

VC-PCR: A PREDICTION METHOD BASED ON VARIABLE SELECTION AND CLUSTERING

Rebecca Marion, Johannes Lederer,
Bernadette Govaerts, Rainer von Sachs

REPRINT | 2024 / 23

ISBA

Voie du Roman Pays 20 - L1.04.01

B-1348 Louvain-la-Neuve

Email : lidam-library@uclouvain.be

<https://uclouvain.be/en/research-institutes/lidam/isba/publication.html>

VC-PCR: A prediction method based on variable selection and clustering

Rebecca Marion¹ | Johannes Lederer²  |
Bernadette Goevarts¹ | Rainer von Sachs¹

¹Institute of Statistics, Biostatistics and Actuarial Sciences, LIDAM, Louvain-la-Neuve, Belgium

²Department of Mathematics, Computer Science, and Natural Sciences, University of Hamburg, Hamburg, Germany

Correspondence

Johannes Lederer, Department of Mathematics, Computer Science, and Natural Sciences, University of Hamburg, Bundesstraße 55, 20146 Hamburg, Germany.
Email:
johannes.lederer@uni-hamburg.de

Funding information

Belgian Fund for Scientific Research (F.R.S.-FNRS, FRIA grant)

Sparse linear prediction methods suffer from decreased prediction accuracy when the predictor variables have cluster structure (e.g., highly correlated groups of variables). To improve prediction accuracy, various methods have been proposed to identify variable clusters from the data and integrate cluster information into a sparse modeling process. But none of these methods achieve satisfactory performance for prediction, variable selection and variable clustering performed simultaneously. This paper presents Variable Cluster Principal Component Regression (VC-PCR), a prediction method that uses variable selection and variable clustering in order to solve this problem. Experiments with real and simulated data demonstrate that, compared to competitor methods, VC-PCR is the only method that achieves simultaneously good prediction, variable selection, and clustering performance when cluster structure is present.

KEYWORDS

dimensionality reduction, latent variables, nonnegative matrix factorization, sparsity, variable clustering

1 | INTRODUCTION

A common objective of transcriptomics is to predict health (cancer status, time to death, etc.) from genomic variables. For the sake of knowledge discovery and interpretability, it is also very

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial-NoDerivs](https://creativecommons.org/licenses/by-nc-nd/4.0/) License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

© 2024 The Author(s). *Statistica Neerlandica* published by John Wiley & Sons Ltd on behalf of Netherlands Society for Statistics and Operations Research.

important to identify the variables that are most important for predicting the health variable (i.e., perform variable selection). However, for complex diseases, it is common for clusters of genes (i.e., gene pathways) to contribute to the disease jointly, rather than on an individual basis (Ma, Song, & Huang, 2007). In these settings, the task of identifying clusters of predictive variables is as important as prediction and variable selection.

Transcriptomics data pose two major challenges to predicting, selecting variables, and clustering variables performed simultaneously, that is, in a single attempt. First, the data is high-dimensional, that is, there is a large number of variables (e.g., genes) with respect to the number of observations. In this case, spurious correlations between relevant and irrelevant variables are common, and as a result, models tend to select too many variables. This overselection results in false positives (Daye & Jeng, 2009), that is, irrelevant variables that are identified as relevant. The second challenge stems from the underlying cluster structure among predictor variables: genes within a given pathway generally share similar functions and have highly correlated expression profiles. This high correlation between predictor variables is known to increase the variance of model coefficients, thereby decreasing prediction accuracy (Park, Hastie, & Tibshirani, 2007).

To address these challenges, various methods have been proposed for estimating a sparse linear model that directly accounts for cluster structure among predictor variables. These methods mostly fall into one of two categories: embedded methods, which simultaneously cluster variables and estimate a sparse model, or two-step methods, which cluster variables and then estimate a sparse model based on these clusters. However, in our observation, none of these existing methods are able to satisfactorily perform the task of simultaneous prediction, variable selection, and variable clustering. Embedded methods tend to be computationally expensive and are generally limited to regression settings, whereas two-step methods tend to identify clusters that are “diluted” with irrelevant variables (Yengo, Jacques, & Biernacki, 2014). Methods from both categories tend to treat the task of variable clustering as a means to improve prediction accuracy rather than as an objective in itself. However, correctly identifying variable clusters (e.g. gene pathways in transcriptomics) is often instrumental to the interpretation of the model, as it allows researchers to better understand disease mechanisms, as well as develop targeted treatments. For convenience of the reader, we have put an in-depth comparison of various methods in Section A of the Appendix.

The goal of this paper is to address these limitations. In this spirit, we propose a method called Variable Cluster Principal Component Regression (VC-PCR) that is motivated by the transcriptomics setup but is designed in a much more general fashion. Unlike existing methods, it addresses three objectives simultaneously: (i) accurate prediction of the response in high-dimensional regression, (ii) selection of the most relevant predictor variables, and (iii) identification of predictive variable clusters. The first step of VC-PCR consists of supervised variable selection and clustering, which is performed using a weighted version of a method that we call Sparse Orthogonal Semi-Nonnegative Matrix Factorization (SOS-NMF). As with the embedded approaches, supervising the variable clustering allows VC-PCR to identify homogeneous, predictive clusters. Moreover, irrelevant or “inactive” variables can be removed from the analysis during the clustering step, improving clustering performance and the recovery of the underlying model. In the second step of VC-PCR, a matrix of latent variables is calculated based on the cluster assignments from the first step, and a classical linear model is estimated to relate these latent variables to the response. Similar to existing two-step methods, the two-step approach employed by VC-PCR is more computationally efficient than embedded methods and can be easily applied to different prediction tasks (regression, classification, survival analysis, etc.). However, unlike other two-step




FIGURE 1 Schematic representation of variable cluster principal component regression (VC-PCR). A response (green column) and predictors (light blue columns) are used to generate cluster assignments for the predictors. In the cluster assignments bar, predictors in the same cluster are shown in the same color, and predictors assigned to none of the clusters are shown in white. Membership degrees are also estimated for each predictor, with values ranging from no membership (white) to high membership (black). These membership degrees are used to calculate latent variables for each cluster. Then, a linear model is estimated based on these latent variables in order to predict the response.

methods, VC-PCR harnesses the advantages of supervision during variable clustering. VC-PCR is summarized in a schematic representation in Figure 1.

Numerical experiments with real and simulated data show that VC-PCR performs better than competing methods in terms of model error, variable selection and clustering. These competing methods are not designed to achieve all three goals simultaneously, meaning that, unlike for VC-PCR, good clustering performance comes at the cost of poor variable selection and/or prediction performance. A weighted VC-PCR, with weights the coefficients for a Ridge regression model, achieves the best variable selection and clustering performance across all simulation settings tested, and it has the best prediction performance for sparse models in settings where variables have predictive cluster structure.

2 | NOTATION

Matrices are denoted with bold-faced capital letters (e.g., $\mathbf{X}$, $\mathbf{\Theta}$), vectors with bold-faced lower-case letters (e.g., $\mathbf{x}$, $\boldsymbol{\theta}$) and scalars with lower-case letters (e.g., x , θ). Row and column vectors of a matrix are indicated by subscripts. For example, if the columns of $\mathbf{X}$ are indexed by the letter j , then the j th column vector is denoted $\mathbf{x}_j$. If the rows are indexed by the letter i , then the i th row vector is denoted $\mathbf{x}_i$. A single element from a matrix is denoted using a lower-case letter with a subscript for

the row index and a subscript for the column index (e.g., x_{ij}). Observations are indexed with the letter $i : 1, \dots, n$, variables with the letter $j : 1, \dots, p$, and clusters with the letter $k : 1, \dots, K$. The Frobenius norm of an $n \times p$ matrix $\mathbf{X}$ is $\|\mathbf{X}\|_F$, where $\|\mathbf{X}\|_F^2 = \sum_{i=1}^n \sum_{j=1}^p x_{ij}^2$. The ℓ_2 norm of an $n \times 1$ vector $\mathbf{x}$ is $\|\mathbf{x}\|_2 = \sqrt{\sum_{i=1}^n x_i^2}$, and the ℓ_1 norm of $\mathbf{x}$ is $\|\mathbf{x}\|_1 = \sum_{i=1}^n |x_i|$. The indicator function $\mathbb{I}[x]$ returns 1 if x is true and 0 otherwise.

Let $\mathbf{X}$ ($n \times p$) be the matrix of predictor variables and $\mathbf{y}$ ($n \times 1$) be the response vector. We assume that the columns of $\mathbf{X}$ are mean-centered and have a sample variance of 1, that is, $\text{mean}(\mathbf{x}_j) = 0$ (sample mean), $\text{var}(\mathbf{x}_j) = 1$ (sample variance) and $\mathbf{x}_j^\top \mathbf{x}_j = n - 1$, $\forall j$. For regression problems, the vector $\mathbf{y}$ is mean-centered and scaled to unit sample variance, and for two-class classification problems, the vector $\mathbf{y}$ is binary.

For regression problems, we assume that the response is a linear function of the predictor variables plus an error:

$$y_i = \mathbf{x}_i^\top \mathbf{b} + \varepsilon_i, \quad (1)$$

where $\mathbf{b}$ ($p \times 1$) contains the model coefficients relating the predictors to the response, $\mathbf{x}_i$ ($p \times 1$) contains the values of the predictors measured for observation i and $\varepsilon_i \stackrel{\text{iid}}{\sim} N(0, \sigma_\varepsilon^2)$. For two-class classification problems, the elements in $\mathbf{y}$ are binary (e.g. $y_i = 1$ if a patient has a given disease and $y_i = 0$ otherwise). We assume that $y_i \stackrel{\text{iid}}{\sim} \text{Bernoulli}(\pi_i)$, where

$$\pi_i = P(y_i = 1) = \frac{\exp(\mathbf{x}_i^\top \mathbf{b})}{1 + \exp(\mathbf{x}_i^\top \mathbf{b})}. \quad (2)$$

3 | VARIABLE CLUSTER—PCR

Our proposed method, VC-PCR, consists of two steps. In the first step, variables are clustered, generating a matrix $\hat{\mathbf{V}}$ ($p \times K$) of cluster membership degrees. Each element $\hat{v}_{jk}$ in $\hat{\mathbf{V}}$ is nonnegative and represents the estimated degree to which variable j belongs to cluster k . In the second step, $\hat{\mathbf{V}}$ is used to calculate a set of latent variables, $\hat{\mathbf{m}}_1, \dots, \hat{\mathbf{m}}_K$, one for each variable cluster, and a linear regression or classification model is estimated based on these latent variables. This second step is similar in spirit to Principal Component Regression (PCR), hence the name “VC-PCR.”

The first step of VC-PCR (variable clustering) is performed using a method that we call SOS-NMF. This method is summarized in Section 3.1. The two steps of the VC-PCR algorithm are explained in Sections 3.2 and 3.3, and the full algorithm and a discussion of its hyperparameters are presented in Section B of the Appendix.

3.1 | Sparse orthogonal semi-NMF

SOS-NMF is the unsupervised method that we developed to perform dimensionality reduction, variable selection and variable clustering on both nonnegative matrices and matrices with positive and negative elements. The objective of SOS-NMF is to decompose $\mathbf{X}$ into the product of two matrices: $\mathbf{U}$ ($n \times K$), a matrix of K unit-variance latent variables, and $\mathbf{V}$ ($p \times K$), a sparse, nonnegative matrix with orthogonal columns. Thanks to the orthogonality and nonnegativity of $\mathbf{V}$, each

vector $\mathbf{v}_j$ (i.e., row of $\mathbf{V}$) can have at most one nonzero value. As a result, variable clusters can be inferred from $\mathbf{V}$: $C_k := \{j \mid v_{jk} > 0\}$. Moreover, the nonzero elements of $\mathbf{V}$ indicate the degree to which a variable belongs to its assigned cluster, where larger values indicate larger membership degrees. The variable selection is possible because we do not assume that the row sums of $\mathbf{V}$ must be positive.

The solution to our proposed SOS-NMF problem is

$$\begin{aligned} \arg \min_{\mathbf{U}, \mathbf{V}} \quad & \frac{1}{2(n-1)} \|\mathbf{X} - \mathbf{U}\mathbf{V}^\top\|_F^2 + \lambda \sum_{j=1}^p \|\mathbf{v}_j\|_1 \\ \text{s.t. } \quad & v_{jk} \geq 0, \quad \forall j, k \\ \text{s.t. } \quad & \text{var}(\mathbf{u}_k) = 1, \quad \forall k \\ \text{s.t. } \quad & \mathbf{V} \text{ has orthogonal columns,} \end{aligned} \quad (3)$$

where λ is a nonnegative hyperparameter controlling the sparsity in $\mathbf{V}$.

3.1.1 | Intuition

The orthogonality and nonnegativity constraints make $\mathbf{V}$ a naturally sparse matrix. The ℓ_1 penalty in Equation (3) encourages $\mathbf{V}$ to be even sparser, with some vectors $\mathbf{v}_j$ being set to zero (see, e.g., section 2.2 in Lederer, 2022). As a result, some variables (e.g., noise variables) may be assigned to none of the clusters and removed from the analysis. This is desirable because noise variables make it difficult to correctly cluster variables (see Vigneau, 2016; Marion, Govaerts, & von Sachs, 2020).

A given variable j is assigned to cluster k (i.e., $v_{jk} > 0$) if its correlation with latent variable k is sufficiently large. As λ increases, the correlation threshold for cluster assignment increases, resulting in the removal of more and more variables from the analysis. Another necessary condition for being assigned to cluster k is having positive correlation with latent variable k . As a result, variables in the same cluster tend to be positively correlated.

Each latent variable vector $\mathbf{u}_k$ is proportional to the weighted sum of the variables j assigned to cluster k , where the weights are the cluster membership degrees v_{jk} . This weighted sum makes it possible to give greater importance to the variables that are the most representative of the cluster.

3.1.2 | Optimization

The solution for SOS-NMF can be found by alternating between the optimization of $\mathbf{U}$ for fixed $\mathbf{V}$ and of $\mathbf{V}$ for fixed $\mathbf{U}$. The derivation of these solutions can be found in the Appendix. For fixed $\mathbf{V}$, the solution for an unconstrained $\mathbf{U}$ is $\tilde{\mathbf{U}} = \mathbf{X}\mathbf{V}(\mathbf{V}^\top\mathbf{V})^{-1}$. We can account for the unit-variance constraint on the latent variables $\mathbf{u}_k$ by calculating the final solution $\hat{\mathbf{U}}$ as follows for all k :

$$\hat{\mathbf{u}}_k = \tilde{\mathbf{u}}_k / \text{sd}(\tilde{\mathbf{u}}_k). \quad (4)$$

Let $C_k := \{j \mid v_{jk} > 0\}$. For fixed $\mathbf{U}$, if $j \in C_k$, then

$$\hat{v}_{jk} = \text{cor}(\mathbf{u}_k, \mathbf{x}_j) - \lambda. \quad (5)$$

The problem then becomes one of determining when $j \in C_k$. There are two necessary conditions for variable j to belong to cluster k : (i) $\text{cor}(\mathbf{u}_k, \mathbf{x}_j) > \lambda$ and (ii) $k = \arg \max_{\ell} \text{cor}(\mathbf{u}_{\ell}, \mathbf{x}_j)$.

3.2 | VC-PCR: cluster memberships

First, a matrix of cluster memberships $\mathbf{V}$ must be estimated. As one of our objectives is to predict a target $\mathbf{y}$, we propose making the clustering step dependent on the target. We do this by integrating information about the target into the SOS-NMF problem using a weight matrix $\mathbf{W}$.

Let $\mathbf{W}$ ($p \times p$) be a diagonal matrix whose entries w_{jj} encode (i) the prior relationship between variable j and the target $\mathbf{y}$ (e.g., positive vs. negative association), as well as (ii) the importance of variable j in the prediction task. For example, the weights could be the coefficients for a Ridge regression model (see Section B for other suggestions). When each variable j is weighted by w_{jj} , the solution to SOS-NMF becomes

$$\arg \min_{\mathbf{U}, \mathbf{V}} \frac{1}{2(n-1)} \|\mathbf{X}\mathbf{W} - \mathbf{U}\mathbf{V}^T\|_F^2 + \lambda \sum_{j=1}^p \|\mathbf{v}_j\|_1 \quad (6)$$

such that $v_{jk} \geq 0 \forall j, k$; $\text{var}(\mathbf{u}_k) = 1 \forall k$; $\mathbf{V}$ has orthogonal columns,

where λ is a nonnegative hyperparameter controlling the sparsity in $\mathbf{V}$. We call this problem Weighted SOS-NMF.

3.2.1 | Intuition

The weight matrix $\mathbf{W}$ allows the clustering to be supervised by the target $\mathbf{y}$. For example, variables that have a similar importance for the prediction problem (i.e., similar weights in absolute value) are more likely to be clustered together. In addition, variables with smaller weights (in absolute value) are more likely to be removed by the ℓ_1 penalty.

In contrast to the original SOS-NMF problem, variables may be assigned to cluster k even if they are negatively correlated with latent variable k . Suppose there are two variables j and ℓ such that $\text{cor}(\mathbf{u}_k, \mathbf{x}_j) = -\text{cor}(\mathbf{u}_k, \mathbf{x}_{\ell}) = \rho$ and $w_{jj} = -w_{\ell\ell} = \nu$. If variable j is assigned to cluster k , then variable ℓ will be as well. The higher the absolute correlation between variables, and the more similar their weights (in absolute value), the more likely it is that they will be clustered together.

3.2.2 | Optimization

The solution for Weighted SOS-NMF can be found by using $\mathbf{X}\mathbf{W}$ as the input to SOS-NMF rather than $\mathbf{X}$. For fixed $\mathbf{V}$, the solution for an unconstrained $\mathbf{U}$ is $\tilde{\mathbf{U}} = \mathbf{X}\mathbf{W}\mathbf{V}(\mathbf{V}^T\mathbf{V})^{-1}$. The solution for each unit-variance latent variable $\mathbf{u}_k$ can be calculated by scaling $\tilde{\mathbf{u}}_k$ (see Equation 4). For fixed $\mathbf{U}$, a given variable j is assigned to cluster k if two conditions are met: (i) $w_{jj}\text{cor}(\mathbf{u}_k, \mathbf{x}_j) > \lambda$ and (ii) $k = \arg \max_{\ell} w_{jj}\text{cor}(\mathbf{u}_{\ell}, \mathbf{x}_j)$. For all $j \in C_k$, the solution $\hat{v}_{jk} = w_{jj}\text{cor}(\mathbf{u}_k, \mathbf{x}_j) - \lambda$.

3.3 | VC-PCR: ordinary least squares

Once the matrix $\mathbf{V}$ of cluster memberships is estimated, we calculate a latent variable matrix $\widehat{\mathbf{M}} = f(\mathbf{X}, \widehat{\mathbf{V}})$ and estimate a linear model predicting $\mathbf{y}$ based on $\widehat{\mathbf{M}}$. Note that the latent variable matrix $\widehat{\mathbf{U}}$ from the first part is not used in this part. This is because $\widehat{\mathbf{U}}$ is a function of $\mathbf{XW}$, a modified version of $\mathbf{X}$. We are interested in relating the original matrix $\mathbf{X}$ to the target $\mathbf{y}$, therefore, the latent variable matrix $\widehat{\mathbf{M}}$ is calculated $\widehat{\mathbf{M}} = \mathbf{X}\widehat{\mathbf{V}}$.

The prediction problem requires solving

$$\arg \min_{\mathbf{a}} L(\widehat{\mathbf{M}}, \mathbf{y}; \mathbf{a}), \quad (7)$$

where $\mathbf{a}$ ($K \times 1$) is a vector of model coefficients and $L(\widehat{\mathbf{M}}, \mathbf{y}; \mathbf{a})$ is a loss function that depends on the modeling task. For linear regression tasks, the loss function $L(\widehat{\mathbf{M}}, \mathbf{y}; \mathbf{a}) = \|\mathbf{y} - \widehat{\mathbf{M}}\mathbf{a}\|_2^2$ and $\mathbf{a}$ can be found using Ordinary Least Squares: $\widehat{\mathbf{a}} = (\widehat{\mathbf{M}}^\top \widehat{\mathbf{M}})^{-1} \widehat{\mathbf{M}}^\top \mathbf{y}$. Estimating a sparse regression model is not necessary because variable selection is already performed in the first part: variables are removed from the analysis thanks to the sparsity penalty in Equation (6). For classification tasks with binary y_i , the likelihood-based loss function $L(\widehat{\mathbf{M}}, \mathbf{y}; \mathbf{a}) = -\sum_{i=1}^n y_i \log(\pi_i) + (1 - y_i) \log(1 - \pi_i)$, where

$$\pi_i = P(y_i = 1) = \frac{\exp(\widehat{\mathbf{m}}_i^\top \mathbf{a})}{1 + \exp(\widehat{\mathbf{m}}_i^\top \mathbf{a})}. \quad (8)$$

The model coefficients $\mathbf{a}$ can be found using Logistic Regression. Again, sparse regression is not necessary here because variable selection has already been performed.

For both regression and classification tasks, a vector $\widehat{\mathbf{b}}$ ($p \times 1$) of coefficients for the original variables in $\mathbf{X}$ can be calculated as a function of $\widehat{\mathbf{V}}$ and $\widehat{\mathbf{a}}$. The predicted value $\widehat{y}_i$ is a function of $\widehat{\mathbf{m}}_i^\top \widehat{\mathbf{a}} = \mathbf{x}_i^\top \widehat{\mathbf{V}} \widehat{\mathbf{a}}$. Therefore, the coefficients $\widehat{\mathbf{b}}$ relating the original variables to the target can be calculated as $\widehat{\mathbf{b}} = \widehat{\mathbf{V}} \widehat{\mathbf{a}}$.

4 | EXPERIMENTS

Experiments with real and simulated data are performed in the regression setting to illustrate the performance of the proposed method, VC-PCR, with related methods from the literature. Performance in terms of model error and complexity is evaluated for all methods and datasets. Variable selection and clustering performance is also evaluated for the simulated datasets, for which the true model support and variable clusters are known by design.

4.1 | Methods compared

Cluster Elastic Net (CEN) (Witten, Shojaie, & Zhang, 2014) is chosen to represent the coefficient-grouping methods because of its computational efficiency compared to the other methods in this category. Split Regularized Regression (SRR) (Christidis, Lakshmanan, Smucler, & Zamar, 2020), the only ensemble method that can be used to cluster variables, is also included. Of the different two-step approaches, Cluster Representative Lasso (CRL) (Bühlmann,

Rütimann, van de Geer, & Zhang, 2013) is tested because it is the most similar to the proposed method. For CRL, three clustering algorithms are tested: hierarchical agglomerative clustering (HAC) based on Ward's criterion (CRL-hclust), HAC based on canonical correlation, as seen in Bühlmann et al. (2013) (CRL-hclustCC), and K-means clustering (CRL-kmeans). For VC-PCR, three choices of weight matrices $\mathbf{W}$ are tested: a diagonal matrix of Lasso regression coefficients (VC-PCR-Lasso), a diagonal matrix of Ridge regression coefficients (VC-PCR-Ridge) and an identity matrix (VC-PCR-Identity). This last choice serves as a benchmark for assessing the added-value of variable-weighting in VC-PCR.

4.2 | Evaluation criteria

The performance of each method is evaluated with respect to different characteristics: the model error and the model complexity, and for the simulation study also the recovery of the model support and the underlying clusters. The model error is quantified with the Mean Squared Error of Prediction (MSEP). Given a vector $\hat{\mathbf{y}} (n_{test} \times 1)$ of predicted response values,

$$\text{MSEP} = \frac{1}{n_{test}} \sum_{i=1}^{n_{test}} (y_i - \hat{y}_i)^2. \quad (9)$$

The model complexity is quantified as the number s of active variables in the model. For methods that estimate one coefficient $\hat{b}_j$ per variable j (e.g., CEN, SRR, etc.),

$$s = \|\hat{\mathbf{b}}\|_0 = \sum_{j=1}^p \mathbb{I}[\hat{b}_j \neq 0]. \quad (10)$$

For methods that estimate one coefficient $\hat{a}_k$ per cluster k (e.g., CRL, etc.),

$$s = \sum_{k=1}^K \hat{p}_k \mathbb{I}[\hat{a}_k \neq 0], \quad (11)$$

where $\hat{p}_k$ is the number of variables in $\hat{\mathbf{C}}_k$.

The Matthews Correlation Coefficient (MCC) (Matthews, 1975) is used in the simulation study to quantify the recovery of the true model support and underlying clusters. Compared to other existing criteria for evaluating binary classification, MCC is considered to be less biased and is more robust to class imbalance (Chicco, 2017; Chicco & Jurman, 2020; Chicco, Tötsch, & Jurman, 2021). Given the number of true positives (TP), true negatives (TN), false positives (FP), and false negatives (FN),

$$\text{MCC} = \frac{(\text{TP})(\text{TN}) - (\text{FP})(\text{FN})}{\sqrt{(\text{TP} + \text{FP})(\text{TP} + \text{FN})(\text{TN} + \text{FP})(\text{TN} + \text{FN})}}, \quad (12)$$

with values in $[-1, 1]$, where a score of 1 represents perfect classification and a score of 0 is equivalent to random classification.

When evaluating model support recovery, variables with nonzero coefficients are assigned to the “positive” class and variables with null coefficients are assigned to the “negative class.”

Cluster recovery is evaluated for all unique pairs of variables. Pairs of variables assigned to the same cluster are in the “positive class” and all other pairs of variables are in the “negative” class.

4.3 | Simulated data

The simulated datasets vary with respect to three factors: the type of correlation structure among predictors (the configuration), the magnitude of correlation between correlated predictors (ρ) and the number of observations (n). The simulated data are generated as follows. For each observation $i : 1, \dots, n$, a vector $\mathbf{x}_i$ is generated from the distribution $N_p(\mathbf{0}, \mathbf{\Sigma})$ and a model error ε_i is generated from $N(0, \sigma_\varepsilon^2)$, where $\mathbf{\Sigma}$ is the covariance matrix for $\mathbf{X}$ and σ_ε^2 is the model noise variance. The response y_i is calculated as $y_i = \mathbf{x}_i^\top \mathbf{b} + \varepsilon_i$, where $\mathbf{b}$ is the vector of true model coefficients. The number of observations $n \in \{25, 50\}$ and p , the number of variables, is fixed to 200.

The vector of true coefficients is fixed for all experiments, meaning that the set of relevant or “active” variables $S := \{j \mid b_j \neq 0\}$ is the same for all datasets. The true vector of coefficients is

$$\mathbf{b} = (\underbrace{1, \dots, 1}_5, \underbrace{0, \dots, 0}_5, \underbrace{-1, \dots, -1}_5, \underbrace{0, \dots, 0}_5, \underbrace{1, \dots, 1}_5, \underbrace{0, \dots, 0}_5, \underbrace{-1, \dots, -1}_5, \underbrace{0, \dots, 0}_5, \underbrace{0, \dots, 0}_{160})^\top. \quad (13)$$

Thus, four blocks of variables are active in the model.

Various different covariance matrices $\mathbf{\Sigma}$ are used to assess how correlation structure among predictor variables impacts performance. The theoretical variance of all variables is set to one. Active variables are split into R sets $G_1, \dots, G_R$ such that for each pair of variables $j, \ell \in S, j \neq \ell$,

$$\sigma_{j\ell} = \begin{cases} \rho & \exists r \text{ s.t. } j, \ell \in G_r; \\ 0 & \exists r \neq r' \text{ s.t. } j \in G_r, \ell \in G_{r'}, \end{cases} \quad (14)$$

where $\rho \in \{0.3, 0.6\}$. As a result, variables belonging to different sets G_r and $G_{r'}$ are uncorrelated and variables belonging to the same set have a correlation of ρ .

Previous works in the literature have tested different types of correlation structure for inactive variables, but they have not compared these different types in a single study. For inactive variables (i.e. $b_j = 0$), we test three common configurations:

1. Some inactive variables are correlated (correlation of ρ) with both active variables and other inactive variables (setup seen in Witten et al., 2014)
2. Some inactive variables are correlated (correlation of ρ) with each other but not with active variables (setup seen in Shen, Huang, & Pan, 2012)
3. Inactive variables are uncorrelated with all other variables (setup seen in Sharma, Bondell, & Zhang, 2013)

These three configurations are summarized in Figure 2 for the setting where $\rho = 0.6$. While the true cluster assignments and coefficients are the same for all configurations, the correlation structure, as seen in the heatmaps of the covariance matrices, varies from one configuration to the next.

It is common practice to consider that all inactive variables belong to the same cluster. Moreover, two active variables j and ℓ are assumed to be in the same cluster if their absolute correlation




FIGURE 2 Experimental configurations for the simulation study. For each configuration, the first and second horizontal bars depict the true cluster assignments and coefficients for the $p = 200$ variables. The square box below is a heatmap of the covariance matrix, Σ . The colors for the coefficients and covariance heatmap represent continuous values (see the legend), and the colors for cluster assignments are arbitrary, representing categories, that is, the different clusters $k \in \{1, \dots, K\}$.

TABLE 1 Summary of simulation parameters.

Type	Parameter	Values	Meaning
Varying	Config	{1, 2, 3}	Type of correlation structure
Varying	ρ	{0.3, 0.6}	Correlation between correlated variables
Varying	n	{25, 50}	Number of observations
Fixed	p	200	Number of variables
Fixed	K	5	Number of true variable clusters
Fixed	SNR	10	Model signal-to-noise ratio

is large and their coefficients are similar in absolute value. Therefore, in this experimental set-up, there are $K = R + 1$ true variable clusters C_k , one cluster for each set G_r of active variables and one cluster for the set of inactive variables. The true clusters are the same for each configuration, as seen in Figure 2.

For each configuration, the variance of the model error, σ_ϵ^2 , is fixed such that all configurations have the same model signal-to-noise ratio (SNR). The SNR is defined as

$$\text{SNR} = \frac{\mathbf{b}^\top \Sigma \mathbf{b}}{\sigma_\epsilon^2}. \quad (15)$$

In addition to the true model coefficients and clusters, several parameters are fixed for all simulation settings. These fixed parameters, as well as the three varying parameters, are summarized in Table 1. A total of $3 \times 2 \times 2 = 12$ different datasets, one for each unique combination of simulation parameters, is generated.

For each dataset and method, all but one of the hyperparameters is tuned using the cross-validation strategy explained in Section C.2 (using especially MSE_P), and the average out-of-sample performance is calculated for each value of the remaining hyperparameter. The latter hyperparameter is the one that most impacts model complexity: δ for the CRL and embedded methods and λ for the VC-PCR methods. For methods requiring cluster initializations, performance is further averaged over all random initializations. The exact numerical implementation and other details of the experiments are provided in the Appendix.

For each configuration of the simulated data, sample size n and correlation level ρ , the model error, cluster recovery and support recovery are assessed for models with different complexities s (see Equations 10 and 11). Figure 3 depicts the performance results for the setting where $n = 50$ and $\rho = 0.6$. The average out-of-sample performance is plotted with respect to s , where each point in the lines corresponds to a different value of the fixed hyperparameter (δ or λ). The vertical black line corresponds to the true number of active variables, 20. For VC-PCR-Lasso, the plotted line does not cover the entire domain of model sizes s . This is because VC-PCR-Lasso depends on two sparsity parameters: δ for the Lasso regression used to generate the weight matrix for Weighted SOS-NMF and λ for the sparsity penalty in Weighted SOS-NMF. For the plots depicted in Figure 3, only one of these hyperparameters is fixed, while the other is tuned using cross-validation. Clearly, hyperparameters yielding larger models are never selected.

For the setting where $n = 50$ and $\rho = 0.6$ (Figure 3), VC-PCR-Ridge outperforms the other methods tested for all of the configurations and performance criteria. For VC-PCR-Ridge, the best solutions in terms of clustering and support recovery are models with a complexity s close to the true number of active variables. Moreover, unlike the other two-step methods (i.e., CRL), the model error for VC-PCR-Ridge is robust to the type of configuration. VC-PCR-Lasso demonstrates a similar level of robustness to configuration type, but its performance is slightly inferior to VC-PCR-Ridge. In contrast to these two methods, the model error for VC-PCR-Identity greatly depends on the configuration type. In addition, the clustering and support recovery are much worse. This demonstrates the added-value of using nonuniform weights in Weighted SOS-NMF, the first step of VC-PCR.

As explained in more detail in the Appendix A, the superiority of VC-PCR compared to the presented competing methods—which do not really supervise via clusters—can be explained as follows. CEN (Witten et al., 2014) which simultaneously optimizes the regression coefficients and cluster assignments and is hence computationally attractive, turns out to have severe difficulties to simultaneously maximize the quality of variable selection and variable clustering: as the penalized solution becomes sparser, less importance is also given to the variable clustering task, resulting in weaker clustering performance for sparser models, regardless of the true underlying model complexity.

CRL in all of its versions (Bühlmann et al., 2013) (as a two-step method) is computationally fast but limited by the fact that their clustering step is unsupervised. As a result, clusters may contain both relevant and irrelevant variables, which is undesirable.

Finally, for SRR (Christidis et al., 2020), as one embedded method inherently different from CEN, like for many of the coefficient-grouping methods, cluster assignment is more of a side effect than a principal objective. Moreover, the best solutions in terms of prediction error tend to be complex models with extensive cluster overlap. As a result, variables are assigned to clusters with less certainty, and the clustering performance is less optimal.

The prediction, variable selection, and clustering performance of VC-PCR generally improves as the sample size n or the correlation ρ between correlated variables increases. VC-PCR-Ridge achieves the best variable selection and clustering performance for all configurations, sample




FIGURE 3 Average out-of-sample performance for $n = 50$ and $\rho = 0.6$, plotted with respect to the average model size s . Each point in the plotted lines corresponds to a value of the hyperparameter (HP) δ or λ ; all other HPs were tuned using k -fold cross-validation. The vertical line indicates the true number of active variables, 20.

sizes and correlation levels tested. This top performance occurs for models with a complexity close to the true complexity, $s = 20$ variables. The prediction performance of VC-PCR is most impacted by the correlation ρ between variables, with weaker performance occurring when there is weak correlation between variables (i.e., weak clustering structure). If ρ is sufficiently high (0.6 in this study), VC-PCR-Ridge has better prediction performance than other methods for all sample sizes, configurations and model complexities tested. This suggests using VC-PCR to test hypotheses about the data's cluster structure: if prediction error is not improved using VC-PCR, then the variables analyzed may not contain predictive cluster structure.

4.4 | Application to a typical genomics data set

We also show that our method can be readily applied to genomics data. We consider the `Liver` dataset from Ghazalpour et al. (2006), which was retrieved from <https://horvath.genetics.ucla.edu/html/CoexpressionNetwork/Rpackages/WGCNA/Tutorials/>. The 3600 predictor variables in the original dataset are liver gene expression levels for 135 mice that were subjected to a Western diet over a period of 16 weeks. The response variable is the weight of the mice divided by their length. Given that this dataset has been analyzed many times already, we do not add another deep analysis but rather aim to illustrate how our method can benefit gene analyses in general.

We perform two simple preprocessing steps that are very common in this type of application: variable screening and normalization. First, all genes with missing values are excluded, as well as one outlier observation ("F2_221") and all observations with missing weights or lengths. Among the remaining observations, $n = 50$ are randomly selected for analysis. The normalized predictor matrix containing these observations is denoted $\tilde{\mathbf{X}}$ and the normalized response vector is denoted $\mathbf{y}$. The variables are then further filtered based on Partial Least Squares (PLS) applied to $\tilde{\mathbf{X}}$ and $\mathbf{y}$ with the number of components $m = 5$. The resulting loadings matrix for $\tilde{\mathbf{X}}$, denoted $\mathbf{Q}$ ($p \times m$), is used to assess the potential importance of variables: for each component $r : 1, \dots, m$, the first 40 variables j with the largest absolute loadings, $|q_{jr}|$, are selected as candidate predictors. The variables with the largest loadings in absolute value are the ones that are most important for reconstructing $\tilde{\mathbf{X}}$, therefore, they contain the most important information from the original dataset. For a given component r , the selected variables tend to be highly correlated, whereas for different components, the selected variables tend to be less correlated, although some correlation and selection overlap is observed. In total, $p = 195$ unique predictors are selected for the final matrix $\mathbf{X}$. The sample correlation matrix for $\mathbf{X}$ (presented in absolute value) is shown in Figure 4. Note that PLS is not used for preprocessing; the matrix $\mathbf{X}$ is simply a subset of the original dataset. Of course, the results of the analysis are affected by these preprocessing steps. The goal of this specific preprocessing is simply to mimic what is usually done in the field (namely variable screening) and to look for structures that match the framework discussed in this paper (potentially relevant clusters) so that we can assess how our method performs when the data fit this framework.

For the `Liver` dataset, the true model support and variable clusters are unknown, so it is not possible to calculate the MCC for variable selection and clustering. Figure 5 depicts the average out-of-sample prediction error (MSEP) plotted with respect to the average model complexity (s). As with the simulated data, VC-PCR-Ridge has the best prediction performance for sparse models (about 20 or fewer variables).

Note finally that our analysis presumes a predictive cluster structure. Our method is not suitable for data where such a structure is not at all realistic, or when the preprocessing is not commensurate with such structures. More research seems to be in order in these cases. The results in




FIGURE 4 Heatmap of the absolute sample correlations between the variables in the *Liver* dataset.




FIGURE 5 Average out-of-sample mean squared error of prediction (MSEP) for the *Liver* dataset, plotted with respect to the average model size s . Each point in the plotted lines represents a value of the hyperparameter (HP) δ or λ ; all other HPs were tuned using k -fold cross-validation.

this section simply support our claim that if there are relevant clusters, our method can perform very well.

5 | DISCUSSION AND CONCLUSIONS

This paper presented a method for prediction and supervised variable clustering, VC-PCR. VC-PCR first identifies variable clusters and performs variable selection via Weighted SOS-NMF. The weights used for Weighted SOS-NMF—which encode the prior relationship between the predictors and the response—make it possible to supervise both the clustering and the variable

selection. VC-PCR then uses the cluster membership degrees to generate latent variables, and a linear model is estimated to relate the latent variables to the response.

Unlike known two-step methods, the clustering in VC-PCR is supervised. As a result, the identified clusters tend to be more homogeneous: active and inactive variables are less likely to be mixed together. VC-PCR is computationally more efficient than embedded methods, and it can be easily applied to a variety of tasks (regression, classification, survival analysis, etc.).

Across all considered configurations, VC-PCR with Ridge-regression weights (VC-PCR-Ridge) yielded the best variable selection and clustering performance. For data with predictive cluster structure, especially in sparse cases, VC-PCR also achieved superior predictions. Unlike related methods from the literature, VC-PCR is able to simultaneously improve prediction, variable selection and clustering performance.

A potential direction for further research is the incorporation of prior knowledge, such as known biological pathways. This, however, would require further investigations.

FUNDING INFORMATION

RM was supported by the Belgian Fund for Scientific Research (F.R.S.-FNRS, FRiA grant).

ACKNOWLEDGMENTS

We thank the editors and the reviewers for their insightful comments and suggestions, which helped us to improve the paper. Open Access funding enabled and organized by Projekt DEAL.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are available from the corresponding author upon reasonable request.

ORCID

Johannes Lederer  <https://orcid.org/0000-0002-5369-3053>

REFERENCES

- Bondell, H. D., & Reich, B. J. (2008). Simultaneous regression shrinkage, variable selection, and supervised clustering of predictors with OSCAR. *Biometrics*, 64(1), 115–123.
- Bühlmann, P., Rütimann, P., van de Geer, S., & Zhang, C. H. (2013). Correlated variables in regression: Clustering and sparse estimation. *Journal of Statistical Planning and Inference*, 143(11), 1835–1858.
- Chicco, C. D. (2017). Ten quick tips for machine learning in computational biology. *Biodata Mining*, 10(1), 1–17.
- Chicco, D., & Jurman, G. (2020). The advantages of the Matthews correlation coefficient (MCC) over F1 score and accuracy in binary classification evaluation. *BMC Genomics*, 21(1), 1–13.
- Chicco, D., Tötsch, N., & Jurman, G. (2021). The Matthews correlation coefficient (MCC) is more reliable than balanced accuracy, bookmaker informedness, and markedness in two-class confusion matrix evaluation. *Biodata Mining*, 14(1), 1–22.
- Christidis, A. A., Lakshmanan, L., Smucler, E., & Zamar, R. (2020). Split regularized regression. *Technometrics*, 62(3), 330–338.
- Dai, L. Y., Liu, J. X., Zhu, R., Kong, X. Z., Hou, M. X., & Yuan, S. S. (2018). *Sparse orthogonal nonnegative matrix factorization for identifying differentially expressed genes and clustering tumor samples*. In *2018 IEEE international conference on bioinformatics and biomedicine (BIBM)* (pp. 1332–1337). Madrid: IEEE.
- Daye, Z. J., & Jeng, X. J. (2009). Shrinkage and model selection with correlated variables via weighted fusion. *Computational Statistics & Data Analysis*, 53(4), 1284–1298.
- Ding, C., Li, T., Peng, W., & Park, H. (2006). *Orthogonal nonnegative matrix t-factorizations for clustering*. In *Proceedings of the 12th ACM SIGKDD international conference on knowledge discovery and data mining* (pp. 126–135). New York, NY: Association for Computing Machinery.

- Ding, C. H., Li, T., & Jordan, M. I. (2008). Convex and semi-nonnegative matrix factorizations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(1), 45–55.
- Ghazalpour, A., Doss, S., Zhang, B., Wang, S., Plaisier, C., Castellanos, R. ... Horvarth, S. (2006). Integrating genetic and network analysis to characterize genes related to mouse weight. *PLoS Genetics*, 2(8), e130.
- Kim, H., & Park, H. (2007). Sparse non-negative matrix factorizations via alternating non-negativity-constrained least squares for microarray data analysis. *Bioinformatics*, 23(12), 1495–1502.
- Kim, S., & Xing, E. P. (2009). Statistical estimation of correlated genome associations to a quantitative trait network. *PLoS Genetics*, 5(8), e1000587.
- Lederer, J. (2022). *Fundamentals of high-dimensional statistics: With exercises and R labs Springer Texts in Statistics*. Cham: Springer.
- Ma, S., Song, X., & Huang, J. (2007). Supervised group lasso with applications to microarray data analysis. *BMC Bioinformatics*, 8(1), 1–17.
- Marion, R., Govaerts, B., & von Sachs, R. (2020). AdaCLV for interpretable variable clustering and dimensionality reduction of spectroscopic data. *Chemometrics and Intelligent Laboratory Systems*, 206, 104169.
- Matthews, B. W. (1975). Comparison of the predicted and observed secondary structure of T4 phage lysozyme. *Biochimica et Biophysica Acta (BBA)-Protein Structure*, 405(2), 442–451.
- Park, M. Y., Hastie, T., & Tibshirani, R. (2007). Averaged gene expressions for regression. *Biostatistics*, 8(2), 212–227.
- Sharma, D. B., Bondell, H. D., & Zhang, H. H. (2013). Consistent group identification and variable selection in regression with correlated predictors. *Journal of Computational and Graphical Statistics*, 22(2), 319–340.
- She, Y. (2010). Sparse regression with exact clustering. *Electronic Journal of Statistics*, 4, 1055–1096.
- Shen, X., Huang, H. C., & Pan, W. (2012). Simultaneous supervised clustering and feature selection over a graph. *Biometrika*, 99(4), 899–914.
- Tibshirani, R., Saunders, M., Rosset, S., Zhu, J., & Knight, K. (2005). Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(1), 91–108.
- Vigneau, E. (2016). Dimensionality reduction by clustering of variables while setting aside atypical variables. *Electronic Journal of Applied Statistical Analysis*, 9(1), 134–153.
- Witten, D. M., Shojaie, A., & Zhang, F. (2014). The cluster elastic net for high-dimensional regression with unknown variable grouping. *Technometrics*, 56(1), 112–122.
- Yang, S., Yuan, L., Lai, Y. C., Shen, X., Wonka, P., & Ye, J. (2012). *Feature grouping and selection over an undirected graph*. In *Proceedings of the 18th ACM SIGKDD international conference on knowledge discovery and data mining* (pp. 922–930). Beijing: ACM.
- Yengo, L., Jacques, J., & Biernacki, C. (2014). Variable clustering in high dimensional linear regression models. *Journal de la Société Française de Statistique*, 155(2), 38–56.
- Yuan, M., & Lin, Y. (2006). Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1), 49–67.
- Zass, R., & Shashua, A. (2007). *Nonnegative sparse PCA*. In *Advances in neural information processing systems* (pp. 1561–1568). Vancouver, B.C: Citeseer.

How to cite this article: Marion, R., Lederer, J., Goevarts, B., & Sachs, R. (2024). VC-PCR: A prediction method based on variable selection and clustering. *Statistica Neerlandica*, 1–25. <https://doi.org/10.1111/stan.12358>

APPENDIX A. RELATED WORK

This section presents methods related to the proposed SOS-NMF and VC-PCR methods. In Section A.1, several NMF methods are summarized and compared to SOS-NMF. In Sections A.2–A.4, three different classes of prediction methods are described and compared

to VC-PCR: coefficient-grouping methods (Section A.2), ensemble methods (Section A.3) and two-step methods (Section A.4).

A.1 NMF methods

SOS-NMF is most closely related to the problem of NMF. While certain methods exist for addressing aspects of the SOS-NMF problem, no method satisfactorily addresses all aspects simultaneously (i.e., sparsity, orthogonality and semi-nonnegativity). This section explores the relationship between SOS-NMF and several closely related NMF methods from the literature. Comparisons to other NMF methods (Semi-NMF, Ding, Li, & Jordan, 2008; Orthogonal NMF Ding, Li, Peng, & Park, 2006; and Sparse NMF (Kim & Park, 2007)) can be found in the Appendix. The objective functions and constraints associated with the NMF methods studied in this paper are summarized in Table A1.

Sparse Orthogonal NMF (SONMF) (Dai et al., 2018) constrains $\mathbf{V}$ to be sparse, nonnegative and orthogonal. In contrast to SOS-NMF, however, the input $\mathbf{Z}$ to SONMF must be nonnegative, making SONMF less flexible. Moreover, the scale of the columns of $\mathbf{V}$ is constrained, rather than the scale of the columns of $\mathbf{U}$. These scaling differences have important implications for the estimation of $\mathbf{V}$. For SOS-NMF, the solution for $\mathbf{v}_j$ can be solved independently for each variable j because the scale of each vector $\mathbf{v}_k$ is not fixed. For SONMF, the solution for v_{jk} depends on all other entries in $\mathbf{V}$ due to the orthonormality constraint. As a result, the solution for each vector $\mathbf{v}_j$ cannot be found independently, and the scale of v_{jk} for variables $j \in C_k$ depends on the size of cluster k . The more variables that are assigned to cluster k , the more likely it is that these variables will have low memberships v_{jk} , to satisfy the constraint $\mathbf{v}_k^T \mathbf{v}_k = 1$. This unfairly penalizes large clusters, as the sparsity penalty makes small values v_{jk} shrink to zero.

SOS-NMF is also related to Nonnegative Sparse Principal Component Analysis (Zass & Shashua, 2007). In contrast to SOS-NMF, the columns of $\mathbf{V}$ in NSPCA are only quasi-orthogonal, unless $\gamma = \infty$. This means that variables may be assigned to multiple clusters (i.e., $\mathbf{v}_j$ may have multiple nonzero entries). For the special case where $\gamma = \infty$, NSPCA differs from SOS-NMF with respect to the scaling constraints on $\mathbf{U}$ and $\mathbf{V}$: SOS-NMF constrains the scale of $\mathbf{U}$ whereas NSPCA constrains the scale of $\mathbf{V}$ (see the Appendix for more details). This difference in scaling has important practical implications. By constraining the scale of $\mathbf{V}$ and encouraging its sparsity, NSPCA unfairly penalizes large clusters (see the explanation above). Moreover, to the best of our knowledge, no optimization strategies have been proposed for solving the NSPCA problem when $\gamma = \infty$, referred to here as Orthonormal NSPCA.

Although it was not originally presented as an NMF method, Clustering around Latent Variables $K + 1$ (CLV $K + 1$, local mode) (Vigneau, 2016) can be rewritten as an NMF problem that is similar to SOS-NMF. In contrast to SOS-NMF, the cluster memberships v_{jk} are constrained to be binary, making them less informative. In addition, the sparsity penalty is weighted by the standard-deviation of the input variables. As a result, cluster assignments only depend on the correlation between the latent variables and the input variables. This means that modifying the scale of the input variables, as is proposed for clustering-part of VC-PCR (see Section 3.2), does not affect the relative importance of variables during clustering. Therefore, unlike for SOS-NMF, prior weighting cannot be used to supervise the variable clustering.

A.2 Coefficient-grouping methods

Most of the supervised variable clustering methods in the literature rely on a regression framework, identifying variable clusters by pushing regression coefficients towards each other in

TABLE A1 Summary of nonnegative matrix factorization (NMF) methods.

Method	Objective function	Constraints
Semi-NMF (Ding et al., 2008)	$\frac{1}{2(n-1)} \ \mathbf{X} - \mathbf{UV}^T \ ^2_F$	$v_{jk} \geq 0, \forall j, k$
Orthogonal NMF (Ding et al., 2006)	$\frac{1}{2(n-1)} \ \mathbf{Z} - \mathbf{UV}^T \ ^2_F$	$v_{jk} \geq 0, \forall j, k$ $u_{ik} \geq 0, \forall i, k$ $\mathbf{V}^T \mathbf{V} = \mathbf{I}_K$
Sparse NMF (Kim & Park, 2007)	$\frac{1}{2(n-1)} \ \mathbf{Z} - \mathbf{UV}^T \ ^2_F + \delta \sum_{i=1}^n \ \mathbf{u}_i \ ^2_2 + \lambda \sum_{j=1}^p \ \mathbf{v}_j \ ^2_1$	$v_{jk} \geq 0, \forall j, k$ $u_{ik} \geq 0, \forall i, k$
Sparse orthogonal NMF (Dai et al., 2018)	$\frac{1}{2(n-1)} \ \mathbf{Z} - \mathbf{UV}^T \ ^2_F + \lambda \sum_{j=1}^p \ \mathbf{v}_j \ _1$	$v_{jk} \geq 0, \forall j, k$ $u_{ik} \geq 0, \forall i, k$ $\mathbf{V}^T \mathbf{V} = \mathbf{I}_K$
Nonnegative Sparse PCA (Zass & Shashua, 2007)	$\frac{1}{2(n-1)} \ \mathbf{XV} \ ^2_F - \gamma \sum_{j=1}^p \ \mathbf{v}_j \ _1 - \gamma \ \mathbf{V}^T \mathbf{V} - \mathbf{I}_K \ ^2_F$	$v_{jk} \geq 0, \forall j, k$
Clustering around Latent variables $K + 1$ (Vigneau, 2016)	$\frac{1}{2(n-1)} \ \mathbf{X} - \mathbf{UV}^T \ ^2_F + \lambda \sum_{j=1}^p \text{sd}(\mathbf{x}_j) \ \mathbf{v}_j \ _1$	$\mathbf{V} \in \{0, 1\}^{p \times K}$ $\mathbf{V}$ has orthogonal columns $\text{var}(\mathbf{u}_k) = 1$
Sparse orthogonal Semi-NMF	$\frac{1}{2(n-1)} \ \mathbf{X} - \mathbf{UV}^T \ ^2_F + \lambda \sum_{j=1}^p \ \mathbf{v}_j \ _1$	$v_{jk} \geq 0, \forall j, k$ $\mathbf{V}$ has orthogonal columns $\text{var}(\mathbf{u}_k) = 1$

Note: Method proposed in this paper.

T A B L E A2 Summary of coefficient-grouping methods.

Method	HPs	Sparsity term	Grouping term	b_j, b_ℓ grouping iff	Grouping type
Flasso Tibshirani, Saunders, Rosset, Zhu, and Knight (2005)	2	$\delta \sum_{j=1}^p b_j $	$\lambda \sum_{j=2}^p b_j - b_{j-1} $	Adjacent	NonAbsVal
CLasso She (2010)	2	$\delta \sum_{j=1}^p b_j $	$\lambda \sum_{1 \leq j < \ell \leq p} b_j - b_\ell $		NonAbsVal
OSCAR Bondell and Reich (2008)	2	$\delta \sum_{j=1}^p b_j $	$\lambda \sum_{1 \leq j < \ell \leq p} \frac{1}{2} b_j - b_\ell + \frac{1}{2} b_j + b_\ell $		AbsVal
PACS Sharma et al. (2013)	2+	$\lambda \sum_{j=1}^p \omega_j b_j $	$\lambda \sum_{1 \leq j < \ell \leq p} \omega_{j\ell}^{(+)} b_j - b_\ell + \sum_{1 \leq j < \ell \leq p} \omega_{j\ell}^{(-)} b_j + b_\ell $		AbsVal
WF Daye and Jeng (2009)	3	$\lambda \sum_{j=1}^p b_j $	$\lambda \sum_{1 \leq j < \ell \leq p} \omega_{j\ell} (b_j - \text{sign}(r_{j\ell}) b_\ell)^2$		AbsVal
SSCFS Shen et al. (2012)	3+	$\delta \sum_{j=1}^p P_\tau(b_j)$	$\lambda \sum_{i,\ell \in \mathcal{E}} P_\tau(b_j - b_\ell)$	Shared edge	NonAbsVal
GFLasso Kim and Xing (2009)	2+	$\delta \sum_{j=1}^p b_j $	$\lambda \sum_{i,\ell \in \mathcal{E}} b_j - \text{sign}(r_{j\ell}) b_\ell $	Shared edge and $r_{j\ell} > 0$	AbsVal
NCFGs Yang et al. (2012)	2+	$\delta \sum_{j=1}^p b_j $	$\lambda \sum_{j,\ell \in \mathcal{E}} b_j - b_\ell $	Shared edge	AbsVal
NCTFGs Yang et al. (2012)	3+	$\delta \sum_{j=1}^p P_\tau(b_j)$	$\lambda \sum_{i,\ell \in \mathcal{E}} P_\tau(b_j - b_\ell)$	Shared edge	AbsVal
CEN Witten et al. (2014)	3	$\delta \sum_{j=1}^p b_j $	$\lambda \sum_{k=1}^K \sum_{j \in C_k} \mathbf{x}_j b_j - \frac{1}{ C_k } \sum_{\ell \in C_k} \mathbf{x}_\ell b_\ell _2^2$	$ r_{j\ell} \rightarrow 1$ and $j, \ell \in C_k$	AbsVal

Notes: $r_{j\ell} = \text{cor}(\mathbf{x}_j, \mathbf{x}_\ell)$; $\mathcal{E}$ is the set of edges in a given graph; $P_\tau(x) = \min(\frac{x}{\tau}, 1)$ for some $\tau > 0$; $\omega_j, \omega_{j\ell}^{(+)}$ and $\omega_{j\ell}^{(-)}$ are weights that depend on the chosen PACS approach; $\omega_{j\ell} = \frac{|r_{j\ell}|^\gamma}{1 - |r_{j\ell}|}$ for some $\gamma > 0$; C_k is the set of variables assigned to cluster k . The grouping type for methods that push coefficients toward each other in absolute value (resp. real value) is “AbsVal” (resp. “NonAbsVal”). For methods involving graphs, the number of hyperparameters (HPs) is presented with a + to indicate that more HPs may be necessary for generating the required graph.

(absolute) value. These *coefficient-grouping methods* (which are listed in Table A2 for the sake of completeness) generally optimize a least-squares criterion with two penalties: a sparsity penalty, which encourages null coefficients, and a grouping penalty, which encourages coefficients for certain variables to take the same (absolute) value. The objective function for these methods usually takes the form

$$J(\mathbf{b}) = \frac{1}{2} \|\mathbf{y} - \mathbf{X}\mathbf{b}\|_2^2 + \text{Penalty}_S(\mathbf{b}; \delta) + \text{Penalty}_G(\mathbf{b}; \lambda), \quad (\text{A1})$$

where $\lambda \geq 0$ and $\delta \geq 0$ are hyperparameters, $\text{Penalty}_S(\mathbf{b})$ is a sparsity penalty and $\text{Penalty}_G(\mathbf{b})$ is a coefficient-grouping penalty. With the exception of CEN (Witten et al., 2014), these methods identify variable clusters after the vector of coefficients $\mathbf{b}$ is estimated: all variables with the same estimated coefficient (in absolute or real value, depending on the method) are assigned to the same cluster.

Some of the coefficient-grouping methods rely on prior graphs. If the graph is not known in advance, a fully connected graph containing all $p(p-1)/2$ possible edges may be used. However, this strategy can be inefficient, even for relatively small p , as the computational time can increase exponentially with the number of edges (Yang et al., 2012). Nongraph-based methods that penalize sums over all pairs of variables also suffer from efficiency problems. Thus, computation time is one of the major challenges for coefficient-grouping methods.

While in the same spirit as other coefficient-grouping methods, CEN (Witten et al., 2014) approaches the problem of supervised variable clustering in a unique way. The grouping penalty for CEN contains the objective function for K-means, where the points to cluster are the variables weighted by their respective regression coefficients. At each iteration of the CEN algorithm, pairs of variables in the same cluster have their coefficients pushed toward each other in absolute value, especially if they have high absolute correlation.

CEN differs from the other methods with respect to two important aspects. First, the grouping penalty does not require summation over pairs of variables, simplifying the calculations necessary. Second, CEN simultaneously optimizes the regression coefficients and cluster assignments, rather than defining the clusters after estimating the coefficients.

Let the variable partition $\mathcal{C} := \{C_1, \dots, C_K\}$. The objective function for CEN is

$$\begin{aligned} J(\mathbf{b}, \mathcal{C}) &= \frac{1}{2} \|\mathbf{y} - \mathbf{X}\mathbf{b}\|_2^2 + \text{Penalty}_S(\mathbf{b}; \delta) + \text{Penalty}_G(\mathbf{b}, \mathcal{C}, \mathbf{X}; \lambda) \\ &= \frac{1}{2} \|\mathbf{y} - \mathbf{X}\mathbf{b}\|_2^2 + \delta \sum_{j=1}^p |b_j| + \lambda \sum_{k=1}^K \sum_{j \in C_k} \|\mathbf{x}_j b_j - \underbrace{\frac{1}{p_k} \sum_{\ell \in C_k} \mathbf{x}_\ell b_\ell}_{\text{centroid } k}\|_2^2, \end{aligned} \quad (\text{A2})$$

where p_k is the number of variables in cluster k .

For fixed $\mathbf{b}$, the CEN problem reduces to minimizing the grouping penalty function with respect to $\mathcal{C}$. This is equivalent to clustering the coefficient-weighted variables $\mathbf{x}_1 b_1, \dots, \mathbf{x}_p b_p$ using K-means. As we will demonstrate below, this weighted K-means problem is closely related to our proposed Weighted SOS-NMF problem. Let $\mathbf{H}$ ($p \times K$) be a binary matrix with orthogonal columns, and let $C_k := \{j \mid h_{jk} = 1\}$. Given a vector $\mathbf{b}$ of regression coefficients, the CEN solution for $C_1, \dots, C_K$ can also be found by estimating $\mathbf{H}$ and inferring the clusters from it. The solution

for $\mathbf{H}$ is given by

$$\begin{aligned} \hat{\mathbf{H}} &= \arg \min_{\mathbf{H}} \|\mathbf{X} \text{diag}(\mathbf{b}) - \mathbf{X} \text{diag}(\mathbf{b}) \mathbf{H} (\mathbf{H}^\top \mathbf{H})^{-1} \mathbf{H}^\top\|_2^2 \\ \text{s.t. } \mathbf{H} &\in \{0, 1\}^{p \times K} \\ \text{s.t. } \mathbf{H} &\text{ has orthogonal columns,} \end{aligned} \quad (\text{A3})$$

where $\text{diag}(\mathbf{b})$ is a diagonal matrix whose diagonal is composed of the vector $\mathbf{b}$ of coefficients. The CEN solution for $C_1, \dots, C_K$ can be found by setting $\hat{C}_k := \{j \mid \hat{h}_{jk} = 1\}, \forall k$.

Clearly, the problem in Equation (A3) is similar to the Weighted SOS-NMF problem, with a few exceptions: (i) the cluster membership matrix $\mathbf{H}$ is binary and (ii) the vector $\mathbf{b}$ used to weight the variables in $\mathbf{X}$ is the minimizer of Equation (A2) for fixed C , not a fixed input. For fixed C and $b_\ell, \ell \neq j$ and $j \in C_k$, the minimizer $\hat{b}_j$ of Equation (A2) is

$$\hat{b}_j = \frac{\text{Soft}(\mathbf{e}_j^\top \mathbf{x}_j + \frac{\lambda}{p_k} \sum_{\ell \in C_k, \ell \neq j} b_\ell \mathbf{x}_j^\top \mathbf{x}_\ell, \delta)}{\mathbf{x}_j^\top \mathbf{x}_j (1 + \frac{\lambda(p_k - 1)}{p_k})}, \quad (\text{A4})$$

where $\mathbf{e}_j = \mathbf{y} - \sum_{\ell \neq j} \mathbf{x}_\ell b_\ell$ and $\text{Soft}(z, \delta) = \text{sign}(z)(|z| - \delta)_+$.

As λ decreases, the vector $\mathbf{b}$ becomes sparser. However, less importance is also given to the variable clustering task, resulting in weaker clustering performance. Therefore, unlike Weighted SOS-NMF, CEN generally achieves better clustering performance for less sparse models, regardless of the true underlying model complexity. This makes it difficult for CEN to simultaneously maximize the quality of variable selection and variable clustering.

A.3 Ensemble methods and SRR

Another approach to the supervised variable clustering problem is to simultaneously estimate K models, one for each cluster, then predict the target based on the average coefficient vector across clusters. SRR (Christidis et al., 2020) takes this “ensemble-method” approach. Rather than estimating a single vector of coefficients, SRR estimates multiple sparse coefficient vectors and encourages these vectors to be as diverse as possible. Although SRR does not explicitly search for variable clusters, they may be inferred from the coefficient vectors: a variable belongs to a given cluster if its coefficient in the corresponding coefficient vector is nonzero. This approach can identify overlapping clusters and does not force coefficients of variables in the same cluster to tend towards the same value.

The objective function for SRR is

$$\begin{aligned} J(\mathbf{b}_1, \dots, \mathbf{b}_K) &= \sum_{k=1}^K \left\{ \frac{1}{2n} \|\mathbf{y} - \mathbf{X} \mathbf{b}_k\|_2^2 + \delta \left[\alpha \sum_{j=1}^p |b_{jk}| + (1 - \alpha) \sum_{j=1}^p b_{jk}^2 \right] + \right. \\ &\quad \left. + \lambda \sum_{g \neq k} \sum_{j=1}^p |b_{jk}| |b_{jg}| \right\}. \end{aligned} \quad (\text{A5})$$

The matrix $\mathbf{B} = [\mathbf{b}_1 \dots \mathbf{b}_K]$ is essentially a cluster membership matrix, where variable j belongs to cluster k if $b_{jk} \neq 0$. Variables can belong to multiple clusters, but “hard” clusters $C_1, \dots, C_K$ can also be defined such that $C_k := \{j \mid k = \arg \max_{\ell} |b_{j\ell}|\}$. Maximal diversity is

achieved when the rows of $\mathbf{B}$ contain only one nonzero element and thus each variable belongs to a single cluster (i.e., $|b_{jk}||b_{jg}| = 0 \forall j, k, g$).

The final vector of regression coefficients used for prediction is an average across all vectors $\mathbf{b}_k$:

$$\bar{\mathbf{b}} = \frac{1}{K} \sum_{k=1}^K \mathbf{b}_k. \quad (\text{A6})$$

Like many of the coefficient-grouping methods, cluster assignment is more of a side effect than a principal objective of SRR. Moreover, the best solutions in terms of prediction error tend to be complex models with extensive cluster overlap. As a result, variables are assigned to clusters with less certainty, and the clustering performance is less optimal.

A.4 Two-step methods

The methods described in Sections A.2 and A.3 are *embedded* methods because model estimation and clustering are performed simultaneously. As a result, the tasks of clustering and modeling are inherently dependent. For two-step methods, these tasks are separated: first, clusters are identified using an unsupervised clustering technique (Step 1), then a model is estimated based on the outputs of the clustering algorithm (Step 2).

One advantage of two-step methods is their computational efficiency compared to embedded methods. It is also easier to adapt them to different types of prediction tasks, such as regression, classification and survival analysis. Indeed, the outputs of the clustering algorithm can simply be used as inputs to existing prediction algorithms. However, unlike VC-PCR, most two-step methods perform clustering without any supervision by the target. As a result, they cannot distinguish variables that are relevant for prediction from irrelevant ones. This is a significant limitation for these methods, as relevant and irrelevant variables may be mixed into the same cluster, thereby diminishing subsequent prediction performance.

Different unsupervised clustering algorithms have been proposed for the first step of the two-step approach, the most common being K-means (used by e.g., Ma et al., 2007) and HAC (used by e.g. Bühlmann et al., 2013; Park et al., 2007). A variety of different strategies for the second step have also been proposed. These strategies are summarized below.

For CRL (Bühlmann et al., 2013; Park et al., 2007), Lasso is used to estimate a sparse model in the second step, where the cluster centroids $\hat{\mathbf{m}}_1, \dots, \hat{\mathbf{m}}_K$ from first step are the predictors. The Lasso solution is given by

$$\arg \min_{\mathbf{a}} L(\hat{\mathbf{M}}, \mathbf{y}; \mathbf{a}) + \delta \|\mathbf{a}\|_1, \quad (\text{A7})$$

where $L(\hat{\mathbf{M}}, \mathbf{y}; \mathbf{a})$ is a loss function, $\mathbf{a} (K \times 1)$ contains the model coefficients and δ is a nonnegative hyperparameter for the Lasso penalty.

Performing regression or classification based on the K cluster centroids, rather than the p variables in $\mathbf{X}$, can be advantageous because K is typically much smaller than p . This decreases the dimensionality of the problem, making it less prone to the curse of dimensionality.

For Cluster Group Lasso (CGL) (Bühlmann et al., 2013), Group Lasso (Yuan & Lin, 2006) is used in the second step. The variables in $\mathbf{X}$, rather the cluster centroids, are used as predictors, and the clusters $\hat{C}_1, \dots, \hat{C}_K$ are used to define the groups used in the grouping penalty. The Group

Lasso solution is given by

$$\arg \min_{\mathbf{b}} L(\mathbf{y}, \mathbf{X}, \mathbf{b}) + \lambda \sum_{k=1}^K \sqrt{\sum_{j \in \hat{C}_k} b_j^2}, \quad (\text{A8})$$

where $L(\mathbf{y}, \mathbf{X}, \mathbf{b})$ is a loss function, $\mathbf{b}$ ($p \times 1$) contains the model coefficients, λ is a nonnegative hyperparameter for the Group Lasso penalty and $\hat{C}_k$ is the k^{th} cluster identified in the first step. For each cluster $\hat{C}_k$, the Group Lasso penalty results in either the selection of all of the variables in the cluster (i.e., they all have nonzero coefficients) or none of them (i.e., they all have null coefficients).

An alternative to CRL and CGL is to estimate independent Lasso models for each cluster identified in Step 1. This strategy, which allows for feature selection within clusters, is used for Supervised Group Lasso (SGLasso) (Ma et al., 2007). SGLasso also performs a third step where these sparse clusters are used to define the groups in Group Lasso (Yuan & Lin, 2006). The Group Lasso step is necessary for selecting the clusters that are important for predicting the response.

All of the previously mentioned methods are limited by the fact that their clustering step is unsupervised. As a result, clusters may contain both relevant and irrelevant variables, which is undesirable.

APPENDIX B. ALGORITHM AND HYPERPARAMETERS

This section summarizes the VC-PCR algorithm and offers guidance with respect to the choice of the hyperparameters λ and K , as well as the weight matrix $\mathbf{W}$.

B.1 Algorithm

The entire VC-PCR algorithm is summarized in Figure B1. At the beginning of the algorithm, $\mathbf{V}$ is initialized based on an initial partition $C_1, \dots, C_K$ of K variable clusters, where $v_{jk} = \mathbb{I}[j \in C_k]$.

B.2 Choice of weights

As mentioned in Section 3.2, the entries w_{jj} in $\mathbf{W}$ should indicate (i) the prior relationship between variable j and the target $\mathbf{y}$ (e.g., positive vs. negative association) and (ii) the importance of variable j in the prediction task. We propose two potential weighting schemes based on the coefficients of Lasso and Ridge models. Given $\mathbf{X}$, $\mathbf{y}$ and a hyperparameter δ , a vector of coefficients $\mathbf{w}$ is estimated by solving

$$\arg \min_{\mathbf{w}} L(\mathbf{X}, \mathbf{y}; \mathbf{w}) + \text{Penalty}(\mathbf{w}; \delta), \quad (\text{B1})$$

where $L(\mathbf{X}, \mathbf{y}; \mathbf{w})$ is a loss function and $\text{Penalty}(\mathbf{w}; \delta)$ is a penalty function. For Lasso coefficients, this penalty is $\delta \sum_{j=1}^p |w_j|$ and for Ridge coefficients, the penalty is $\delta \sum_{j=1}^p w_j^2$.

When Lasso weights are used, prior variable selection is performed, as variables with Lasso coefficients equal to zero are essentially excluded from the variable clustering task. The remaining variables are weighted by a value that reflects their relative importance in the prediction task, as well as the direction of their correlation with the target. Ridge weights encode similar information but do not preselect variables: in most cases, all variables are given a nonzero weight. However, variables with high absolute correlation are more likely to have similar weights in absolute value. This encourages SOS-NMF to cluster these variables together.

Require: A predictor matrix $\mathbf{X}$ ($n \times p$) with mean-centered, unit-variance columns, a target vector $\mathbf{y}$ ($n \times 1$) (binary for classification, centered and scaled for regression), a diagonal weight matrix $\mathbf{W}$ ($p \times p$), an initial partition $C_1, \dots, C_K$ and a hyperparameter $\lambda > 0$

Ensure: A sparse cluster membership matrix $\hat{\mathbf{V}}$, a latent variable matrix $\hat{\mathbf{M}}$ and model coefficients $\hat{\mathbf{a}}$

► Initialization

```

1: for each cluster  $k \in 1, \dots, K$  do
2:   for each variable  $j \in 1, \dots, p$  do
3:      $v_{jk} = \mathbb{I}[j \in C_k]$ ;
4:   end for each
5: end for each

```

► Cluster memberships

```

6: while not converged do

```

► Update latent variable matrix

```

7:   if there are columns of zeroes in  $\mathbf{V}$  then
8:     remove columns of zeroes from  $\mathbf{V}$ ;
9:   end if
10:   $\hat{\mathbf{U}} = \mathbf{X}\mathbf{W}\mathbf{V}(\mathbf{V}^T\mathbf{V})^{-1}$ ;
11:  for each cluster  $k$  in  $1, \dots, K$  do
12:     $\mathbf{u}_k = \hat{\mathbf{u}}_k / \text{sd}(\hat{\mathbf{u}}_k)$ ;
13:  end for each

```

► Update cluster membership matrix

```

14:  for each variable  $j$  in  $1, \dots, p$  do
15:     $\hat{k} = \arg \max_k w_{jk} \text{cor}(\mathbf{u}_k, \mathbf{x}_j)$ ;
16:     $v_{j\hat{k}} = (w_{j\hat{k}} \text{cor}(\mathbf{u}_{\hat{k}}, \mathbf{x}_j) - \lambda)_+$ ;
17:  end for each
18: end while
19:  $\hat{\mathbf{V}} = \mathbf{V}$ 

```

► OLS

```

20:  $\hat{\mathbf{M}} = \mathbf{X}\hat{\mathbf{V}}$ ;

```

```

21: Calculate  $\hat{\mathbf{a}}$  by solving Eq. (7);

```

FIGURE B1 Variable cluster principal component regression (VC-PCR).

BB.3 Choice of hyperparameters

The hyperparameters are discussed later in Section C.2.

APPENDIX C. DETAILS ON THE EXPERIMENTS

C.1 Implementation

The experiments were programmed in R. Weighted SOS-NMF, VC-PCR, and SRR were implemented using in-house code. CEN was performed using code provided by the authors. Lasso and Ridge regression was performed using the function *glmnet* from the *glmnet* package. For the CRL methods, Ward's hierarchical clustering and K-means clustering were performed using the functions *hclust* and *kmeans* from the *stats* package. Hierarchical clustering based on canonical correlation was performed using the function *hclustvar* from the *ClustOfVar* package.

Of the methods compared, CRL with K-means clustering (CRL-kmeans) and hierarchical clustering based on Ward's criterion (CRL-hclustCor) were the fastest methods (faster than VC-PCR by a factor of 16 and 7, respectively). This is likely due to the fact that the *kmeans* function from the *stats* package is partly coded in C and the *hclust* function is partly coded in Fortran. Compared to VC-PCR, CEN was four times slower, CRL with hierarchical clustering based on canonical correlation (CRL-hclustCC) was seven times slower and SRR was several 100 times slower.

C.2 Hyperparameter selection

The methods compared in this study depend on the choice of several hyperparameters. The VC-PCR algorithm, as well as the embedded method SRR, depends on three hyperparameters: K , the number of clusters in the initial partitions, δ , the hyperparameter used to estimate the weights in $\mathbf{W}$, and λ , which determines the degree of sparsity in $\mathbf{V}$. In practice, these hyperparameters are tuned using k -fold cross-validation, using a criterion such as the MSEP for regression problems or the Matthews Correlation Coefficient (MCC) (Matthews, 1975) for classification problems. For larger data, k -fold cross-validation might become computationally challenging. We would then expect methods such as AV (Lederer, 2022, Chap. 4) to be more appropriate. We now discuss the selection of these hyperparameters in a bit more detail.

For the sparsity hyperparameter λ , the smallest value leading to $\mathbf{V} = \mathbf{0}$ can be determined based on the first iteration of the while loop in Figure B1. Let $\mathbf{u}_k^{(1)}$ be the k th latent variable at the first iteration. All elements v_{jk} are equal to zero if λ is greater than or equal to $w_{jj}\text{cor}(\mathbf{u}_k^{(1)}, \mathbf{x}_j)$, for all j, k . Therefore, the largest value of λ to consider is

$$\lambda_{\max} = \max_{j,k} w_{jj}\text{cor}(\mathbf{u}_k^{(1)}, \mathbf{x}_j). \quad (\text{C1})$$

For SRR, the fourth hyperparameter α is fixed to 1, resulting in a Lasso penalty (see Equation A5).

For the hyperparameters δ and λ , 10 values are tested for each method. For the simulated dataset, three values of K are tested: $\{4, 5, 6\}$ (the true value is 5). For the real dataset, five values of K are tested: $\{10, 20, 40, 50, 60\}$. Some methods, such as CRL-kmeans, CEN and VC-PCR, require an initial set of clusters for each value of K . Five random initializations are tested for these methods, and the resulting performance is subsequently averaged (see Section S1).

Concerning (k -fold) cross-validation to select hyperparameters via evaluation of the average out-of-sample performance, we note that this involves k -fold cross-validation at two levels: an inner and an outer loop. The dataset is split into 10 outer folds. Then, for each iteration m of the outer loop, the samples in outer fold m are set aside and fivefold cross-validation (the nested inner loop) is performed on the remaining samples. The latter samples are normalized before analysis: the target and predictors are centered and scaled to unit variance. The hyperparameters yielding the model with the smallest average MSEP across the inner folds are then selected and a final model is trained using all samples in the inner folds, which are also normalized. The performance of this model is evaluated using the samples in outer fold m , which were not used during model selection or estimation. These samples are normalized using the same parameters used to normalize the samples that trained the model. After iterating through all of the outer folds, the average out-of-sample performance is then calculated.