Data-Driven Sparse Partial Least Squares

Hadrien Lorenzo, Olivier Cloarec, Rodolphe Thiébaut, Jérôme Saracco

► To cite this version:

Hadrien Lorenzo, Olivier Cloarec, Rodolphe Thiébaut, Jérôme Saracco. Data-Driven Sparse Partial Least Squares. Statistical Analysis and Data Mining, 2021, 10.1002/sam.11558 . hal-03368956

HAL Id: hal-03368956

<https://hal.inria.fr/hal-03368956>

Submitted on 7 Oct 2021

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.

Data-Driven Sparse Partial Least Squares

Hadrien Lorenzo*, ASTRAL, Inria BSO, Talence, France

Olivier Cloarec, Corporate Research Advanced Data Analytics, Sartorius, Aubagne, France

Rodolphe Thiébaut, SISTM, Inria BSO/INSERM - U1219 BPH, Bordeaux cedex, France
Vaccine Research Institute, Créteil cedex, France

Jérôme Saracco, ASTRAL, Inria BSO, Talence, France

October 7, 2021

Abstract

In the supervised high dimensional settings with a large number of variables and a low number of individuals, variable selection allows a simpler interpretation and more reliable predictions. That subspace selection is often managed with supervised tools when the real question is motivated by variable prediction. We propose a Partial Least Square (PLS) based method, called data-driven sparse PLS (ddsPLS), allowing variable selection both in the covariate and the response parts using a single hyper-parameter per component. The subspace estimation is also performed by tuning a number of underlying parameters. The ddsPLS method is compared to existing methods such as classical PLS and two well established sparse PLS methods through numerical simulations. The observed results are promising both in terms of variable selection and prediction performance. This methodology is based on new prediction quality descriptors associated with the classical R^2 and Q^2 and uses bootstrap sampling to tune parameters and select an optimal regression model.

Keywords PLS regression, Supervised learning, Variable selection, Soft thresholding, Multi-block data

1 Introduction

While recent biotechnologies offer more and more descriptors (denoted as p) for monitoring biological processes, the number of observations (denoted as n) does not increase, which implies the ill-posed problem $n \ll p$ in supervised settings (such as regression framework). Different methods have been developed in the past century to tackle this high dimensional problem and, Partial Least Squares (PLS), particularly gathers attention since its introduction by Wold (1966). Indeed, this methodology has firstly become a standard in chemometrics calibration Martens et al. (1992) and is now used in biology Zoong Lwe et al. (2020), animal genetics Kautt et al. (2020), in human genetics Rechten et al. (2017) and in many other areas. However, to keep the ease of regression models' interpretation as p grows further, variable selection has also started to be needed by PLS users.

Following the L_1 -penalization of the regression coefficients introduced by Tibshirani (1996), models that handle variable selection solutions (denoted as *sparse* methodologies in the following) and produce model regularization allowing to keep confidence in the built models have been introduced by Lê Cao et al. (2008); Chun and Keleş (2010). Even if Chun and Keleş (2010) proved that PLS consistency is lost as n shrinks and p grows, the PLS and sparse PLS methods have allowed authors to analyze real datasets meaningfully. The PLS estimation goes through $R \in \mathbb{N}^*$ identical processes which can be generally described as follows.

- (a) Estimate the covariance matrix between covariate ($\mathbf{x}$) and response ($\mathbf{y}$) parts.
- (b) Estimate the singular-space associated with the largest singular-value of the previous matrix.
- (c) Project the covariate and response parts on the previously defined subspace.
- (d) Estimate the linear regression matrix of the response part on the covariate part in the subspace.
- (e) Remove the information carried by the current subspace from the covariate and response parts.

The step (a) corresponds to the estimation of a covariance matrix between the response and the covariate parts. The chosen estimator is very often $\mathbf{M} = 1/n \sum_{i=1}^n (\mathbf{y}_i - \bar{\mathbf{y}})(\mathbf{x}_i - \bar{\mathbf{x}})'$, for a data-set $(x_i, y_i)_{i=1..n}$ of size n with empirical means $\bar{\mathbf{x}}$ and $\bar{\mathbf{y}}$. This estimator is very sensitive to the curse of dimensionality, when the sample size is small and/or the number of variables is large, see for example Bickel and Levina (2008). However, the works previously mentioned only apply regularization methods in step (b), trying to overcome the failure to estimate the covariance matrix.

*Corresponding author, hadrien.lorenzo@inria.fr

Therefore, the present work focuses on a regularized estimation of the covariance matrix directly in step (a). More precisely, the considered solution and its associated optimization problem are

$$S_\lambda(\mathbf{M}) = \arg \min_{\mathbf{\Sigma} \in \mathbb{R}^{q \times p}} \|\mathbf{M} - \mathbf{\Sigma}\|^2 + 2\lambda \|\mathbf{\Sigma}\|, \quad (1)$$

where the soft-thresholding operator, an element-wise matrix operator, is defined as

$$\forall \mathbf{X} \in \mathbb{R}^{p \times q}, S_\lambda(\mathbf{X}) = (S_\lambda(x_{i,j}))_{(i,j) \in \llbracket 1, p \rrbracket \times \llbracket 1, q \rrbracket} = (\text{sign}(x_{i,j}) \max(0, |x_{i,j}| - \lambda))_{(i,j) \in \llbracket 1, p \rrbracket \times \llbracket 1, q \rrbracket},$$

$\|\bullet\|$ and $|\bullet|$ are respectively the Frobenius and the L_1 norms. Note that if $\lambda = 0$, the chosen estimator is the empirical estimator of the variance-covariance matrix.

So far, no theoretical development has been conducted to study the behavior of this estimator but for the case of a variance matrix, that is when $\mathbf{x} = \mathbf{y}$. Indeed, this case has been widely analyzed for the soft-thresholding operator but also for different thresholding operators. Johnstone and Lu (2009) have explored the pertinence of thresholding the diagonal of the empirical covariance matrices in the objective of building sparse versions of the Principal Components Analysis (PCA). Bickel and Levina (2008) have put forward theoretical and numerical arguments to various types of thresholding operators. Also, Deshpande and Montanari (2016) have shown the interest of such an idea where the whole empirical covariance matrix is thresholded, and the diagonal is removed.

In the PLS context, the number of components must be tuned, which is generally performed using prediction error. For sparse models, other parameters must also be adjusted by cross-validation and, more specifically, the leave-one-out method in the context of PLS. Recently, a sparse PLS approach (denoted as sgPLS for sparse group PLS and introduced in Sutton et al. (2018); Lique et al. (2016)) has used bootstrap approach in Broc et al. (2021) to efficiently overcome the difficulty of the “large p , small n ” multi-block framework to select optimal models.

Recently, a meta-analysis of sparse PLS methodologies was performed by Mehmood et al. (2020) over 16 different types of strategies divided in 3 classes which are “filter”, “wrapper” and “embedded” methods. The first class “filter” corresponds to a filter which is applied on the output of a PLS model (on the regression coefficients for example) and allows to classify variables according to their importance. The second class “wrapper” considers methodologies where the result of a “filter” is used to refit a PLS model. This operation is produced a certain amount of times. The most classical methods in this context are the backward (and also forward even if this has not been studied in the PLS case yet) variable selection methodologies. It also takes into account randomized methodologies regarding to the studied variables or sub samples. This list is non exhaustive and we encourage the reader to refer to the section “3.2 Wrapper methods” of Mehmood et al. (2020) to appreciate the richness of the solutions devised. The third class “embedded” corresponds to context where variable selection is performed in the PLS model building.

In the following, the proposed sparse PLS model, based on empirical covariance thresholding, soft-thresholding more specifically, is studied. The thresholds and the number of components are tuned thanks to bootstrap operations. Following the classification of sparse PLS methods described in Mehmood et al. (2020), the proposed methodology is close to the third class “embedded”. Yet this methodology modifies the PLS algorithm working on the covariance matrix before the PLS model is actually built, this closeness to the data gave its name: *data-driven sparse PLS* (ddsPLS). This ddsPLS methodology opens a new class of sparse PLS methods.

In Section 2, the underlying statistical model associated with latent variables is introduced. Section 3 describes the PLS (NIPALS-PLS/PLS2) algorithm and details the existing sparse PLS solutions. Section 4 introduces the proposed ddsPLS solution and the associated model selection algorithm. Based on numerical studies, Section 5 reports the quality of the proposed methodology and compares it to the above-presented methods through simulation designs. Finally, Section 6 provides some concluding remarks.

2 A latent variable model

Let $\mathbf{x}$ (resp. $\mathbf{y}$) be a p -dimensional (resp. q -dimensional) random variable. Let assume that the variables of $\mathbf{x}$ and of $\mathbf{y}$ are centered and of unit variance. Let $n \in \mathbb{N}^*$ be the sample size. Let $\mathbf{X}$ and $\mathbf{Y}$ be $n \times p$ and $n \times q$ random matrices where, for each row $i \in \llbracket 1, n \rrbracket$, $\mathbf{x}_i \sim \mathbf{x}$ and $\mathbf{y}_i \sim \mathbf{y}$. Let us further assume that the $\mathbf{x}_i$ ’s, respectively the $\mathbf{y}_i$ ’s, are independent to each other. Different methods such as PCA and PLS are based on the principle of generative latent vector common to $\mathbf{x}$ and $\mathbf{y}$, denoted by ϕ hereafter, a $\mathcal{R}$ -dimensional random variable. The corresponding latent variable model can be written as

$$\begin{cases} \mathbf{x} = \mathbf{A}'\phi + \epsilon \text{ where } \mathbf{A} = [\mathbf{a}_1, \dots, \mathbf{a}_{\mathcal{R}}]' \in \mathcal{M}_{\mathcal{R} \times p}(\mathbb{R}) \text{ with } \|\mathbf{a}_r\| \neq 0, \\ \mathbf{y} = \mathbf{D}'\phi + \xi \text{ where } \mathbf{D} = [\mathbf{d}_1, \dots, \mathbf{d}_{\mathcal{R}}]' \in \mathcal{M}_{\mathcal{R} \times q}(\mathbb{R}) \text{ with } \|\mathbf{d}_r\| \neq 0, \\ \text{with } \text{var}(\phi) = \mathbb{I}_{\mathcal{R}}, \mathbb{E}\phi = \mathbf{0}_{\mathcal{R}}, \mathbb{E}\epsilon = \mathbf{0}_p, \mathbb{E}\xi = \mathbf{0}_q, \\ \text{cov}(\phi, \epsilon) = \mathbf{0}_{\mathcal{R} \times p}, \text{cov}(\phi, \xi) = \mathbf{0}_{\mathcal{R} \times q}, \text{cov}(\epsilon, \xi) = \mathbf{0}_{p \times q}. \end{cases} \quad (2)$$

This model implies that $\text{var}(\mathbf{x}) = \mathbf{A}'\mathbf{A} + \text{var}(\epsilon)$ and $\text{var}(\mathbf{y}) = \mathbf{D}'\mathbf{D} + \text{var}(\xi)$. Furthermore, if the components of the random errors ϵ and ξ are independent and share the same variance, denoted by σ^2 , it comes $\text{var}(\mathbf{x}) = \mathbf{A}'\mathbf{A} + \sigma^2\mathbb{I}$ and $\text{var}(\mathbf{y}) = \mathbf{D}'\mathbf{D} + \sigma^2\mathbb{I}$. In all cases, the relation holds $\text{cov}(\mathbf{y}, \mathbf{x}) = \mathbf{D}'\mathbf{A}$.

Remark 1 *Helland and Almøy (1994) have studied the number R of latent variables shared by $\mathbf{y}$ and $\mathbf{x}$. They specified in which measure R is different from $\mathcal{R}$ in the general case. They have shown that the dimension R is equal to the largest number of eigen-vectors of $\mathbf{A}'\mathbf{A} = \text{var}(\mathbf{x})$, which have non-null projections on $\mathbf{A}'\mathbf{D} = \text{cov}(\mathbf{x}, \mathbf{y})$. In the PLS framework, R is the dimension of interest and is retrieved in the methodology as the estimated number of components.*

The underlying latent variable model (2) intuitively implies the following regression model

$$\mathbf{y} = \mathbf{B}'\mathbf{x} + \mathbf{e}, \quad (3)$$

where $\mathbf{B} \in \mathbb{R}^{p \times q}$ is a non-stochastic matrix, $\mathbf{e} \in \mathbb{R}^q$ is a centered residual q -dimensional random vector, and $\mathbf{e}$ and $\mathbf{x}$ are independent. Matrix $\mathbf{B}$ must satisfy

$$(\mathbf{AB})'\phi = \mathbf{D}'\phi \xRightarrow{\mathbb{E} \cdot \phi'} \mathbf{AB} = \mathbf{D}. \quad (4)$$

In the general case, without any additional hypothesis, $\mathbf{B}$ is not unique, but for the trivial cases or when $\mathbf{AA}'$ is invertible, the ordinary least-squares estimator of $\mathbf{B}$ only uses Equation (3), minimizing the l_2 -norm of $\mathbf{e}$, and provides $\hat{\mathbf{B}} = \mathbf{X}^+\mathbf{Y}$, which is known to be unstable in the case of multicollinearity in the $\mathbf{X}$ matrix. Let us notice that the latent variable model (2) is also not identifiable. Indeed, for any matrix $\mathbf{G} \in \mathcal{GL}_{\mathcal{R}}(\mathbb{R})$, denoting by $\mathbf{t} = \mathbf{G}'\phi$, $\mathbf{P} = \mathbf{G}^{-1}\mathbf{A}$ and $\mathbf{C} = \mathbf{G}^{-1}\mathbf{D}$, model (2) implies the following latent decomposition used by most of the authors.

$$\begin{cases} \mathbf{x} = \mathbf{P}'\mathbf{t} + \epsilon, \\ \mathbf{y} = \mathbf{C}'\mathbf{t} + \xi, \\ \text{with } \text{var}(\mathbf{t}', \epsilon', \xi')' \text{ diagonal,} \end{cases} \quad (5)$$

where $\mathbf{t}$ is called the score and $\mathbf{P}$ and $\mathbf{C}$ are called the loadings. Since the score and loadings are not identifiable, the PLS method concentrates on building a subspace that covers the space drawn by ϕ through an unknown matrix $\mathbf{G}$. Note that, only when numerical simulations are carried out, Equation (4) allows to screen the quality of structural reconstruction of the current estimator $\hat{\mathbf{B}}$ of $\mathbf{B}$ thanks to the metrics (for example)

$$\frac{\|\hat{\mathbf{AB}} - \mathbf{D}\|^2}{\|\mathbf{D}\|^2}, \quad (6)$$

which is as close to 0 as the estimated model is close to the theoretical one.

The PLS algorithm is a multi-linear regression method able to estimate regression matrices even in the context of degenerate datasets, i.e. when $\mathbf{X}'\mathbf{X}$ is singular. The principle of this algorithm is to iteratively estimate, here for the r^{th} -iteration, the first right and left singular vectors, $\mathbf{u}_r$ and $\mathbf{v}_r$, of the empirical covariance matrix $\mathbf{M}^{(r)} = \mathbf{Y}^{(r)'}\mathbf{X}^{(r)}/(n-1)$ where $\mathbf{Y}^{(r)}$ and $\mathbf{X}^{(r)}$ are the residual matrices of the previous steps, see technical details in next section.

3 NIPALS-PLS and sparse PLS solutions

This section recalls three different PLS algorithms. First the original PLS-Nipals algorithm, denoted by NIPALS-PLS Wold (1966) hereafter, is presented. Then the sparse PLS algorithm using L_1 -weights regularization, denoted by sPLS Lê Cao et al. (2008) hereafter, is detailed. Finally, a modification of the previous sparse solution to control the non-convex part of the optimization problem, denoted by SPLS Chun and Keleş (2010) hereafter, is provided.

3.1 NIPALS-PLS (PLS2) solution

The NIPALS-PLS Wold (1966) algorithm, often referred to as PLS2 algorithm in the literature, solves the optimization problem given at the first row of Table 1 The r -th iteration of this algorithm is described in (7): using the initializations $\mathbf{X}^{(1)} = \mathbf{X}$ and $\mathbf{Y}^{(1)} = \mathbf{Y}$,

$$\forall r \in \llbracket 1, R \rrbracket \quad \left\{ \begin{array}{l} \text{(a)} \quad \mathbf{u}_r = \overrightarrow{\text{RSV}}(\mathbf{M}^{(r)}), \quad \mathbf{v}_r = \overrightarrow{\text{RSV}}(\mathbf{M}^{(r)'}), \\ \text{(b)} \quad \mathbf{t}_r = \mathbf{X}^{(r)}\mathbf{u}_r, \\ \text{(c)} \quad \mathbf{p}_r = \mathbf{X}^{(r)'}\mathbf{t}_r/\mathbf{t}_r'\mathbf{t}_r, \\ \text{(d)} \quad \mathbf{c}_r = \mathbf{Y}^{(r)'}\mathbf{t}_r/\mathbf{t}_r'\mathbf{t}_r, \\ \text{(e)} \quad \mathbf{X}^{(r+1)} = \mathbf{X}^{(r)} - \mathbf{t}_r\mathbf{p}_r', \quad \mathbf{Y}^{(r+1)} = \mathbf{Y}^{(r)} - \mathbf{t}_r\mathbf{c}_r', \end{array} \right. \quad (7)$$

where $\overrightarrow{\text{RSV}}$ (for Right-Singular-Vector) is the function that returns the first right singular vector of its argument. In this algorithm, step (a) estimates the weights, step (b) estimates the scores, steps (c) and (d) estimate the linear regression matrices of $\mathbf{X}^{(r)}$ on $\mathbf{t}_r$ and $\mathbf{Y}^{(r)}$ on $\mathbf{t}_r$, and finally step (e) applies deflation. The regression matrix estimation is performed using the relation

$$\hat{\mathbf{B}} = \mathbf{U}(\mathbf{P}'\mathbf{U})^{-1}\mathbf{C}' \quad (8)$$

where $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_R]$, $\mathbf{P} = [\mathbf{p}_1, \dots, \mathbf{p}_R]$ and $\mathbf{C} = [\mathbf{c}_1, \dots, \mathbf{c}_R]$. If one wants to build the scores from an $\mathbf{x}$ -matrix, potentially different from the one used to build the model, it must not use the matrix $\mathbf{U}$ directly but the matrix $\mathbf{U}(\mathbf{P}'\mathbf{U})^{-1}$: $\mathbf{XU}(\mathbf{P}'\mathbf{U})^{-1} = (\mathbf{t}_1, \dots, \mathbf{t}_R) = (\mathbf{X}^{(1)}\mathbf{u}_1, \dots, \mathbf{X}^{(R)}\mathbf{u}_R)$.

Remark 2 In its original form, the PLS2/NIPALS-algorithm details the step (a) alternating estimations of $\mathbf{u}_r$ and $\mathbf{v}_r$ until convergence of $\mathbf{u}_r$ such as

$$\begin{aligned} \mathbf{u}_r &= \mathbf{X}^{(r)'} \mathbf{s}_r / \|\mathbf{X}^{(r)'} \mathbf{s}_r\| \\ \mathbf{t}_r &= \mathbf{X}^{(r)} \mathbf{u}_r / (\mathbf{u}_r' \mathbf{u}_r) \\ \mathbf{c}_r &= \mathbf{Y}^{(r)'} \mathbf{t}_r / (\mathbf{t}_r' \mathbf{t}_r) \\ \mathbf{s}_r &= \mathbf{Y}^{(r)} \mathbf{c}_r / (\mathbf{c}_r' \mathbf{c}_r) \end{aligned},$$

where $\mathbf{s}_r$ is initialised as the first column of $\mathbf{Y}^{(r)}$, see Tenenhaus (1998) for example. This alternating algorithm and the $\overrightarrow{RSV}$ operators both build $\mathbf{u}_r$ and the form associated to the operator will be kept in the following due to its compact aspect.

3.2 Two sparse PLS solutions

Variable selection can be performed during the computation of the components used to construct the PLS models. Based on L_1 -penalization, two sparse PLS solutions have been developed, denoted sPLS Lê Cao et al. (2008) and SPLS Chun and Keleş (2010) hereafter. More precisely, the NIPALS-PLS maximization criterion is then modified such as detailed in the second and third rows, respectively, of Table 1

sPLS Lê Cao et al. (2008) introduces sparsity both in the $\mathbf{X}$ and in the $\mathbf{y}$ parts through $2R+1$ parameters. The underlying idea is to construct a solution using the soft threshold operator $S_\lambda(\cdot)$. The r -th step of the associated algorithm is detailed in (9): using the initializations $\mathbf{X}^{(1)} = \mathbf{X}$ and $\mathbf{Y}^{(1)} = \mathbf{Y}$,

$$\forall r \in \llbracket 1, R \rrbracket \left\{ \begin{array}{l} \text{(a)}^\dagger \text{ Initialize } \mathbf{v}_r \text{ normalized, iterate until convergence of the two steps :} \\ \quad \text{(i)}^\dagger \quad \mathbf{u}_r = S_{\lambda_u^{(r)}} \left(\mathbf{M}^{(r)'} \mathbf{v}_r \right), \quad \mathbf{u}_r = \mathbf{u}_r / \|\mathbf{u}_r\|, \\ \quad \text{(ii)}^\dagger \quad \mathbf{v}_r = S_{\lambda_v^{(r)}} \left(\mathbf{M}^{(r)} \mathbf{u}_r \right), \quad \mathbf{v}_r = \mathbf{v}_r / \|\mathbf{v}_r\|, \\ \text{(b)} \quad \mathbf{t}_r = \mathbf{X}^{(r)} \mathbf{u}_r, \\ \text{(c)} \quad \mathbf{p}_r = \mathbf{X}^{(r)'} \mathbf{t}_r / \mathbf{t}_r' \mathbf{t}_r, \\ \text{(d)} \quad \mathbf{c}_r = \mathbf{Y}^{(r)'} \mathbf{t}_r / \mathbf{t}_r' \mathbf{t}_r, \\ \text{(e)} \quad \mathbf{X}^{(r+1)} = \mathbf{X}^{(r)} - \mathbf{t}_r \mathbf{p}_r', \quad \mathbf{Y}^{(r+1)} = \mathbf{Y}^{(r)} - \mathbf{t}_r \mathbf{c}_r', \end{array} \right. \quad (9)$$

where the regression matrix estimation is performed using the relation (8). One might notice that $\lambda_v^{(r)}$ fixes the sparsity of $\mathbf{v}_r$. However, since Equation (e) in Algorithm (7) does not explicitly impose sparsity in the $\mathbf{y}$ part, then the estimated matrix $\hat{\mathbf{B}}$ is not constrained to be sparse in $\mathbf{Y}$. In the implementation, the search for parameters $(\lambda_u^{(r)}, \lambda_v^{(r)})$ is driven, in an equivalent way, by two new parameters $(keep_X, keep_Y)_r$ which respectively correspond to the number of selected variables in both $\mathbf{x}$ and $\mathbf{y}$ parts for the r^{th} component.

SPLS Chun and Keleş (2010) concentrates on the $\mathbf{x}$ part sparsity and needs only 4 parameters to be tuned. Looking at the optimization problem of SPLS Chun and Keleş (2010), the sparse solution $\mathbf{c}$, relying on a L_1 -penalization via the parameter λ_1 , is not directly sought. Instead the algorithm defines $\mathbf{c}$ as a direction relatively close to $\mathbf{w}$, the “associated PLS-based solution”. This distance is measured according to the $\mathbf{S}^{X,(r)}$ -metric and the importance of each of these two terms in the overall optimization problem is tuned by a parameter $\kappa \in [0, 1]$. If κ is high (equal 1 at the limit), the problem yields to the PLS Wold (1966) solution, while a low κ means that $\mathbf{c}$ and $\mathbf{w}$ must be very close to each other but not necessarily to the classical PLS solution. According to Chun and Keleş (2010), λ_2 can be tuned to $+\infty$ as to preserve the soft-thresholding solution and $\kappa < 1/2$ avoids local solution issues. There are now only 3 parameters (R, κ, λ_1) left to be tuned. Algorithm (10) is the associated algorithm in the case of so-called NIPALS deflation, as described by the authors where NIPALS-PLS(.) is a function that takes a covariate and a response matrix for the two first arguments and a number of components in third and sends the NIPALS-PLS regression matrix such as described in Algorithm (7). In the following, the regression matrix is estimated as $\hat{\mathbf{B}} = \mathbf{B}^{(R)}$, also $\mathcal{A}$ is the index set of active variables, updated at each iteration, such as

$$\forall r \in \llbracket 1, R \rrbracket : \left\{ \begin{array}{l} \mathbf{Z}^{(r)} = \mathbf{X}' \mathbf{Y}^{(r)} \quad \text{and} \quad \mathbf{S}^{X,(r)} = \mathbf{Z}^{(r)} \mathbf{Z}^{(r)'} \\ \text{Initialize } \mathbf{c}^{(r)}, \text{ iterate until convergence of the two steps :} \\ \quad \mathbf{w}^{(r)} = \arg \min_{\mathbf{w}, \mathbf{w}' \mathbf{w} = 1} -\kappa \mathbf{w}' \mathbf{S}^{X,(r)} \mathbf{w} + (1 - \kappa) (\mathbf{c}^{(r)} - \mathbf{w})' \mathbf{S}^{X,(r)} (\mathbf{c}^{(r)} - \mathbf{w}), \\ \quad \mathbf{c}^{(r)} = \arg \min_{\mathbf{c}} \left(\mathbf{Z}^{(r)'} \mathbf{c} - \mathbf{Z}^{(r)'} \mathbf{w}^{(r)} \right)' \left(\mathbf{Z}^{(r)'} \mathbf{c} - \mathbf{Z}^{(r)'} \mathbf{w}^{(r)} \right) + \lambda_1 |\mathbf{c}|, \\ \mathcal{A} = \left\{ i \in \llbracket 1, p \rrbracket \mid \mathbf{w}_i^{(r)} \neq 0 \vee \mathbf{B}_{i,\cdot}^{(r-1)} = \mathbf{0} \right\}, \\ \mathbf{B}_{\mathcal{A}}^{(r)} = \text{NIPALS-PLS} \left(\mathbf{X}_{\mathcal{A}}, \mathbf{Y}^{(r)}, r \right) \text{ and } \mathbf{B}_{-\mathcal{A}}^{(r)} = \mathbf{0}, \\ \text{(e')} \quad \mathbf{Y}^{(r+1)} = \mathbf{Y}^{(r)} - \mathbf{X} \mathbf{B}^{(r)}. \end{array} \right. \quad (10)$$

The deflation in sPLS Lê Cao et al. (2008) is performed as in the classical NIPALS-PLS algorithm (see row (e) of Algorithm (7)), while only the $\mathbf{Y}$ part is deflated in the SPLS Chun and Keleş (2010) algorithm (using row (e') of Algorithm (10), instead of row (e) of Algorithm (7)).

Table 1: Core optimization problem of each PLS-based, where $\mathbf{M}^{(r)} = \mathbf{Y}^{(r)'}\mathbf{X}^{(r)}/n - 1$ and $\mathbf{S}^{X,(r)} = \mathbf{X}'\mathbf{Y}^{(r)}\mathbf{Y}^{(r)'}\mathbf{X}$.

Method	Optimization Problem	Constraints	Parameters
NIPALS-PLS Wold (1966)	$\max_{\mathbf{u}, \mathbf{v}} \mathbf{v}'\mathbf{M}^{(r)}\mathbf{u}$	$\mathbf{u}'\mathbf{u} = \mathbf{v}'\mathbf{v} = 1$	R
sPLS Lê Cao et al. (2008)	$\min_{\mathbf{u}, \mathbf{v}} \left\ (n-1)\mathbf{M}^{(r)} - \mathbf{v}\mathbf{u}' \right\ ^2 + \lambda_u^{(r)} \mathbf{u} + \lambda_v^{(r)} \mathbf{v} $	$\mathbf{u}'\mathbf{u} = \mathbf{v}'\mathbf{v} = 1$	$R, (\lambda_u^{(r)}, \lambda_v^{(r)})_r$
SPLS Chun and Keleş (2010)	$\min_{\mathbf{w}, \mathbf{c}} -\kappa \mathbf{w}\mathbf{S}^{X,(r)}\mathbf{w} + (1-\kappa)(\mathbf{c}-\mathbf{w})'\mathbf{S}^{X,(r)}(\mathbf{c}-\mathbf{w}) + \lambda_1 \mathbf{c} + \lambda_2\ \mathbf{c}\ $	$\mathbf{w}'\mathbf{w} = 1$	$R, \kappa, \lambda_1, \lambda_2$
ddsPLS	$\arg \max_{\mathbf{u}, \mathbf{v}} \mathbf{v}'S_{\lambda_r}(\mathbf{M}^{(r)'})\mathbf{u}$	$\mathbf{u}'\mathbf{u} = \mathbf{v}'\mathbf{v} = 1$	$R, (\lambda_r)_r$

The method proposed in the next section is a new sparse PLS methodology based on a soft-thresholding operation of the empirical covariance matrix, where the associated hyperparameters are ajusted via bootstrap.

4 Data-driven sparse PLS

This section details the ddsPLS method. First, Section 4.1 provides the underlying algorithm. Then Section 4.3 introduces the main quality descriptors generally used in PLS regression and their adaptations to the bootstrap context with slight modifications. Section 4.4 explains how to tune the different hyperparameters of the ddsPLS method, and an illustration based on an “easy” example is given in Section 4.6. Finally Section 4.2 provides the extension of the ddsPLS approach to the case of multi-block $\mathbf{x}$ datasets.

Remark 3 *Since the underling latent variable model (2) assumes centered and scaled $\mathbf{x}$ and $\mathbf{y}$ variables, the data are (marginally) standardized on the overall dataset before applying the ddsPLS methodology. So data are standardized by subtracting the empirical mean and dividing by the empirical standard deviation. This standardization step makes the ddsPLS methodology sensitive to potential extreme observations’ adverse effects and thus to outliers.*

4.1 ddsPLS algorithm

As noticed in Mehmood et al. (2020), sPLS Lê Cao et al. (2008) and SPLS Chun and Keleş (2010) are embedded sparse methodologies. But those methodologies are based on the penalization of the $\mathbf{x}$ parts from the eigendecomposition of $\mathbf{Y}'\mathbf{X}$ (and also of the $\mathbf{y}$ part for sPLS Lê Cao et al. (2008), but through a different regularization coefficient). The idea of the ddsPLS method is to soft-threshold directly the empirical covariance matrix right, using the $S_\lambda(\cdot)$ operator, before building the weights. This operation naturally tends to remove non-interesting variables since their associated soft-thresholded empirical covariance coefficients are equal to 0, for a threshold λ large enough, and the associated weight goes to 0. This regularization of the “data itself” (actually the empirical covariance matrix) and not of the associated weights (which are first eigenvectors of the empirical covariance matrices) led to the name “data-driven sparse PLS (ddsPLS)” to the proposed method. The corresponding optimization problem is provided in the last row of Table 1 One might notice that if $\lambda_r = 0$, it corresponds to the usual NIPALS-PLS Wold (1966) solution. The associated algorithm of the ddsPLS method is detailed below. Note that, to introduce sparsity in the response in the regression matrix $\hat{\mathbf{B}}$, the step (d) of the Algorithm (7) is modified such as

$$\mathbf{c}_r = \left(\mathbf{Y}^{(r)}\mathbf{\Pi}_r \right)' \mathbf{t}_r / \mathbf{t}_r' \mathbf{t}_r$$

where $\mathbf{\Pi}_r = \text{diag}(\{\delta_{\neq 0}(\mathbf{v}_r)_j\}_{j \in [1, q]})$ is sparse since “ $\mathbf{v}_r$ ” is sparse for λ large enough. Here also, if $\lambda = 0$, the result is equal to the solution offered by the NIPALS-PLS Wold (1966) solution.

The algorithm of ddsPLS is then

$$\forall r \in [1, R] \quad \left\{ \begin{array}{l} \text{(a)* } \mathbf{u}_r = \overrightarrow{\text{RSV}} \left(S_{\lambda^{(r)}} \left(\mathbf{M}^{(r)} \right) \right), \mathbf{v}_r = \overrightarrow{\text{RSV}} \left(S_{\lambda^{(r)}} \left(\mathbf{M}^{(r)'} \right) \right), \\ \text{(b) } \mathbf{t}_r = \mathbf{X}^{(r)}\mathbf{u}_r, \\ \text{(c) } \mathbf{p}_r = \mathbf{X}^{(r)'}\mathbf{t}_r / \mathbf{t}_r' \mathbf{t}_r, \\ \text{(d)* } \left\{ \begin{array}{l} \mathbf{\Pi}_r = \text{diag}(\{\delta_{\neq 0}(\mathbf{v}_r)_j\}_{j \in [1, q]}) \\ \mathbf{c}_r = \arg \min_{\mathbf{V}} \left\| \mathbf{Y}^{(r)}\mathbf{\Pi}_r - \mathbf{t}_r \mathbf{V}' \right\|^2 = \left(\mathbf{Y}^{(r)}\mathbf{\Pi}_r \right)' \mathbf{t}_r / (\mathbf{t}_r' \mathbf{t}_r). \end{array} \right. \\ \text{(e) } \mathbf{X}^{(r+1)} = \mathbf{X}^{(r)} - \mathbf{t}_r \mathbf{p}_r', \mathbf{Y}^{(r+1)} = \mathbf{Y}^{(r)} - \mathbf{t}_r \mathbf{c}_r', \end{array} \right. \quad (11)$$

where the regression matrix estimation is also performed using the relation (8).

Notice that $R + 1$ parameters need to be tuned to obtain an optimal solution: the number R of components and the associated regularization coefficient $\lambda^{(r)}$ for each component. Based on different well-known quality descriptors used in PLS and on new appealing ones introduced in our ddsPLS context (see next subsection), a way to calibrate these different parameters is provided in Section 4.4.

4.2 Adaptation to the multi-block structure

A multi-block generalization of the latent variable model (2) is

$$\forall t \in \llbracket 1, T \rrbracket, \quad \mathbf{x}_t = \mathbf{A}'_t \boldsymbol{\phi} + \boldsymbol{\epsilon}_t \quad \text{and} \quad \mathbf{y} = \mathbf{C}' \boldsymbol{\phi} + \boldsymbol{\xi}, \quad (12)$$

where T is the total number of blocks. The corresponding underlying regression model can be written as

$$\mathbf{y} = \sum_{t=1}^T \mathbf{B}'_t \mathbf{b}_t + \mathbf{e}. \quad (13)$$

Let's introduce the following notation:

$$\mathbf{x} = (\mathbf{x}'_1, \dots, \mathbf{x}'_T)' \quad , \quad \mathbf{A} = [\mathbf{A}_1, \dots, \mathbf{A}_T] \quad \text{and} \quad \mathbf{B} = [\mathbf{B}'_1, \dots, \mathbf{B}'_T]'. \quad (14)$$

Model (12) can be rewritten in the form of Model (2), thus it is possible to use the ddsPLS methodology as previously in this multi-block framework.

Remark 4 *The adaptation of the PLS method to the multi-block context has firstly been described in Wold (1984) and finally defined in Vivien (2002) to take the form*

$$\begin{aligned} \max_{\mathbf{w}_t, \mathbf{v}, c_t} \quad & \mathbf{v}' \mathbf{Y}^{(r)'} \sum_{t=1}^T \left(c_t \mathbf{X}_t^{(r)} \mathbf{w}_t \right), \\ \text{s.t.} \quad & \|\mathbf{w}_t\| = \|\mathbf{v}\| = 1 \\ \text{and} \quad & \sum_{t=1}^T c_t^2 = 1. \end{aligned} \quad (15)$$

If $\mathbf{u}$ is associated to a solution of any of the previous algorithms (except for SPLS Chun and Keleş (2010)) in steps (a), (i') or (a'), and $\mathbf{u}_t$ is the extracted matrix from $\mathbf{u}$ associated with block t , then by identification

$$\mathbf{w}_t = \mathbf{u}_t / \|\mathbf{u}_t\| \quad \text{and} \quad c_t = \|\mathbf{u}_t\|. \quad (16)$$

Remark 5 *For the deflation steps (see steps (e) and (e')), the "mbPLS" algorithm (for multi-block PLS) uses the super-score deflation where the super-score gathers information from the T different blocks: $\mathbf{t}_r = \sum_{t=1}^T c_t \mathbf{X}_t^{(r)} \mathbf{w}_t$. Hence, in the following, when NIPALS-PLS Wold (1966) or sPLS Lê Cao et al. (2008) are used in the multi-block framework, the super-score deflation is always used. Moreover, for the multi-block version of the ddsPLS algorithm, the super-score deflation strategy is also used.*

Note different authors consider that deflating each block on the super-score might mix block information and, therefore, might also drive to senseless conclusions. For example, Westerhuis and Smilde (2001) present specific designs where no deflation on the $\mathbf{X}$ blocks provides better prediction results. This corresponds to the solution chosen in the SPLS Chun and Keleş (2010) algorithm.

Remark 6 *When the response variable $\mathbf{y}$ has multi-blocks, the proposed methodology remains the same where the different blocks are concatenated, as done in this section if $\mathbf{x}$ has multi-blocks, see (14), and the ddsPLS methodology can again be applied on the resulting data-set.*

4.3 Model quality descriptors

The PLS methodology is known to provide the correct common subspaces between the $\mathbf{x}$ and the $\mathbf{y}$ parts if there is indeed information to catch. Whereas, if this cross structure is poor, the methodology sometimes over-fits focusing on the $\mathbf{x}$ part variance in particular when the dimension p of $\mathbf{x}$ is large, see for example Wakeling and Morris (1993); Cloarec (2014). Its side effect artificially increases the R^2 statistics, also denoted as explained variance or determination coefficient, which only describes the goodness of fit on the train dataset:

$$R^2 = 1 - \frac{\sum_{j=1}^q \sum_{i=1}^n (y_{i,j} - \hat{y}_{i,j})^2}{\sum_{j=1}^q \sum_{i=1}^n (y_{i,j} - \bar{y}_j)^2}, \quad (17)$$

where $\bar{y}_j$ is the empirical mean of the j^{th} -variable of the $\mathbf{y}$ part and $\hat{y}_{i,j}$ is the estimation of $y_{i,j}$ thanks to the current model. The numerator is often denoted as the RSS (Residual Sum of Squares) and the denominator as the TSS (Total Sum of Squares) and thus $R^2 = 1 - \text{RSS}/\text{TSS}$. This statistics can actually be interpreted as an estimator of

$$\gamma(\mathbf{p}) = 1 - \frac{\sum_{j=1}^q \text{var}(y_j - y_j^{(\mathbf{p})})}{\sum_{j=1}^q \text{var}(y_j)}, \quad (18)$$

built on the (potentially stochastic) prediction model $\mathbf{p}$ where $y_j^{(\mathbf{p})}$ is an estimator of y_j based on the prediction model $\mathbf{p}$. This metrics is as close to 1 as the prediction error is low. Its interest is to take into account the different scales of the variables y_j . As an interpretation, asymptotically (when $n \rightarrow +\infty$), if this metric is equal to 0, then the associated prediction model performs equivalently to the mean prediction model, and if it is below 0, it performs worse. In the oracle context (a prediction model which knows the true relations between the informational variables), the estimator of $\mathbf{y}$ by the oracle predictor model $\mathbf{p}^*$ is $\mathbf{y}^* = \mathbf{A}' \boldsymbol{\phi}$ according to Model (2), and the previous metrics can be written as

$$\gamma^* = 1 - \frac{\sum_{j=1}^q \text{var}(\epsilon_j)}{\sum_{j=1}^q \text{var}(y_j)}. \quad (19)$$

Remark 7 The two previous statistics are global in the sense of the q -dimensional response variable $\mathbf{y}$ and therefore do not allow for a marginal critique of the model. Another interesting statistics would be $\gamma_j(\mathbf{p}) = 1 - \text{var}(y_j - y_j^{(\mathbf{p})}) / \text{var}(y_j)$, for $j = 1 \dots q$, which exhibits the marginal effects of each variable on the response variable. However, the response variables have been standardized in order to be processed on the same scale (see Remark 3), which rationalizes the definition of the γ . Although this does not make it possible to identify marginally the variables of $\mathbf{y}$ that can potentially degrade the criterion, their deleterious effects remain observable. Moreover, in the case of marginal γ_j , it would have been necessary to define a transformation aggregating the q marginal criteria making it possible to define a decision algorithm on the conservation or not of the current component. The definition of such a transformation is not trivial and, therefore, will not be considered in this work. For all these reasons, the global statistics γ is used in the proposed methodology.

The R^2 estimator is subject to over-fitting, this is due to the fact that the quality of the model is assessed on the whole sample used to build the model.

To circumvent this over-fitting drawback, the criterion Q^2 has been introduced by Stone (1974); Geisser (1974). It is based on F -folds cross-validation such as

$$Q^2 = 1 - \frac{\sum_{j=1}^q \sum_{f=1}^F \sum_{i \notin \text{cv}_f} (y_{i,j} - \hat{y}_{i,j}^{(\text{cv}_f)})^2}{\sum_{j=1}^q \sum_{f=1}^F \sum_{i \notin \text{cv}_f} (y_{i,j} - \bar{y}_j)^2}, \quad (20)$$

where $\hat{y}_{i,j}^{(\text{cv}_f)}$ is the estimation of $y_{i,j}$ thanks to the current model where the observations in “cv $_f$ ” are used as the train sample for fold “ f ”. The numerator is often denoted as the PRESS (Prediction Error Sum of Squares) and thus $Q^2 = 1 - \text{PRESS}/\text{TSS}$.

Thanks to R^2 and Q^2 , different rules have been used to select “optimal” models. In the general prediction context, the most popular maximizes Q^2 , which is equivalent to minimize the Mean Squared Error in Prediction (MSEP) since the denominator of the right-hand side term of Q^2 does not use the model’s prediction (but the mean estimator). In the context of PLS, a rule of thumb is to favor models for which the difference between R^2 and Q^2 is minimum, which is supposed to correspond to a minimum of over-fitting, see for example Cloarec (2014).

Those metrics R^2 and Q^2 are helpful to evaluate the quality of a complete model. Other metrics have been introduced to evaluate the quality of a current r^{th} -component of the model and are denoted as R_r^2 and Q_r^2 , see for example Tenenhaus (1998). Let us suppose that the first $(r-1)$ components have already been built, then the quality of the r^{th} -component can be screened thanks to

$$R_r^2 = 1 - \frac{\sum_{j=1}^q \sum_{i=1}^n (y_{i,j} - (\hat{y}_{i,j}^{(r)} - \hat{y}_{i,j}^{(r-1)}) - \bar{y}_j)^2}{\sum_{j=1}^q \sum_{i=1}^n (y_{i,j} - \bar{y}_j)^2}, \quad Q_r^2 = 1 - \frac{\sum_{j=1}^q \sum_{f=1}^F \sum_{i \notin \text{cv}_f} (y_{i,j} - \hat{y}_{i,j}^{(\text{cv}_f, r)})^2}{\sum_{j=1}^q \sum_{f=1}^F \sum_{i \notin \text{cv}_f} (y_{i,j} - \hat{y}_{i,j}^{(\text{cv}_f, r-1)})^2}, \quad (21)$$

where $\hat{y}_{i,j}^{(r)}$, resp. $\hat{y}_{i,j}^{*,(r)}$, is the prediction of the i^{th} -observation and the j^{th} -response variable thanks to model based on components the first r and all observations, resp. all observations except the i^{th} one. By convention, $\hat{y}_{i,j}^{(0)} = \bar{y}_j$ and $\hat{y}_{i,j}^{*,(0)} = \frac{1}{n-1} \sum_{i' \neq i} y_{i',j}$. The metrics R_r^2 represents the proportion of variability which is explained by the new component r and not by the $(r-1)$ previous ones, while the Q_r^2 does the same but based on out-of-fold samples from the cross-validation fold.

Rules of thumb exist to validate the current component based on the value of Q_r^2 , for example $Q_r^2 \geq 0.0975$ following for instance Tenenhaus (2005); Lê Cao et al. (2008), or $Q_r^2 \geq 0.05$ according to Pérez-Enciso and Tenenhaus (2003); Tenenhaus (2004) or $Q_r^2 \geq 0$ (or ≥ 0.05 if $n \leq 100$) according to Thévenot (2016).

Based on these descriptors Q_r^2 , another metrics which describes the overall quality of a R -components model is the cumulative Q^2 criterion, denoted as $Q_{R,(cum)}^2$ hereafter, and defined as an aggregation of the R criteria $(Q_1^2, \dots, Q_R^2)$ such as

$$Q_{R,(cum)}^2 = 1 - \prod_{r=1}^R (1 - Q_r^2).$$

This metric $Q_{R,(cum)}^2$ is used in Tenenhaus (1998) with the rule that the current component R is retained if $Q_{R,(cum)}^2 \gg Q_{R-1,(cum)}^2$.

In the previous statistics Q_r^2 (and thus in the metric $Q_{R,(cum)}^2$), the performances of numerous prediction models are simultaneously mixed, over the F folds, in the numerator and in the denominator. One consequence of this is that model selection is driven by prediction models for which the errors are the most important. To circumvent this drawback, for F -folds cross-validation, one way is first to consider the F different “marginal” Q^2 -like statistics, denoted as

$$1 - \frac{\sum_{j=1}^q \sum_{i \notin \text{cv}_f} (y_{i,j} - \hat{y}_{i,j}^{(\text{cv}_f)})^2}{\sum_{j=1}^q \sum_{i \notin \text{cv}_f} (y_{i,j} - \bar{y}_j^{(\text{cv}_f)})^2}, \quad \text{for } f = 1, \dots, F, \quad (22)$$

and then to combine them a posteriori using the empirical mean. To the best of our knowledge, this approach has never been proposed in the PLS literature. However, when the sample size n is small, using F -folds cross-validation is not relevant and the use of bootstrap is interesting to enrich the underlying information of the available data. This is the strategy chosen for the continuation of this work.

Based on B bootstrap samples, the bootstrapped versions of R^2 and Q^2 are given by

$$\bar{R}_B^2 = \frac{1}{B} \sum_{b=1}^B R_b^2 \quad \text{and} \quad \bar{Q}_B^2 = \frac{1}{B} \sum_{b=1}^B Q_b^2 \quad (23)$$

with, for the current bootstrap sample b ,

$$R_b^2 = 1 - \frac{\sum_{j=1}^q \sum_{i \in \text{IN}(b)} (y_{i,j} - \hat{y}_{i,j}^b)^2}{\sum_{j=1}^q \sum_{i \in \text{IN}(b)} (y_{i,j} - \bar{y}_j^b)^2}, \quad Q_b^2 = 1 - \frac{\sum_{j=1}^q \sum_{i \in \text{OOB}(b)} (y_{i,j} - \hat{y}_{i,j}^b)^2}{\sum_{j=1}^q \sum_{i \in \text{OOB}(b)} (y_{i,j} - \bar{y}_j^b)^2}, \quad (24)$$

where $\text{IN}(b)$ and $\text{OOB}(b)$ are respectively the set of the in-bag observations and the set of the out-of-bag observations, $\hat{y}_{i,j}^b$ is the prediction of the i^{th} -observation and the j^{th} -response variable and $\bar{y}_j^b$ is the empirical mean of the j^{th} -response variable estimated on the bootstrap sample b .

In the same way, bootstrapped versions of R_r^2 and Q_r^2 are given by

$$\bar{R}_{B,r}^2 = \frac{1}{B} \sum_{b=1}^B R_{b,r}^2 \quad \text{and} \quad \bar{Q}_{B,r}^2 = \frac{1}{B} \sum_{b=1}^B Q_{b,r}^2 \quad (25)$$

where

$$R_{b,r}^2 = 1 - \frac{\sum_{j=1}^q \sum_{i \in \text{IN}(b)} \left(y_{i,j} - \left(\hat{y}_{i,j}^{b,(r)} - \hat{y}_{i,j}^{b,(r-1)} \right) - \bar{y}_j^b \right)^2}{\sum_{j=1}^q \sum_{i \in \text{IN}(b)} (y_{i,j} - \bar{y}_j^b)^2}, \quad Q_{b,r}^2 = 1 - \frac{\sum_{j=1}^q \sum_{i \in \text{OOB}(b)} \left(y_{i,j} - \hat{y}_{i,j}^{b,(r)} \right)^2}{\sum_{j=1}^q \sum_{i \in \text{OOB}(b)} \left(y_{i,j} - \hat{y}_{i,j}^{b,(r-1)} \right)^2}. \quad (26)$$

Remark 8 Those metrics, $\bar{R}_B^2$, $\bar{Q}_B^2$, $\bar{R}_{B,r}^2$ and $\bar{Q}_{B,r}^2$, share the same philosophies as those discussed for (17), (20) and (21), but they introduce two innovative ideas (except for the bootstrap analysis itself). Firstly, they are empirical means over all bootstrap replications. This allows not mixing bootstrap quadratic errors, which is a good feature since estimations can vary from one bootstrap sample to another. Secondly, the “ R^2 ”-spirit metrics are based directly on the b^{th} bootstrap replication but use the associated in-bag sample, while the “ Q^2 ”-spirit metrics use the associated out-of-bag sample. This implies that the estimators R_b^2 and $R_{b,r}^2$ are even more subject to over-fitting and lead the difference “ $\bar{R}_B^2 - \bar{Q}_B^2$ ” to be more discriminant and informative in the presence of over-fitting case.

The following section details how these metrics are used to adjust the hyperparameters of the ddsPLS method by following the spirit of their cross-validation versions.

4.4 Selection of the hyperparameters

The ddsPLS algorithm detailed in (11) and the other sparse PLS methods require hyperparameters to be adjusted. In the case of ddsPLS, the following regularization parameters, R and $\Lambda_R = \{\lambda_r, r = 1, \dots, R\}$, must be optimally tuned. Due to the variance constraints of the variables $\mathbf{x}$ and $\mathbf{y}$, we know that $\lambda_r \in [0, 1]$ where

- $\lambda_r = 0$, corresponds to the NIPALS-PLS solution,
- $\lambda_r = 1$, leads to no variable selected and thus to the prediction model equal to empirical mean estimation.

This optimal hyperparameter selection is based on the minimization of $R_r^2 - Q_r^2$ (that might correspond to a minimum in over-fitting for the r^{th} -component). To this end, the proposed ddsPLS strategy considers $\bar{R}_{B,r}^2 - \bar{Q}_{B,r}^2$, the bootstrapped version of the corresponding metrics defined in (23).

More precisely, the selected model must satisfy that, for a current model built on $(r-1)$ components with estimated values for regularization parameters $\hat{\Lambda}_{r-1} = \{\hat{\lambda}_1, \dots, \hat{\lambda}_{r-1}\}$, the r^{th} -component is acceptable if

1. the r^{th} -component performs better than the mean estimation (i.e. $\bar{Q}_{B,r}^2 > 0$),
2. the model built on r components performs better than the one built on $(r-1)$ components (i.e. $\bar{Q}_B^2$ is increasing with the number of components),
3. the bootstrapped explained variance is close to the bootstrapped cross-validated explained variance (i.e. $\bar{R}_B^2 - \bar{Q}_B^2$ is minimum on the set of hyperparameters).

This hyperparameter selection procedure is illustrated on a toy example in Section 4.6.

4.5 A theoretical lower bound for the regularisation coefficients

Among the sparse covariance estimators the solution developed by Cai and Liu (2011) uses a fully data-driven estimator for each element of the covariance matrix. Even if they consider the variance case (where $\mathbf{y} = \mathbf{x}$) and assume that random parts are normally distributed, they use the adaptive threshold

$$\lambda_{i,j} = \delta \sqrt{\frac{\hat{\theta}_{i,j} \log p}{n}}, \quad (27)$$

where δ can be whether taken equal to 2, whether empirically chosen through cross-validation and $\theta_{i,j} = \text{var}((\mathbf{x}_i - \mathbb{E}\mathbf{x}_i)(\mathbf{x}_j - \mathbb{E}\mathbf{x}_j))$, the variance of the covariance coefficient between $\mathbf{x}_i$ and $\mathbf{x}_j$.

As n decreases/ p increases, the probability for a covariance coefficient to be larger than a given value increases. The value of $\lambda_{i,j}$ estimates the variability of the covariance coefficient below which non null coefficient can be considered as null with a high probability.

Inspired from the previous work, let us introduce the following value

$$\begin{aligned} \lambda_r^{(0)} &= \frac{1}{pq} \sum_{j=1}^q \sum_{i=1}^p \sqrt{\frac{\theta_{(j,i)}^{(r)} \log(\max(p, q))}{n}} \\ \text{where } \theta_{(j,i)}^{(r)} &= \frac{1}{n} \sum_{k=1}^n \left(x_{(k,i)}^{(r)} y_{(k,j)}^{(r)} - m_{(j,i)}^{(r)} \right)^2, \end{aligned}$$

where $x_{(k,i)}^{(r)}$, $y_{(k,j)}^{(r)}$ and $m_{(j,i)}^{(r)}$ are the elements of the matrices $\mathbf{X}^{(r)}$, $\mathbf{Y}^{(r)}$ and $\mathbf{M}^{(r)}$ respectively at positions (k, i) , (k, j) and (j, i) respectively. This value $\lambda_r^{(0)}$ is the lowest accessible value by the algorithm detailed in the previous section. It prevents from building components in which a large number of variables would be selected. This lower bound for λ is implemented in the package.

4.6 A toy example

Consider the latent variable model (2) with the following “easy” data structure

$$\mathbf{A} = \sqrt{1 - \sigma^2} \begin{pmatrix} \mathbf{1}'_{50} & \mathbf{0}'_{950} \end{pmatrix}, \quad \mathbf{D} = \sqrt{1 - \sigma^2} \begin{pmatrix} 1 \end{pmatrix}, \quad (28)$$

where $1 - \sigma^2 = 0.9025$ and

$$\boldsymbol{\psi} = (\phi', \boldsymbol{\epsilon}'_{1 \dots 50} / \sigma, \boldsymbol{\epsilon}'_{51 \dots 1000}, \xi' / \sigma)' \sim \mathcal{N}(\mathbf{0}, \mathbb{I}_{1+1000+1}).$$

The dependent variable $\mathbf{y}$ is only associated with the first 50 variables of $\mathbf{x}$ while the last 950 variables of $\mathbf{x}$ are not linked to $\mathbf{y}$. According to Remark 1, the true value for the number of components is $R = 1$. Thus, in this simulation framework, one of the objectives of the ddsPLS method is to properly retain one component for which the first 50 variables of $\mathbf{x}$ are selected in the final model. Moreover the corresponding prediction squared error is expected around “ $\gamma = 1 - \sigma^2 = 0.9025$ ”, for a sample with a reasonable number n of observations.

Figure 1 gathers results of 3 simulations based on 3 sample sizes, respectively, $n = 50, 100, 200$. For the bootstrap-based selection criterion, $B = 50$ bootstrap samples are used. For each of these samples, the proposed algorithm keeps only one component. The 3 graphs at the top of Figure 1 provide the plot of $\hat{R}_{B,1}^2$ and $\hat{Q}_{B,1}^2$ versus λ , for each sample size. The “optimal” $\hat{\lambda}_1$ ’s (for which $\hat{R}_{B,1}^2 - \hat{Q}_{B,1}^2$ are minimum) are also located in these graphs. One might notice that selecting model considering only $\hat{Q}_{B,1}^2$ would lead to bad conclusions, where its maximum (see the colored star on the curve of $\hat{Q}_{B,1}^2$) is around $\lambda = 0.8$ (for $n = 50$ or 200) or 0.3 (for $n = 100$) which is an over-fitting region according to the difference $\hat{R}_{B,1}^2 - \hat{Q}_{B,1}^2$.

The graph at the bottom left of Figure 1 shows the number of selected variables for each value of λ and each of the 3 samples ($n = 50, 100$ and 200). These curves are decreasing until a first loss of the 950 non-informative variables of $\mathbf{x}$ which occurs naturally at lower λ values when n is large. Then one can observe a large plate, and finally, the curves become decreasing again, resulting in the loss of the 50 relevant variables of $\mathbf{x}$ in the model. Indeed, as n increases, the minimum correlation to remove the 950 non-informative variables of $\mathbf{x}$ is small, and the maximum correlation keeping the 50 informative variables is high. For the 3 selected models (with the optimal $\hat{\lambda}_1$ values), only the 50 relevant variables of $\mathbf{x}$ are selected. The right bottom graph of Figure 1 shows boxplots of the corresponding first 50 regression coefficients for the 3 sample sizes. As the variables share the same distribution and are independent, the regression coefficient associated with the first 50 variables of $\mathbf{x}$ should naturally be close to $1/50$, while the last 950 coefficients regression coefficient associated with the last 950 variables of $\mathbf{x}$ are equal to 0. One might notice from the three boxplots, the coefficients of the first 50 relevant variables get closer to this expected value $1/50$ as n increases.

In the following, let’s define the True Positive Rate (TPR) and the False Positive Rate (FPR) of variable (of $\mathbf{x}$) selection as $\text{TPR} = \text{TP}/P$ and $\text{FPR} = \text{FP}/N$ where TP is the number of selected variables among those that should be selected, P is the total number of variables that should be selected, FP is the number of selected variables among those that should not be selected, and N is the total number of variables that should not be selected. Naturally, the optimal target is to have $\text{TPR} = 1$ and $\text{FPR} = 0$. In this simulation

Figure 1: Simulation results for the toy example.

study, the TPR and FPR of variables of $\mathbf{x}$ selection are optimal for the 3 sample sizes. Note that the TPR and FPR are also optimal for the variable of $\mathbf{y}$ selection since the one-dimensional dependant variable is always retained in the final selected models.

Finally, according to Equation (6), the metrics $\|\mathbf{A}\hat{\mathbf{B}} - \mathbf{D}\|^2 / \|\mathbf{D}\|^2$ can be calculated in this simulation framework: its value is approximately equal to $0.002 \ll 1$ and so shows reliable results.

This “toy” example illustrated the principle of the selection criterion for the ddsPLS method. In Section 5, others simulations will lead to estimate its good numerical behavior in more complex contexts.

5 Simulation study

Different synthetic structures have been used to evaluate the numerical behavior of the proposed ddsPLS methodology. Since the multi-block framework can be reduced to a mono-block framework, only mono-block structure in the $\mathbf{x}$ part are considered in this simulation study via three specific designs.

The first design corresponds to two groups of variables in the $\mathbf{x}$ part, which are not correlated, and a 3-dimensional response variable $\mathbf{y}$. The second design simulates spectroscopic data and will be detailed below.

For each design, 100 replications will be considered.

In this simulation study, various methodologies are compared with the proposed “ddsPLS” approach. The different methodologies are based on what has already been proposed by other authors and adapted to the bootstrap versions of γ . More precisely,

- “sPLS 1”, “sPLS 2” and “sPLS 3” build sparse PLS models as detailed in Lê Cao et al. (2008). The additional notation “1” refers to the solution minimizing the difference $\bar{R}_{B,r}^2 - \bar{Q}_{B,r}^2$ for each component r . The notation “2” refers to the solution maximizing $\bar{Q}_{B,r}^2$ for each component r , which is the closest solution to what is done currently using this model, corresponding to minimizing the cross-validation error (here replaced with maximizing the out-of-bag normalized error). Finally the notation “3” refers to model which builds variables while there exists $(keep_X, keep_Y)$ such as $\bar{Q}_{R,(cum)}^2 > 0.0975$ (based on CV results) where the couple $(keep_X, keep_Y)_r$ is then choosing maximizing $\bar{Q}_{R,(cum)}^2$. This is what the most popular methodology for this method, see for example Tenenhaus (1998).
- “NIPALS-PLS” corresponds to classical PLS algorithm where the number of components is chosen minimizing $\bar{R}_{B,r}^2 - \bar{Q}_{B,r}^2$ for each component.
- “SPLS” builds sparse PLS models as detailed in Chun and Keleş (2010). The parameter κ is left to the value “ $\kappa = 0.5$ ” (knowing that $\kappa \in [0, 1]$) which provides a balanced compromise between sparsity (low κ) and prediction results close to the classical PLS problem (high κ). Using a grid of values corresponding to $\eta = \{0.1, 0.2, \dots, 0.9\}$ and $R = \{1, 2, \dots, 10\}$, the considered best model corresponds to the minimum of the mean squared prediction error.

For “SPLS” and “sPLS 3”, the maximum number of possible components is fixed to 10. Also, for those two methodologies, the metric Q^2 , based on the F -folds cross-validation, is used, with a number F of folds equal to $F = (25, 50, 100, 100, 100)$ for $n = (25, 50, 100, 200, 400)$ respectively. Note that, when $n = (25, 50, 100)$, leave-one-out cross-validation is performed.

For the other approaches (“**sPLS 1**”, “**sPLS 2**”, “**NIPALS-PLS**” and “**ddsPLS**”), Bootstrap versions of Q^2 or R^2 metrics are used with a number B of bootstrap samples adapted to the sample size n : more specifically $B = (1000, 500, 300, 100, 100)$ for $n = (25, 50, 100, 200, 400)$ respectively.

5.1 Design 1: simulated data

The first design considers

$$\mathbf{A} = \sqrt{1 - \sigma^2} \begin{pmatrix} \alpha_3 \mathbf{1}_{3,50} & \mathbf{0}_{3,50} & \mathbf{0}_{3,900} \\ \mathbf{0}_{2,50} & \alpha_2 \mathbf{1}_{2,50} & \mathbf{0}_{2,900} \end{pmatrix}_{(5,1000)} \quad \text{and} \quad \mathbf{D} = \sqrt{1 - \sigma^2} \begin{pmatrix} \alpha_3 \mathbf{1}_{3,1} & \mathbf{0}_{3,1} & \mathbf{0}_{3,1} \\ \mathbf{0}_{2,1} & \alpha_2 \mathbf{1}_{2,1} & \mathbf{0}_{2,1} \end{pmatrix}_{(5,3)} \quad (29)$$

where $\forall k \in \mathbb{N}^*$, $\alpha_k = 1/\sqrt{k}$, $\sqrt{1 - \sigma^2} = 0.99$ and

$$\boldsymbol{\psi} = (\boldsymbol{\phi}', \boldsymbol{\epsilon}'_{1 \dots 100}/\sigma, \boldsymbol{\epsilon}'_{101 \dots 1000}, \boldsymbol{\xi}'_{1 \dots 2}/\sigma, \xi_3)' \sim \mathcal{N}(\mathbf{0}, \mathbb{I}_{5+1000+3}).$$

The corresponding latent variable model is then:

$$x_j = \begin{cases} \sqrt{1 - \sigma^2}(\phi_1 + \phi_2 + \phi_3)/\sqrt{3} + \epsilon_j & \text{for } j = 1 \dots 50 \\ \sqrt{1 - \sigma^2}(\phi_4 + \phi_5)/\sqrt{2} + \epsilon_j & \text{for } j = 51 \dots 100 \\ \epsilon_j & \text{for } j = 101 \dots 1000 \end{cases} \quad \text{and} \quad \begin{cases} y_1 = \sqrt{1 - \sigma^2}(\phi_1 + \phi_2 + \phi_3)/\sqrt{3} + \xi_1 \\ y_2 = \sqrt{1 - \sigma^2}(\phi_4 + \phi_5)/\sqrt{2} + \xi_2 \\ y_3 = \xi_3 \end{cases} \quad (30)$$

An analysis of matrices $\mathbf{A}$ and $\mathbf{D}$ described in (29) shows that $\mathcal{R} = 5$, but there are multi-collinearities in $\mathbf{A}$ and the real objective for a relevant number of components might be $R = 2$, which is proved according to Helland and Almøy (1994), as noticed in Remark 1. Also, only 100 variables over the 1000 variables of the variable $\mathbf{x}$ should be selected, moreover 2 out of the 3 variables of the dependent variable $\mathbf{y}$ should be selected. The theoretical γ is equal to $\gamma^* = 2(1 - \sigma^2)/3 = 0.6534$ according to (19).

According to Figure 2 and for large n ($=200$ or 400), all the methodologies almost always build the correct number of components, with the exception of the “**SPLS**” approach. For low n ($=25, 50$ or 100), the “**NIPALS-PLS**” and “**sPLS 3**” methodologies mostly build far too many components, which is also the case for the “**sPLS 2**” and “**SPLS**” methodologies but in a mitigate manner with median values equal to 3. The “**ddsPLS**” methodology and “**sPLS 1**” almost always build the correct number of components with a clear advantage for the first-mentioned method.

According to Figure 3 when $n = 400$ (large sample size), all the methodologies correctly provide “ Q^2 ” values close to γ^* . For all the cases, “**ddsPLS**” underestimates the value of Q^2 . When $n = 25$, the bootstrap based versions of Q^2 seem to under-estimate γ^* while the “**sPLS 3**” over-estimates it and “**SPLS**” shows high variability around this value. For “**NIPALS-PLS**”, “**sPLS 3**” and “**SPLS**” methodologies, the bad performances are due to over-fitting because of the high values retained for R for these approaches. However, the proposed Q^2 metric, $\bar{Q}_B^2$, used for “**NIPALS-PLS**” (and for “**sPLS 1**” and “**sPLS 2**”), detects over-fitting providing low Q^2 values, while “**sPLS 3**” provides high values of Q^2 which are often larger than the true expected value, this being essentially due to how Q^2 is calculated and not representative of the actual quality of the model. Finally, the “**sPLS 1**” and “**sPLS 2**” methodologies show slightly better results than “**ddsPLS**” in that low n context. Even if this metric seems worse for “**ddsPLS**”, this might be subject to over-fitting and other descriptors (such as $\|\mathbf{A}\hat{\mathbf{B}} - \mathbf{D}\|^2/\|\mathbf{D}\|^2$, or a prediction error on an independent data-set) should be used. This results have been observed and are accessible on the right part of the Figure 3 and in the Figure 4 Those figures are commented below.

In Figure 3 from the boxplots of $\|\mathbf{A}\hat{\mathbf{B}} - \mathbf{D}\|^2/\|\mathbf{D}\|^2$, for each methodology and each sample size, the methodologies ranked in descending order of performance are: “**ddsPLS**”, “**sPLS 3**”, “**SPLS**”, “**sPLS 1**”, “**sPLS 2**” and “**NIPALS-PLS**” for low n . As n increases, this order is maintained and all the methodologies get better results for this metric.

We define the scaled RMSE (Root Mean-Squared Error) such as

$$\forall j \in [1, q], \quad \text{scRMSE}_j = \sqrt{\frac{\sum_{i=1}^{n_{test}} (y_{i,j} - \hat{y}_{i,j})^2}{\sum_{i=1}^{n_{test}} (y_{i,j} - \bar{y}_j)^2}},$$

estimated on a independent sample of size n_{test} where $\bar{y}_j$ is estimated on the independent test sample and $\hat{y}_{i,j}$ is estimated according to the current prediction model. The Figure 4 shows the scaled RMSE for the different methodologies in all cases, for each response variable separately. Methodologies “**ddsPLS**” and “**sPLS 3**” show the best performances for low y_1 and y_2 while “**ddsPLS**” and “**NIPALS-PLS**” are the best for y_3 . More interestingly, “**sPLS 3**” and “**SPLS**” are especially not precise for y_3 , this is associated with the remark formulated in the previous paragraph associated with the potential over-fitting of those methodologies.

In terms of variable selection, the Figure 5 insures that the True Positive Rate (TPR) of variable of $\mathbf{x}$ selection is mostly equal to 1 for all methodologies, except for “**ddsPLS**”. On the contrary, the False Positive Rate (FPR) is mostly equal to 0 only for “**ddsPLS**” and is closely equal to 0 for “**SPLS**”. Other methodologies based on “sPLS” decrease their FPR with n until n reaches 100. More precisely, the FPR of “**sPLS 1**” indeed slowly decreases with n , “**sPLS 2**” and “**sPLS 3**” increase their FPR for $n \geq 100$. This shows that Q^2 maximization based methods (both for cross-validation and bootstrap versions) seem to

Figure 2: Number of components built for design 1, gray areas correspond to the objective, that is $R = 2$ components.

build sparser models as the number n of individuals grows. On the contrary, $Q^2 - R^2$ minimization based methods decrease their FPR as n grows.

In the y part, regardless of n , all methodologies select always y_1 and y_2 , and so the FPR equal 1 in all cases. According to Figure 5b), all methodologies naturally decreases the selection rate of y_3 as n increases, except for “sPLS 3”. Furthermore, among the bootstrapped methodologies, “ddsPLS” outperforms the others. In terms of variable selection, this can lead to the conclusion that bootstrap-based versions provide better results than those based on cross-validation, and more precisely the one associated with the minimization of $\bar{R}_{B,r}^2 - \bar{Q}_{B,r}^2$ which perform better than those based on the maximization of $\bar{Q}_{B,r}^2$, for both x and y parts.

Figure 3: Main performance metrics for design 1. The horizontal line represents the value of $Q^2 = \gamma^* = 2\varepsilon^2/3 \approx 0.6534$ which corresponds to the objective, where $\varepsilon = \sqrt{1 - \sigma^2}$.

Design 1 with varying q As to evaluate the performance of the proposed methodology if q varies, it has been considered to add extra response variables, as to reach $q = \{10, 50, 100\}$. Each of these extra response variables is not associated nor with the other response variables nor with the other response variables.

Figure 4: Scaled RMSE (Root Mean Squared Errors) for design 1, for each response variable, on an independent test sample of size $n_{test} = 1000$. Horizontal lines represent the objective values $\approx (0.141, 0.141, 1)$.

As previously, $N = 100$ different data sets (with $n = 50$) have been generated to build different prediction models. Different information can be highlighted.

- All the models are built on 2 components.
- All the models select, in the response part, only the two relevant variables (the two first ones).
- The first row of Figure 6 gives the scaled RMSE of the N models, for each scenario and for each response variable. An independent data set of size $n_{test} = 1000$ has been used to estimate those distributions. It is clear that all the models predict well the two selected variables (y_1 and y_2) and do not "predict" for others but returns the empirical mean values.
- The Second row of Figure 6 shows the frequency of selection in the covariate part for each scenario. Only the 100 first covariates, associated with y_1 and y_2 , are selected.

This experiment shows that the proposed "**ddsPLS**" methodology is robust to an increasing number q of response variables.

5.2 Design 2: simulated data inspired by real data

In this design, datasets with structures similar to spectroscopic data have been generated. Using the routine already used by Cloarec (2014), the $\mathbf{x}$ part contains spectra observations and is divided into two blocks. The first one ($\mathbf{X}_1$) corresponds to spectroscopy with large overlapping peaks and a low number of variables, such as Ultraviolet, Visible and Near-Infrared spectrophotometry. The second block ($\mathbf{X}_2$) has narrow peaks and a high number of variables but with much less overlap between peaks. The blocks $\mathbf{X}_t, t = 1, 2$ are built using the following procedure:

$$\mathbf{X}_t = \sqrt{1 - \sigma_t^2} \mathbf{C}_t \mathbf{S}_t + \sigma_t \mathbf{E}_t \quad (31)$$

with

$$\mathbf{S}_t = \left[(s_{t,r,j})_{(r,j)} \right] \in \mathbb{R}^{\mathcal{R} \times p_t}, \quad \mathbf{C}_t = \left[(c_{t,i,r})_{(i,r)} \right] \in \mathbb{R}^{n \times \mathcal{R}} \text{ and } \mathbf{E}_t = [(\mathbf{e}'_{t,i})_i]' \in \mathbb{R}^{n \times p_t}$$

where

$$\forall (r, j) \in \llbracket 1, \mathcal{R} \rrbracket \times \llbracket 1, p_t \rrbracket, \quad s_{t,r,j} = \sum_{k=1}^{K_r} h_{r,k} \exp \left\{ - (j - \pi_{r,k})^2 / (2\omega_{r,k}^2) \right\}, \quad \mathbf{e}_{t,i} \sim \mathcal{N}(\mathbf{0}_{p_t}, \mathbb{I}_{p_t}) \text{ and } \sigma_t = 0.05.$$

The score matrix is here denoted $\mathbf{C}_t$ for "Concentration", and the matrix of the weights is denoted $\mathbf{S}_t$ for "Spectrum". Let p_t be the number of variables in each block t and $\mathcal{R}_t$ the number of latent variables ($\mathcal{R}_t = 10$ for both blocks). For each latent variable, $r \in \llbracket 1, \mathcal{R}_t \rrbracket$, the number K_r of peaks is selected randomly from a uniform discrete random distribution $\mathcal{U}[\llbracket 1, 5 \rrbracket]$ in order to generate, $\forall k \in \llbracket 1, K_r \rrbracket$, a spectrum with the following characteristics:

- position of the peaks is sampled from a discrete uniform distribution $\pi_{r,k} \sim \mathcal{U}[\llbracket 2, p_t \rrbracket]$;
- relative peak heights are sampled from a continuous uniform distribution $h_{r,k} \sim \mathcal{U}(0, 1)$;
- peak widths are sampled from a continuous uniform distribution $\omega_{r,k} \sim \mathcal{U}(2, 3)$ and are proportional to the number of variables.

The previous parameters allow to build the matrix $\mathbf{S}_t$. To build the matrix $\mathbf{C}_t$ such as the intensity of each latent variable in each observation, the following Gaussian distribution is used: $c_{t,i,r} \sim \mathcal{N}(\mu_{t,r}, \sigma_{t,r}^2)$ where

- the expected intensity of each latent variable is randomly sampled from an exponential distribution: $\log(\mu_{t,r}) \sim \mathcal{U}(0, 1)$,

(a) Variable selection performances in the x part.

(b) Variable selection performance in the y part measured by the number of selected y variables where the gray areas correspond to theoretical objectives, that is two selected variables: y_1 and y_2 (which are always selected, providing a TPR equal to 1).

Figure 5: Variable selection performances in the x and y parts for design 1

Figure 6: Scaled RMSE for each response variable and selection frequencies in the covariate part for the three scenarios $q = \{10, 50, 100\}$.

- the variance of the intensity for each latent variable is randomly sampled from a uniform distribution $\sigma_{t,r}^2 \sim \mathcal{U}(0, \sigma_\mu \mu_{t,r})$ where σ_μ , when low enough, allows to generate positive values for $c_{t,i,r}$ with high probability.

Moreover, in order to get the correlation between the two blocks $\mathbf{X}_1$ and $\mathbf{X}_2$, the first five columns of $\mathbf{C}_2$ have been replaced by the first five columns of $\mathbf{C}_1$. Figure 7 provides a visualization of an extract from the x part for a dataset simulated according to the procedure.

The 5-dimensional y part is generated from the model $\mathbf{Y} = \sqrt{1 - \sigma_y^2} \mathbf{C}_y + \sigma_y \mathbf{E}_y$, where each row of $\mathbf{E}_y$

follows $\mathcal{N}(\mathbf{0}, \mathbb{I})$, $\sigma_y = 0.2$ corresponds to the standard deviation of this observation noise, and the matrix $\mathbf{C}_y$ is used to manage the link with the $\mathbf{x}$ part. More precisely, the first 3 columns of $\mathbf{C}_y$ are the standardized first 3 columns of $\mathbf{C}_1$. The last two columns of $\mathbf{Y}$ are generated from a standardized normal distribution. Hence only the first 3 variables of the $\mathbf{y}$ part are related to the $\mathbf{x}$ part.

According to the way the spectra are generated, the objective number R of components of the underlying model is equal to 3 with a high probability.

Figure 7: First five observations of simulated datasets $\mathbf{X}_1$ and $\mathbf{X}_2$ from design 2.

Figures 8(a) and 8(b) respectively show, for 100 replications, the number $\hat{R}$ of components retained in each model and the associated number of selected variables in the $\mathbf{x}$ part. Since the correct dimension is equal to $R = 3$, all the methodologies (apart from “NIPALS-PLS”, which builds far too many components) succeed to get most of times models with three components, and only “ddsPLS” has a clear majority of models built on three components. Moreover, according to Figure 8(b), “sPLS 1” and “sPLS 2” build often non-sparse models with very large numbers of selected variables. “sPLS 3” also builds models with various numbers of selected variables and also builds non sparse models. The methodologies “ddsPLS” and “SPLS” build the sparsest models, with a clear advantage for the first methodology even if “ddsPLS” sometimes builds non sparse models. Note that these graphs do not make it possible to appreciate the quality of the selection procedure and the corresponding estimated model different methodologies; this point is discussed below.

(a) Number $\hat{R}$ of components built for each methodology for design 2. Gray area corresponds to the correct value $R = 3$.

(b) Number of selected variables in the $\mathbf{x}$ part for each methodology for design 2.

Figure 8: Estimated number of components $\hat{R}$ (top) and number of selected variables (bottom) for design 2.

Since the simulated spectra are not sparse, it is not possible to directly evaluate the relevance of the variable selections. However, the 3 spectra, which are common to $\mathbf{X}_1$, $\mathbf{X}_2$ and $\mathbf{Y}$ are known in this simulation framework. It is reasonable to assume that the relevance of a variable is associated with the intensity of this variable in these three spectra of interest. Then, the selected variables are described in terms of TPR and FPR by varying a threshold α over the maximum values of the spectra for each variable along with

the three spectra of interest as to build ROC (Receiver Operating Characteristic) curves and so create a common score of selectivity corresponding to the Area Under the Curve (AUC). For a relevant analysis, a common distribution of α for all methodologies has been used.

The top right graph of the Figure 9 shows the AUC metrics computed for the models built over the 100 replications of datasets. It allows evaluating the performance of the methodologies in terms of variable selection. It seems that “sPLS 1” and “sPLS 2” are the least efficient methodologies to select the correct variables efficiently. The methodologies “sPLS 3” and “SPLS” get better AUC but keep high variability. Finally, the proposed “ddsPLS” approach provides the best AUC.

Top left graphs (density plots) of the Figure 9 detail more precisely the quality of the selection procedure of the five competing methodologies (in columns), via the TPR versus α (first row), the FPR versus α (second row) and the FPR versus the TPR (third row).

- From the first row, one might notice that “sPLS 1” and “sPLS 2” present a huge loss of TPR for very low α and do not succeed to keep TPR close to 1 for low values of α . “sPLS 3” also suffers from a fall although less important and even succeeds to keep TPR close to 1 for low values of α . “SPLS” keeps this tendency, and “ddsPLS” performs even better.
- According to the second row, it seems that “sPLS 1” and “sPLS 2” quickly (in terms of α) fall in FPR, which is encouraging in a certain way, while the three other methodologies keep it high for low values of α . This actually, quite paradoxically, allows considering that “ddsPLS”, “SPLS” and “sPLS 3” are, in fact, better methodologies because they do not tend to remove too quickly variables; feature that is often criticized in sparse methods.
- The third row shows the ROC curves in density and once again confirms the previous observations: “ddsPLS”, “SPLS” and “sPLS 3”, in decreasing order of performance, outperform the other methodologies since the high densities are concentrated at the top of their respective plot, corresponding to high TPR. Unfortunately, the precision is not sufficient to allow a correct differentiation of each distribution for high TPR’s.

The five histograms in Figure 9 (bottom right graphs) show the conditional distribution of the FPR given $\text{TPR} > 0.9$. “ddsPLS” clearly provides the heaviest distribution in that TPR area and gets the heaviest lower distribution tail, which are the two reasonable objectives of a selection method.

Figure 9: Variable selection performances for design 2. Top left graphs: density plots for the five competing methodologies (in columns) of the TPR versus α (first row), of the FPR versus α (second row) and of FPR versus TPR (third row), red areas correspond to densities higher than 4000 for the first two rows and 3000 otherwise. Top right graph: boxplots of AUC for the five competing methodologies. Bottom right graphs: histograms of FPR given $\text{TPR}_{i,0.9}$ for the five competing methodologies.

In terms of Q^2 results (not provided here), “ddsPLS”, “SPLS”, “sPLS 1” and “sPLS 2” are the most efficient methodologies (respectively equal to 0.47, 0.46, 0.47 and 0.43 in median for the all simulated datasets) for $\hat{R} = 3$ components built models. Note that those four methodologies get lower Q^2 than expected ($\gamma = 3\epsilon_y^2/5 = 0.597$ in the considered simulation framework). Only “sPLS 3” gets higher median for Q^2 , equal to 0.55. For more complex models built with $\hat{R} > 3$ components, the results also show that “sPLS 3” give a median Q^2 equal to 0.64 for $\hat{R} = 4$, 0.75 for $\hat{R} = 5$, 0.93 for $\hat{R} = 6$, 0.91 for $\hat{R} = 8$ and 0.99 for $\hat{R} = 10$, while the four other methodologies get decreasing values for median Q^2 . Finally, in the context of that simulation design, it seems again that “sPLS 3” does not succeed to beat over-fitting.

Concerning the selection in the y part (which is only available for “ddsPLS”, “sPLS 1”, “sPLS 2” and “sPLS 3” methodologies), the following results were obtained. Below, the notation (y_j, N_j) means that variable y_j has been selected N_j times over the 100 replications:

- “**ddsPLS**” : $(y_1, 87), (y_2, 87), (y_3, 81), (y_4, 18), (y_5, 19),$
- “**sPLS 1**” : $(y_1, 83), (y_2, 86), (y_3, 74), (y_4, 37), (y_5, 36),$
- “**sPLS 2**” : $(y_1, 88), (y_2, 84), (y_3, 82), (y_4, 53), (y_5, 51),$
- “**sPLS 3**” : $(y_1, 89), (y_2, 91), (y_3, 81), (y_4, 32), (y_5, 33).$

It is clear that “**ddsPLS**” outperforms other methodologies regarding the FPR’s in **y** part selection, while the TPR’s appear to be equivalent for the four methodologies.

6 Conclusion

The proposed ddsPLS methodology is a sparse PLS method, both in **x** and **y**, based on soft-thresholding using bootstrap operations for model selection. This solution fixes sparsity in both **x** and **y** parts using a single parameter, which is also interpretable in terms of minimal empirical covariance allowed to be selected in the model. Moreover this solution shows very good results both in prediction and in selection, through the various simulation models analysed.

While most of the other sparse PLS are based on cross-validation approaches, our solution uses bootstrap. This allows building smooth curves of the chosen criterion even in the ill-conditioned contexts of low n and large p data.

Based on the proposed bootstrap version of sparse PLS, a new relevant definition of the Q^2 and the R^2 statistics denoted as $\bar{Q}_B^2$ and $\bar{R}_B^2$ where B is the number of bootstrap samples, has been introduced. Those new versions of Q^2 and R^2 do not mix the errors of the different folds but aggregate the fold estimations through the bootstrap mean estimator.

Different aspects must be highlighted regarding those new statistics. The newly introduced γ^* is a theoretical limit for the Q^2 and R^2 statistics that our criterion respects while the tested competing methods violate it in the $n \ll p$ context even for simple simulation schemes. The chosen criterion $\bar{R}_B^2 - \bar{Q}_B^2$ allows to efficiently select sparse models.

Finally, the ddsPLS methodology can be applied to different data contexts such as biological data, omics data and chemometrics, where the sample size n is indeed much lower than the number p of variables.

An implementation of the proposed ddsPLS methodology is available at <https://github.com/hlorenzo/ddsPLS2>. A vignette is also available to facilitate the use of the corresponding R package and the associated functions for any user not necessarily specialised in R or statistical modelling. Note that this package uses the **foreach** (v.1.5.0) package to parallelize the computations associated with the different bootstrap samples. These codes are written in C^{++} thanks to the packages **Rcpp** (v.1.0.5) and **RcppEigen** (v.0.3.3.7.0) in order to reduce the computation time and to ensure a portability towards other languages, Python in particular.

Whatsoever, all simulation codes (regarding **Toy Example**, **Design 1** and **Design 2**) are accessible via the repository https://github.com/hlorenzo/simulation_ddspls, where two files (**README.md** and **simulation_ddspls.html**) help to use the different simulation functions.

Acknowledgments

We thank the editor and the two anonymous referees for their comments that led to an improved version of this paper. Hadrien Lorenzo received the support of “Sartorius, Corporate Research Advanced Data Analytics” for this project.

References

- Bickel, P. J. and E. Levina, 2008: Regularized estimation of large covariance matrices. *The Annals of Statistics*, **36**, no. 1, 199 – 227, doi:10.1214/009053607000000758.
URL <https://doi.org/10.1214/009053607000000758>
- Broc, C., T. Truong, and B. Liquet, 2021: Penalized partial least squares for pleiotropy. *BMC bioinformatics*, **22**, no. 1, 1–31.
- Cai, T. and W. Liu, 2011: Adaptive thresholding for sparse covariance matrix estimation. *Journal of the American Statistical Association*, **106**, no. 494, 672–684, doi:10.1198/jasa.2011.tm10560.
URL <https://doi.org/10.1198/jasa.2011.tm10560>
- Chun, H. and S. Keleş, 2010: Sparse partial least squares regression for simultaneous dimension reduction and variable selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **72**, no. 1, 3–25.
- Cloarec, O., 2014: Can we beat over-fitting? *Journal of Chemometrics*, **28**, no. 8, 610–614.
- Deshpande, Y. and A. Montanari, 2016: Sparse PCA via covariance thresholding. *Journal of Machine Learning Research*, **17**, no. 141, 1–41.

- Geisser, S., 1974: A predictive approach to the random effect model. *Biometrika*, **61**, no. 1, 101–107, doi:10.1093/biomet/61.1.101.
URL <https://doi.org/10.1093/biomet/61.1.101>
- Helland, I. S. and T. Almøy, 1994: Comparison of prediction methods when only a few components are relevant. *Journal of the American Statistical Association*, **89**, no. 426, 583–591.
- Johnstone, I. M. and A. Y. Lu, 2009: On consistency and sparsity for principal components analysis in high dimensions. *Journal of the American Statistical Association*, **104**, no. 486, 682–693.
- Kautt, A. F., C. F. Kratochwil, A. Nater, G. Machado-Schiaffino, M. Olave, F. Henning, J. Torres-Dowdall, A. Härer, C. D. Hulsey, P. Franchini, M. Pippel, E. W. Myers, and A. Meyer, 2020: Contrasting signatures of genomic divergence during sympatric speciation. *Nature*, **588**, no. 7836, 106–111.
- Lê Cao, K.-A., D. Rossouw, C. Robert-Granié, and P. Besse, 2008: A sparse PLS for variable selection when integrating omics data. *Statistical applications in genetics and molecular biology*, **7**, no. 1.
- Liquet, B., P. L. De Micheaux, B. P. Hejblum, and R. Thiébaut, 2016: Group and sparse group partial least square approaches applied in genomics context. *Bioinformatics*, **32**, no. 1, 35–42.
- Martens, H., T. Ni, et al., 1992: *Multivariate calibration*. John Wiley & Sons.
- Mehmood, T., S. Sæbø, and K. H. Liland, 2020: Comparison of variable selection methods in partial least squares regression. *Journal of Chemometrics*, **34**, no. 6, e3226, doi:<https://doi.org/10.1002/cem.3226>, e3226 cem.3226.
URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/cem.3226>
- Pérez-Enciso, M. and M. Tenenhaus, 2003: Prediction of clinical outcome with microarray data: a partial least squares discriminant analysis (pls-da) approach. *Human genetics*, **112**, no. 5-6, 581–592.
- Rechtien, A., L. Richert, H. Lorenzo, [...], R. Thiébaut, M. Altfeld, and M. Addo, 2017: Systems vaccinology identifies an early innate immune signature as a correlate of antibody responses to the ebola vaccine rvsv-zebov. *Cell Reports*, **20**, no. 9, 2251–2261, doi:10.1016/j.celrep.2017.08.023.
- Stone, M., 1974: Cross-validatory choice and assessment of statistical predictions. *Journal of the Royal Statistical Society. Series B (Methodological)*, **36**, no. 2, 111–147.
URL <http://www.jstor.org/stable/2984809>
- Sutton, M., R. Thiébaut, and B. Liquet, 2018: Sparse partial least squares with group and subgroup structure. *Statistics in Medicine*, **37**, no. 23, 3338–3356, doi:<https://doi.org/10.1002/sim.7821>.
URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/sim.7821>
- Tenenhaus, M., 1998: *La régression PLS: théorie et pratique*. Editions technip.
- 2004: PLS Regression and PLS Path Modeling for Multiple Table Analysis. *Proceedings in COMPSTAT, Proceedings in Computational Statistics*, J. Antoch, ed., Physica-Verlag HD, Heidelberg, 489–499.
- 2005: La régression logistique PLS. *Proceedings in Modèles statistiques pour données qualitatives*, Editions technip, 263–276.
- Thévenot, E. A., 2016: ropls: Pca, pls (-da) and opls (-da) for multivariate analysis and feature selection of omics data. *Bioconductor*.
- Tibshirani, R., 1996: Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 267–288.
- Vivien, M., 2002: *Approches PLS linéaires et non linéaires pour la modélisation de multi-tableaux. Théorie et applications*. Ph.D. thesis, Université Montpellier I.
- Wakeling, I. N. and J. J. Morris, 1993: A test of significance for partial least squares regression. *Journal of Chemometrics*, **7**, no. 4, 291–304.
- Westerhuis, J. A. and A. K. Smilde, 2001: Deflation in multiblock PLS. *Journal of chemometrics*, **15**, no. 5, 485–493.
- Wold, H., 1966: Estimation of principal components and related models by iterative least squares. *Multivariate analysis*, 391–420.
- Wold, S., 1984: Three PLS algorithms according to SW. *Proceedings in Symposium MULDAST, Multivariate analysis in science and technology*, 26–30.
- Zoong Lwe, Z. S., R. Welte, D. Anco, S. Naveed, S. Rustgi, and S. Narayanan, 2020: Heat stress elicits remodeling in the anther lipidome of peanut. *Scientific Reports*, **10**, no. 1, 22163, doi:10.1038/s41598-020-78695-3.