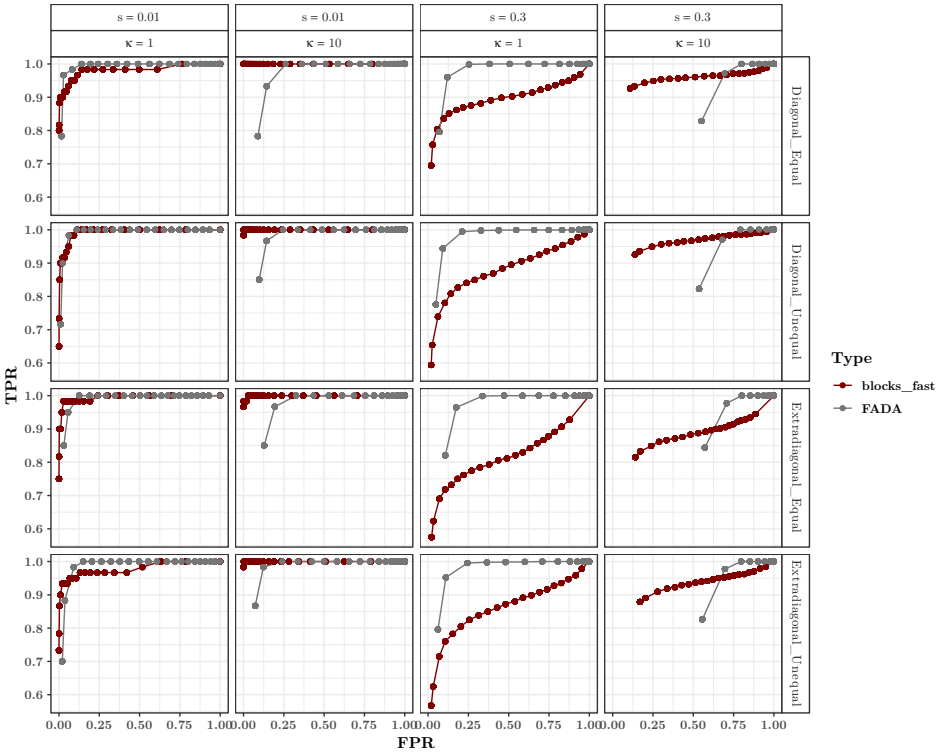

Figure 4.24 – Means of the ROC curves obtained from 100 replications comparing the variables selected by the MultiVarSel strategy using either $\boldsymbol{\Sigma}^{-1/2}$ obtained by BlockCov to remove the dependence or the methodology proposed by FADA methodology. κ is linked to the signal to noise ratio and s denotes the sparsity levels *i.e* the fraction of non-zero elements in $\boldsymbol{B}$.

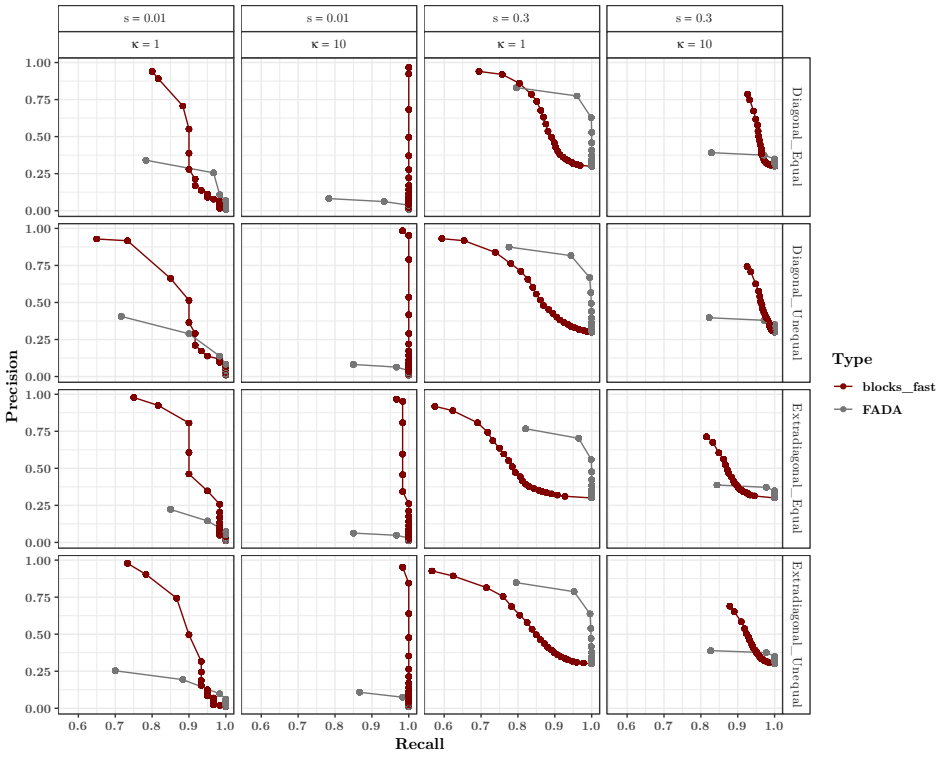

Figure 4.25 – Means of the precision recall curves obtained from 100 replications comparing the variables selected by the MultiVarSel strategy using either $\boldsymbol{\Sigma}^{-1/2}$ obtained by BlockCov to remove the dependence or the methodology proposed by FADA methodology. κ is linked to the signal to noise ratio and s denotes the sparsity levels *i.e* the fraction of non-zero elements in $\boldsymbol{B}$.

4.6.2 Groups of metabolites

Group 1	Group 2	Group 3	Group 4
Alanine	Arginine	Glutamate	beta.Sitosterol
Asparagine	Cystein	alpha.Tocopherol	Campesterol
Aspartate	Gaba	gamma.Tocopherol	Eicosanoate
Glycine	Glutamine	Linolenic.acid	Heptadecanoate
Isoleucine	Tryptophan	H2SO4	Stearic.acid
Leucine	Linoleic.acid	X2.Oxoglutarate	Tetracosanoate
Lysine	Quercetin	Mannitol	BenzoylX.Glucosinolate.3
Phenylalanine	BenzoylGlucosinolate.3Breakdown	Urea	Sulfite
Proline	Nonanenitrile.9.methylthio	Fructose.6.P	U2609.4.361
Serine	UGlucosinolatebreakdown140.5	Digalactosylglycerol	U3122.4.202.I3M.
Threonine	X2.Hydroxyglutarate	Galactinol	dihydroxybenzoate
Tyrosine	Citrate	Galactosylglycerol	beta.indole.3.acetonitrile
Valine	Erythronate	Rhamnose	U1837.6.368
X5..methylthio.pentanenitrile	Galactonate	Stachyose	U1841.9.394
Octanenitrile.8.methylthio	Gluconate	Sucrose	U2003.8.293
UGlucosinolatebreakdown140.4	Glycerate	U1093.6.147	U2371.1.361
Succinate	Malate	U1124.3.140	U2375.6.191
Threonate	Allantoin	U1530.2.314	U2513.2.296
Arabitol	Erythritol	U2053.6.321.1	U2513.2.296.1
myo.Inositol	Ethanolamine	U2109.3.305	U2692.9.361
Glycerol.3.P	Sorbitol	U2197.2.494	U2798.377
myo.Inositol.1.P	Threitol	U2315.2.245	U2942.2.556
Phosphate	Xylitol	U3898.1.204	U3063.0.361
U2206.2.299	Ethylphosphate		U3415.9.498
Fructose	Glucose		
Glucopyranose..H2O.	Mannose		
U1154.3.156	Raffinose		
U1393.172	Ribose		
U1541.8.263	U1127.5.140		
U1647.2.403	U1172.9.281		
U1705.2.319.pentitol.	U1559.4.217		
U1729.0.273	U1628.9.233		
U1816.2.228	U1849.2.285		
U1859.2.246	U1927.0.204		
U2076.9.204	U1939.1.210		
U2170.6.361	U1983.0.217		
U2184.1.299	U1983.0.217.1		
U2251.5.361	U2012.7.361		
U2278.6.361	U2282.4.349		
U2550.7.149	U2400.1.179		
U2731.2.160	U2779.9.361		
U2857.8.342			
U2929.1.297			
U3041.1.361			
U3080.7.361			
U3100.8.361			

Group 5	Group 6	Group 7	Group 8
Quercitrin	X4MTB	BenzoylGlucosinolate.2Breakdown	Maleate
Dehydroascorbate	X5MTP	Hexanenitrile.6methylthio	Pentonate.4
Fumarate	X6MTH	Sinapinate.trans	U1408.4.298
Sinapinate.cis	X7MTH	Anhydroglucose	U1617.8.146
Arabinose	X8MTO	U1125.1.140	U1767.3.243
Galactose	UGlucosinolate140.1	U1290.198	U1904.9.204
U1127.4.169	U1129.9.184	U1371.5.151	U2828.8.361
U1718.0.157	U1270.1.240	U1549.7.130	U2839.3.312
U1931.5.202	U1897.2.327	U1568.5.313	U2882.5.297
U2261.0.218	U2473.361	U1592.8.217	U3008.3.457
U2412.1.157	U2529.8.361	U1700.6.288	U3168.2.290
U2588.9.535	U2756.4.271	U1759.4.331	U3218.5.297
U2688.5.333	U2924.3.361	U1852.0.217	U3910.6.597.Trigalactosylglycerol.
U3213.1.400	U3279.7.361	U1872.1.204.methyl.hexopyranoside.	U2443.7.217
U1380.5.184		U1958.217	
		U2053.6.321	
		U2087.6.321	
		U2150.9.279	
		U2271.6.249	
		U3188.1.361	
		U3701.368	
		U4132.5.575	

4.6.3 Groups of proteins

Proteins present in group 8 : AT1G14170.1 , AT1G20260.1 , AT1G42970.1 , AT1G47980.1 , AT1G55210.1 , AT1G75280.1 , AT2G19900.1 , AT2G22240.1 , AT2G28900.1 , AT2G32920.1 , AT2G37970.1 , AT3G12580.1 , AT3G13930.1 , AT3G26650.1 , AT3G26720.1 , AT3G44300.1 , AT3G47930.1 , AT3G54400.1 , AT3G55800.1 , AT4G16760.1 , AT4G20830.1 , AT4G25740.1 , AT4G34870.1 , AT4G35790.1 , AT5G11880.1 , AT5G12040.1 , AT5G14030.1 , AT5G17380.1 , AT5G22810.1 , AT5G26000.1 , AT5G50370.1 , AT5G66190.1 , AT5G67360.1 , ATCG00480.1.

Chapitre 5

Etude du dialogue entre cellules dendritiques et lymphocytes Th

Production scientifique

Ce chapitre est un résumé de l'article accepté au journal Cell : M. Grandclaudon*, M. Perrot-Dockès*, C. Trichot, O. Mostafa-Abouzid, W. Abou-Jaoudé, F. Berger, P. Hupé, D. Thieffry, L. Sansonnet, J. Chiquet, C. Lévy-Leduc, V. Soumelis *A quantitative multivariate model of human dendritic cell-T helper cell communication.*

** : ces auteurs ont contribué de manière égale à cette publication* Pour plus de détails, un lecteur intéressé peut se référer à l'article complet disponible en Annexe.

5.1 Introduction

La communication entre cellules peut se faire par l'échange de signaux moléculaires produits par une cellule donnée et transmis à une cellule cible qui émettra d'autres signaux en réponse. Pour simplifier le discours nous appellerons « input » les signaux émis par la cellule donnée et « output » les signaux émis en réponse par la cellule cible. Ce mode de communication requiert l'analyse de signaux multiples émis par les cellules. Cependant, la plupart des études sont univariées et se concentrent sur l'effet d'un input ou d'un groupe d'inputs sur un output. En négligeant la diversité des inputs et des outputs elles ne permettent pas d'approcher un contexte réaliste de communication cellulaire.

L'objectif de ce projet est de s'approcher de ce contexte réaliste pour mieux comprendre la communication entre les cellules dendritiques (DC) et les lymphocytes T-helper (Th). Une brève introduction à l'immunologie et au dialogue entre les cellules dendritiques et les lymphocytes Th est disponible en introduction de cette thèse dans la partie 1.4. Une vision simple de ce dialogue est que lorsqu'une cellule dendritique rencontre un agent pathogène elle transmet divers signaux à un lymphocyte Th, dit naïf, qui en réponse va se différencier en un profil donné et émettre d'autres signaux spécifiques. Ces signaux varient en fonction du type d'agent pathogène

rencontré par la cellule dendritique. Plus précisément, en présence d'un pathogène intracellulaire, donc d'une maladie auto-immune, la DC va sécréter de l'interleukine 12 (IL12), une fois capté par le lymphocyte T naïf celui-ci va se différencier en lymphocytes Th1 et sécréter notamment de l'interféron gamma (IFNg) pour combattre cette maladie auto-immune. De même, en présence d'un parasite extra-cellulaire, les DC vont sécréter des signaux qui vont amener le lymphocyte T naïf à se différencier en un lymphocyte Th2 qui va sécréter des interleukines 4, 5, 13 (IL4, IL5, IL13). De nombreuses études ont mis en avant d'autres profils tels que le profil Th17 induit par la présence d'IL6, TNFa, IL23, TGFb qui sécrètent IL17A et IL17F pour répondre à la présence de bactéries et de champignons extérieurs (voir Park et al., 2005). La figure 1.4 qui est une version simplifiée de la figure 1 de Leung et al. (2010) montre différents profils Th décrits.

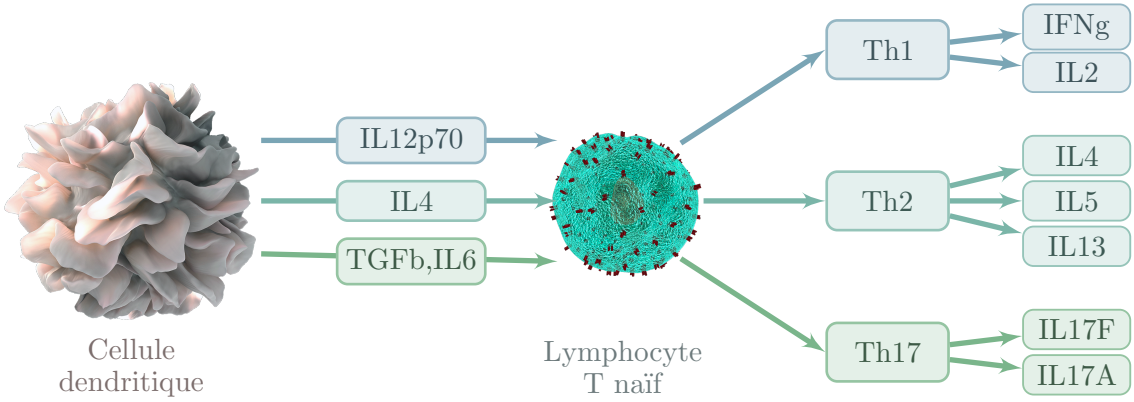

FIGURE 5.1 – Les différents profils Th

5.2 Description des données

5.2.1 Protocole expérimental

Dans le cadre d'une collaboration avec Vassili Soumelis et Maximilien Grandclaudon de l'Institut Curie nous avons eu accès à des données mesurant la réponse des lymphocytes Th aux signaux des DC sous diverses conditions de perturbations visant à reproduire in vitro, l'environnement in situ et in vivo. En pratique, un signal qui agit sur l'activité de la DC (ci-après appelé perturbateur) est injecté à un groupe de DC qui va sécréter des signaux en réponse, 24h après l'insertion du perturbateur dans la culture des cellules dendritiques 36 signaux sécrétés par les DC sont mesurés. Tous les signaux sécrétés par les DC sont alors mis en présence d'un groupe de lymphocyte Th qui va sécréter à nouveau des signaux. Six jours plus tard, 17 des signaux sécrétés par les lymphocytes Th sont mesurés ainsi que l'expansion cellulaire. Pour simplifier le discours nous appellerons par la suite « inputs » les

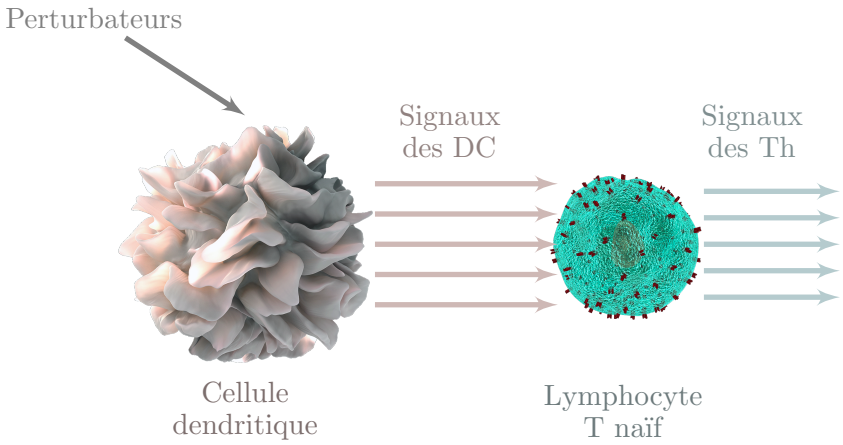

FIGURE 5.2 – Représentation du protocole expérimental et du dialogue entre DC et lymphocytes Th

signaux des DC et « outputs » les signaux des lymphocytes Th et l'expansion cellulaire. Cette expérience a été répétée sur 428 couples groupe de DC / groupe de lymphocytes Th . La figure 5.2 résume ce protocole.

5.2.2 Une grande diversité

Afin de mieux comprendre la communication entre les DC et les lymphocytes Th les données visent à capter une grande diversité dans la sécrétion des DC et donc probablement une grande diversité dans les signaux sécrétés par les lymphocytes Th. Pour cela, on fait varier le perturbateur injecté dans les DC ainsi que le type de DC dans le protocole expérimental décrit section 5.2.1. Plus précisément on utilise des DC provenant du sang et des DC dérivées des monocytes, ainsi que différentes combinaisons et doses de 14 perturbateurs. Les 428 échantillons sont finalement obtenus à l'aide de 82 couples types de DC / types de perturbateurs. La diversité que cela engendre sur les signaux des cellules dendritiques (resp. des lymphocytes Th) est représentée par une heatmap des moyennes sur plusieurs réplicats des différents signaux dans les 82 conditions de perturbation (voir à gauche (resp. à droite) de la figure 5.3).

5.3 Modélisation

Nous appliquons ensuite la méthode décrite dans les chapitres 2 et 3 et implémentée dans le package `MultiVarSel` pour modéliser les valeurs des signaux des lymphocytes Th en fonction des signaux des DC. Cette méthode est fondée sur le modèle 1.2 :

$$\boldsymbol{Y} = \boldsymbol{X}\boldsymbol{B} + \boldsymbol{E},$$

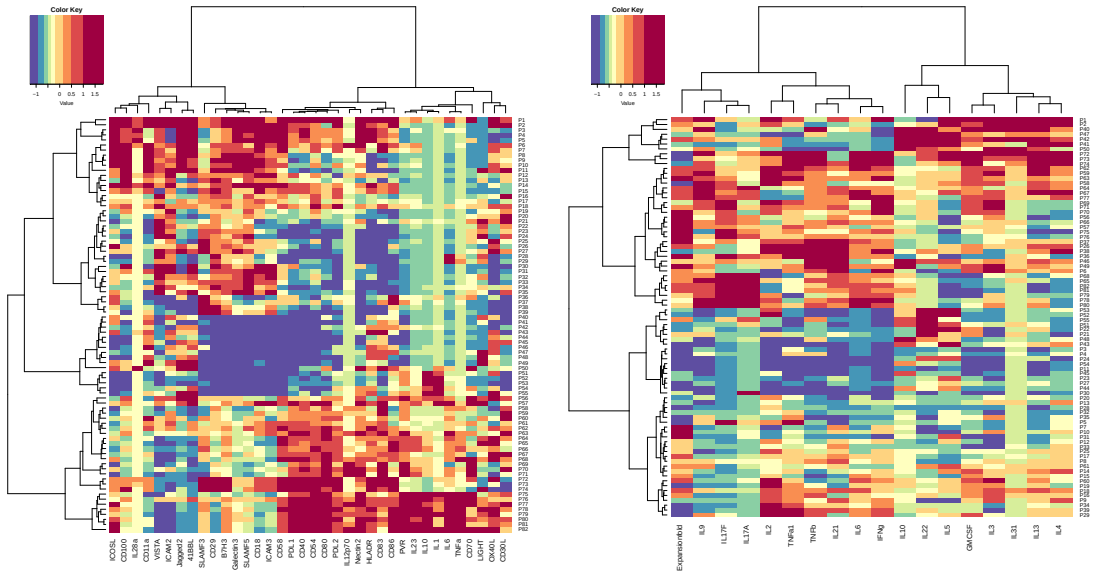

FIGURE 5.3 – Heatmap représentant la moyenne des différents signaux sur les 82 conditions de perturbations obtenues à partir de plusieurs réplicats. Les signaux des DC sont représentés à gauche et ceux des Th à droite.

où la matrice $\boldsymbol{Y}$ (resp. la matrice $\boldsymbol{X}$) contient les concentrations des $q = 18$ signaux des lymphocytes Th (resp. $p = 36$ signaux des DC) pour les 428 échantillons. Dans ce cas $n = 428 \gg q = 18$, la matrice de covariance empirique de $\boldsymbol{E}$ est donc utilisée pour modéliser la matrice de covariance des lignes de $\boldsymbol{E}$. En prenant en compte des connaissances immunologiques et afin de garder de nombreuses associations entre les signaux des lymphocytes Th et des cellules dendritiques le seuil de la *stability selection* est ici fixé à 0.65. Une fois les variables sélectionnées les coefficients sont alors ré-estimés en utilisant la méthode des moindres carrés ordinaires. Les valeurs des coefficients ainsi obtenus sont disponibles dans la figure 5.4. En appliquant un clustering hiérarchique sur les valeurs de ces coefficients ainsi obtenues on retrouve des groupes contenant les signaux des lymphocytes Th caractéristiques des différents profils (voir figure 5.1), ce qui montre la cohérence de nos résultats car il est attendu que des signaux qui appartiennent au même profil Th soient expliqués par les mêmes signaux de DC.

5.4 Validation biologique

Une prédiction intéressante du modèle est l'association positive entre IL12p70 et IL17F. En effet IL12p70 est connu comme promoteur caractéristique d'INFg et donc des Th1 mais n'est pas connu comme promoteur d'IL17F et des Th17. De plus certaines études ont même montré une absence d'impact voir un impact négatif d'IL12p70 sur les cellules Th17. Deux types d'expériences ont alors été mises en

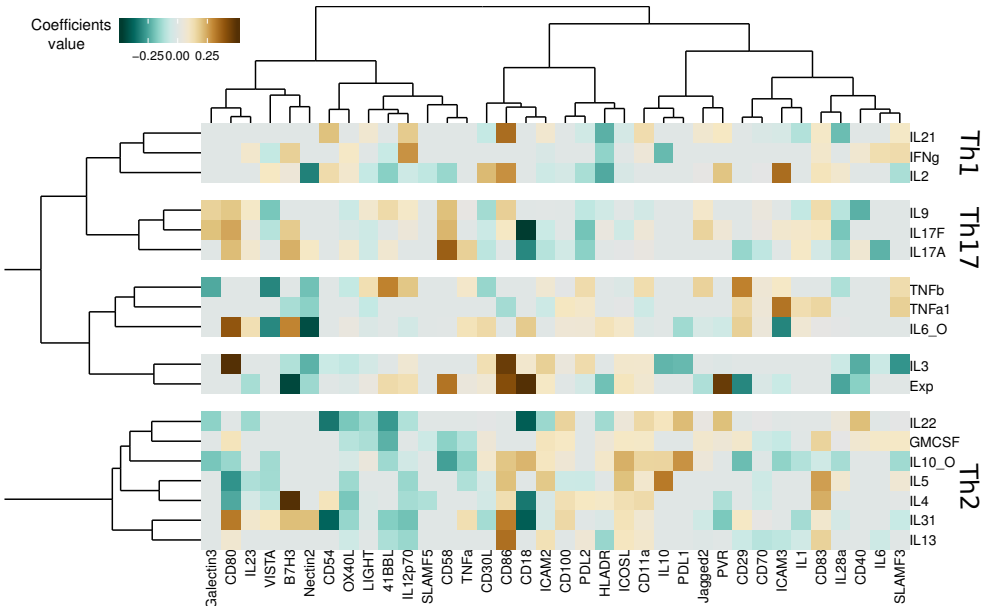

FIGURE 5.4 – Heatmap représentant les coefficients obtenus lors de la modélisation des signaux des lymphocytes Th par les signaux des cellules dendritiques en utilisant la stratégie décrite dans les chapitres 2 et 3 avec un seuil de 0.65.

place afin de valider cette hypothèse.

La première consiste à perturber des cellules dendritiques à l'aide d'un des agents pathogènes présent dans le jeu de données (ici Zymosan (concentré à 10g/ml)) puis de les mettre en contact avec des lymphocytes Th en bloquant ou non l'effet d'IL12p70. Cette expérience a été répétée sur 7 échantillons indépendants. Au risque de 5% cette expérience a mis en évidence une diminution d'IL17F lorsque l'on bloque IL12p70. Elle a cependant aussi mis en évidence une baisse du niveau d'IL17A lorsque l'on bloque IL12p70 ce qui n'était pas prévu par le modèle.

La seconde expérience consiste à donner directement IL12p70 à un lymphocyte Th naïf. Sur cette expérience aucune différence significative n'a été mise en évidence. L'effet observé par le modèle et par la première expérience est peut être dû à l'effet de l'IL12p70 en présence d'un autre input. Nous proposons alors un modèle permettant de mettre en évidence ce type d'effet. Ce modèle émet l'hypothèse d'une association entre IL12p70 et IL17F en présence de l'input IL1. Cette association a été validée par des expériences biologiques qui montrent en effet qu'un lymphocyte Th produit plus d'IL17F en présence d'IL12p70 et IL1 que en présence d'IL1 seul.

Cette association entre IL12p70, IL1 et IL17F est un exemple de validation biologique obtenue à l'aide du modèle. Pour plus de précision sur les différentes associations validées biologiquement le lecteur pourra se référer à l'article soumis dans la revue à comité de lecture Cell et disponible en annexe.

Conclusion et perspectives

Dans cette thèse nous avons mis en place une méthode, adaptée à la grande dimension, permettant de faire de la sélection de variables dans le modèle linéaire multivarié. Plus précisément, dans le chapitre 2, nous proposons un estimateur parcimonieux des coefficients dans le modèle linéaire général et établissons des conditions sous lesquelles la consistance en signe de cet estimateur est vérifiée. Cet estimateur nécessite l'estimation de la covariance existant entre les différentes variables réponses. Dans le cas simple où la matrice de covariance est celle d'un processus autorégressif d'ordre 1 et la matrice de design est celle d'une ANOVA équilibrée à 1 facteur, nous avons montré que les conditions pour obtenir la consistance en signe de notre estimateur sont vérifiées. Nous avons ensuite développé des méthodes pour l'estimation de matrices de covariance en grande dimension lorsqu'elle est supposée être une matrice de Toeplitz symétrique (dans le chapitre 3) et lorsqu'elle est supposée être une matrice par blocs (dans le chapitre 4). Ces différentes méthodes ont été appliquées à des problématiques de protéomique, de métabolomique et d'immunologie.

Dans ce dernier chapitre nous suggérons, dans un premier temps, des pistes qui permettraient d'élargir nos résultats de consistance à d'autres types de matrices de covariance. Dans un second temps, nous proposons d'autres estimateurs pénalisés issus du modèle linéaire univarié que nous généralisons au cas multivarié afin de prendre en compte des spécificités dues à des problématiques biologiques spécifiques décrites à la fin de ce chapitre.

6.1 Vers d'autres cas vérifiant les conditions de consistance en signe de notre estimateur

Les conditions du théorème 2.5 portent sur la matrice de covariance $\mathbf{\Sigma}$ mais aussi sur son inverse. Pour montrer que ces conditions sont vérifiées dans le cas de l'AR(1) nous avons bénéficié du fait que ces deux matrices ont une forme explicite simple. Pour généraliser à d'autres matrices de covariance en utilisant la même stratégie, il faudrait idéalement avoir des formes explicites de ces matrices ainsi que de leurs inverses.

Dans le cas des ARMA(p,q) il n'existe pas de formes explicites directes de la matrice de covariance $\mathbf{\Sigma}_{p,q}$ ni de son inverse. Cependant, Haddad (2004) propose des relations de récurrence sur le paramètre q pour obtenir $\mathbf{\Sigma}_{p,q}$ et $\mathbf{\Sigma}_{p,q}^{-1}$, ainsi que des relations de récurrence sur le paramètre p pour obtenir $\mathbf{\Sigma}_{p,0}$ et $\mathbf{\Sigma}_{p,0}^{-1}$, ce qui permet d'obtenir $\mathbf{\Sigma}_{p,q}$ et $\mathbf{\Sigma}_{p,q}^{-1}$ à partir des matrices explicites $\mathbf{\Sigma}_{1,0}$ et $\mathbf{\Sigma}_{1,0}^{-1}$.

De même, lorsque la matrice de covariance est une matrice de covariance par blocs diagonaux son inverse est aussi une matrice de covariance par blocs diagonaux, chaque bloc étant inversé indépendamment. Dans le cas **Diagonal-Equal** décrit dans le chapitre 4, la matrice de covariance est une matrice avec L blocs diagonaux où pour tout l dans $[\![1, L]\!]$ le bloc l est de la forme :

$$a_l \boldsymbol{I}_{\boldsymbol{q_l}} + b_l \boldsymbol{J}_{\boldsymbol{q_l}},$$

où q_l est la taille du bloc l, a_l et b_l sont des réels (avec a_l non nul), $\boldsymbol{I}_{\boldsymbol{q_l}}$ est la matrice identité de $\mathbb{R}^{q_l}$ et $\boldsymbol{J}_{\boldsymbol{q_l}}$ est une matrice de taille $q_l \times q_l$ qui a tous ses éléments égaux à 1. L'inverse de telles matrices a une forme explicite :

$$\frac{1}{a_l}(\boldsymbol{I}_{\boldsymbol{q_l}} - \frac{b_l}{a_l + q_l b_l}\boldsymbol{J}_{\boldsymbol{q_l}}).$$

Cependant, dans ce cas, l'estimation de $\boldsymbol{\Sigma}$ ne dépend pas de un mais de plusieurs paramètres. De plus, les erreurs de prédiction de la matrice $\boldsymbol{\Sigma}$ et de son inverse, ne viennent pas seulement de l'erreur de l'estimation des paramètres mais aussi de l'erreur faite pour retrouver les blocs. Ces deux points compliquent alors la vérification des hypothèses (A9) et (A10) du théorème 2.5.

6.2 Sélection de variables à l'aide d'autres régressions pénalisées

Les chapitres 2 et 3 de cette thèse présentent une méthode permettant de faire de la sélection de variables dans le modèle linéaire général dans des cas de grande dimension. Cependant, ces méthodes ne permettent pas de prendre en compte une connaissance biologique ou une problématique spécifique indiquant qu'un groupe de variables réponses ou un groupe de variables explicatives ont plus de chance, voire sont forcées, d'avoir des effets similaires. Dans cette section nous nous intéresserons à des méthodes qui permettent de prendre en compte de telles contraintes.

6.2.1 Introduction à d'autres régressions pénalisées

Le Lasso permet de faire de la sélection de variables dans le modèle linéaire univarié. Cependant, dans certains cas, les variables sont structurées en L groupes et il peut être intéressant de les sélectionner selon les groupes. Pour cela, Yuan & Lin (2006) proposent l'estimateur du `group lasso` :

$$\widehat{b}_G = \text{Argmin}_{b_1,\dots,b_L} \left\{ ||y - \sum_{1 \leq l \leq L} \boldsymbol{X}_{(l)} b_{(l)}||_2^2 + \lambda \sum_{1 \leq l \leq L} \sqrt{p_l}||b_l||_2 \right\}, \tag{6.1}$$

où $\boldsymbol{X}_{(l)}$ est la matrice contenant les variables appartenant au groupe l et p_l est le nombre de variables appartenant au groupe l.

Considérons maintenant qu'il existe des raisons pour lesquelles les coefficients liés à certaines colonnes de $\boldsymbol{X}$ soient similaires. Pour cela, Xin et al. (2014) proposent l'estimateur du `fused lasso généralisé` :

$$\widehat{b}_F = \text{Argmin}_b ||y - \boldsymbol{X}b||_2^2 + \left\{ \lambda_1 \sum_{(i,j)\in\mathcal{G}} |b_i - b_j| + \lambda_2 ||b||_1 \right\}, \tag{6.2}$$

où $\mathcal{G}$ est l'ensemble des couples (i, j), indices des coefficients qu'on présume similaires. Hoefling (2010) propose un algorithme pour calculer $\widehat{b}_F$. Inciter deux ou plusieurs variables à avoir des coefficients similaires sera appelé fusion par la suite.

6.2.2 Extensions au cadre multivarié

Dans cette section nous adaptons les pénalités utilisées dans le `group lasso` et le `fused lasso` au cas multivarié. Ceci est fait simplement en utilisant des arguments similaires à ceux utilisés dans la section 1.2.2 et dans le chapitre 2 (plus précisément à l'équation (1.13)). Après une telle transformation il est possible d'appliquer les algorithmes du `group lasso` et du `fused lasso`, en faisant varier les groupes selon la problématique. En effet, dans le cadre multivarié, il peut être intéressant de sélectionner une variable explicative, ou un groupe de variables explicatives, pour toutes les réponses, par exemple. De même il peut être intéressant d'inciter la fusion de coefficients liés à une ou plusieurs variables explicatives pour toutes les réponses. Nous avons implémenté ces différents modèles dans le package R `VariSel` qui permet de sélectionner une association (similaire ou différente) d'une ou plusieurs variables explicatives sur une ou toutes les variables réponses. Ce package utilise les algorithmes du `lasso`, du `group Lasso` et du `fused Lasso` implémentés respectivement dans les packages `glmnet` (Friedman et al., 2010b), `gglasso` (Yang & Zou, 2017) et `FusedLasso` (Hoefling, 2014). Le package contient par ailleurs des fonctions permettant de comparer les modèles obtenus en utilisant ces différentes pénalités, par bootstrap ou par cross-validation par exemple.

La figure 6.1 présente un exemple simple avec deux variables réponses Y_1 et Y_2 qu'on essaie d'expliquer en fonction de quatre variables explicatives $X_1^{G_1}$, $X_2^{G_1}$, $X_3^{G_2}$, $X_4^{G_2}$, ces variables sont divisées en deux groupes : G_1 composé de $X_1^{G_1}, X_2^{G_1}$ et G_2 composé de $X_3^{G_2}, X_4^{G_2}$. Dans cette figure nous représentons six modèles. Dans la première ligne on représente trois modèles de type « `group-lasso` » qui sélectionnent soit un groupe de variables explicatives sur une réponse (à gauche), soit une variable explicative (sans prendre en compte les groupes) sur toutes les réponses (au centre), soit un groupe de variables explicatives pour toutes les réponses (à droite). Dans la seconde ligne on représente trois modèles de type « `fused-lasso` » qui

influencent soit un groupe de variables explicatives à avoir le même coefficient sur une réponse, soit une variable explicative à avoir un coefficient similaire pour toutes les réponses, soit un groupe de variables explicatives à avoir des coefficients similaires, aussi bien entre eux que sur toutes les réponses. On notera que les matrices représentées dans la figure 6.1 ne sont pas réellement estimées par les modèles mais, elles permettent de mettre en avant leurs différences. En effet, dans les modèles de type « `group-lasso` » les coefficients appartenant au même groupe sont forcés à être sélectionnés ensemble mais ne sont pas forcés à être égaux, alors que dans le cas du « `fused-lasso` » les coefficients sont incités à être égaux sans être forcés à être sélectionnés ensemble.

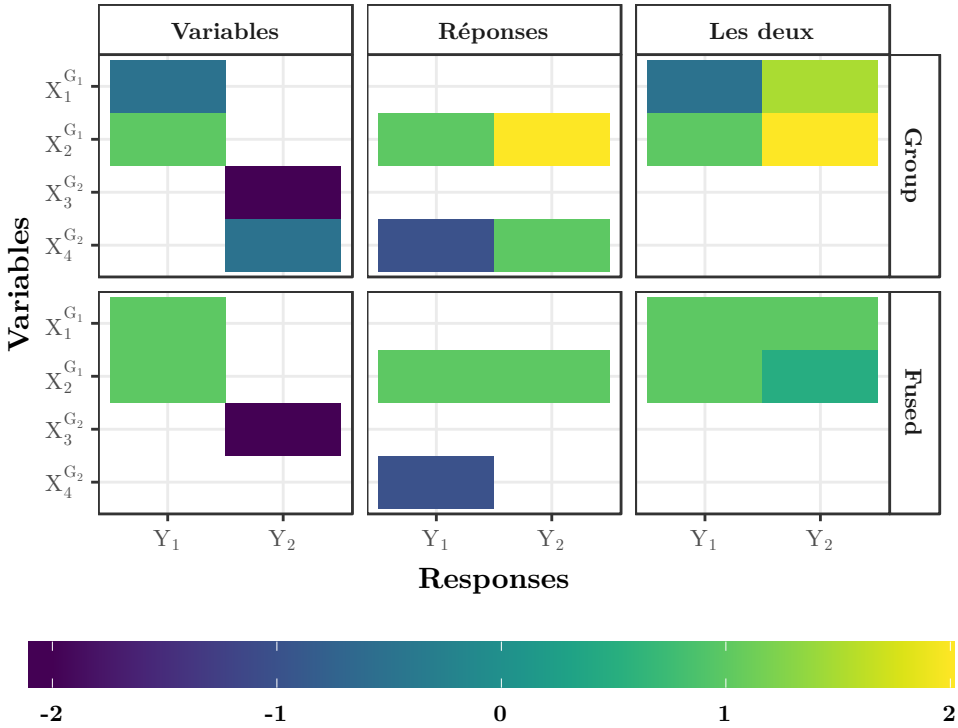

FIGURE 6.1 – Matrices de coefficients pouvant être estimées en utilisant les pénalités de groupe (6.1) ou de fusion (6.2) aussi bien sur les réponses que sur les groupes ou les deux. Ici la variable X_1 est liée à la variable X_2 et la variable X_3 est liée à la variable X_4.

Il serait intéressant de montrer la consistance en signe de ces estimateurs sous certaines conditions. Pour cela, en mimant notre démarche pour montrer le théorème 2.5, on pourrait s'inspirer des propriétés théoriques de ces estimateurs dans le cas univarié. Par exemple Bach (2008) proposent des hypothèses sous lesquelles la consistance en signe du groupe lasso est vérifiée. De même Rinaldo et al. (2009) proposent des résultats de consistance, adaptés au cas du `fused-lasso`, lorsque la matrice $\boldsymbol{X}$ est l'identité.

6.2.3 Applications

Ces méthodes peuvent être utiles dans de nombreuses applications biologiques. Nous avons commencé à nous y intéresser lors d'une collaboration avec Charlotte Brault, Agnès Doligez et Timothée Flutre des unités AGAP et GQE de l'INRA. Le but de cette collaboration est de chercher des plants de vignes qui sont peu affectés par la sécheresse. Pour cela nous proposons de chercher des régions d'ADN qui peuvent être liées à un ou plusieurs gènes et qui seraient spécifiques de traits caractérisants la résistance à la sécheresse. Ces régions d'ADN sont appelées des locus de caractères quantitatifs (QTL pour quantitative trait loci en anglais). Ici nous nous concentrerons sur les régions caractéristiques des allèles des différents gènes. Une colonne $\boldsymbol{X}_i^{(a)}$ d'une matrice $\boldsymbol{X}$ de QTL est ici le nombre de réplicats de l'allèle a d'un gène i donné pour les n échantillons. Afin de pouvoir sélectionner, en même temps, tous les allèles liés à un gène donné nous proposons d'utiliser des modèles de type « `group-lasso` ». La résistance à la sécheresse étant caractérisée par différentes variables réponses il peut aussi être intéressant de sélectionner tous les allèles d'un même gène pour toutes les réponses en même temps. De même, il peut être intéressant d'inciter les coefficients liant les différentes variables réponses à un allèle d'un gène à être similaires, cela peut ce faire à l'aide de modèles de type « `fused-lasso` ».

Enfin, ces modèles peuvent nous permettre d'étudier plus spécifiquement les données décrites dans le chapitre 5. Dans ce chapitre on étudie le lien entre des signaux sécrétés par des cellules dendritiques (inputs) et les signaux sécrétés en réponse par des lymphocytes Th (outputs). Ces données sont obtenues à partir de deux types de cellules dendritiques qu'il serait intéressant de prendre en compte dans la modélisation. Ainsi, pour chaque couple input / output on aurait deux coefficients, un pour chaque types de cellules dendritiques. Étudier ces données à l'aide d'un modèle de type « `fused-lasso` » en incitant les coefficients liant un couple input /output aux deux types de cellules dendritiques à être égaux permettrait de mettre en évidence des similarités et des différences sur l'effet des inputs sur les outputs en fonction du type de cellules dendritiques.

Le passage du génotype au phénotype est certes beaucoup moins obscur que du temps des « spermatistes » et des « ovistes » mais n'est pas totalement explicite pour autant. Être capable de manipuler des données de grande dimension peut contribuer à clarifier ce problème, c'est dans ce cadre que s'inscrit cette thèse qui cherche à proposer des outils adaptés. Nous espérons que les méthodes décrites dans cette thèse seront utiles pour résoudre des problématiques biologiques mais avons conscience de tout le chemin qu'il reste à parcourir pour percer le « mystère » du passage du génotype au phénotype.

Bibliographie

Akaike, H. (1970). Statistical predictor identification. Annals of the institute of Statistical Mathematics, 22(1) :203–217.

Alquier, P. and Doukhan, P. (2011). Sparsity considerations for dependent variables. Electron. J. Statist., 5 :750–774.

Audoin, C., Cocandeau, V., Thomas, O., Bruschini, A., Holderith, S., and Genta-Jouve, G. (2014). Metabolome consistency : Additional parazoanthines from the mediterranean zoanthid parazoanthus axinellae. Metabolites, 4 :421–432.

Bach, F. R. (2008). Consistency of the group lasso and multiple kernel learning. Journal of Machine Learning Research, 9(Jun) :1179–1225.

Bai, B., Novák, O., Ljung, K., Hanson, J., and Bentsink, L. (2018). Combined transcriptome and translatome analyses reveal a role for tryptophan-dependent auxin biosynthesis in the control of DOG1-dependent seed dormancy. New Phytologist, 217(3) :1077–1085.

Bai, B., Peviani, A., Horst, S., Gamm, M., Bentsink, L., and Hanson, J. (2017). Extensive translational regulation during seed germination revealed by polysomal profiling. New Phytologist, 214(1) :233–244.

Banerjee, O., Ghaoui, L. E., and D'aspremont, A. (2008a). Model selection through sparse maximum likelihood estimation for multivariate gaussian or binary data. Journal of Machine Learning Research, 9 :485–516.

Banerjee, O., Ghaoui, L. E., and d'Aspremont, A. (2008b). Model selection through sparse maximum likelihood estimation for multivariate gaussian or binary data. Journal of Machine learning research, 9(Mar) :485–516.

Bates, D. and Maechler, M. (2017). Matrix : Sparse and Dense Matrix Classes and Methods. R package version 1.2-8.

Belloni, A., Chernozhukovand, V., and Wang, L. (2011). Square-root lasso : pivotal recovery of sparse signals via conic programming. Biometrika, 98(4) :791–806.

Bickel, P. J. and Levina, E. (2008). Covariance regularization by thresholding. Ann. Statist., 36(6) :2577–2604.

Bien, J. and Tibshirani, R. J. (2011). Sparse estimation of a covariance matrix. Biometrika, 98(4) :807–820.

Blum, Y., Houee-Bigot, M., and Causeur, D. (2016a). FANet : Sparse Factor Analysis model for high dimensional gene co-expression Networks. R package version 1.1.

Blum, Y., Houée-Bigot, M., and Causeur, D. (2016b). Sparse factor model for co-expression networks with an application using prior biological knowledge. Statistical applications in genetics and molecular biology, 15(3) :253—272.

Boccard, J. and Rudaz, S. (2016). Exploring omics data from designed experiments using analysis of variance multiblock orthogonal partial least squares. Analytica Chimica Acta, 920 :18 – 28.

Brockwell, P. and Davis, R. (1991). Time Series : Theory and Methods. Springer Series in Statistics. Springer-Verlag New York.

Brockwell, P. J. and Davis, R. A. (1990). Time Series : Theory and Methods. Springer-Verlag New York, Inc., New York, NY, USA.

Cai, T., Liu, W., and Luo, X. (2011). A constrained l1 minimization approach to sparse precision matrix estimation. Journal of the American Statistical Association, 106(494) :594–607.

Cai, T. T. and Yuan, M. (2012). Adaptive covariance matrix estimation through block thresholding. Ann. Statist., 40(4) :2014–2042.

Cattell, R. B. (1966). The scree test for the number of factors. Multivariate behavioral research, 1(2) :245–276.

Chen, Y., Du, P., and Wang, Y. (2014). Variable selection in linear models. Wiley Interdisciplinary Reviews : Computational Statistics, 6(1) :1–9.

Chiquet, J., Mary-Huard, T., and Robin, S. (2016). Structured regularization for conditional Gaussian graphical models. Statistics and Computing, 27(3) :789–804.

Desboulets, L. D. D. (2018). A Review on Variable Selection in Regression Analysis. working paper or preprint.

Devijver, E. and Gallopin, M. (2018). Block-diagonal covariance selection for high-dimensional gaussian graphical models. Journal of the American Statistical Association, 113(521) :306–314.

Dieterle, F., Ross, A., Schlotterbeck, G., and Senn, H. (2006). Probabilistic quotient normalization as robust method to account for dilution of complex biological mixtures. application in 1h nmr metabonomics. Analytical Chemistry, 78(13) :4281–4290.

Dobriban, E. (2018). Permutation methods for factor analysis and PCA. arXiv :1710.00479.

Draper, N. R. and Smith, H. (1998). Applied Regression Analysis. Wiley.

Durand, T. C., Cueff, G., Godin, B., Valot, B., Clément, G., Gaude, T., and Rajjou, L. (2019). Combined proteomic and metabolomic profiling of the arabidopsis thaliana vps29 mutant reveals pleiotropic functions of the retromer in seed development. International journal of molecular sciences, 20(2) :362.

Eckart, C. and Young, G. (1936). The approximation of one matrix by another of lower rank. Psychometrika, 1(3) :211–218.

Fan, J. and Li, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. Journal of the American Statistical Association, 96(456) :1348–1360.

Fan, J., Yuan, L., and Han, L. (2016). An overview of the estimation of large covariance and precision matrices. The Econometrics Journal, 19(1) :C1–C32.

Faraway, J. J. (2004). Linear Models with R. Chapman & Hall/CRC.

Field, B., Cardon, G., Traka, M., Botterman, J., Vancanneyt, G., and Mithen, R. (2004). Glucosinolate and amino acid biosynthesis in arabidopsis. Plant Physiology, 135(2) :828–839.

Friedman, J., Hastie, T., and Tibshirani, R. (2008). Sparse inverse covariance estimation with the graphical lasso. Biostatistics, 9(3) :432.

Friedman, J., Hastie, T., and Tibshirani, R. (2010a). Regularization paths for generalized linear models via coordinate descent. Journal of Statistical Software, 33(1) :1–22.

Friedman, J., Hastie, T., and Tibshirani, R. (2010b). Regularization paths for generalized linear models via coordinate descent. Journal of Statistical Software, 33(1) :1–22.

Galland, M., Huguet, R., Arc, E., Cueff, G., Job, D., and Rajjou, L. (2014). Dynamic proteomics emphasizes the importance of selective mrna translation and protein turnover during arabidopsis seed germination. Molecular & Cellular Proteomics, 13(1) :252–268.

Gianola, D., Perez-Enciso, M., and Toro, M. A. (2003). On marker-assisted prediction of genetic value : beyond the ridge. Genetics, 163(1) :347–365.

Giraud, C. (2014). Introduction to High-Dimensional Statistics. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. Taylor & Francis.

Haddad, J. N. (2004). On the closed form of the covariance matrix and its inverse of the causal arma process. Journal of Time Series Analysis, 25(4) :443–448.

Harville, D. (2001). Matrix Algebra : Exercises and Solutions : Exercises and Solutions. Springer New York.

Hastie, T., Tibshirani, R., and Friedman, J. (2001). The Elements of Statistical Learning. Springer Series in Statistics. Springer New York Inc., New York, NY, USA.

Heinze, G., Wallisch, C., and Dunkler, D. (2018). Variable selection - A review and recommendations for the practicing statistician. Biom J, 60(3) :431–449.

Higham, N. J. (2002). Computing the nearest correlation matrix - a problem from finance. IMA Journal of Numerical Analysis, 22(3) :329–343.

Hoefling, H. (2010). A path algorithm for the fused lasso signal approximator. Journal of Computational and Graphical Statistics, 19(4) :984–1006.

Hoefling, H. (2014). FusedLasso : Solves the generalized Fused Lasso. R package version 1.0.6.

Horn, J. L. (1965). A rationale and test for the number of factors in factor analysis. Psychometrika, 30(2) :179–185.

Horn, R. A. and Johnson, C. R. (1986). Matrix Analysis. Cambridge University Press, New York, NY, USA.

Hossain, Z., Amyot, L., McGarvey, B., Gruber, M., Jung, J., and Hannoufa, A. (2012). The translation elongation factor eef-1bβ1 is involved in cell wall biosynthesis and plant development in arabidopsis thaliana. PLoS One. e30425.

Hosseini, M. J. and Lee, S.-I. (2016). Learning sparse gaussian graphical models with overlapping blocks. In Lee, D. D., Sugiyama, M., Luxburg, U. V., Guyon, I., and Garnett, R., editors, Advances in Neural Information Processing Systems 29, pages 3808–3816. Curran Associates, Inc.

Hrydziuszko, O. and Viant, M. R. (2012). Missing values in mass spectrometry based metabolomics : an undervalued step in the data processing pipeline. Metabolomics, 8(1) :161–174.

Huang, Z., Footitt, S., and Finch-Savage, W. E. (2014). The effect of temperature on reproduction in the summer and winter annual arabidopsis thaliana ecotypes bur and cvi. Annals of botany, 113(6) :921–929.

Jacquard, A. (1978). Eloge de la différence. La génétique et les hommes.

Johnson, R. A. and Wichern, D. W., editors (1988). Applied Multivariate Statistical Analysis. Prentice-Hall, Inc., Upper Saddle River, NJ, USA.

Karahalil, B. (2016). Overview of systems biology and omics technologies. Current medicinal chemistry, 23(37) :4221–4230.

Kendall, S. L., Hellwege, A., Marriot, P., Whalley, C., Graham, I. A., and Penfield, S. (2011). Induction of dormancy in arabidopsis summer annuals requires parallel regulation of dog1 and hormone metabolism by low temperature and cbf transcription factors. The Plant Cell, 23(7) :2568–2580.

Kerdaffrec, E. and Nordborg, M. (2017). The maternal environment interacts with genetic variation in regulating seed dormancy in swedish arabidopsis thaliana. PloS one, 12(12). e0190242.

Kirwan, J., Broadhurst, D., Davidson, R., and Viant, M. (2013). Characterising and correcting batch variation in an automated direct infusion mass spectrometry (dims) metabolomics workflow. Analytical and Bioanalytical Chemistry, 405(15) :5147–5157.

Kuhl, C., Tautenhahn, R., Boettcher, C., Larson, T. R., and Neumann, S. (2012). Camera : an integrated strategy for compound spectra extraction and annotation of liquid chromatography/mass spectrometry data sets. Analytical Chemistry, 84 :283–289.

Lê Cao, K.-A., Boitard, S., and Besse, P. (2011). Sparse pls discriminant analysis : biologically relevant feature selection and graphical displays for multiclass problems. BMC Bioinformatics, 12(1) :253.

Ledoit, O. and Wolf, M. (2004). A well-conditioned estimator for large-dimensional covariance matrices. Journal of Multivariate Analysis, 88(2) :365 – 411.

Lee, W. and Liu, Y. (2012). Simultaneous Multiple Response Regression and Inverse Covariance Matrix Estimation via Penalized Gaussian Maximum Likelihood. J. Multivar. Anal., 111 :241–255.

Leung, S., Liu, X., Fang, L., Chen, X., Guo, T., and Zhang, J. (2010). The cytokine milieu in the interplay of pathogenic Th1/Th17 cells and regulatory T cells in autoimmune disease. Cell. Mol. Immunol., 7(3) :182–189.

Lütkepohl, H. (2005). New introduction to multiple time series analysis. Springer, Berlin.

MacGregor, D. R., Kendall, S. L., Florance, H., Fedi, F., Moore, K., Paszkiewicz, K., Smirnoff, N., and Penfield, S. (2015). Seed production temperature regulation of primary dormancy occurs through control of seed coat phenylpropanoid metabolism. New Phytologist, 205(2) :642–652.

Mallows, C. L. (1973). Some comments on c p. Technometrics, 15(4) :661–675.

Mardia, K., Kent, J., and Bibby, J. (1979). Multivariate analysis. Probability and mathematical statistics. Academic Press.

Mardia, K. V., Kent, J. T., and Bibby, J. M. (1980). Multivariate Analysis (Probability and Mathematical Statistics). Academic Press.

Mehmood, T., Liland, K. H., Snipen, L., and Saebo, S. (2012). A review of variable selection methods in partial least squares regression. Chemometrics and Intelligent Laboratory Systems, 118 :62 – 69.

Meinshausen, N. and Bühlmann, P. (2010). Stability selection. Journal of the Royal Statistical Society : Series B (Statistical Methodology), 72(4) :417–473.

Meinshausen, N. and Buhlmann, P. (2010). Stability selection. Journal of the Royal Statistical Society, 72(4) :417–473.

Meng, C., Kuster, B., Culhane, A. C., and Gholami, A. M. (2014). A multivariate approach to the integration of multi-omics datasets. BMC Bioinformatics, 15(1) :162.

Molstad, A. J., Weng, G., Doss, C. R., and Rothman, A. J. (2018). An explicit mean-covariance parameterization for multivariate response linear regression.

Mosmann, T. R., Cherwinski, H., Bond, M. W., Giedlin, M. A., and Coffman, R. L. (1986). Two types of murine helper t cell clone. i. definition according to profiles of lymphokine activities and secreted proteins. The Journal of immunology, 136(7) :2348–2357.

Mosmann, T. R. and Coffman, R. (1989). Th1 and th2 cells : different patterns of lymphokine secretion lead to different functional properties. Annual review of immunology, 7(1) :145–173.

Muller, K. E. and Stewart, P. W. (2006). Linear Model Theory : Univariate, Multivariate, and Mixed Models. John Wiley & Sons.

Nicholson, J. K., Lindon, J. C., and Holmes, E. (1999). 'metabonomics' : understanding the metabolic responses of living systems to pathophysiological stimuli via multivariate statistical analysis of biological nmr spectroscopic data. Xenobiotica, 29(11) :1181–1189. PMID : 10598751.

Nishii, R. (1984). Asymptotic properties of criteria for selection of variables in multiple regression. The Annals of Statistics, pages 758–765.

Park, H., Li, Z., Yang, X. O., Chang, S. H., Nurieva, R., Wang, Y.-H., Wang, Y., Hood, L., Zhu, Z., Tian, Q., et al. (2005). A distinct lineage of cd4 t cells regulates tissue inflammation by producing interleukin 17. Nature immunology, 6(11) :1133.

Perrot-Dockès, M., Lévy-Leduc, C., Chiquet, J., Sansonnet, L., Brégère, M., Étienne, M.-P., Robin, S., and Genta-Jouve, G. (2018). A variable selection approach in the multivariate linear model : an application to lc-ms metabolomics data. Statistical applications in genetics and molecular biology, 17(5).

Perrot-Dockès, M., Lévy-Leduc, C., Sansonnet, L., and Chiquet, J. (2018). Variable selection in multivariate linear models with high-dimensional covariance matrix estimation. Journal of Multivariate Analysis, 166 :78 – 97.

Perrot-Dockès, M., Lévy-Leduc, C., and Chiquet, J. (2019). MultiVarSel : Variable Selection in a Multivariate Linear Model. R package version 1.1.3.

Perthame, E., Friguet, C., and Causeur, D. (2016). Stability of feature selection in classification issues for high-dimensional correlated data. Statistics and Computing, 26(4) :783–796.

Perthame, E., Friguet, C., and Causeur, D. (2019). FADA : Variable Selection for Supervised Classification in High Dimension. R package version 1.3.4.

Pourahmadi, M. (2013). High-Dimensional Covariance Estimation. Wiley Series in Probability and Statistics.

Provart, N. J., Alonso, J., Assmann, S. M., Bergmann, D., Brady, S. M., Brkljacic, J., ..., and Dangl, J. (2016). 50 years of arabidopsis research : highlights and future directions. New Phytologist, 209(3) :921–944.

R Core Team (2017). R : A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria.

Rajjou, L., Gallardo, K., Debeaujon, I., Vandekerckhove, J., Job, C., and Job, D. (2004). The effect of α-amanitin on the arabidopsis seed proteome highlights the distinct roles of stored and neosynthesized mrnas during germination. Plant physiology, 134(4) :1598–1613.

Raven, P., Singer, S., Johnson, G., Mason, K., Losos, J., Bouharmont, J., Masson, P., and Van Hove, C. (2017). Biologie. Biologie. De Boeck Supérieur.

Ren, S., Hinzman, A. A., Kang, E. L., Szczesniak, R. D., and Lu, L. J. (2015). Computational and statistical analysis of metabolomics data. Metabolomics, 11(6) :1492–1513.

Riekeberg, E. and Powers, R. (2017). New frontiers in metabolomics : from measurement to insight. F1000Research, 6.

Rinaldo, A. et al. (2009). Properties and refinements of the fused lasso. The Annals of Statistics, 37(5B) :2922–2952.

Rothman, A. J. (2012). Positive definite estimators of large covariance matrices. Biometrika, 99(3) :733–740.

Rothman, A. J., Bickel, P. J., Levina, E., and Zhu, J. (2008). Sparse permutation invariant covariance estimation. Electron. J. Statist., 2 :494–515.

Rothman, A. J., Levina, E., and Zhu, J. (2010). Sparse multivariate regression with covariance estimation. Journal of Computational and Graphical Statistics, 19(4) :947–962.

Saccenti, E., Hoefsloot, H. C. J., Smilde, A. K., Westerhuis, J. A., and Hendriks, M. M. W. B. (2013). Reflections on univariate and multivariate analysis of metabolomics data. Metabolomics, 10(3) :361–374.

Schwarz, G. et al. (1978). Estimating the dimension of a model. The annals of statistics, 6(2) :461–464.

Smith, C., Want, E., O'Maille, G., Abagyan, R., and Siuzdak, G. (2006). XCMS : Processing mass spectrometry data for metabolite profiling using Nonlinear peak alignment, matching, and identification. Analytical Chemistry, 78(3) :779–787.

Smith, R., Mathis, A., and Prince, J. (2014). Proteomics, lipidomics, metabolomics : a mass spectrometry tutorial from a computer scientist's point of view. BMC Bioinformatics, 15.

Stuart, J. M., Segal, E., Koller, D., and Kim, S. K. (2003). A gene-coexpression network for global discovery of conserved genetic modules. Science, 302(5643) :249–255.

Thompson, M. L. (1978). Selection of variables in multiple regression : Part i. a review and evaluation. International Statistical Review/Revue Internationale de Statistique, pages 1–19.

Tibshirani, R. (1996). Regression shrinkage and selection via the Lasso. J. Royal. Statist. Soc B., 58(1) :267–288.

Varah, J. (1975). A lower bound for the smallest singular value of a matrix. Linear Algebra and its Applications, 11(1) :3 – 5.

Verdegem, D., Lambrechts, D., Carmeliet, P., and Ghesquière, B. (2016). Improved metabolite identification with midas and magma through ms/ms spectral dataset-driven parameter optimization. Metabolomics, 12(6) :1–16.

Vogel, J. T., Cook, D., Fowler, S. G., and Thomashow, M. F. (2006). The cbf cold response pathways of arabidopsis and tomato. Cold Hardiness in Plants : Molecular Genetics, Cell Biology and Physiology, pages 11–29.

Volpe, E., Touzot, M., Servant, N., Marloie-Provost, M.-A., Hupé, P., Barillot, E., and Soumelis, V. (2009). Multiparametric analysis of cytokine-driven human th17 differentiation reveals a differential regulation of il-17 and il-22 production. Blood, 114(17) :3610–3614.

von Luxburg, U. (2007). A tutorial on spectral clustering. Statistics and Computing, 17(4) :395–416.

Walker, A. (1964). Asymptotic properties of least-squares estimates of parameters of the spectrum of a stationary non-deterministic time-series. Journal of the Australian Mathematical Society, 4(3) :363–384.

Wen, F., Yang, Y., Liu, P., and Qiu, R. C. (2016). Positive definite estimation of large covariance matrix using generalized nonconvex penalties. IEEE Access, 4 :4168–4182.

Wittenburg, D., Teuscher, F., Klosa, J., and Reinsch, N. (2016). Covariance between genotypic effects and its use for genomic inference in half-sib families. G3 : Genes, Genomes, Genetics, 6(9) :2761–2772.

Wittstock, U. and Halkier, B. A. (2002). Glucosinolate research in the arabidopsis era. Trends in plant science, 7(6) :263–270.

Xin, B., Kawahara, Y., Wang, Y., and Gao, W. (2014). Efficient generalized fused lasso and its application to the diagnosis of alzheimer's disease. In Twenty-Eighth AAAI Conference on Artificial Intelligence.

Yang, J., Benyamin, B., McEvoy, B. P., Gordon, S., Henders, A. K., Nyholt, D. R., Madden, P. A., Heath, A. C., Martin, N. G., Montgomery, G. W., et al. (2010). Common snps explain a large proportion of the heritability for human height. Nature genetics, 42(7) :565.

Yang, Y. and Zou, H. (2017). gglasso : Group Lasso Penalized Learning Using a Unified BMD Algorithm. R package version 1.4.

Yuan, M. and Lin, Y. (2006). Model selection and estimation in regression with grouped variables. Journal of the Royal Statistical Society Series B, 68 :49–67.

Yuan, M. and Lin, Y. (2007). Model selection and estimation in the gaussian graphical model. Biometrika, 94(1) :19–35.

Zhang, A., Sun, H., Wang, P., Han, Y., and Wang, X. (2012). Modern analytical techniques in metabolomics analysis. Analyst, 137 :293–300.

Zhang, H., Zheng, Y., Yoon, G., Zhang, Z., Gao, T., Joyce, B., Zhang, W., Schwartz, J., Vokonas, P., Colicino, E., Baccarelli, A., Hou, L., and Liu, L. (2017). Regularized estimation in sparse high-dimensional multivariate regression, with application to a DNA methylation study. Stat Appl Genet Mol Biol, 16(3) :159–171.

Zhao, P. and Yu, B. (2006). On model selection consistency of lasso. Journal of Machine Learning Research, 7 :2541–2563.

Zheng, X. and Loh, W.-Y. (1995). Consistent variable selection in linear models. Journal of the American Statistical Association, 90(429) :151–156.

Annexe

Cette annexe présente l'article soumis au journal Cell :
M. Grandclaudon*, M. Perrot-Dockès*, C. Trichot, O. Mostafa-Abouzid, W. Abou-Jaoudé, F. Berger, P. Hupé, D. Thieffry, L. Sansonnet, J. Chiquet, C. Lévy-Leduc, V. Soumelis
A quantitative multivariate model of human dendritic cell-T helper cell communication.
** : ces auteurs ont contribué de manière égale à cette publication*

Un résumé est disponible dans le chapitre 5.

Article

Cell

A Quantitative Multivariate Model of Human Dendritic Cell-T Helper Cell Communication

Graphical Abstract

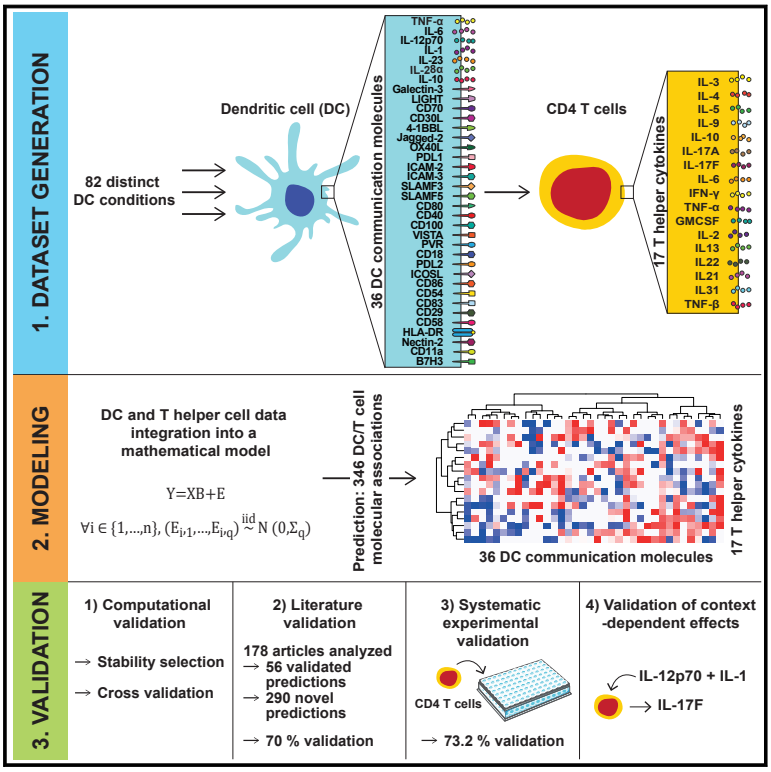

Highlights

- 428 protein-level measurements of 36 DC communication molecules and 17 Th cytokines
- Data-driven quantitative model of DC-T cell communication extensively validated
- Systematic and unbiased predictions of context-dependent mechanisms
- Validation of a new context-dependent role of IL-12p70 in Th17 differentiation

Authors

Maximilien Grandclaudon, Marie Perrot-Dockès, Coline Trichot, ..., Julien Chiquet, Céline Lévy-Leduc, Vassili Soumelis

Correspondence

vassili.soumelis@curie.fr

In Brief

Grandclaudon et al. show that combinatorial rules that explain communication between dendritic cells and T helper cells can be helpful in vaccine design and immunotherapy.

Grandclaudon et al., 2019, Cell *179*, 432–447
October 3, 2019 © 2019 Elsevier Inc.
https://doi.org/10.1016/j.cell.2019.09.012

CellPress

Cell

Article

A Quantitative Multivariate Model of Human Dendritic Cell-T Helper Cell Communication

Maximilien Grandclaudon,[1,2,8] Marie Perrot-Dockès,[3,8] Coline Trichot,[1,2,8] Léa Karpf,[1,2] Omar Abouzid,[1,2] Camille Chauvin,[1,2] Philémon Sirven,[1,2] Wassim Abou-Jaoudé,[4] Frédérique Berger,[1,5,6] Philippe Hupé,[1,6,7] Denis Thieffry,[4] Laure Sansonnet,[3] Julien Chiquet,[3] Céline Lévy-Leduc,[3] and Vassili Soumelis[1,2,9,*]

[1]Institut Curie, Centre de Recherche, PSL Research University, 75005 Paris, France
[2]INSERM U932, Immunity and Cancer, 75005 Paris, France
[3]UMR MIA-Paris, AgroParisTech, INRA—Université Paris-Saclay, 75005 Paris, France
[4]Computational Systems Biology Team, Institut de Biologie de l'École Normale Supérieure, Centre National de la Recherche Scientifique UMR8197, INSERM U1024, École Normale Supérieure, PSL Université, 75005 Paris, France
[5]Institut Curie, PSL Research University, Unit of Biostatistics, 75005 Paris, France
[6]Institut Curie, PSL Research University, INSERM U900, 75005 Paris, France
[7]Mines Paris Tech, 77305 Cedex Fontainebleau, France
[8]These authors contributed equally
[9]Lead Contact
*Correspondence: vassili.soumelis@curie.fr
https://doi.org/10.1016/j.cell.2019.09.012

SUMMARY

Cell-cell communication involves a large number of molecular signals that function as words of a complex language whose grammar remains mostly unknown. Here, we describe an integrative approach involving (1) protein-level measurement of multiple communication signals coupled to output responses in receiving cells and (2) mathematical modeling to uncover input-output relationships and interactions between signals. Using human dendritic cell (DC)-T helper (Th) cell communication as a model, we measured 36 DC-derived signals and 17 Th cytokines broadly covering Th diversity in 428 observations. We developed a data-driven, computationally validated model capturing 56 already described and 290 potentially novel mechanisms of Th cell specification. By predicting context-dependent behaviors, we demonstrate a new function for IL-12p70 as an inducer of Th17 in an IL-1 signaling context. This work provides a unique resource to decipher the complex combinatorial rules governing DC-Th cell communication and guide their manipulation for vaccine design and immunotherapies.

INTRODUCTION

Cell-cell communication involves the exchange of molecular signals produced by a given cell and transmitting an effect through specific receptors expressed on target cells. This process requires integration of multiple communication signals of different nature during homeostatic or stress-related responses. For example, differentiation of pluripotent hematopoietic stem cells into mature myeloid or lymphoid blood cells requires the collective action of multiple cytokines, growth factors, and Notch ligands (Balan et al., 2018). In the context of stress, multiple signals need to be integrated by innate and adaptive immune cells, including cytokines, growth factors, inflammatory mediators, and immune checkpoints (Chen and Flies, 2013; Macagno et al., 2007). In most studies, these communication molecules have been studied as individual stimuli to a target cell by gain- and loss-of-function experiments. This provides important knowledge regarding the downstream effects of the signals but prevents us from widely addressing their function in various contexts of other co-expressed communication signals.

Context dependency is an important aspect of verbal language communication that can directly affect the meaning of individual words but also modify the logic of syntactic rules (Cariani and Rips, 2017; Kintsch and Mangalath, 2011). Similarly, context dependencies may dramatically affect the function of biologically active communication signals. For example, we have shown that 90% of the transcriptional response to type I interferon in human CD4 T cells depends on the cytokine context (T helper 1 [Th1], Th2, or Th17; Touzot et al., 2014). Other studies have identified major context-dependent functions of immune checkpoints, such as OX40-ligand (Ito et al., 2005), and regulatory cytokines, such as transforming growth factor β (TGF-β) (Ivanov et al., 2006; Manel et al., 2008; Volpe et al., 2008). These studies suggest that communication molecules function as words of a complex language with grammar defining combinatorial rules of co-expression and mutual influence of one signal over the function (meaning) of another signal.

Three levels of biological complexity need to be integrated to decipher those combinatorial rules: (1) the multiplicity of input communication signals to include as many possible contextual effects; (2) communication signals at their naturally occurring concentrations; and (3) a large number of output responses in target cells, reflecting the effect of cell-cell communication quantitatively and qualitatively. Those three levels create a bottleneck in deciphering cell-cell communication.

Here we developed an integrative approach combining (1) coupled protein-level measurement of multiple communication signals and output response molecules in target cells; (2) a multivariate mathematical modeling strategy enabling us to infer the input-output relationships for individual signals, taking into account the context and configuration of all other signals; and (3) experimental validation of model-derived hypotheses. We applied this framework to decipher human dendritic cell (DC)-Th cell communication, which potentially involves over 70 individual molecular stimuli (Chen and Flies, 2013), including cytokines, tumor necrosis factor (TNF) family members, integrins, nectins, Notch ligands, and galectins (Tindemans et al., 2017; Zhu et al., 2010; Zygmunt and Veldhoen, 2011). These molecules can all be expressed by DCs and function as communication signals to T cells (hereafter called Th stimuli). They can be measured at the protein level by highly specific assays to optimize biological relevance.

By using this unbiased data-driven approach, we could capture the simultaneous effects of large numbers of DC-T cell communication signals in naturally occurring patterns and expression levels. Our systems-level model revealed novel emergent and context-dependent mechanisms controlling Th cell differentiation. A similar framework can be applied to systematically decipher the communication of other cell types.

RESULTS

Generation of a Unique Multivariate Dataset of Human DC-Th Cell Communication

To induce a broad range of DC molecular states expressing various patterns of communication signals, human monocyte-derived DCs (MoDCs) and primary blood CD11c$^+$ DCs (bDCs), were activated for 24 h with a diversity of DC-modulating signals (hereafter called DC perturbators). These included 14 distinct stimuli that were grouped in three categories reflecting various physiopathological contexts: (1) the endogenous factors interferon β (IFN-β), GM-CSF, TSLP, and PGE2; (2) the Toll-like receptor ligands lipopolysaccharide (LPS) (a Toll-like receptor 4 [TLR4] agonist), PAM3CSK4 (a TLR1 and 2 agonist), Curdlan (a Dectin1 agonist), zymosan (a TLR2 and Dectin1 agonist), R848 (a TLR7 and 8 agonist), poly(I:C) (a TLR3 agonist), and aluminum potassium sulfate (Alum, an NLRP3 inflammasome inducer); and (3) the whole pathogens heat-killed *Candida albicans* (HKCA), heat-killed *Listeria monocytogenes* (HKLM), heat-killed *Staphylococcus aureus* (HKSA), heat-killed *Streptococcus pneumoniae* (HKSP), and influenza virus (flu). These 14 DC perturbators were used in distinct doses and combinations to further increase the diversity of DC communication molecules and downstream functional effects (Table S1). In each independent experiment, we included a medium condition as a negative control and LPS (100 ng/mL) and/or zymosan (10 μg/mL) as positive controls. A total of 66 perturbators were used on MoDCs and 16 on bDCs, totaling 82 distinct "DC conditions" (C1–C82; Table S1).

Under each DC condition, we measured 36 DC-expressed molecules that influence Th cell differentiation in at least one published study (STAR Methods) and can be measured with a highly specific antibody-based assay. Twenty-nine were measured by fluorescence-activated cell sorting (FACS) at the DC surface (Figure S1A), and 7 were measured in the 24-h DC culture supernatant (STAR Methods).

Following 24-h culture under each of the 82 DC perturbation conditions, the same DC batch was used to stimulate naive CD4 T cells in a heterologous co-culture system. On day 6 of co-culture, we measured Th cell expansion fold (Exp Fold) and a total of 17 distinct Th cytokines broadly representing the spectrum of Th cell output responses (STAR Methods). In total, we produced a unique dataset of coupled measurements of DC-derived Th stimuli and Th response cytokines from 428 independent observations from 44 independent donors (Figure 1A; Table S2).

Variability and Specificity of DC Communication Signals

We asked whether our systematic DC stimulation strategy could generate important variations in the expression of individual DC-derived Th stimuli. All Th stimuli were expressed over at least three logs (Figure 1B) with high coefficients of variation (>0.44; Figure 1C). Interleukins had higher variability (10^4–10^5) and high coefficients of variation from 2.72 for interleukin-12 (IL-12) p70 (IL-12) to 1.43 for IL-6. CD11a had a wide expression range (10^4) but the smallest coefficient of variation (0.44), with values distributed around the mean (Figure 1C). Hence, we were able to generate highly variable expression patterns for all Th stimuli.

We sought to identify conserved and specific patterns of Th stimuli in response to standard DC perturbators. We compared the expression levels of DC-derived Th stimuli under three conditions belonging to distinct classes of microbes—LPS (100 ng/mL, bacteria), zymosan (10 μg/mL, fungi), and flu (1×, Viruses)—that were used across at least 17 MoDC biological replicates (Figure 1D). Medium MoDCs (negative control) expressed lower levels of activation-associated communication molecules (Figures 1D and S1B). We confirmed previous findings, validating our experimental system: (1) zymosan induced specifically IL-10 and IL-23, (2) flu induced a large amount of IL-28α, and (3) LPS and zymosan induced a large amount of IL-12 (Figures 1D and S1B). In addition, we identified novel specific inductions of DC-derived Th stimuli: zymosan-treated MoDCs expressed the highest levels of CD54 and PVR, flu-treated MoDCs specifically induced ICOSL, and LPS-treated MoDCs induced the highest levels of CD30L and CD83 (Figure 1D). Specificity of expression of a given signal for a given DC stimulation was determined using Wilcoxon statistical test (Figure S1B). Hence, standard DC perturbators induced specific patterns of Th stimuli.

Defining the Spectrum of DC Communication States

Next we aimed to assess the spectrum of DC communication states, as defined by their expression pattern of communication signals, across the 82 DC conditions. We computed the mean expression of biological replicates for each DC-derived Th stimulus and performed unsupervised hierarchical clustering to identify classes of the most similar conditions (C1–C82, y axis) and DC-derived Th stimuli (x axis) (Figure 2A). This revealed five groups of DC conditions (Figure 2B). Each of the four standard DC conditions (Figure 1D) belonged to a different group (Figure 2A).

Group 1 was defined by high expression of adhesion molecules such as CD18, ICAM-2, ICAM-3, and CD29 and low levels

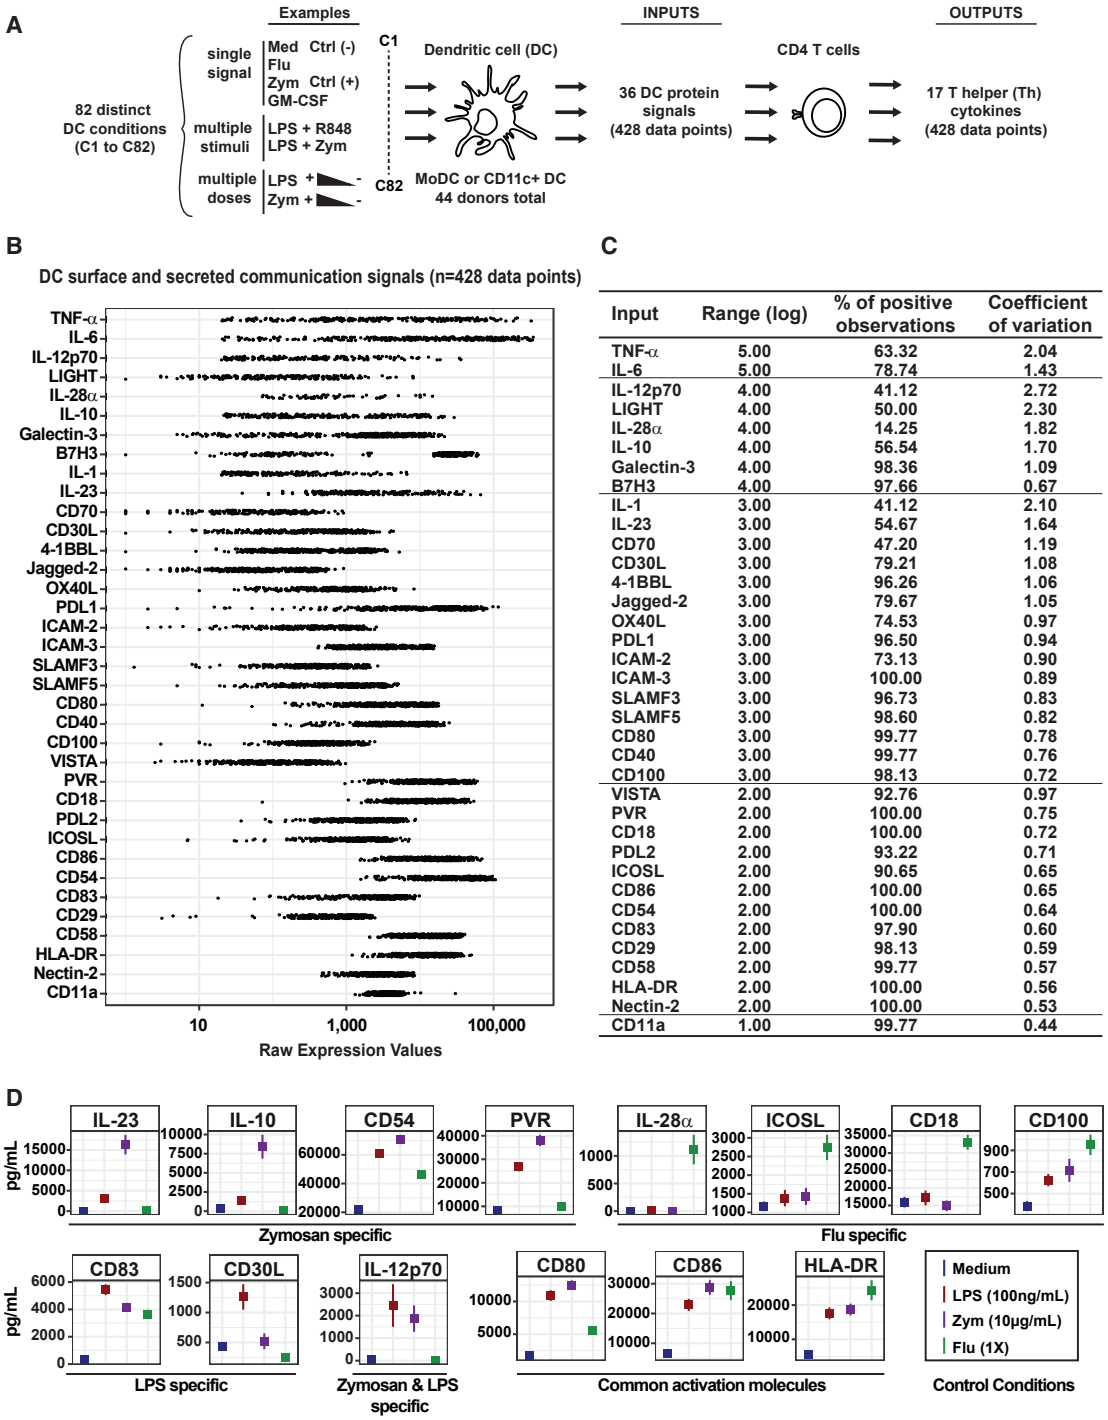

Input	Range (log)	% of positive observations	Coefficient of variation
TNF-α	5.00	63.32	2.04
IL-6	5.00	78.74	1.43
IL-12p70	4.00	41.12	2.72
LIGHT	4.00	50.00	2.30
IL-28α	4.00	14.25	1.82
IL-10	4.00	56.54	1.70
Galectin-3	4.00	98.36	1.09
B7H3	4.00	97.66	0.67
IL-1	3.00	41.12	2.10
IL-23	3.00	54.67	1.64
CD70	3.00	47.20	1.19
CD30L	3.00	79.21	1.08
4-1BBL	3.00	96.26	1.06
Jagged-2	3.00	79.67	1.05
OX40L	3.00	74.53	0.97
PDL1	3.00	96.50	0.94
ICAM-2	3.00	73.13	0.90
ICAM-3	3.00	100.00	0.89
SLAMF3	3.00	96.73	0.83
SLAMF5	3.00	98.60	0.82
CD80	3.00	99.77	0.78
CD40	3.00	99.77	0.76
CD100	3.00	98.13	0.72
VISTA	2.00	92.76	0.97
PVR	2.00	100.00	0.75
CD18	2.00	100.00	0.72
PDL2	2.00	93.22	0.71
ICOSL	2.00	90.65	0.65
CD86	2.00	100.00	0.65
CD54	2.00	100.00	0.64
CD83	2.00	97.90	0.60
CD29	2.00	98.13	0.59
CD58	2.00	99.77	0.57
HLA-DR	2.00	100.00	0.56
Nectin-2	2.00	100.00	0.53
CD11a	1.00	99.77	0.44

Figure 1. Variability and Specificity of DC Communication Signals

(A) Experimental strategy.

(B) Raw expression values of the 36 DC communication signals (n = 428 data points).

(C) Statistical descriptors of the 36 DC communication signals: expression range (log magnitude), percentage of positive observations among the 428 datapoints, and coefficient of variation.

(D) Average expression values and SD for the four indicated DC signals for MoDCs.

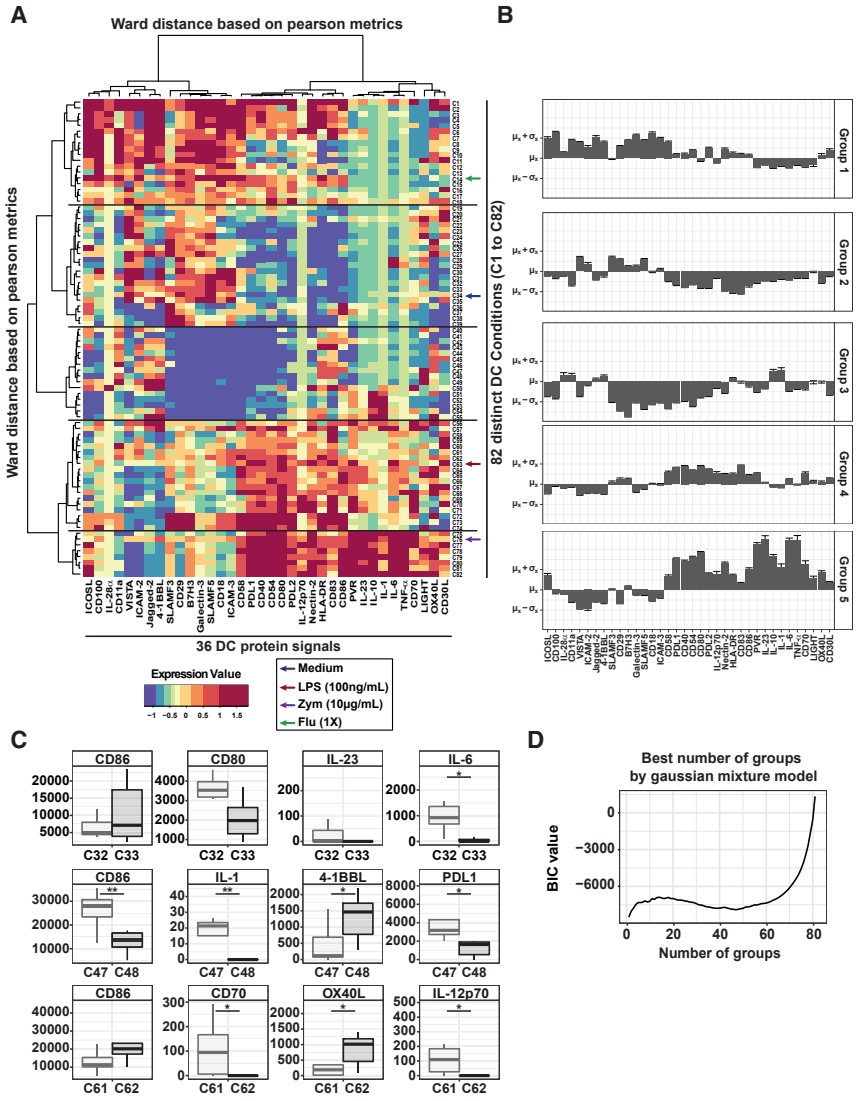

Figure 2. The Diversity of DC States Is Defined by Unique Combinations of Communication Molecules

(A) Heatmap showing expression values of each 36 DC-derived signal, performed with hierarchical clustering on Pearson metrics for the DC signals and Euclidian distances for the 82 DC conditions.

(B) Expression profiles (mean and SD) of the 36 communication molecules within the five groups of DC conditions, defined by hierarchical clustering. Expression data were logged and scaled so μ represents the mean and σ the SD of the expression of a given DC signal across the whole dataset.

(C) Boxplot of selected DC signals for pairs of stimulatory conditions defined as being the most correlated within our dataset by Pearson correlation (t test).

(D) Best number of groups by Gaussian mixture model, determined using the 428 points of the 36 DC parameters.

of co-stimulatory molecules and cytokines with the exception of high IL-28a. Group 2 showed low expression for most DC-derived Th stimuli but high levels of integrins, VISTA and B7H3, suggesting a capacity to interact with T cells and transmit co-inhibitory signals. Group 3 showed a complementary pattern, lack of group 1- and group 2-specific molecules, and intermediate or high levels of co-stimulatory molecules such as CD83, CD86, HLA-DR, 4-1BBL, and OX40L. This suggested potent T cell stimulating functions. Group 4 exhibited high levels of molecules from the B7 and TNF superfamilies, such as CD80, CD86, PDL1, PDL2, and CD40, but intermediate or low cytokine levels. In contrast, group 5 showed the highest level of cytokines and molecules of the B7 and TNF superfamilies (Figure 2B).

Next we sought to analyze intra-cluster heterogeneity. We selected three pairs of perturbators most closely related as defined by Euclidian distance (C32 [MoDC HKLM, MOI 1] and C33 [MoDC HKCA, MOI 1], C47 [bDC LPS, 100 ng/mL] and C48 [bDC HLKM, MOI 1], and C61 [MoDC R848, 1 µg/mL] and C62 [MoDC PAM3, 10 µg/mL]) and compared them regarding expression of the 36 DC-derived Th stimuli (Figure 2C). C32 and C33 did not exhibit significant differences in CD80 and CD86 expression, reflecting equal levels of DC activation. They were statistically different only for IL-6, with levels ranging from complete absence in C33 to over 1 ng/mL in C32 (Figure 2C). In contrast, the pairs C47/C48 and C61/C62 showed significant differences for multiple Th stimuli. C47 expressed significantly more CD86, PDL1, and IL-1 than C48. On the contrary, C48 expressed higher levels of 4-1BBL. C61 and C62 showed marked differences in CD70 and IL-12 (higher in C61) and OX40L (higher in C62) levels. Hence, each DC condition expressed unique combinations of DC-derived Th stimuli, suggesting different communication potential with CD4 T cells.

An unsupervised Gaussian mixture model showed that the highest Bayesian information criterion (BIC) value corresponded to 82 groups, confirming that each DC condition induced a unique profile of the 36 communication molecules (Figure 2D).

Using principal-component analysis (PCA), we showed that neither the date of the experiment nor the donor batch had a major effect on clustering (Figure S1C; STAR Methods).

The Heterogeneity of DC-Induced Th Cytokine Responses

We characterized the diversity of CD4 T cell output responses, as assessed by Th cytokine profiles, following co-culture of naive CD4 T cells with activated DCs across the 82 conditions described previously. Th cytokines exhibited important variations across the 428 observations (Figure 3A). Some cytokines, such as IL-2, TNF-α, GM-CSF, TNF-β, and IL-3, were always detected (Figure S2A).

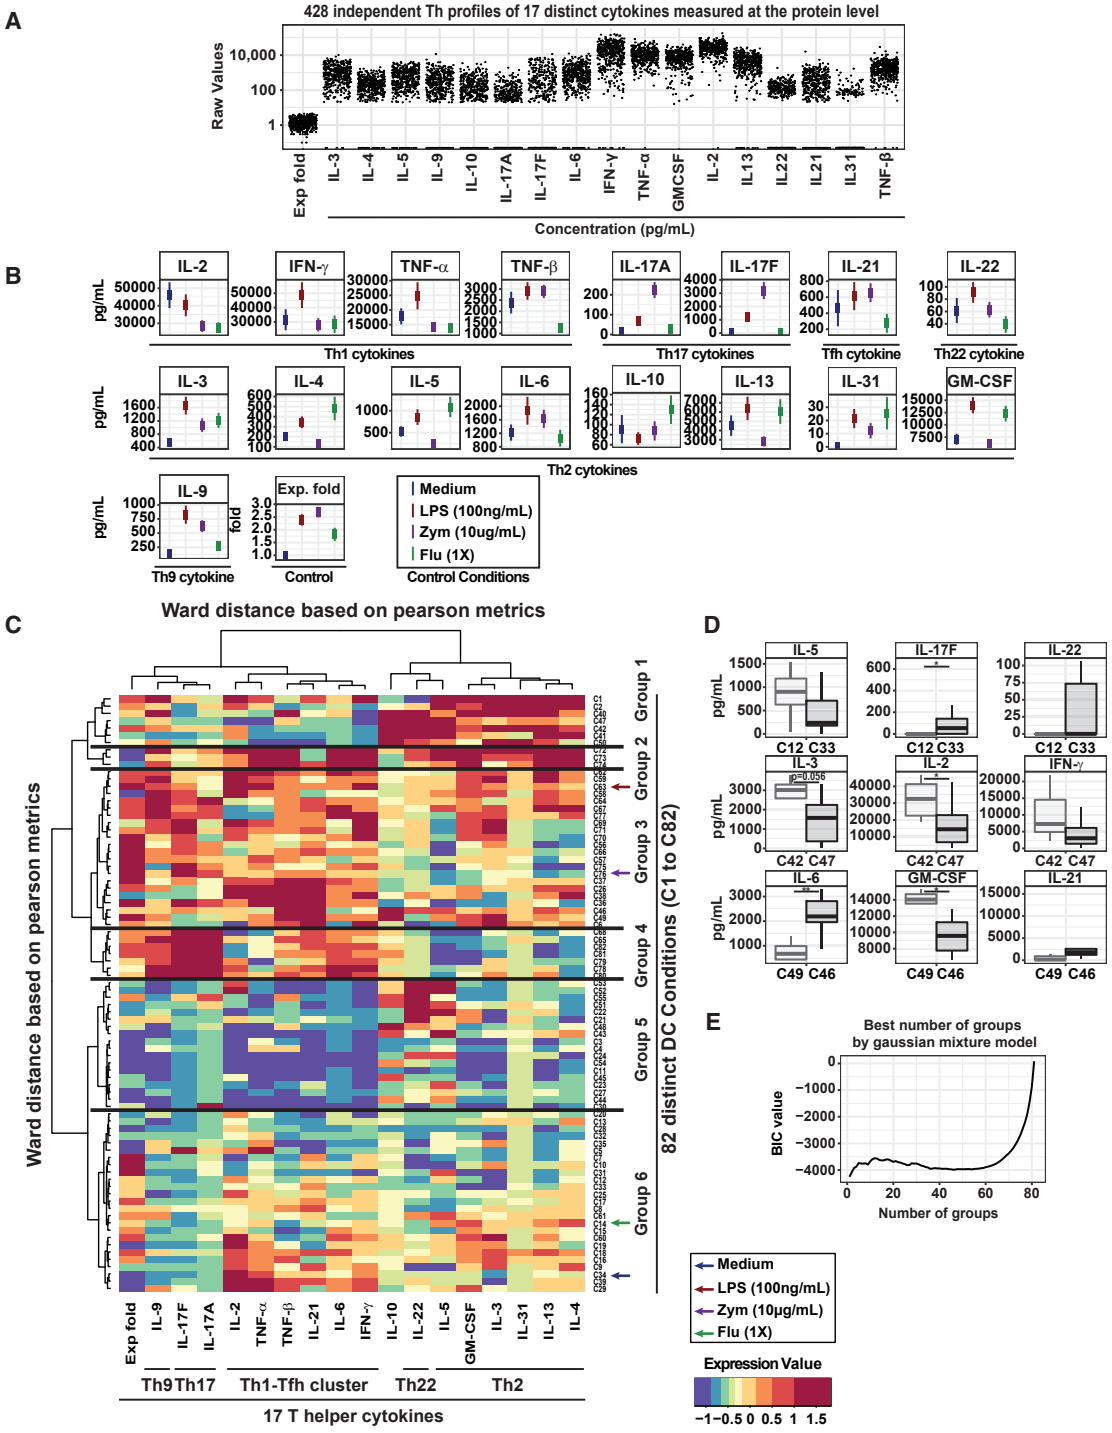

Figure 3. Th Cytokine Responses Mirror the Variability in DC Communication States

(A) Raw expression values of each of the 18 Th-derived parameters (n = 418 data points).

(B) Average expression values and SD for all Th-derived signals under the MoDC conditions medium, LPS, zymosan, and flu.

(C) Heatmap of expression values of each 18 Th parameters, performed with hierarchical clustering on Pearson metrics for the DC signals and Euclidian distances for the T cell conditions.

(D) Boxplot of Th signals for pairs of conditions selected as being the most correlated within our dataset by Pearson correlation (t test).

(E) Best number of groups by Gaussian mixture model, determined only using the 428 points of the 18 Th parameters.

To identify Th subset signatures, we compared cytokine expression under our four standard conditions: medium (negative control), LPS, zymosan, and flu. The Th17 cytokines IL-17A and IL-17F were induced predominantly in zymosan MoDCs. LPS MoDCs induced mixed Th1, Th2, and Th9 responses characterized by high IFN-γ, IL-13, IL-3, and IL-9 compared with medium. Flu MoDCs induced the Th2 cytokines IL-4, IL-5, and IL-31 (Figures 3B and S2B). These results indicate that, under the LPS, zymosan, and flu conditions, each DC state induced a distinct set of Th cytokine responses corresponding to prototypical Th signatures or mixed Th profiles.

Th Cytokine Responses Mirror the Variability in DC Communication States

We asked whether Th cytokine responses would reveal distinct patterns or a continuum of responses mirroring each of the DC communication states (Figure 2A). We performed hierarchical Pearson clustering on our 18 distinct Th-derived variables across the entire 82 DC-activating conditions (Figure 3C). This revealed 6 distinct groups, although intra-group heterogeneity was evident in almost all groups. Interestingly, DC perturbation conditions (C1–C82) did not appear in the same order compared with DC communication signal clustering (Figure 2A), indicating that closely related patterns of DC-derived Th stimuli did not necessarily induce the closest patterns in Th cytokine responses.

Group 1 was dominated by production of IL-10, IL-22, IL-5, GM-CSF, IL-3, IL-31, IL-13, and IL-4 (Figure S2C). Group 2 was the most heterogeneous and included the inflammatory cytokines TNF-α and IL-6 co-expressed with variable levels of the Th1 (IFN-γ) and Th2 (IL-4 and IL-13) cytokines (Figure S2C). Group 3 expressed IL-21, IFN-γ, and IL-17F but no or low IL-17A, suggesting the possibility of differential regulatory mechanisms (Figure S2C). Group 4 was dominated by the Th17 cytokines IL-17A and IL-17F, group 5 by IL-22, and group 6 by IL-2. Distinct sets of DC perturbation conditions and, hence, patterns of DC-derived communication molecules were associated with each of these groups (Figure 3C). This was the first suggestion of specific rules underlying input-output relationships in DC-Th communication.

Because of intra-group heterogeneity, we asked whether most correlated conditions within the same cluster would differ from each other (Figure 3D). C12 and C33 were associated to different levels in IL-17F, whereas C42 and C47 were different in IL-2 and C46 and C49 were different in IL-6 and GM-CSF levels (Figure 3D). As for the DC dataset, we found that 82 was the best number of groups in our Th-derived dataset, based on a Gaussian mixture model (Figure 3E). This suggested that a single DC profile of communication molecules would induce a unique set of Th cytokines.

A Data-Driven Lasso Penalized Regression Model Predicts Th Cytokine Responses from Combinations of DC-Derived Th Stimuli

Having generated distinct patterns of DC-derived communication signals associated with a diversity of induced CD4 T cell cytokine responses, the question of their relationship appeared to be critical to decipher DC-Th communication. Given the complexity of the dataset and the lack of clear hypotheses concerning the majority of DC-derived Th stimuli, we applied an unsupervised mathematical modeling strategy (Figure 4A).

The MultiVarSel strategy with stability selection performed similarly as the internal positive control and better than other methodologies tested (Figure S3A; STAR Methods). Therefore, we applied MultiVarSel to the modeling of our experimental data (Figure 4A). This methodology takes into account the dependencies that may exist among Th cell cytokines and combines Lasso criterion and stability selection to select associations between DC-derived signals (inputs) and Th cytokines (outputs) (STAR Methods).

Our multivariate model identified a large number of significant positive (red) and negative (blue) associations of the 36 DC-derived Th stimuli with the 17 Th-derived cytokines (Figure 4B). White squares represent the absence of significant association (Figure 4B). The frequency of selection obtained for each input-output association is provided in Figure S3B.

Our mathematical model revealed (1) the effect of each DC communication signal on Th output responses and (2) the critical regulators for each Th cytokine. For example, negative regulators of IL-10 were OX40L, 4-1BBL, IL-12, TNF-α, CD58, VISTA, Galectin-3, CD80, CD29, IL-1, ICAM-3, SLAMF3, IL-28α, and CD83, and positive regulators were Jagged-2, PDL1, IL-10, CD11a, HLA-DR, ICOSL, CD100, CD30L, CD18, ICAM-2, and CD86 (Figure 4B). Hence, the model can predict IL-10 production by responding Th cells for any DC, given the expression level of these molecules. It allows simulating loss or gain of function of an input. Similar insight can be obtained for each of the 17 Th cytokine responses, which may be explained by a combination of DC-derived communication signals.

We used computational cross-validation to evaluate the error of prediction of our model (Figure 4C). For all Th cytokines, the multivariate outperformed the best univariate model (Figure S3C). We ranked Th cytokines based on their prediction errors; the Th variables best explained by our model were IL-6, IL-17F, Exp Fold, and IL-3 (Figure 4C).

To address DC type specificity in model performance, we calculated the cross-validation error for each Th output of the MoDC and bDC dataset, respectively. Our model predicted equally well the majority of the outputs for the two DC types (Figure S3D). For a few outputs, mostly IL-22 and TNF-β, the model was more error prone in bDCs than MoDCs (Figure S3D). Interestingly, a higher prediction error was found for TNF-β in 5 of 118 observations (Figure S3E), where TNF-β levels were very high (range, 6.7–22.2). This suggested that a TNF-β-promoting input signal might be involved in those 5 cases but not included in our model. For IL-22, more observations had a higher prediction error in bDCs compared with MoDCs, but the prediction error range and distributions were similar, suggesting that the input-output relationship was conserved (Figure S3E).

We performed hierarchical clustering for both DC and T cell-derived variables to identify co-regulations between Th outputs. We retrieved relevant clusters of Th cytokines belonging classically to the same Th subset (Figure 4B). The Th2-related cytokines IL-13, IL-31, IL-5, IL-4, IL-10, and GM-CSF were found in the same cluster, suggesting that their induction would be controlled by similar mechanisms. IL-17A and IL-17F were also

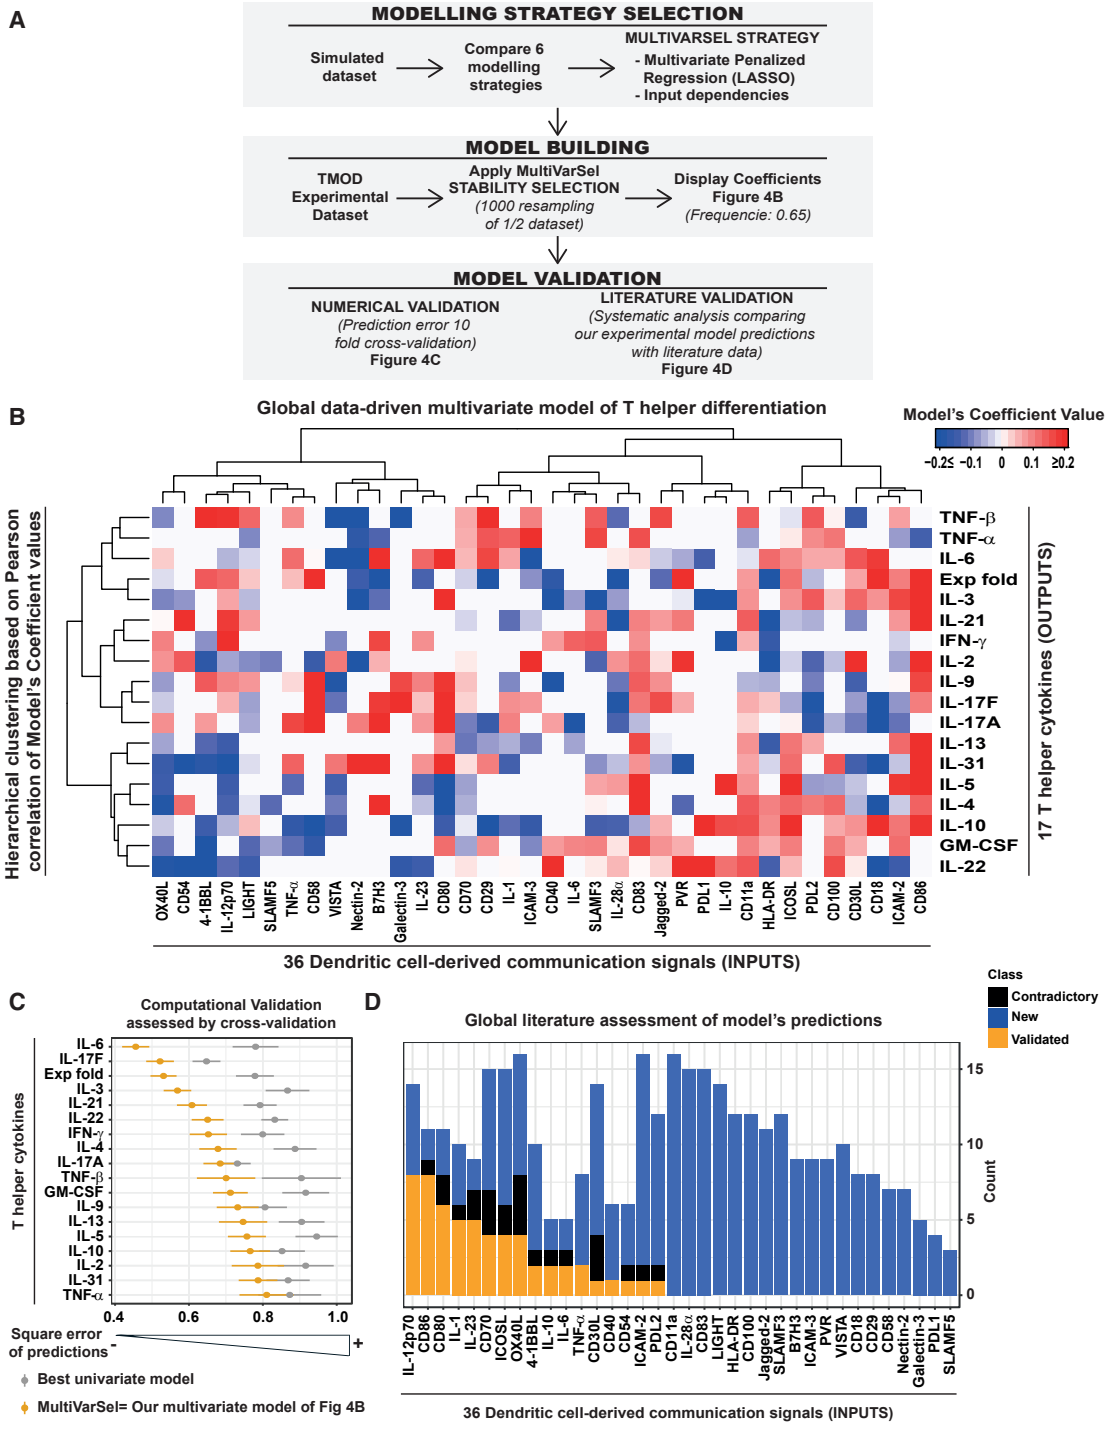

Figure 4. A Data-Driven Lasso Penalized Regression Model Predicts Th Differentiation Outcomes from DC-Derived Communication Signals

(A) Mathematical modeling strategy.

(B) Heatmap of the model's coefficient values of the MultiVarSel-derived model, explaining the 18 Th parameters based on the 36 DC-derived signals (Pearson correlation-based hierarchical clustering).

(C) Mean and SE of prediction error values obtained by 10-fold cross-validation for Th parameters using the multivariate model (yellow) and the best univariate model (gray) within the 36 DC signals.

(D) Literature-based validation score. For each DC signal, all predicted associations with Th cytokines were categorized as "new," "validated," or "contradictory."

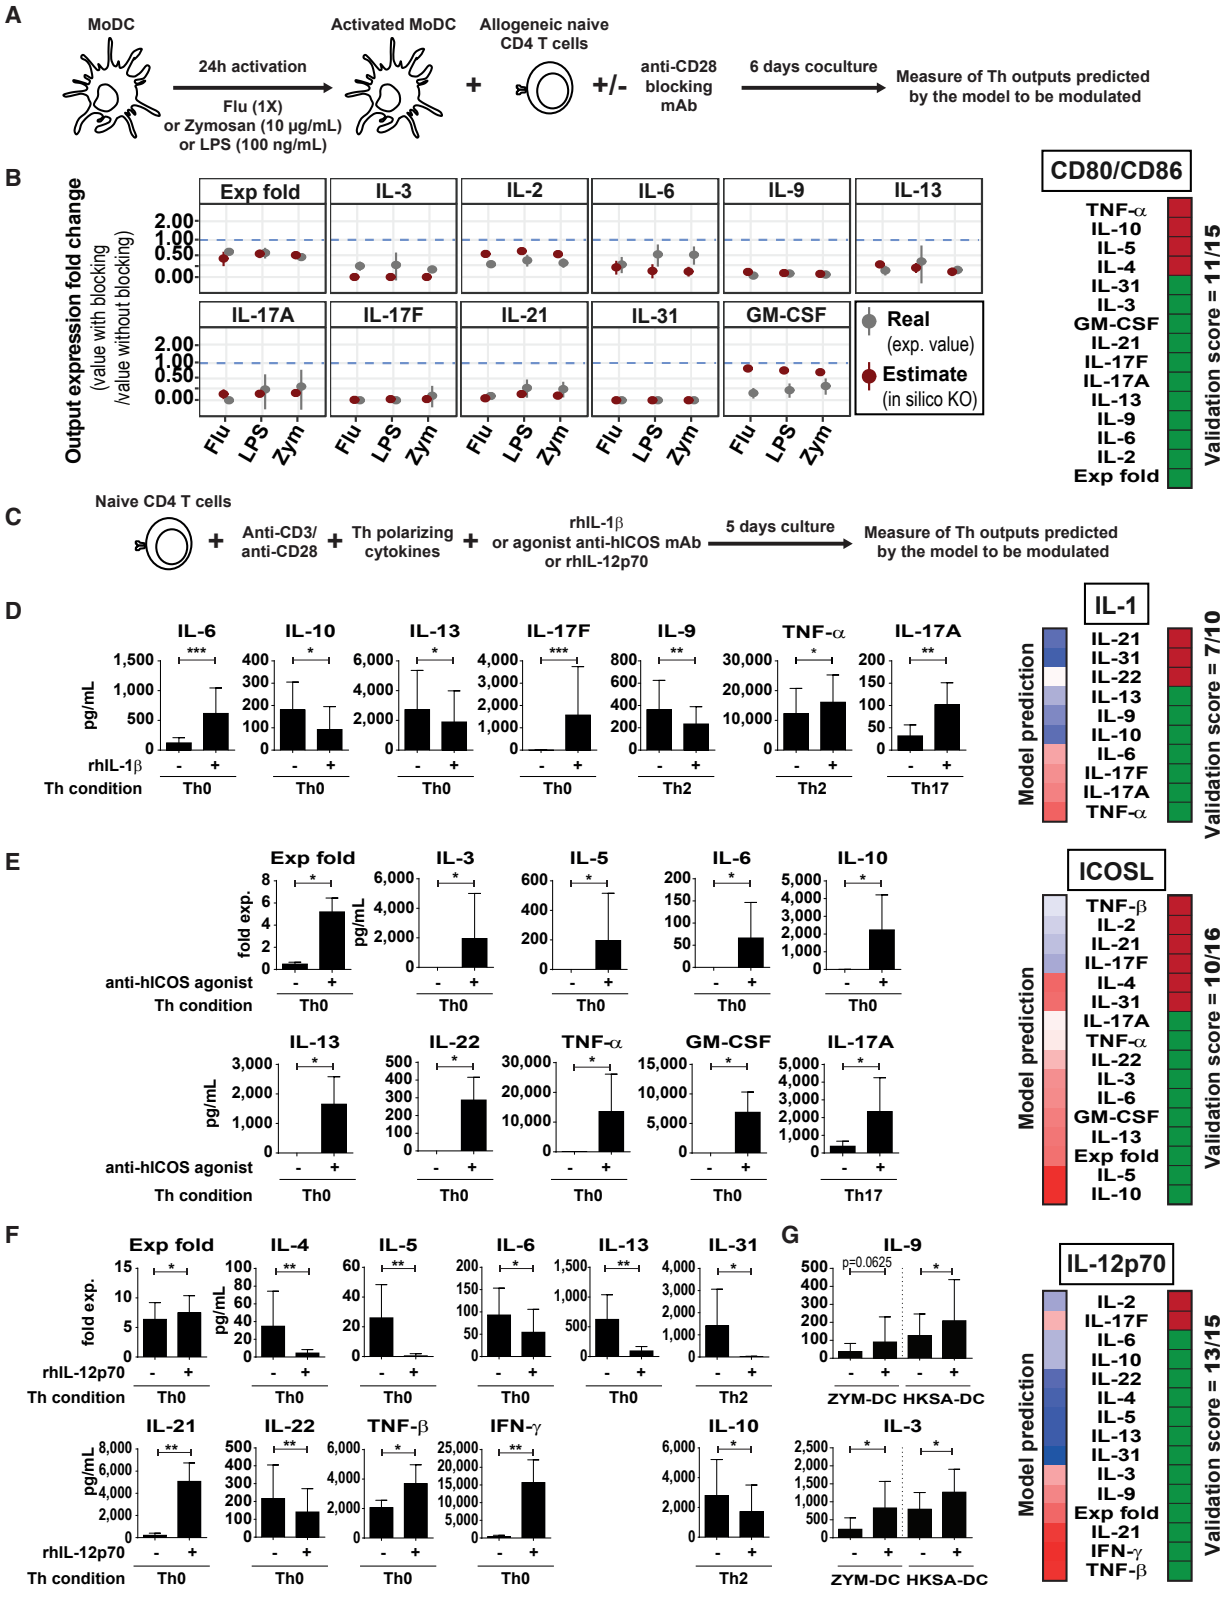

(legend on next page)

in the same cluster, implying that the model associated them with closely related DC communication signals (Figure 4B). Surprisingly, our model closely related IL-9 expression to IL-17A and IL-17F, suggesting common regulators. It also clustered IL-22 closer to the Th2 than to the Th17 cytokines. IL-21 was associated with the Th1 cytokines IL-2 and IFN-γ (Figure 4B).

The Multivariate DC-Th Model Reveals Novel Regulators of Th Cytokine Responses

We systematically compared our model results with the literature as a knowledge-based validation but also novelty assessment. We screened 178 relevant articles (STAR Methods) and extracted information regarding specific molecular control of a given Th cytokine by DC-derived signals measured in our model (Table S3). We computed a validation score based on the number of articles identifying the same associations than our model (STAR Methods). IL-12 ranked as the top DC communication signal for which our model predictions globally recapitulated existing knowledge (8 of 13 predicted associations). Among other known associations, IL-23 was positively associated with IL-17A and IL-17F, IL-10 was positively associated with IL-10 and negatively with IFN-γ, and CD40 was positively associated with IFN-γ.

However, the model also predicted 290 associations that were not described previously. Putative novel regulators were identified for all Th outputs (Table S4). The robustness of each prediction could be estimated by the value of the coefficient and by the frequency of detection of the association (Table S4). Examples of high scores were B7H3 and CD83 association with IL-4, 4-1BBL association with IL-9, ICOSL association with IL-13, and OX40L negative association with IL-22 (Table S4). Overall, literature knowledge was retrieved for 80 distinct input-output relationships presented in our model (Figure 4B); 56 were in agreement with our model, representing a global literature validation score of 70%.

Systematic and Independent Experimental Validation of Model's Predictions

We performed systematic experimental validation by selecting a subset of target inputs and systematically measuring the Th outputs selected by our model. We assessed the novelty of each validated prediction (Table S3).

First we addressed systematic validations of model predictions by blocking experiments (Figure 5A). We performed double *in silico* knockout for CD80 and CD86 under the three conditions—LPS (100 ng/mL), flu (1×), and zymosan (10 μg/mL) MoDCs—in which CD80 and CD86 were highly expressed and predicted an effect on 15 distinct Th outputs (Figure 5B), 11 of which were successfully experimentally validated (STAR Methods). The positive role of CD80 and CD86 on IL-3 and IL-31, to our knowledge, have not been described elsewhere. The predictions we failed to validate were for IL-4, IL-5, IL-10, and TNF-α (Figure S4A), all predicted to be decreased by CD80/CD86.

Then we validated the effects of three additional inputs: IL-1, ICOSL, and IL-12 used as exogenous factors (Figure 5C). First we gave the selected input together with anti CD3/CD28 signals (Th0) and systematically measured all Th outputs predicted by the model to be influenced by that input. In the absence of any effect, we gave the selected input under a Th2 (IL-4) or Th17 (IL-6, IL-1β, IL-23, and TGF-β) condition to detect additional synergistic or inhibitory effects required to validate the predicted effect. For example, it is not possible to validate the inhibition of a Th2 cytokine without significant production of this cytokine at baseline.

We focused on the ten predictions made by our model for IL-1 (Figure 5D). By adding IL-1β to the Th0 condition, we were able to detect significant upregulation of IL-6 and IL-17F and significant downregulation of IL-10 and IL-13. IL-10 downregulation and IL-6 upregulation were also significant in the Th2 context (Figure S4B). Under a Th2 condition, we validated significant upregulation of TNF-α and downregulation of IL-9 by IL-1β (Figure S4B), not seen in Th0 (Figure S4B). Under a Th17 condition, we observed a positive effect of IL-1β on IL-17A. We could not validate the predictions regarding IL-21, IL-31, and IL-22 (Figure S4B). In total, 7 of 10 predicted effects of IL-1 were validated. Interestingly, the positive role of IL-1β on induction of IL-6 by Th cells was not known (Table S3) and may suggest new biology and amplification loops in an inflammatory context.

We used a similar strategy to validate predictions regarding ICOSL using an anti-ICOS agonistic antibody. Overall, we validated 10 of 16 predictions (Figure 5E and S4C; STAR Methods). Interestingly, five of the 10 validated predictions were novel (Table S3; IL-5, IL-13, IL-3, GM-CSF, and IL-22), suggesting common pathways to induce IL-22 and Th2 responses.

Finally, we experimentally tested the predictions regarding IL-12 (Figure 5F). Adding IL-12 to the Th0 condition validated an induction of IFN-γ, IL-21, Exp Fold, and TFN-β. We also validated the inhibitory role of IL-12 on Th2 cytokine (IL-4, IL-5, and IL-13), IL-6, and IL-22 production. Using the Th2 condition, we further validated the inhibitory role of IL-12 on IL-10 and IL-31. The effects of IL-12 on TNF-β, IL-31, and IL-6 have not been described previously (Table S3).

Because our anti-CD3/CD28 system did not allow validating IL-12 effects on IL-2, IL-17F, IL-3, and IL-9 (Figure S4D), we wondered whether DC-dependent factors could affect the role

Figure 5. Independent and Systematic Experimental Validation of the Model's Prediction

(A) CD28 blocking experimental design in DC-T co-culture.

(B) Comparison of the predicted versus observed fold change following CD28 blocking; n = 6 donors.

(C) Experimental scheme of the "adding" validation procedure used in (D)–(F).

(D) DC-free validation experiment studying the effect of adding IL-1β in Th0, Th2, and Th17. Naive T cells were stimulated by anti-CD3/CD28 beads; n = 6 donors.

(E) DC-free validation experiment studying the effect of adding ICOS in Th0 and Th17. Naive T cells were stimulated by coated anti-CD3 and ICOS antibodies and soluble anti-CD28; n = 6 donors.

(F) IL-12 validation experiments in the DC-free system. Naive T cells were stimulated by anti-CD3/CD28 beads under Th0 and Th2 conditions; n = 8 donors.

(G) Validation of IL-12 predictions regarding IL-3 and IL-9. bDCs were cultured with naive CD4 T cells. IL-12 at 10 ng/mL was added for 6 days; n = 6 donors.

For (B) and (D)–(G), each panel shows the mean and SD of cytokine concentration, measured on restimulated Th supernatants (Wilcoxon test).

of IL-12 on these cytokines. We selected DC conditions with very low production of IL-12 (C51 and C55; Figure 2A) and performed a co-culture with naive T cells, adding or not adding IL-12. As a positive control, IL-12 was able to induce IFN-γ in both zymosan and HKSA conditions (Figure S4E). We did not validate the role of IL-12 on IL-2 or IL-17F regulation (data not shown). However, we validated that IL-3 was induced by IL-12 in both zymosan DCs (C51) and HKSA DCs (C54) (Figure 5G), whereas IL-9 was significantly upregulated only in HKSA DCs. Overall, we were able to experimentally validate 13 of 15 predictions regarding IL-12.

Our systematic strategy established a validated prediction of the input-output relationship in 41 of 56 cases (73.2%), 13 representing new mechanisms identified by the model. This number is similar to or higher than the computational cross-validation (Figure 4C). Predictions with higher stability selection frequencies were more validated than those with low stability selection (Figure S4F). However, the value of the model's coefficients was not statistically different between the two groups (Figure S4F), indicating that the model efficiently captured associations with low coefficient values.

Although IL-12 was the input best explained by our model, we could not validate the predicted association between IL-12 and IL-17F (Figure S4D), neither in the literature nor in our systematic experimental validation. Previous studies have shown either no effect (Volpe et al., 2008) or a negative effect (Acosta-Rodriguez et al., 2007) of IL-12 on Th17 differentiation. We hypothesized that context-dependent effects may lead to new functions of IL-12, not accomplished by IL-12 as a single agent.

A Context-Dependent Model Reveals a Role of IL-12 in Th17 Differentiation

We designed a strategy to capture context-dependent effects of one input on any given output by integrating new composite variables into the model (Figure 6A). These new input variables were based on the co-occurrence of a given input with other DC-derived communication signals (i.e., contexts). They adopted the value of the given input (for instance, IL-12) in each observation where the co-expressed DC signal was present, and they took a zero value when the co-signal was absent. We could derive 455 context-dependent variables.

The model including all context-dependent variables performed less well (higher error of prediction) than our classical MultiVarSel strategy (Figure S5A), most likely because of overfitting issues dependent on the dataset size, with a number of input variables exceeding the number of data points used to fit the model. Therefore, we derived 36 models, each one integrating the context dependencies of one input (Table S5). For each of these models, we reported the coefficient and the stability selection frequencies of each input (Table S5). To globally estimate the influence of context dependencies within our data, we quantified the number of times an input variable was selected, either "alone" or "with" another one. We derived percentages of context dependencies and represented the results either per input (Figure S5B) or per output (Figure S5C). The inputs most likely to present "context-dependent" functions were PDL1 and SLAMF3, whereas CD11a and CD70 were mostly context-independent (Figure S5B). When analyzing the outputs, the models revealed that all cytokines could be regulated by context-dependent mechanisms with relatively similar percentages (range, 0.13–0.22) (Figure S5C).

We used this strategy to explain the role of IL-12 in the control of Th17 differentiation through identification of context-dependent effects. We found that adding context-dependent variables for IL-12 improved the model predictions for IL-17F and performed equally well for IL-17A (Figure 6B). We then focused on DC-derived signals that were kept significant by the model and observed distinct associations of the new IL-12 context-dependent variables with IL-17A and IL-17F (Figure 6C), including some differentially associated with IL-17A and IL-17F, respectively. Among various contexts, we found that IL-12 in the context of IL-1, ICAM-2, or Jagged-2 was associated with IL-17F, whereas IL-12 in the context of CD70, IL-23, or LIGHT was associated with IL-17A.

As a first level of *in silico* validation, we selected a DC condition under which IL-12 was co-expressed with many of these contexts, and DC-derived signals induced IL-17A and IL-17F by responder Th cells. Zymosan (10 μg/mL) on MoDCs fulfilled these criteria (Figures 1D and 3C). To study the specific effects of IL-12 in the context of all other DC communication signals induced by zymosan, we performed *in silico* IL-12 knockout in the IL-12 context-dependent model. We compared predicted values for IL-17A and IL-17F when IL-12 was kept or not kept in the model (Figure 6D). *In silico* knockout of IL-12 diminished the production of both IL-17A and IL-17F under the zymosan (10 μg/mL) condition. As experimental validation, we performed independent DC/T cell co-culture experiments using MoDCs treated with 10 μg/mL zymosan in the presence and absence of IL-12-neutralizing antibody (Figure 6E). Blocking IL-12 significantly decreased the production of IL-17A and IL-17F, as predicted (Figure 6E), and inhibited IFN-γ production (Figure S5D). The same model predicted no effect of blocking IL-12 in Curdlan MoDCs (Figure S5E), which we validated experimentally (Figure S5F).

Synergistic Interaction between IL-12 and IL-1 Explains Induction of IL-17F without IL-17A

Our model predicted distinct roles of IL-12 on IL-17A and IL-17F production depending on the context in which IL-12 is expressed. Interestingly, IL-12, IL-1, and CD80 were the top variables almost systematically selected by the model to explain the differences between IL-17A and IL-17F (Figure 7A). This corroborated the results in Figure 6C, where we found that IL-12 in the context of IL-1 was associated with IL-17F but not IL-17A. The model estimate for a stability selection of less than 0.8 indicated that IL-12, IL-1, and CD80 were positive contributors to the differences between IL-17A and IL-17F (Figure S6A). Consequently, we hypothesized that the combination of IL-12 with IL-1 would induce IL-17F independent of IL-17A.

To experimentally validate our hypothesis, we used a DC-free Th polarization assay, allowing us to specifically study the interaction between IL-12 and IL-1 regardless of any other molecular context. Naive CD4 T cells were polyclonally activated with anti-CD3/CD28 beads and put in distinct cytokine treatments: Th0 (no cytokine) and Th2 (IL-4) as negative controls; Th17 (IL-1β+IL-23+IL-6+TGF-β) as a positive control, IL-12, IL-1β, and IL-12+IL-1β. IL-12 alone induced IFN-γ and IL-21 and inhibited Th2-related cytokines, as expected (Figure S6B). IL-12 alone induced neither IL-17F nor IL-17A, but combining IL-12

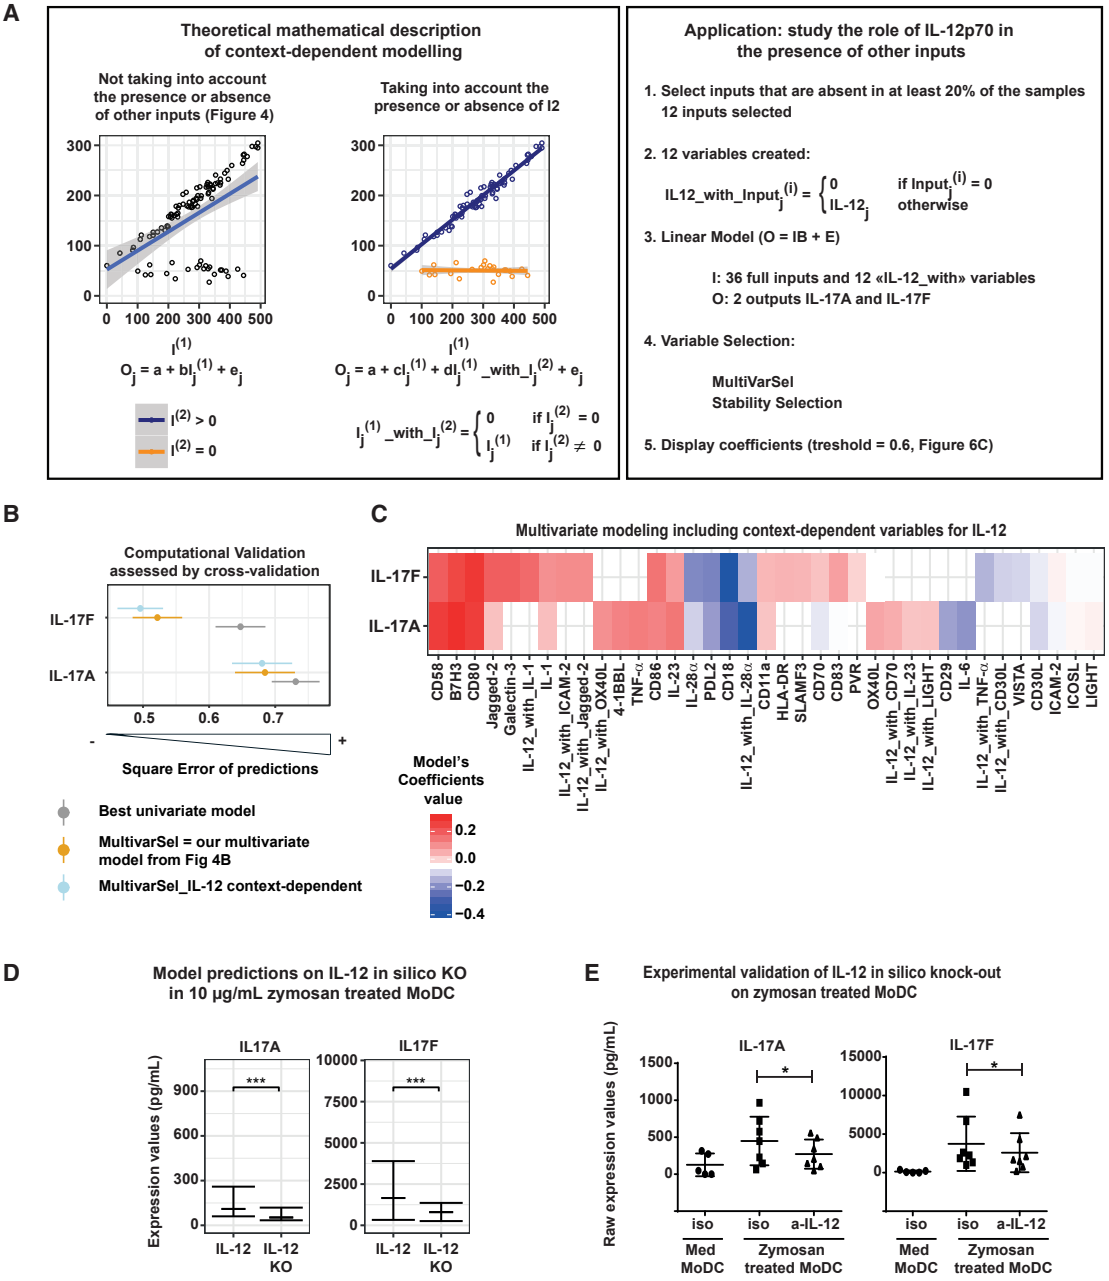

Figure 6. A Context-Dependent Model Reveals a Role of IL-12 in Th17 Differentiation

(A) Context-dependent modeling and application to IL-12. I, input; O, output.

(B) Error of prediction values obtained by 10-fold cross-validation for IL-17A and IL-17F, comparing the best univariate model (gray), MultiVarSel (yellow), and MultiVarSel with context dependencies (blue).

(C) Heatmap of the model's coefficient value of the context-dependent multivariate model explaining IL-17A and IL-17F.

(D) Model predictions regarding IL-12 *in silico* knockout (KO) under the zymosan MoDC condition for IL-17A and IL-17F values (blue) compared with experimental values in the presence of IL-12 (yellow); paired t test.

(E) Concentrations of IL-17A and IL-17F produced by Th cells after differentiation with zymosan MoDCs in the presence of anti-IL-12 neutralizing antibody or a matched isotype; n = 6 donors, paired t test.

with IL-1β dramatically induced IL-17F at levels comparable with the positive control, without a detectable amount of IL-17A, which fully validated the model predictions (Figure 7B).

This effect was specific to the IL-12+IL-1β combination IL-6, IL-23, or TGF-β alone or combined with IL-12 could not induce IL-17F expression (Figure S6C). The exact same pattern of Th

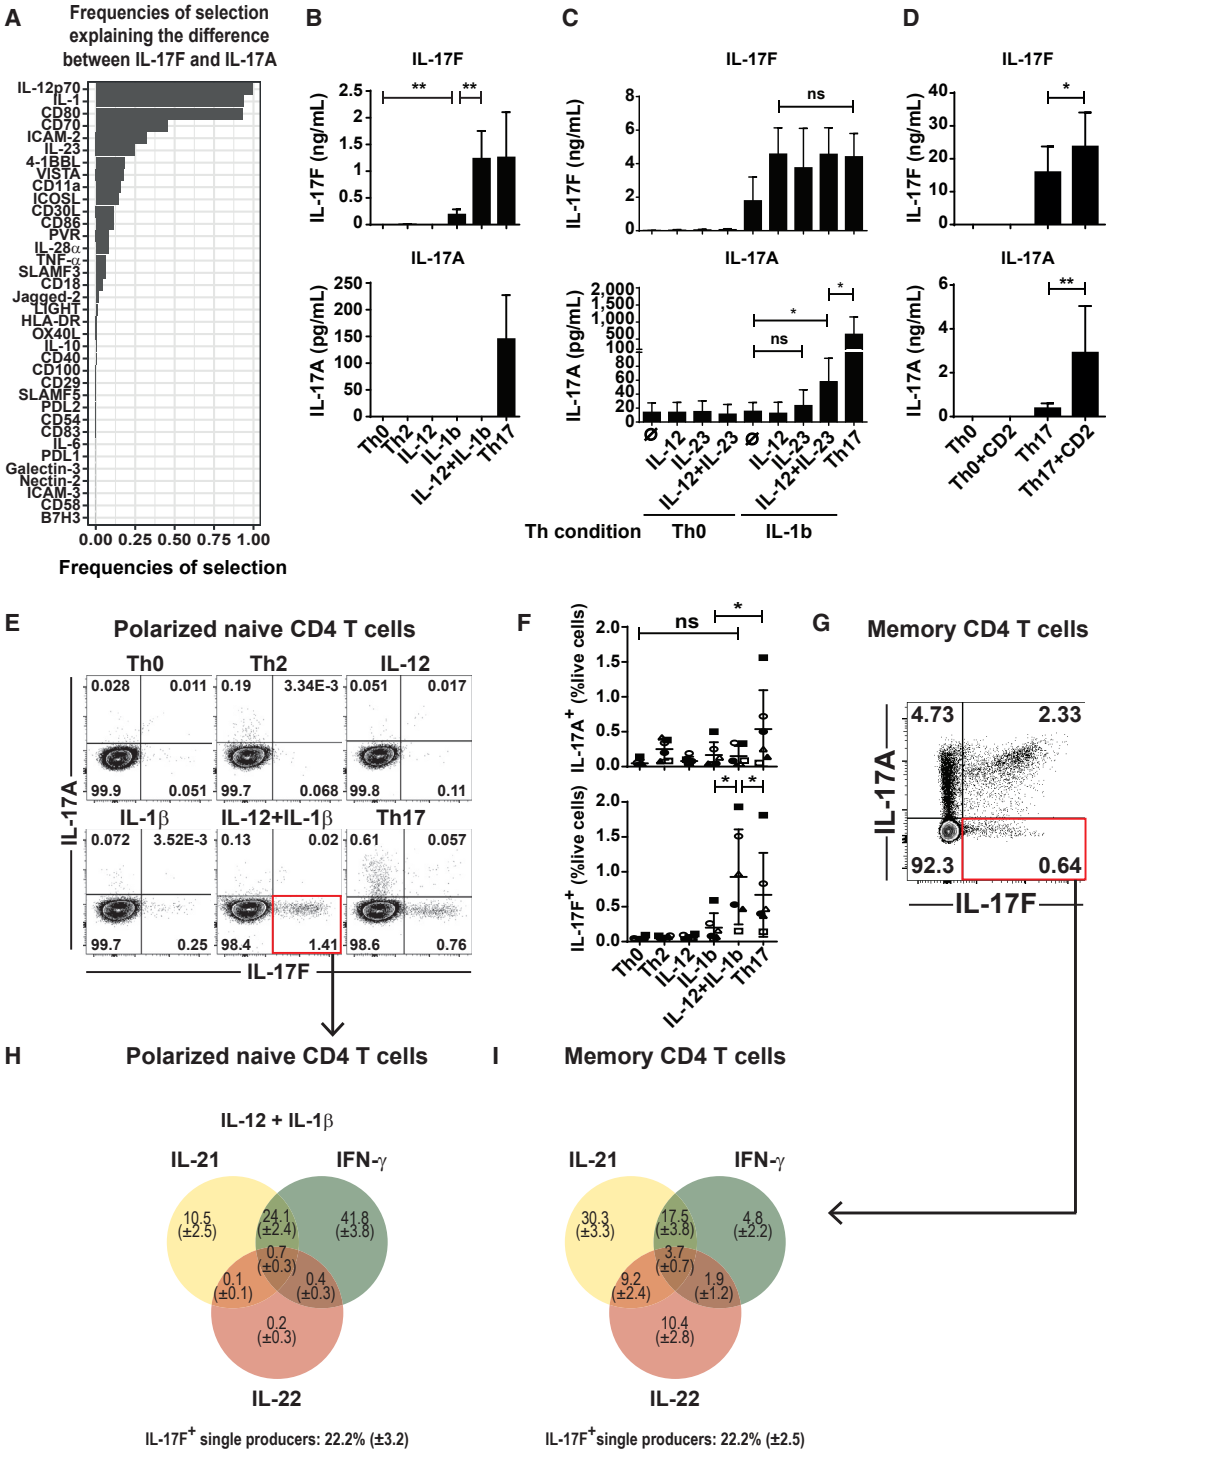

Figure 7. Synergistic Interaction of IL-12 and IL-1 Promotes IL-17F without IL-17A

(A) Stability selection frequencies of the different DC signals by a multivariate model, explaining the difference between IL-17F and IL-17A.

(B) Concentration of cytokines measured on restimulated Th supernatants. Naive CD4 T cells were differentiated for 5 days with anti-CD3/CD28 beads under the indicated conditions; n = 6 donors, paired t test.

(C) The same experimental design as in (B), with conditions as annotated; n = 6 donors, Wilcoxon test.

(D) Coated anti-CD2 and anti-CD3 together with soluble anti-CD28 were given for 5 days to naive CD4 T cells under Th0 or Th17 conditions. Cytokine concentrations were measured after 24-h restimulation on day 5. Mean and SD are shown; n = 8, Wilcoxon test.

(legend continued on next page)

cytokine expression was obtained by combining IL-1α or IL-1β with IL-12, which fit model predictions because those two variables were highly correlated (Figure S6D). The capacity of IL-12+IL-1β to induce IL-17F was resistant to the presence of other Th differentiation factors, such as IL-4 (Figure S6E). Using CellTrace Violet (CTV; Figure S6F), we could show that the production of IL-17F could not be attributed to the distinct proliferation capacity of Th cells under the IL-12+IL-1β condition.

Next we questioned whether Th cells generated under the IL-12+IL-1β condition would express transcription factors classically associated with Th17 differentiation. We measured 63 RNA transcripts by qPCR under Th0, Th2, IL-1β, IL-12, IL-12+IL-1β, and Th17 conditions (Table S6). The 63 genes included master regulators of the Th1 and Th2 subsets, such as T-bet and GATA3, respectively, and Th17 regulators, such as RORc, STAT3, BATF, and SATB1 (Ciofani et al., 2012). IL-17A and IL-17F regulation at the mRNA level mirrored the protein level (Figure S6H). IL-12+IL-1β induced significantly more RORc, BATF, and Bcl6 than IL-12 or IL-1β alone (Figure S6H), which could explain the induction of IL-17F and IL-21. Still, the levels of RORc and Bcl6 were lower in IL-12+IL-1β than under the Th17 condition (Figure S6H). T-bet was highly induced in IL-12+IL-1β in comparison with the IL-12 or Th17 conditions, indicating that Th1 differentiation was maintained and that T-bet did not inhibit IL-17F production. IL-12Rb2, a Th1 marker, was downregulated by IL-1β when added to IL-12, whereas IL-12, IL-12+IL-1β, and Th17 conditions all induced the IL-23 receptor (Figure S6H). SATB1 was specifically upregulated in IL-12+IL-1β in comparison with Th17 or IL-1β alone (Figure S6H), suggesting that it could play a role in the specific upregulation of IL-17F.

To globally assess the expression of the various Th lineage-specific factors, across IL-12- and IL-1-containing conditions, we performed a principal-component analysis (PCA) including all 63 mRNA variables (Figure S7A). Cells from the IL-12+IL-1β condition had an intermediate expression pattern between the IL-12 (Th1) and Th17 conditions. By decomposing the PCA space into vectors for each variable, we found that IL-17F, IL-23R, ICOS, and T-bet projected predominantly along the IL-12+IL-1β condition (Figure S7B), again pointing to mixed Th1/Th17 features.

We then addressed the link between IL-12 and IL-17A in various contexts. IL-12 with IL-23 was predicted to induce IL-17A but not IL-17F (Figure 6C). In a DC-free Th polarization assay, we used IL-12, IL-23, or IL-12+IL-23 and found that none of these conditions induced IL-17A (Figure 7C). We hypothesized that a third input could explain the positive link between "IL-12_with_IL-23" and IL-17A. Using an unsupervised analysis, we found IL-1 as a top variable with the highest correlation (Figure S7C). In addition, IL-12 and IL-17A positive correlation was significant specifically in the group of data points where IL-23 and IL-1 were expressed (Figures S7D and S7E) and was lost when only IL-1 or IL-23 was expressed with IL-12 (Figure S7D). Therefore, we tested whether IL-12+IL-23 would induce IL-17A in the presence of IL-1β. We validated a significant induction of IL-17A with no effect on IL-17F when IL-12 and IL-23 were given in the presence of IL-1β compared with IL-12 or IL-23 (Figure 7C). We measured IL-17A and IL-17F by qPCR and retrieved the same induction pattern (Figure S7F). Last, we could show that RORc was higher in IL-12+IL-23+IL-1β than in IL-12+IL-1β (Figure S7F).

Finally, we observed that our modeling strategy always identified CD58 as a main Th17 inducer because it positively affected both IL-17A and IL-17F (Figures 4B and 6C), an association that we had not seen during our systematic literature review (Figure 4D; Table S3). To test this hypothesis, we used an agonist anti-CD2 antibody that mimics the presence of CD58 (STAR Methods). As predicted, IL-17A and IL-17F were not induced by anti-CD2 alone under the Th0 condition. However, anti-CD2 significantly induced production of IL-17A and IL-17F under Th17 conditions (Figure 7D), which was confirmed by intracellular FACS staining (Figures S7H and S7I), with IL-17F upregulation restricted to IL-17A-positive cells (Figure S7I).

To establish the cytokine co-expression profiles of IL-12+IL-1β-treated Th cells at the single-cell level, we performed intracellular cytokine staining (Figure 7E). We confirmed that IL-12+IL-1β induced significantly more IL-17F-positive Th cells without co-production of IL-17A (Figure 7F). In naive CD4 T cells polarized with the Th17 cytokine cocktail (IL-1β, IL-6, TGF-β, and IL-23) we mainly found two subsets of Th17 cells producing either IL-17A or IL-17F, with very few cells co-producing both cytokines. To check for *in vivo* existence of those IL-17A and IL-17F single producers, we analyzed the human CD4 T cell memory compartment by intracellular FACS in healthy donor peripheral blood mononuclear cells (PBMCs). We could identify a small fraction of Th cells expressing only IL-17F in the absence of IL-17A, suggesting that this phenotype constitutes a differentiation endpoint (Figure 7G).

To gain more insight into the functional properties of these "Th17F" cells, we studied their co-production with IL-21, IFN-γ, and IL-22, all relevant to the Th17 and/or IL-12 pathways, *in vitro* (Figure S7J) and *ex vivo* (Figure S7K). Among IL-17F$^+$IL-17A$^-$ cells generated with IL-12 and IL-1β, the majority co-produced IFN-γ (41.8%), IL-21 (10.5%), or both (24.1%) (Figure 7H), reflecting a dominant role of IL-12. IL-17F$^+$/IL-17A$^-$ memory CD4 cells preferentially co-expressed IL-21 (30.3%) and IL-21 together with IFN-γ (17.5%) (Figure 7I), which matched the *in vitro* differentiated CD4 T cells. In addition, the percentage of IL-17F$^+$/IL-17A$^-$/IL-22$^-$/IL-21$^-$/IFN-γ$^-$ cells between *in vitro* IL-12+IL-1β stimulation and the *ex vivo* restimulated memory

(E) Day 5 intracellular FACS analysis of Th cells differentiated as in (B). Dot plots show a representative donor.
(F) Quantification of live total CD4 T cells producing either IL-17A or IL-17F; n = 6 donors, paired t test.
(G) Representative donor of CD4 memory T cells with intracellular FACS staining for IL-17A versus IL-17F.
(H) Venn Diagrams of IL-17F$^+$/IL17A$^-$ Th cells co-producing IL-22, IFN-γ, and IL-21 of naive CD4 T cells under the IL-12+IL-1β. IL-12+IL; mean percentage and confidence interval, n = 6 donors.
(I) Venn Diagrams of IL-17F$^+$/IL17A$^-$ Th cells co-producing IL-22, IFN-γ, and IL-21 of memory CD4 T cells stimulated for 5 h with PMA and ionomycin; mean percentage of 6 donors with confidence interval.

compartment was similar (22.2%), which indirectly supported that IL-12+IL-1β induced the emergence of IL-17F single producers.

Taken together, our results demonstrate a synergy between IL-12 and IL-1 in inducing IL-17F single-producing Th cells with possible physiopathological relevance.

DISCUSSION

Cell-cell communication may involve several tens of communication signals functioning concomitantly and possibly interacting with each other. These signals, in turn, modify many molecular and functional parameters in target cells. Such complexity cannot be captured and formalized without an integrated mathematical modeling approach. Theoretical models of Th cell differentiation have already been established (Abou-Jaoudé et al., 2015; Naldi et al., 2010) and include a large number of possible inputs to T cells. However, they suffer from three limitations: (1) they include input signals that may be expressed by diverse cell types in different anatomical locations; (2) they do not recapitulate combinations of input signals in their naturally occurring patterns and concentrations; and (3) they use prior knowledge to infer input-output relationships, which does not integrate possible context-dependencies and interactions. In parallel, data-driven models have been developed in response to predefined stimuli, such as Th17 (Yosef et al., 2013) or Th1/Th2 (Antebi et al., 2013), which do not recapitulate the integration of multiple communication signals. In our study, we applied an unbiased data-driven approach specifically designed to model DC-Th communication. Combinations and concentrations of input communication signals were measured as naturally determined by their intrinsic biological regulation. Subsequently, the input-output relationships were learned from the experimental data and integrated into any underlying context dependency and interaction, even when not described previously. This maximizes the relevance of the model and the potential for novel discoveries.

Cells can change state in response to environmental cues, a concept defined as plasticity (da Silva-Diz et al., 2018; Liu et al., 2001). Each cell state may be associated with different communication potential; i.e., different expression patterns of communication signals (Soumelis et al., 2002; Wang et al., 2014). To broadly cover the possible DC states, we used various DC-stimulatory conditions (cytokines, viruses, bacteria, fungi) at various doses and combinations and across a large number of observations (>400). This prevented us from biasing our observations toward certain quantitatively or qualitatively extreme behaviors. After the model has learned the rules from such an extended range of observations, we anticipate that it should be able to predict behaviors in situations not necessarily covered in our original dataset, as confirmed in our computational and experimental validations. This opens possibilities of application in many areas of immunology, inflammation, and immunotherapy.

RNA sequencing (RNA-seq) has offered a means of capturing the expression of many communication signals and their receptors to infer cell-cell communication between various cell types (Vento-Tormo et al., 2018). However, the RNA-to-protein correlation can be rather low (Liu et al., 2016) and varies a lot depending on the gene (Edfors et al., 2016). Consequently, RNA copies of a gene cannot be associated with a given functional output, preventing quantitative mathematical modeling. Functional response in target cells can only be estimated indirectly through surrogate activation markers, which is most often not performed. In our approach, all measurements of communication signals and output variables were done at the protein level, hence directly measuring the bioactive communication molecules with a direct link to a specific response in target cells. This ensures robustness of the modeling strategy, as evidenced by our model's ability to recapitulate most of the known relationships in DC-Th cell communication.

Modeling complex biological behaviors in a quantitative manner is challenging. In data-driven models, it relies in large parts on the choice of explanatory (input) variables, which drive the induction or regulation of output variables. Here we selected DC-derived communication molecules through exhaustive literature mining. The model was able to integrate 36 input and 18 output variables in a quantitative manner, which makes it a reference in the field. We have been able to describe patterns of DC communication molecules in a way that goes beyond the classical view of immature versus mature DCs (Banchereau and Steinman, 1998; Guermonprez et al., 2002). In fact, we showed that almost every DC stimulatory condition leads to a distinct DC state. This is a first step in defining general combinatorial rules of DC-derived communication molecules: co-expressed molecules form the basis of putative context-dependent effects. Through the large number of variables handled by the model, we identified 290 novel associations explaining major immunoregulatory cytokines, which may lead to the discovery of novel functions of known DC molecules and suggest novel therapeutic targets.

Going further into the complexity of communication, we explored context dependencies of communication signals. In verbal communication, the context may dramatically alter the meaning of an individual word. Currently, there is no systematic way to search for context dependencies in biological communication. In our modeling strategy, we devised a method that introduces context-dependent variables for a given molecule. This allows unbiased identification of context-dependent functions that would have been missed by classical regression models. For example, we identified a new function for IL-12 in promoting IL-17F production by Th cells, which was completely unexpected based on prior knowledge (Korn et al., 2009). Identifying such context dependencies before therapeutic targeting of a DC-Th communication molecule may improve the prediction of its effect.

Given that DC-Th communication is central to a large number of physiopathological conditions (Keller, 2001), we can foresee multiple applications for the model. Based on expression pattern of DC molecules, the model can predict the induced Th cytokine profile. Quantitative measurements of DC communication molecules in a given disease or in an individual patient *ex vivo* can be used to simulate the corresponding Th response. Depending on the outcome, strategies may be devised to re-orient the response toward a protective or less pathogenic profile, again through model-based predictions. Alternatively, starting from a Th profile (cytokine or groups of cytokines), the appropriate molecular targets can be manipulated through gain- or loss-of-function experiments to amplify or inhibit a given Th cytokine. Last, the model can help predict the most appropriate vaccine adjuvant to obtain

protective immunity against some microbes or to re-orient a pathogenic Th response. For example, all DC molecules positively associated in the model to Th2 responses are potential targets to decrease pathogenic Th2 allergic inflammation (Ito et al., 2005; Nakayama et al., 2017; Soumelis et al., 2002).

Using DC-Th communication as a model, we established a framework that can now be applied to other types of cell-cell communication following 5 major steps: (1) systematic perturbation of the ''sender'' cell to generate a diversity of communication states; (2) broad, quantitative, and protein-level measurement of communication molecules relevant to the sender cell; (3) systematic quantitative assessment of the response in ''receiver'' or target cells; (4) MultiVarSel modeling of the input-output relationship, which defines communication rules; (5) *in silico* and experimental validation. Currently, we believe that cell type specificities in expression of communication molecules and in their function would prevent us from generalizing our DC-Th model to other cell types. Comparing different quantitative models of cell-cell communication will ultimately tell us whether cells speak the same language (i.e., whether they express similar patterns of communication molecules) and whether the same communication molecule has the same meaning (function) when expressed by two different cell types.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- LEAD CONTACT AND MATERIALS AVAILABILITY
- EXPERIMENTAL MODEL AND SUBJECT DETAILS
  - Human subjects
- METHOD DETAILS
  - PBMCs purification
  - MoDC generation and activation
  - Blood dendritic cells purification
  - $CD4^+$ T lymphocytes purification
  - Paired protein measurement in DC/T coculture
  - IL-12 blocking experiment
  - CD28 blocking experiment
  - Addition of rhIL-12p70 during DC/T coculture
  - DC-free Th cell polarization
  - ICOS agonism
  - CD2 agonism
  - Flow cytometry analysis
  - Cytokine quantification
  - Gene expression quantification
  - Anti-human ICOS monoclonal blocking antibody
- QUANTIFICATION AND STATISTICAL ANALYSIS
  - Dataset quality control – batch effect
  - Dataset quality control – T cell expansion
  - Statistical tests
  - Statistical analysis
  - Model comparison and ROC Curves
  - Modeling strategy
  - Systematic literature review
- DATA AND CODE AVAILABILITY

SUPPLEMENTAL INFORMATION

Supplemental Information can be found online at https://doi.org/10.1016/j.cell.2019.09.012.

ACKNOWLEDGMENTS

We thank the Institut Curie Cytometry platform for cell sorting, P. Gestraud and F. Coffin for advice regarding statistical analyses, and O. Lantz and N. Manel for important discussions. We wish to thank L. Pattarini, B. Fould, W. Cohen, O. Geneste, B. Lockhart, V. Blanc, and their teams from the Institut de Recherche Servier for having produced and generously shared the anti-ICOS antibody. This work was supported by funding from Agence Nationale de la Recherche: ANR-16-CE15-0024-01, ANR-10-IDEX-0001-02 PSL*, ANR-11-LABX-0043 and ANR-17-CE14-0025-02 from Center of Clinical Inverstigation IGR-Curie: CIC IGR-Curie 1428 and from Ligue contre le Cancer: EL2016.LNCC/VaS. M.G. was funded by ANRS and ARC.

AUTHOR CONTRIBUTIONS

M.G., C.T., L.K., C.C., and P.S. performed the experiments. M.G. and V.S. designed the experiments. M.P.-D. performed the statistical analysis. O.A., W.A.-J., and M.G. performed literature mining. P.H., D.T., F.B., L.S., J.C., C.L.-L., and V.S. participated in and supervised the statistical analysis. M.G. and V.S. wrote the manuscript. V.S. supervised the study.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: February 19, 2019
Revised: June 20, 2019
Accepted: September 9, 2019
Published: October 3, 2019

REFERENCES

Abou-Jaoudé, W., Monteiro, P.T., Naldi, A., Grandclaudon, M., Soumelis, V., Chaouiya, C., and Thieffry, D. (2015). Model checking to assess T-helper cell plasticity. Front. Bioeng. Biotechnol. *2*, 86.

Acosta-Rodriguez, E.V., Napolitani, G., Lanzavecchia, A., and Sallusto, F. (2007). Interleukins 1beta and 6 but not transforming growth factor-beta are essential for the differentiation of interleukin 17-producing human T helper cells. Nat. Immunol. *8*, 942–949.

Alculumbre, S., and Pattarini, L. (2016). Purification of Human Dendritic Cell Subsets from Peripheral Blood. Methods Mol. Biol. *1423*, 153–167.

Antebi, Y.E., Reich-Zeliger, S., Hart, Y., Mayo, A., Eizenberg, I., Rimer, J., Putheti, P., Pe'er, D., and Friedman, N. (2013). Mapping differentiation under mixed culture conditions reveals a tunable continuum of T cell fates. PLoS Biol. *11*, e1001616.

Balan, S., Arnold-Schrauf, C., Abbas, A., Couespel, N., Savoret, J., Imperatore, F., Villani, A.C., Vu Manh, T.P., Bhardwaj, N., and Dalod, M. (2018). Large-Scale Human Dendritic Cell Differentiation Revealing Notch-Dependent Lineage Bifurcation and Heterogeneity. Cell Rep. *24*, 1902–1915.e6.

Banchereau, J., and Steinman, R.M. (1998). Dendritic cells and the control of immunity. Nature *392*, 245–252.

Cariani, F., and Rips, L.J. (2017). Conditionals, Context, and the Suppression Effect. Cogn. Sci. (Hauppauge) *41*, 540–589.

Chen, L., and Flies, D.B. (2013). Molecular mechanisms of T cell co-stimulation and co-inhibition. Nat. Rev. Immunol. *13*, 227–242.

Ciofani, M., Madar, A., Galan, C., Sellars, M., Mace, K., Pauli, F., Agarwal, A., Huang, W., Parkhurst, C.N., Muratet, M., et al. (2012). A validated regulatory network for Th17 cell specification. Cell *151*, 289–303.

da Silva-Diz, V., Lorenzo-Sanz, L., Bernat-Peguera, A., Lopez-Cerda, M., and Muñoz, P. (2018). Cancer cell plasticity: Impact on tumor progression and therapy response. Semin. Cancer Biol. *53*, 48–58.

Edfors, F., Danielsson, F., Hallström, B.M., Käll, L., Lundberg, E., Pontén, F., Forsström, B., and Uhlén, M. (2016). Gene-specific correlation of RNA and protein levels in human cells and tissues. Mol. Syst. Biol. *12*, 883.

Guermonprez, P., Valladeau, J., Zitvogel, L., Théry, C., and Amigorena, S. (2002). Antigen presentation and T cell stimulation by dendritic cells. Annu. Rev. Immunol. *20*, 621–667.

Ito, T., Wang, Y.H., Duramad, O., Hori, T., Delespesse, G.J., Watanabe, N., Qin, F.X., Yao, Z., Cao, W., and Liu, Y.J. (2005). TSLP-activated dendritic cells induce an inflammatory T helper type 2 cell response through OX40 ligand. J. Exp. Med. *202*, 1213–1223.

Ivanov, I.I., McKenzie, B.S., Zhou, L., Tadokoro, C.E., Lepelley, A., Lafaille, J.J., Cua, D.J., and Littman, D.R. (2006). The orphan nuclear receptor RORgammat directs the differentiation program of proinflammatory IL-17+ T helper cells. Cell *126*, 1121–1133.

Keller, R. (2001). Dendritic cells: their significance in health and disease. Immunol. Lett. *78*, 113–122.

Kintsch, W., and Mangalath, P. (2011). The construction of meaning. Top. Cogn. Sci. *3*, 346–370.

Korn, T., Bettelli, E., Oukka, M., and Kuchroo, V.K. (2009). IL-17 and Th17 Cells. Annu. Rev. Immunol. *27*, 485–517.

Liu, Y.J., Kanzler, H., Soumelis, V., and Gilliet, M. (2001). Dendritic cell lineage, plasticity and cross-regulation. Nat. Immunol. *2*, 585–589.

Liu, Y., Beyer, A., and Aebersold, R. (2016). On the Dependency of Cellular Protein Levels on mRNA Abundance. Cell *165*, 535–550.

Macagno, A., Napolitani, G., Lanzavecchia, A., and Sallusto, F. (2007). Duration, combination and timing: the signal integration model of dendritic cell activation. Trends Immunol. *28*, 227–233.

Manel, N., Unutmaz, D., and Littman, D.R. (2008). The differentiation of human T(H)-17 cells requires transforming growth factor-beta and induction of the nuclear receptor RORgammat. Nat. Immunol. *9*, 641–649.

Meinshausen, N., and Bühlmann, P. (2010). Stability selection. J. R. Stat. Soc. Series B. Stat. Methodol. *72*, 417–473.

Nakayama, T., Hirahara, K., Onodera, A., Endo, Y., Hosokawa, H., Shinoda, K., Tumes, D.J., and Okamoto, Y. (2017). Th2 Cells in Health and Disease. Annu. Rev. Immunol. *35*, 53–84.

Naldi, A., Carneiro, J., Chaouiya, C., and Thieffry, D. (2010). Diversity and plasticity of Th cell types predicted from regulatory network modelling. PLoS Comput. Biol. *6*, e1000912.

Perrot-Dockès, M., Lévy-Leduc, C., Chiquet, J., Sansonnet, L., Brégère, M., Étienne, M.P., Robin, S., and Genta-Jouve, G. (2018a). A variable selection approach in the multivariate linear model: an application to LC-MS metabolomics data. Stat. Appl. Genet. Mol. Biol. *17*, /j/sagmb.2018.17.issue-5/sagmb-2017-0077/sagmb-2017-0077.xml.

Perrot-Dockès, M., Lévy-Leduc, C., Sansonnet, L., and Chiquet, J. (2018b). Variable selection in multivariate linear models with high-dimensional covariance matrix estimation. J. Multivariate Anal. *166*, 78–97.

Soumelis, V., Reche, P.A., Kanzler, H., Yuan, W., Edward, G., Homey, B., Gilliet, M., Ho, S., Antonenko, S., Lauerma, A., et al. (2002). Human epithelial cells trigger dendritic cell mediated allergic inflammation by producing TSLP. Nat. Immunol. *3*, 673–680.

Tibshirani, R. (1996). Regression Shrinkage and Selection via the Lasso. J. R. Stat. Soc. Series B Stat. Methodol. *58*, 267–288.

Tindemans, I., Peeters, M.J.W., and Hendriks, R.W. (2017). Notch Signaling in T Helper Cell Subsets: Instructor or Unbiased Amplifier? Front. Immunol. *8*, 419.

Touzot, M., Grandclaudon, M., Cappuccio, A., Satoh, T., Martinez-Cingolani, C., Servant, N., Manel, N., and Soumelis, V. (2014). Combinatorial flexibility of cytokine function during human T helper cell differentiation. Nat. Commun. *5*, 3987.

Vento-Tormo, R., Efremova, M., Botting, R.A., Turco, M.Y., Vento-Tormo, M., Meyer, K.B., Park, J.E., Stephenson, E., Polański, K., Goncalves, A., et al. (2018). Single-cell reconstruction of the early maternal-fetal interface in humans. Nature *563*, 347–353.

Volpe, E., Servant, N., Zollinger, R., Bogiatzi, S.I., Hupé, P., Barillot, E., and Soumelis, V. (2008). A critical function for transforming growth factor-beta, interleukin 23 and proinflammatory cytokines in driving and modulating human T(H)-17 responses. Nat. Immunol. *9*, 650–657.

Wang, Y., Chen, X., Cao, W., and Shi, Y. (2014). Plasticity of mesenchymal stem cells in immunomodulation: pathological and therapeutic implications. Nat. Immunol. *15*, 1009–1016.

Yosef, N., Shalek, A.K., Gaublomme, J.T., Jin, H., Lee, Y., Awasthi, A., Wu, C., Karwacz, K., Xiao, S., Jorgolli, M., et al. (2013). Dynamic regulatory network controlling TH17 cell differentiation. Nature *496*, 461–468.

Zhu, J., Yamane, H., and Paul, W.E. (2010). Differentiation of effector CD4 T cell populations (*). Annu. Rev. Immunol. *28*, 445–489.

Zygmunt, B., and Veldhoen, M. (2011). T helper cell differentiation more than just cytokines. Adv. Immunol. *109*, 159–196.

université PARIS-SACLAY | École doctorale de mathématiques Hadamard (EDMH)

Titre : Méthodes régularisées pour l'analyse de données multivariées en grande dimension: théorie et applications.

Mots clés : Méthodes régularisées, données multivariées, grande dimension

Résumé : Dans cette thèse nous nous intéressons au modèle linéaire général (modèle linéaire multivarié) en grande dimension. Nous proposons un nouvel estimateur parcimonieux des coefficients de ce modèle qui prend en compte la dépendance qui peut exister entre les différentes réponses. Cet estimateur est obtenu en estimant dans un premier temps la matrice de covariance des réponses puis en incluant cette matrice de covariance dans un critère Lasso.
Les propriétés théoriques de cet estimateur sont étudiées lorsque le nombre de réponses peut tendre vers l'infini plus vite que la taille de l'échantillon. Plus précisément, nous proposons des conditions générales que doivent satisfaire les estimateurs de la matrice de covariance et de son inverse pour obtenir la consistance en signe des coefficients.
Nous avons ensuite mis en place des méthodes, adaptées à la grande dimension, pour l'estimation de matrices de covariance qui sont supposées être des matrices de Toeplitz ou des matrices avec une structure par blocs, pas nécessairement diagonaux.
Ces différentes méthodes ont enfin été appliquées à des problématiques de métabolomique, de protéomique et d'immunologie.

Title : Regularized methods to study multivariate data in high dimensional settings: theory and applications.

Keywords : Regularized methods, multivariate data, high dimension

Abstract : In this PhD thesis we study general linear model (multivariate linear model) in high dimensional settings. We propose a novel variable selection approach in the framework of multivariate linear models taking into account the dependence that may exist between the responses. It consists in estimating beforehand the covariance matrix of the responses and to plug this estimator in a Lasso criterion, in order to obtain a sparse estimator of the coefficient matrix.
The properties of our approach are investigated both from a theoretical and a numerical point of view. More precisely, we give general conditions that the estimators of the covariance matrix and its inverse have to satisfy in order to recover the positions of the zero and non-zero entries of the coefficient matrix when the number of responses is not fixed and can tend to infinity.
We also propose novel, efficient and fully data-driven approaches for estimating Toeplitz and large block structured sparse covariance matrices in the case where the number of variables is much larger than the number of samples without limiting ourselves to block diagonal matrices.
These approaches are applied to different biological issues in metabolomics, in proteomics and in immunology.

Université Paris-Saclay
Espace Technologique / Immeuble Discovery
Route de l'Orme aux Merisiers RD 128 / 91190 Saint-Aubin, France