Supervised dimension reduction for ordinal predictors

Liliana Forzani, Pamela Llop, Diego Tomassi

Facultad de Ingeniería Química, Universidad Nacional del Litoral

Researchers of CONICET

and

Rodrigo García Arancibia

Instituto de Economía Aplicada Litoral (FCE-UNL)

Abstract

In applications involving ordinal predictors, common approaches to reduce dimensionality are either extensions of unsupervised techniques such as principal component analysis, or variable selection procedures that rely on modeling the regression function. In this paper, a supervised dimension reduction method tailored to ordered categorical predictors is introduced. It uses a model-based dimension reduction approach, inspired by extending sufficient dimension reductions to the context of latent Gaussian variables. The reduction is chosen without modeling the response as a function of the predictors and does not impose any distributional assumption on the response or on the response given the predictors. A likelihood-based estimator of the reduction is derived and an iterative expectation-maximization type algorithm is proposed to alleviate the computational load and thus make the method more practical. A regularized estimator, which simultaneously achieves variable selection and dimension reduction, is also presented. Performance of the proposed method is evaluated through simulations and a real data example for socioeconomic index construction, comparing favorably to widespread use techniques.

1. Introduction.

Regression models with ordinal predictors are common in many applications. For instance, in economics and the social sciences, ordinal variables are used to predict phenomena like income distribution, poverty, consumption patterns, nutrition, fertility, healthcare decisions, and subjective well-being, among others [e.g. 3, 44, 39, 27, 36, 20]. In marketing research, customer preferences are used to create automatic recommendation systems, as in the case of Netflix [e.g. 2, 43], where the ratings for unseen movies can be predicted using the user's previous ratings and information about the consumer preferences for the whole database.

In this context, when the number of predictors is large, it is of interest to reduce the dimensionality of the space by combining them into a few variables in order to get efficiency in the estimation as well as an understanding of the model. The commonly used dimension reduction techniques for ordinal variables are adaptations of standard principal component analysis (PCA)

[34, 28]. For example, in the case of the *Índice de Focalización de Pobreza* (a socio-economic index commonly used in Latin America), the first normalized principal component is used to predict poverty status, even if this outcome variable was never used to estimate the scaling. It is clear, however, that ignoring the response when building such an index can lead to a loss of predictive power compared to the full set of predictors.

A different approach to dimensionality reduction is to perform variable selection on the original set of predictors. A method adapted to ordinal predictors is proposed in [21]. Despite the fact that this method uses information from the response to achieve variable selection, it performs simultaneously regression modeling by assuming a parametric model for the response as a function of predictors.

For regression and classification tasks it is widely accepted that supervised dimension reduction is a better alternative than PCA-like approaches. Sufficient dimension reduction (SDR), in particular, has gained interest in recent years as a principled methodology to achieve dimension reduction on the predictors $\mathbf{X} \in \mathbb{R}^p$ without losing information about the response Y. Formally, for the regression of $Y|\mathbf{X}$, SDR amounts to finding a transformation $R(\mathbf{X}) \in \mathbb{R}^d$, with $d \leq p$, such that the conditional distribution of $Y|\mathbf{X}$ is identical to that of $Y|R(\mathbf{X})$. Nevertheless, there is no need to assume a distribution for Y or for $Y|\mathbf{X}$. Thus, the obtained reductions can subsequently be used with any prediction rule. Moreover, when the reduced space has low dimension, it is feasible to plot the response versus the reduced variables. This can play an important role in facilitating model building and understanding [8, 9].

Most of the methodology in SDR is based on the inverse regression of $\mathbf{X}$ on Y, which translates a p-dimensional problem of regressing $Y|\mathbf{X}$ into p (easier to model) one-dimensional regressions corresponding to $\mathbf{X}|Y$. Estimation in SDR was developed originally for continuous predictors and was based on the moments of the conditional distribution of $\mathbf{X}|Y$ (SIR, [33]; SAVE, [16]; pHd, [32]; PIR [4]; MAVE, [49, 31, 15, 51, 14]; DR, [30], see also [13, 17, 7, 15, 50] and [9], where much of its terminology was introduced). Later, [10] introduced the so called *model-based* inverse regression of $\mathbf{X}|Y$ (see also [11, 12]). The main advantage of this approach is that provides an estimator of the sufficient reduction that contains all the information in $\mathbf{X}$ that is relevant to Y, allowing maximum likelihood estimators which are optimal in terms of efficiency and $\sqrt{n}$-consistent under mild conditions when the model holds.

Along the lines of the model-based SDR approach, the up to date methodology is for predictors belonging to a general exponential family of distributions (See [5]). Then, when attempting to apply SDR to ordinal predictors, a first approach could be to treat them as polythomic variables, ignoring their natural order. Then a multinomial distribution can be postulated over them, which can be treated as member of the exponential family. However, ordered variables usually do not follow a multinomial distribution and the order information is lost when treating them as multinomial. There have been attempts to use dummy variables to deal with ordinal data, but this procedure has been shown to introduce spurious correlations [28]. Another approach is to treat the ordered predictors as a discretization of some underlying continuous random variable. This technique is commonly used in the social sciences and is known as the *latent variable model*. In this context, the latent variables are usually modeled as normally or logistically distributed, obtaining the so-called ordered probit and logit models, respectively [22, 35]. While for each scientific phenomenon the latent variable can take a particular meaning (e.g., utility in economic choice problems, liability to diseases in genetics, or tolerance to a drug in toxicology), a general

interpretation of a latent variable may be the *propensity* to observe a certain value j of an ordered categorical variable [45]. Regardless of their philosophical meaning and the criticisms about their real existence, latent variables are very useful for generating distributions for modeling, hence their widespread use.

In this paper, we develop a supervised dimension reduction method for ordinal predictors, based on the SDR for the regression of the response given the underlying normal latent variables. Under this context, we present a maximum likelihood estimator of the reduction and we propose an approximate expectation maximization (EM) algorithm for its practical computation, which is close to recent developments in graphical models for ordinal data [23] and allows for computationally efficient estimation without losing accuracy.

The rest of this paper is organized as follows. In Section 2 we describe the inverse regression model for ordinal data and its dimensionality reduction. In Section 3 we derive the Maximum likelihood estimates of the reduction and we also present a variable selection method. Section 4 is dedicated to developing a permutation test for choosing the dimension for the reduction. Simulation results are presented in Section 5. Section 6 contains a socio-economic application using the methodology developed in this paper to create a socio-economic status (SES) index from ordinal predictors. Finally, a concluding discussion is given in Section 7. All proofs and other supporting material are given in the appendices. Matlab codes for the algorithm and simulations are available at `http://www.fiq.unl.edu.ar/pages/investigacion/` `investigacion-reproducible/grupo-de-estadistica-statistics-group.php`.

2. Model

Let us consider the regression of a response $Y \in \mathbb{R}$ on a predictor $\mathbf{X} = (X_1, X_2, \ldots, X_p)^T$, where each X_j, $j = 1, \ldots, p$ is an ordered categorical variable, i.e., $X_j \in \{1, \ldots, G_j\}$, $j = 1, \ldots, p$. To state a dimension reduction of $\mathbf{X}$ inspired by the model-based SDR approach (see [10]), we should model the inverse regression of $\mathbf{X}$ on Y. However, as we stated in the introduction, the model-based SDR techniques deal with continuous predictors. Therefore, in order to frame our problem in that context, we will assume the existence of a p-dimensional vector of unobserved underlying continuous latent variables $\mathbf{Z} = (Z_1, Z_2, \ldots, Z_p)^T$, with $E(\mathbf{Z}) = 0$, such that each observed X_j is a discretizing of Z_j as follows. There exists a set of thresholds $\mathbf{\Theta}^{(j)} = \{\theta_0^{(j)}, \theta_1^{(j)}, \ldots, \theta_{G_j}^{(j)}\}$, that split the real line in disjoints intervals $-\infty = \theta_0^{(j)} < \theta_1^{(j)} < \cdots < \theta_{G_j-1}^{(j)} < \theta_{G_j}^{(j)} = +\infty$ and

$$X_j = \sum_{g=1}^{G_j} g \mathbb{I}(\theta_{g-1}^{(j)} \leq Z_j < \theta_g^{(j)}), \tag{1}$$

where $\mathbb{I}(A)$ is the indicator function of the set A. Therefore, $X_j = g \Leftrightarrow Z_j \in [\theta_{g-1}^{(j)}, \theta_g^{(j)})$ and $P(X_j = g) = P(\theta_{g-1}^{(j)} \leq Z_j < \theta_g^{(j)})$.

In the framework of model-based inverse regression, we adopt, following [11] that the variable $\mathbf{Z}$ given Y is normal with mean depending on Y and constant variance, i.e.

$$\mathbf{Z}|Y = \boldsymbol{\mu}_Y + \boldsymbol{\epsilon}, \tag{2}$$

3

where $\boldsymbol{\mu}_Y = E(\mathbf{Z}|Y)$ and the error $\boldsymbol{\epsilon}$ is independent of Y, normally distributed with mean $\mathbf{0}$ and covariance (positive definite) matrix $\boldsymbol{\Delta}$. As usual in latent variable models for ordinal data (see [26]), additionally to $E(\mathbf{Z}) = 0$, we set the diagonal $[\boldsymbol{\Delta}]_{jj} \doteq \delta_j = 1$ in order to allow for model identification.

Since $E(\mathbf{Z}|Y)$ depends on Y we could model that dependence as a function of $\mathbf{f}_Y \in \mathbb{R}^r$ a vector of r known functions with $E((\mathbf{f}_Y - E(\mathbf{f}_Y))(\mathbf{f}_Y - E(\mathbf{f}_Y))^T)$. Under this model, each coordinate of $\mathbf{Z}|Y$ follows a linear model with predictor vector $\mathbf{f}_Y$ and therefore, when Y is quantitative, we can use inverse plots to get information about the choice of $\mathbf{f}_y$, which is not possible in the regression of Y on $\mathbf{X}$. When Y is continuous, $\mathbf{f}_y$ usually will be a flexible set of basis functions, like polynomial terms in Y, which may also be used when it is impractical to apply graphical methods to all of the predictors. When Y is categorical and takes values $\{C_1, \ldots, C_h\}$, we can set $r = h - 1$ and specify the jth element of $\mathbf{f}_y$ to be $\mathbb{I}(y \in C_j)$, $j = 1, \ldots, h$. When Y is continuous, we can also *slice* its values into h categories $\{C_1, \ldots, C_h\}$ and then specify the jth coordinate of $\mathbf{f}_y$ as for the case of a categorical Y. For more details see [1]. As a consequence, model (2) can be expressed as

$$\mathbf{Z}|Y = \boldsymbol{\Gamma}\{\mathbf{f}_Y - E(\mathbf{f}_Y)\} + \boldsymbol{\epsilon}, \tag{3}$$

where $\boldsymbol{\epsilon}$ is independent of Y, normally distributed with mean $\mathbf{0}$ and covariance (positive definite) matrix $\boldsymbol{\Delta}$.

For the regression of Y on the continuous latent variable $\mathbf{Z}$, under model (3) the minimal SDR is $R(\mathbf{Z}) = \boldsymbol{\alpha}^T\mathbf{Z}$, with $\boldsymbol{\alpha}$ a basis for $\boldsymbol{\Delta}^{-1}\text{span}(\boldsymbol{\Gamma})$ by [11], Theorem 2.1. Note that if $R(\mathbf{Z}) = \boldsymbol{\alpha}^T\mathbf{Z}$ is a sufficient reduction, then $R(\mathbf{Z}) = \mathbf{A}\boldsymbol{\alpha}^T\mathbf{Z}$ is a sufficient reduction for any invertible $\mathbf{A} \in \mathbb{R}^{d \times d}$ [9]. Therefore what is identifiable is the span of $\boldsymbol{\alpha}$, not $\boldsymbol{\alpha}$ itself. In the SDR literature, the identifiable parameter $\text{span}(\boldsymbol{\alpha})$ is called a *sufficient reduction subspace*. If $\dim(\text{span}(\boldsymbol{\Gamma})) = d \leq \min\{r, p\}$, (3) can be re-written as

$$\mathbf{Z}|Y = \boldsymbol{\Delta}\boldsymbol{\alpha}\boldsymbol{\xi}\{\mathbf{f}_Y - E(\mathbf{f}_Y)\} + \boldsymbol{\epsilon}, \tag{4}$$

where $\boldsymbol{\alpha} \in \mathbb{R}^{p \times d}$ with $d \leq p$ is a semi-orthogonal matrix whose columns form a basis for the d-dimensional subspace $\boldsymbol{\Delta}^{-1}\text{span}(\boldsymbol{\Gamma})$, $\boldsymbol{\xi} \in \mathbb{R}^{d \times r}$ is a full rank d matrix with $r \geq d$ (see [11], [1]).

Coming back to our problem of interest, in order to propose a supervised dimension reduction for the regression of $Y|\mathbf{X}$ let us observe that, since $\mathbf{X}$ is a function of $\mathbf{Z}$, $\boldsymbol{\alpha}^T\mathbf{Z}$ will be also the sufficient dimension reduction for $Y|\mathbf{X}$, i.e. $Y \perp\!\!\!\perp \mathbf{X}|\boldsymbol{\alpha}^T\mathbf{Z}$ (see Proposition 4.5 in [9]). However, since $\mathbf{Z}$ is unobservable, and the only information available is $\mathbf{X}$, we take the conditional expectation of $\boldsymbol{\alpha}^T\mathbf{Z}$ given $\mathbf{X}$, instead of $\boldsymbol{\alpha}^T\mathbf{Z}$ for the reduction of $Y|\mathbf{X}$ since it is the best predictor of $\boldsymbol{\alpha}^T\mathbf{Z}$ in terms of minimum Mean Square Error. Therefore, for the regression of Y on X, the proposed supervised dimension reduction will be

$$R(\mathbf{X}) = E(\boldsymbol{\alpha}^T\mathbf{Z}|\mathbf{X}). \tag{5}$$

Remark 1. *Observe that in this case, regardless of the encoding of $\mathbf{X}$, the reduction $R(\mathbf{X})$ is completely identified since for each $j = 1, \ldots, p$, $E(Z_j) = 0$ and $X_j = g \Leftrightarrow Z_j \in [\theta_{g-1}^{(j)}, \theta_g^{(j)})$ and therefore, whatever the coding of X_j is, the underlying (and as a consequence the thresholds) does not change.*
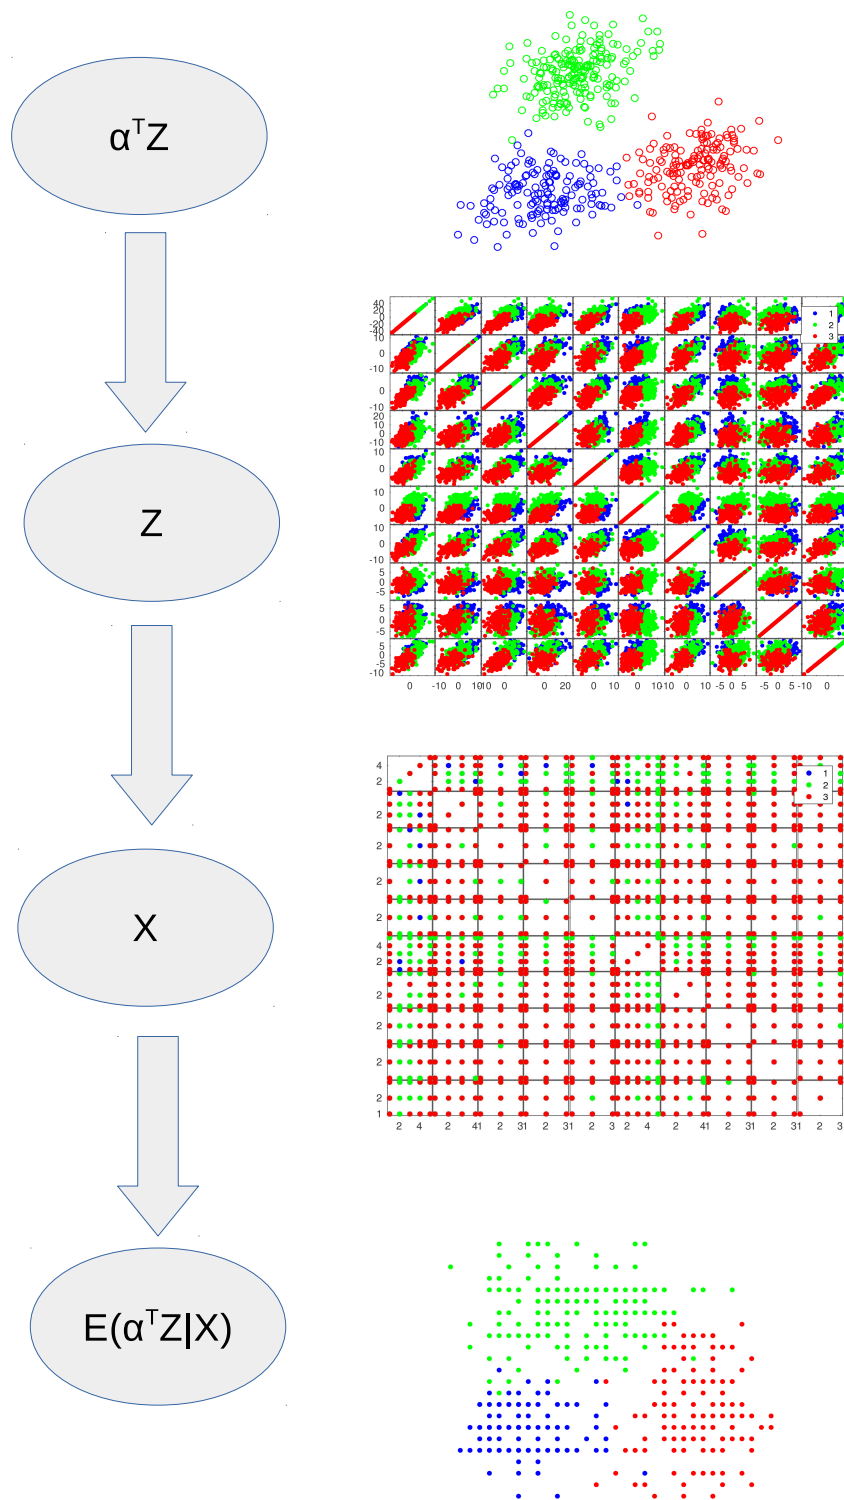
Figure 1: Data generation and processing chain according to the assumed model and proposed dimension reduction scheme.

Figure 2: Residual dependence between the response and the predictors, given the reduction. Hilbert-Schmidt condicional independence criterion (HSCIC) is used as measure of conditional dependence.

Figure 1 helps to understand how the proposed method works, showing an example corresponding to a categorical outcome with three nominal values $Y \in \{1, 2, 3\}$. Suppose we are able to observe the underlying continuous variables $\mathbf{Z}$ and that their distribution follows model (4), with $\boldsymbol{\alpha} \in \mathbb{R}^{p \times 2}$. It means that the characteristic information needed to discriminate between the three groups lies actually in the two-dimensional subspace spanned by the columns of $\boldsymbol{\alpha}$. Top panel of Figure 1 shows such information, plotting the coordinates of $\boldsymbol{\alpha}^T \mathbf{Z}$. For each group indexed by Y, the data fall in a cluster well separated from the others. If we are not allowed to observe this reduced subspace but the complete underlying predictors $\mathbf{Z}$, we would have a situation as described by the scatter plots between pairs of predictors depicted in the second panel of Figure 1. Despite we can still see some separation between clusters, it is not as clear as in the sufficient low-dimensional subspace. In real scenarios with ordinal data, according to the assumed model we do not have access to observe $\mathbf{Z}$ either, but a discretized version $\mathbf{X}$ which is a function of the underlying $\mathbf{Z}$ throught the set of fixed thresholds $\boldsymbol{\Theta}$. This situation is illustrated in the third panel. It is clear from the figure that the continuous values of $\mathbf{Z}$ collapse into a discrete set of values in $\mathbf{X}$ and it is now much harder to discriminate between the groups indexed by Y. Nevertheless, the dimension reduction approach proposed in this paper projects the data again onto a 2-dimensional subspace, as shown in panel at the bottom of Figure 1. Note that clear separation between clusters is recovered; indeed, the information available in this subspace closely resembles that in the characteristic subspace spanned by $\boldsymbol{\alpha}$ (compare the first and last panels).

It is interesting to see also how well the proposed reduction captures the information about Y that is available in $\mathbf{X}$. We know that for the underlying variables, $\mathbf{Z} \perp\!\!\!\perp Y \mid \boldsymbol{\alpha}^T \mathbf{Z}$. If we can measure the residual dependence between $\mathbf{Z}$ and Y given $\boldsymbol{\alpha}^T \mathbf{Z}$ using a suitable measure of statistical independence $\rho(\mathbf{Z}, Y \mid \boldsymbol{\alpha}^T \mathbf{Z})$, we will have $\rho(\mathbf{Z}, Y \mid \boldsymbol{\alpha}^T \mathbf{Z}) = 0$ in the population. In practice, for a finite random sample of $(Y, \mathbf{Z}, \boldsymbol{\alpha}^T \mathbf{Z})$, this value will be greater than cero. This is shown in the left-most boxplot in Figure 2. The boxplot was computed using 100 simulated data sets and choosing for the measure ρ the Hilbert-Schmidt condicional independence criterion introduced in [41]. Isotropic Gaussian kernels are used to embed the observations into a RKHS. Kernel bandwiths are set to the median of the pairwise distances between samples. Since $\boldsymbol{\alpha}^T \mathbf{Z}$ is

still a sufficient reduction for the regression $Y|\mathbf{X}$, in the population we would also have $\rho(\mathbf{X}, Y \mid \boldsymbol{\alpha}^T\mathbf{Z}) = 0$. A sample estimate of this quantity, using the same data and dependence measure as before, is shown in the second boxplot of the figure. Finally, if we compute the same empirical measure, with the same random sample but for the practical reduction $R(\mathbf{X}) = E(\boldsymbol{\alpha}^T\mathbf{Z} \mid \mathbf{X})$ instead of $\boldsymbol{\alpha}^T\mathbf{Z}$, we obtain the boxplot shown on the right of the figure. This figure shows that on average the empricial version of $\rho(\mathbf{X}, Y \mid E(\boldsymbol{\alpha}^T\mathbf{Z}|\mathbf{X}))$ is a little bigger than the empirical version of $\rho(\mathbf{X}, Y \mid \boldsymbol{\alpha}^T\mathbf{Z})$ for the same random sample of $(Y, \mathbf{Z}, \boldsymbol{\alpha}^T\mathbf{Z})$, albeit they are very close. This suggests that, even if we cannot claim sufficiency of the proposed reduction, it is really close to the ideal unattainable reduction $\alpha^T\mathbf{Z}$.

3. Estimation

For the supervised dimension reduction given in (5), we need to estimate the semiorthogonal basis matrix $\boldsymbol{\alpha}$. If $\mathbf{Z}$ were observed, the maximum likelihood estimator of $\boldsymbol{\alpha}$ would be the one derived in [11]. That is, $\widehat{\boldsymbol{\alpha}} = \widetilde{\boldsymbol{\Sigma}}^{-1/2}\mathbf{v}$, where $\mathbf{v}$ are the first d eigenvectors of the symmetric matrix $\widetilde{\boldsymbol{\Sigma}}^{-1/2}\widetilde{\boldsymbol{\Sigma}}_{\text{fit}}\widetilde{\boldsymbol{\Sigma}}^{-1/2}$, $\widetilde{\boldsymbol{\Sigma}}$ is the sample marginal covariance of the predictors and $\widetilde{\boldsymbol{\Sigma}}_{\text{fit}}$ is the sample covariance of the fitted values of the regression $\mathbf{Z}|\mathbf{f}_y$. This estimation procedure is called Principal Fitted Components (PFC). However, $\mathbf{Z}$ is not observed, and therefore the sample covariance matrices (marginal and fitted) cannot be estimated directly. In view of the robustness proven in [11], we could consider applying the methodology directly to $\mathbf{X}$ in a naive way and it still will obtain a $\sqrt{n}$ consistent estimator. This approach will be the initial value of our algorithm to obtain the maximum likelihood estimate under the true model.

For the estimation, let as assume we have a random sample of n points $(y_i, \mathbf{x}_i)$ drawn from the joint distibution of $(Y, \mathbf{X})$ following model (4) and that the dimension d of the reduction is known (later in Section 4 we will consider how to infer it). In what follows, we will call $\boldsymbol{\Theta} \doteq \{\boldsymbol{\Theta}^{(1)}, \ldots, \boldsymbol{\Theta}^{(p)}\} = \{\theta_0^{(1)}, \ldots, \theta_{G_1}^{(1)}, \ldots\ldots, \theta_0^{(p)}, \ldots, \theta_{G_p}^{(p)}\}$ and $C(\mathbf{X}, \boldsymbol{\Theta}) = [\theta_{X_1-1}^{(1)}, \theta_{X_1}^{(1)}) \times \cdots \times [\theta_{X_p-1}^{(p)}, \theta_{X_p}^{(p)})$. In order to obtain the estimator we need to maximize the log-likelihood function of the observed data

$$\sum_{i=1}^{n} \log f_{\mathbf{X}}(\mathbf{x}_i|y_i; \boldsymbol{\Theta}, \boldsymbol{\Delta}, \boldsymbol{\alpha}, \boldsymbol{\xi}). \tag{6}$$

Since $\mathbf{Z}|Y$ is normally distributed, using (1) we can compute, for each i, the truncated unnormalized density $f_{\mathbf{X},\mathbf{Z}}(\mathbf{x}_i, \mathbf{z}_i|y_i; \boldsymbol{\Theta}, \boldsymbol{\Delta}, \boldsymbol{\alpha}, \boldsymbol{\xi})$ as

$$f_{\mathbf{X},\mathbf{Z}}(\mathbf{x}_i, \mathbf{z}_i|y_i; \boldsymbol{\Theta}, \boldsymbol{\Delta}, \boldsymbol{\alpha}, \boldsymbol{\xi}) = (2\pi)^{-p/2}|\boldsymbol{\Delta}|^{-1/2}e^{-\frac{1}{2}\text{tr}(\boldsymbol{\Delta}^{-1}(\mathbf{z}_i - \boldsymbol{\Delta}\boldsymbol{\alpha}\boldsymbol{\xi}\bar{\mathbf{f}}_{y_i})(\mathbf{z}_i - \boldsymbol{\Delta}\boldsymbol{\alpha}\boldsymbol{\xi}\bar{\mathbf{f}}_{y_i})^T)}I_{\{\mathbf{z}_i \in C(\mathbf{x}_i, \boldsymbol{\Theta})\}},$$

where $\bar{\mathbf{f}}_{y_i} \doteq \mathbf{f}_{y_i} - n^{-1}\sum_{i=1}^{n}\mathbf{f}_{y_i}$. Therefore, for each i, the unnormalized marginal density $\mathbf{X}|Y$ will be

$$f_{\mathbf{X}}(\mathbf{x}_i|y_i; \boldsymbol{\Theta}, \boldsymbol{\Delta}, \boldsymbol{\alpha}, \boldsymbol{\xi}) = \int f_{\mathbf{X},\mathbf{Z}}(\mathbf{x}_i, \mathbf{z}_i|y_i; \boldsymbol{\Theta}, \boldsymbol{\Delta}, \boldsymbol{\alpha}, \boldsymbol{\xi})\, d\mathbf{z}_i.$$

As an exact computation of the likelihood is difficult due to the multiple integrals involved, maximum likelihood estimates are often obtained using an iterative expectation-maximization (EM) algorithm. This is a common choice for models with latent variables, since it exploits the reduced complexity of computing the joint likelihood of the complete data $(\mathbf{X}, \mathbf{Z})$. We will follow this approach in the present paper. The corresponding algorithm is described below.

3.1. Algorithm

In this section we present the EM algorithm, closely related to the one given in [23], to estimate the parameters in model (4). Throughout this section, we will use superscripts $A^{(k)}$ to indicate the value of quantity A at the kth iteration of the algorithm. In addition, to make notation easier, let us collect the parameters $\boldsymbol{\Delta}, \boldsymbol{\alpha}, \boldsymbol{\xi}$ into a single parameter vector $\boldsymbol{\Omega} \doteq \{\boldsymbol{\Delta}, \boldsymbol{\alpha}, \boldsymbol{\xi}\}$. The procedure starts with Step 0, where we initialize $\boldsymbol{\Omega}^{(0)}$ using the estimators obtained from PFC applied to $\mathbf{X}$. Then, the algorithm iterates between the following two steps until convergence is reached: Step 1 is devoted to estimating $\boldsymbol{\Theta}^{(k)}$ given $\boldsymbol{\Omega}^{(k-1)}$ and Step 2 to getting $\boldsymbol{\Omega}^{(k)}$ by maximizing the conditional expectation (given $\boldsymbol{\Omega}^{(k-1)}$ and $\boldsymbol{\Theta}^{(k)}$) of the joint log-likelihood (6). This step is properly the EM step.

Step 1: Estimation of $\boldsymbol{\Theta}$: Given $\boldsymbol{\Omega}^{(k-1)}$ from Step 0 or from a previous iteration, let $\boldsymbol{\Psi}^{(k-1)} \doteq \boldsymbol{\Delta}^{(k-1)}\boldsymbol{\alpha}^{(k-1)}\boldsymbol{\xi}^{(k-1)}$. For each $j = 1, \ldots, p$ and $g = 1, \ldots, G_j$ define

$$L_g^{(j)}(\theta) \doteq \#\{i : x_{ij} \leq g\} - \sum_{i=1}^{n} \Phi\left(\frac{\theta - \boldsymbol{\Psi}_j^{(k-1)}\bar{\mathbf{f}}_{y_i}}{\delta_j^{(k-1)}}\right),$$

where Φ is the cumulative distribution function of the standard normal, for each j, $\delta_j^{(k-1)} = (\boldsymbol{\Delta}^{(k-1)})_{jj}$, $\boldsymbol{\Psi}_j^{(k-1)}$ indicates the jth row of $\boldsymbol{\Psi}^{(k-1)}$, x_{ij} is the jth coordinate of $\mathbf{x}_i$, and $\#S$ indicates the cardinality of the set S. Then, take $\hat{\theta}_0^{(j)} = -\infty =$, $\hat{\theta}_{G_j}^{(j)} = +\infty$. For $g = 1, \ldots, G_j - 1$, assign to $\hat{\theta}_g^{(j)}$ the unique solution of the equation $L_g^{(j)}(\theta) = 0$. Set $\boldsymbol{\Theta}^{(k)} = \{\theta_0^{(1)}, \ldots, \theta_{G_1}^{(1)}, \ldots, \theta_0^{(p)}, \ldots, \theta_{G_p}^{(p)}\}$. Here the definition of $L_g^{(j)}$ is based on the normality assumption on the conditional distribution of the underlying continuous variable. More precisely, we define $L_g^{(j)}$ as a *search* function of thresholds using the underlying (normal) cumulative distribution function.

Step 2: Estimation of $\boldsymbol{\Delta}^{(k)}, \boldsymbol{\alpha}^{(k)}, \boldsymbol{\xi}^{(k)}$: Given $\boldsymbol{\Theta}^{(k)}$ computed in Step 1 and $\boldsymbol{\Omega}^{(k-1)}$ from Step 0 or from a previous iteration, we apply the EM algorithm to maximize (6). The EM algorithm consist in finding $\boldsymbol{\Omega}^{(k)}$ that maximize over $\boldsymbol{\Omega}$

$$Q(\boldsymbol{\Omega}|\boldsymbol{\Omega}^{(k-1)}) = \sum_{i=1}^{n} E_{\mathbf{z}_i|y_i;\boldsymbol{\Omega}^{(k-1)}}\left[\log f_{\mathbf{x}_i,\mathbf{z}_i}(\mathbf{x}_i, \mathbf{z}_i|y_i; \boldsymbol{\Omega})\big|y_i; \boldsymbol{\Omega}^{(k-1)}\right]. \tag{7}$$

These produce

$$\begin{aligned}
\boldsymbol{\alpha}^{(k)} &= \mathbf{S}^{-1/2}\hat{\boldsymbol{\zeta}}_d\mathbf{N}, \\
(\boldsymbol{\Delta}^{-1})^{(k)} &= \mathbf{S}^{-1} + \boldsymbol{\alpha}^{(k)}((\boldsymbol{\alpha}^{(k)})^T\mathbf{S}_{\mathrm{res}}\boldsymbol{\alpha}^{(k)})^{-1}(\boldsymbol{\alpha}^{(k)})^T - \boldsymbol{\alpha}^{(k)}((\boldsymbol{\alpha}^{(k)})^T\mathbf{S}\boldsymbol{\alpha}^{(k)})^{-1}(\boldsymbol{\alpha}^{(k)})^T, \\
\boldsymbol{\xi}^{(k)} &= ((\boldsymbol{\alpha}^{(k)})^T\boldsymbol{\Delta}^{(k)}\boldsymbol{\alpha}^{(k)})^{-1}(\boldsymbol{\alpha}^{(k)})^T\mathbf{M}^T\mathbf{F}(\mathbf{F}^T\mathbf{F})^{-1},
\end{aligned}$$

where $\mathbf{N}$ is a matrix such that $(\boldsymbol{\alpha}^{(k)})^T\boldsymbol{\alpha}^{(k)} = \mathbf{I}_d$ and the $\hat{\boldsymbol{\zeta}}_d$ are the first d eigenvectors of $\mathbf{S}^{-1/2}\mathbf{S}_{\mathrm{fit}}\mathbf{S}^{-1/2}$, where the matrices $\mathbf{S} \in \mathbb{R}^{p \times p}$ and $\mathbf{S}_{\mathrm{fit}} \in \mathbb{R}^{p \times p}$ are given by

$$\mathbf{S} = \frac{1}{n}\sum_{i=1}^{n} E_{\mathbf{z}_i|\mathbf{x}_i,y_i;\boldsymbol{\Omega}^{(k-1)}}(\mathbf{z}_i\mathbf{z}_i^T|\mathbf{x}_i, y_i; \boldsymbol{\Omega}^{(k-1)}) \quad \text{and} \quad \mathbf{S}_{\mathrm{fit}} = n^{-1}\mathbf{M}^T\mathbf{F}(\mathbf{F}^T\mathbf{F})^{-1}\mathbf{F}^T\mathbf{M}$$

8

with $\mathbf{F} \in \mathbb{R}^{n \times r}$ and $\mathbf{M} \in \mathbb{R}^{n \times p}$ matrices whose transposes are given by $\mathbf{F}^T = [\bar{\mathbf{f}}_{y_1}, \ldots, \bar{\mathbf{f}}_{y_n}]$ and $\mathbf{M}^T = [E_{\mathbf{z}_1|\mathbf{x}_1, y_1; \mathbf{\Omega}^{(k-1)}}(\mathbf{z}_1|\mathbf{x}_1, y_1; \mathbf{\Omega}^{(k-1)}), \ldots, E_{\mathbf{z}_n|\mathbf{x}_n, y_n; \mathbf{\Omega}^{(k-1)}}(\mathbf{z}_n|\mathbf{x}_n, y_n; \mathbf{\Omega}^{(k-1)}))]$ and the residual matrix $\mathbf{S}_{\text{res}}$ is defined by $\mathbf{S}_{\text{res}} = \mathbf{S} - \mathbf{S}_{\text{fit}}$. Details of the derivation and the EM algorithm are given in Appendix A.

$\underline{\textbf{Step 3:}}$ Check convergence. If it is not reached, go to **Step 1**. We check convergence simply by looking to see whether $Q(\mathbf{\Omega}^{(k)}|\mathbf{\Omega}^{(k-1)})$ stops increasing from one iteration to the next. Specifically, we check whether $(Q(\mathbf{\Omega}^{(k)}|\mathbf{\Omega}^{(k-1)}) - Q(\mathbf{\Omega}^{(k-1)}|\mathbf{\Omega}^{(k-2)}))/Q(\mathbf{\Omega}^{(k-1)}|\mathbf{\Omega}^{(k-2)}) < \epsilon$, with ϵ typically set to 10^{-6}.

3.2. Estimation with variable selection

When we compute the linear combinations implied in (5), we need to include all the original variables. This means that even non-relevant or redundant variables are included in the final model, making it harder to interpret. To overcome this limitation, we can perform variable selection, in this way obtaining linear combinations that include only the active, relevant variables. Following [6], the maximization of (7) is equivalent to finding, in each iteration,

$$\boldsymbol{\alpha}^{(k)} = \arg\min_{\boldsymbol{\alpha}} \left\{ - \text{tr}(\boldsymbol{\alpha}^T \mathbf{S}_{\text{fit}} \boldsymbol{\alpha}) \right\}, \qquad \text{subject to } \boldsymbol{\alpha}^T \mathbf{S} \boldsymbol{\alpha} = \mathbf{I}_d. \qquad (8)$$

To induce variable selection in dimension reduction, we introduce a group-lasso type penalty since, in order not to choose a particular variable X_j, we need to make the whole jth row of $\boldsymbol{\alpha}$, $\boldsymbol{\alpha}_j$, equal to 0. For that, following [6], we use a mixed ℓ_1/ℓ_2 norm, where the inner norm is the ℓ_2 norm of each row of $\boldsymbol{\alpha}$. Adding the penalty term to (8), we get

$$\boldsymbol{\alpha}^{(k)} = \arg\min_{\boldsymbol{\alpha}} \left\{ - \text{tr}(\boldsymbol{\alpha}^T \mathbf{S}^{-1/2} \mathbf{S}_{\text{fit}} \mathbf{S}^{-1/2} \boldsymbol{\alpha}) + \lambda \sum_{i=1}^{p} ||\boldsymbol{\alpha}_i||_2 \right\}, \qquad \text{subject to } \boldsymbol{\alpha}^T \mathbf{S} \boldsymbol{\alpha} = \mathbf{I}_d.$$

The parameter λ can be found using an information criterion, such as Akaike's (AIC) or Bayes' (BIC) criteria. Details are given in [6]. Another approach is to find the value λ^* that minimizes the prediction error via a cross-validation experiment, but this requires adopting a specific prediction rule.

It is interesting to note that this procedure performs at the same time variable selection and dimension reduction without modeling the regression for $Y|E(\mathbf{Z}|\mathbf{X})$ or $Y|\mathbf{X}$. Thus, the obtained reduction can be used later with any prediction rule of choice. This is different, for instance, from the approach proposed in [21], where the variable selection is driven by a particular regression model.

3.3. Computing the reduction

In order to compute the reduction (5), we need to estimate $\boldsymbol{\alpha}$ and $E(\mathbf{Z}|\mathbf{X})$. In the preceding paragraph we focused on computing an estimate of $\boldsymbol{\alpha}$. To estimate $E(\mathbf{Z}|\mathbf{X})$, observe that, when the response Y is discrete, using Bayes' rule we get

$$
\begin{aligned}
E(\mathbf{Z}|\mathbf{X}) &= E_{Y|\mathbf{X}}\big(E(\mathbf{Z}|\mathbf{X}, Y = y)\big) \\
&= \sum_{y \in S_Y} P(Y = y|\mathbf{X}) E(\mathbf{Z}|\mathbf{X}, y) \\
&= \sum_{y \in S_Y} w(y) E(\mathbf{Z}|\mathbf{X}, y),
\end{aligned}
$$

9

where,

$$w(y) = \frac{P(\mathbf{X}|Y = y)P(Y = y)}{\sum_{y \in S_Y} P(\mathbf{X}|Y = y)P(Y = y)},$$

Then, to estimate $E(\mathbf{Z}|\mathbf{X})$ we take: $\widehat{E}(\mathbf{Z}|\mathbf{X}, y) = \mathbf{M}$ from the Step 2 of the EM algoritm (Section 3.1); $\widehat{P}(\mathbf{X}|Y = y) = \int_{C(\mathbf{X},\widehat{\boldsymbol{\Theta}})} f_{\mathbf{Z}|Y}(\mathbf{z}|y; \widehat{\boldsymbol{\Delta}}, \widehat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\xi}})\, d\mathbf{z}$ with $\widehat{\boldsymbol{\Theta}}, \widehat{\boldsymbol{\Delta}}, \widehat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\xi}}$ obtained from Steps 1 and 2 of Section 3.1 or 3.2; and $\widehat{P}(Y = y)$ from the sample. When the response is continuous, we can simply slice Y in h bins and use the previous procedure. Note also that the sample space of $\mathbf{X}$ is finite and so the sample space of $E(\boldsymbol{\alpha}^T \mathbf{Z}|\mathbf{X})$ is finite too. Thus, a priori we can tabulate those values for future use and avoid computations when we need to reduce a new instance of $\mathbf{X}$. Nevertheless, when p is moderately large and there are several ordered categories for each variable X_j, the amount of memory needed to store such a look up table can become too large in practice. For instance, if $p = 20$ and we have $G_j = 3$ for each X_j, to store all the values of $E(\boldsymbol{\alpha}^T \mathbf{Z}|\mathbf{X})$ in double precision we need around 26 GB of memory.

4. Choosing the dimension

Our developments in Section 3 assumed that the dimension $d \le \min(p, r)$ of the reduction subspace was known. In practical settings, this dimension should be inferred from the data. For model-based SDR, which allows for likelihood computation, likelihood-ratio and information criteria such as AIC or BIC have been proposed to drive the selection of d [11]. The accuracy of these methods, however, is not robust to deviations from the assumed model. When the main goal is prediction, a common choice is to assess different values of d according to their performance at predicting out-of-sample cases in a cross validation setting. The value of d picked is the one that achieves the minimum prediction error.

Another way to choose the dimension is via permutation tests, as introduced in [17]. For that, assume that $\boldsymbol{\alpha} \in \mathbb{R}^{p \times m}$ and $(\boldsymbol{\alpha}, \boldsymbol{\alpha}_0) \in \mathbb{R}^{p \times p}$ is unitary. A permutation test relies on the fact that $\mathcal{S}_{\boldsymbol{\alpha}}$ is a sufficient dimension reduction subspace for the regression of Y on $\mathbf{Z}$ whenever $(Y, \boldsymbol{\alpha}^T \mathbf{Z}) \perp\!\!\!\perp \boldsymbol{\alpha}_0^T \mathbf{Z}$. Note that this implies $m \ge d$. For this test, we consider the statistic $\hat{\Lambda}_m = 2(\mathcal{Q}_p(\mathbf{I}_p) - \mathcal{Q}_m(\widehat{\boldsymbol{\alpha}}))$, where $\mathcal{Q}_r$ is given by the Q function in (7), evaluated at the estimator obtained in the EM algorithm given above, for a fixed dimension of $\mathcal{S}_{\boldsymbol{\alpha}}$. Set $m = 0$. A procedure adapted to ordinal data to infer d via permutation testing involves the following steps:

(i) Obtain $\widehat{\boldsymbol{\alpha}}$, the MLE of $\boldsymbol{\alpha}$ and compute $\hat{\Lambda}_m$. Obtain also $\widehat{\boldsymbol{\alpha}}_0$.

(ii) For the data $(Y_i, E(\widehat{\boldsymbol{\alpha}}^T \mathbf{Z}|\mathbf{X}_i), E(\widehat{\boldsymbol{\alpha}}_0^T \mathbf{Z}|\mathbf{X}_i))$, permute the columns corresponding to $E(\widehat{\boldsymbol{\alpha}}_0^T \mathbf{Z}|\mathbf{X}_i)$ to get a new sample $(Y_i, E(\mathbf{Z}|\mathbf{X}_i)^*)$. For the new data, obtain the MLE and compute $\hat{\Lambda}_m^*$.

(iii) Repeat step 2, B times.

(iv) Compute the fraction of the $\hat{\Lambda}_m^*$ that exceed $\hat{\Lambda}_m$. If this value is smaller than the chosen significance level and if $m < \min(r, p)$ set $m = m + 1$ and go to step 1. Otherwise, return $d = m$.

In this way, the inferred d is the smallest m that fails to reject the null hypothesis of independence between $(Y, E(\widehat{\boldsymbol{\alpha}}^T \mathbf{Z}|\mathbf{X}_i))$ and $E(\widehat{\boldsymbol{\alpha}}_0^T \mathbf{Z}|\mathbf{X}_i)$.

5. Simulations

In this section we illustrate the performance of the proposed method using simulated data. A critical aspect of the implementation is the computation of the E-step of the EM algorithm. Exact computation of the truncated moments involved in this step would make the proposed method infeasible in practice even for a dimensionality of the predictors of order $5 < p < 10$, depending on the number of ordered categories. To address this, we implement an approximate estimation method adapted from [23]. The main idea is to use a recursion to iteratively compute the truncated moments of a multivariate normal distribution. The derivation and further details are given in Appendix B. The first step is to validate the proposed approximate method on the E-step by comparing its performance with the exact computation of the truncated moments. Then, we compare the performance of the proposed SDR-based method against standard methodology developed for continuous data. Taking advantage of the computational savings obtained with the approximate E-step computation, we then illustrate the performance of the proposed strategy to infer the dimension of the dimension reduction subspace using permutation testing and cross validation. Finally, we illustrate the performance of the regularized estimator proposed in Section 3.2 in a prediction task.

5.1. Validation of the proposed algorithm for the E-step of the algorithm

The most demanding part of the proposed method is the computation of the truncated moments of the multivariate normal distributions in the E-step. The approximate iterative method proposed for its computation is a main ingredient to allow the application of the methodology in practical settings. In this section we validate this strategy by comparing the approximate computation against the exact computation using the algorithm proposed in [29]. Since the exact computation involves a high computational cost even for moderate dimensions of the predictor vectors, we set $p = 5$ and $n = 100$. In addition, we set $G_j = 4$ for $j = 1, 2, \ldots, 5$. The data was generated according to (4), with $Y \sim N(0, 1)$. For the basis matrix $\boldsymbol{\alpha}$, we set $\sqrt{p}\,\boldsymbol{\alpha} = (\mathbf{1}_p \;\; \text{sign}(\mathbf{e}))$, with $\mathbf{1}_p$ a column vector of ones of size $p = 5$ and $\mathbf{e} \sim N_p(\mathbf{0}, \mathbf{I})$. For the covariance matrix $\boldsymbol{\Delta}$, we set $\boldsymbol{\Delta} = \mathbf{I} + \boldsymbol{\alpha}\mathbf{B}\boldsymbol{\alpha}^T$, with $\mathbf{B}$ a 2×2 symmetric random matrix fixed at the outset. We also chose a polynomial basis for $\mathbf{f}_Y$, with $r = 2$. The same choice of $\mathbf{f}_Y$ was used for the estimation. The experiment was replicated 100 times. In each run, the same training sample was used with both methods.

To assess the accuracy of the estimation, we measured the angle between the subspace spanned by the true $\boldsymbol{\alpha}$ and the one spanned by the estimate $\widehat{\boldsymbol{\alpha}}$. This quantity ranges from 0 degrees if the two subspaces are identical to 90 degrees if they do not share any information. The average angle obtained with the exact method was 10.21 degrees with a standard deviation of 6.37 degrees, whereas for the estimate obtained with the approximate method the average angle was 12.65 degrees with a standard deviation of 5.70 degrees. The 95% confidence interval for the average difference between the angles obtained with both methods is $(2.07, 2.81)$ degrees. These values suggest that the price to pay for the introduction of the approximate computation is very small.

It is also illustrative to see the impact of the approximation on prediction. Using plain linear regression for $Y | E(\widehat{\boldsymbol{\alpha}}^T \mathbf{Z} | \mathbf{X})$, the MSE of the residuals averaged over the 100 runs is 0.863 when estimating the reduction using the approximate E-step, whereas it is 0.811 when using the exact method. In both cases the standard deviation of the averaged MSE is 0.04. Thus, the difference in the MSE obtained with the approximate method represents less than 1.5% of the average MSE

11

obtained with the exact method. The importance of the approximate method for the E-step is better understood when noting the big difference in computing time. Using plain MATLAB implementations for both methods, computation with the exact method takes on average 1.13×10^2 seconds for each run, while with the approximate method this time was reduced to 0.26 seconds, a difference of three orders of magnitude. Overall, these results show that the approximate method to compute the truncated moments is a viable alternative: it reduces the computing time in practical applications without a significant loss in accuracy.

5.2. Performance of the proposed method

In this section we assess the performance of the proposed method on simulated data. For this example, we set $p = 20$, $d = 2$ and a polynomial basis with $r = 2$ for $\mathbf{f}_Y$. As in Section 5.1, we generate the data according to (4), with $Y \sim N(0, 1)$, $\sqrt{p}\,\boldsymbol{\alpha} = (\mathbf{1}_p \ \text{sign}(\mathbf{e}))$, with $\mathbf{e} \sim N_p(\mathbf{0}, \mathbf{I})$, and covariance matrix $\boldsymbol{\Delta} = \mathbf{I} + \boldsymbol{\alpha}\mathbf{B}\boldsymbol{\alpha}^T$, with $\mathbf{B}$ a $d \times d$ symmetric random matrix fixed at the outset. The values of G_j in this case ranged from 3 to 5. For the estimation we used a polynomial basis with $r = 2$. To evaluate the performance, we computed the angle between the true reduction $\boldsymbol{\alpha}^T\mathbf{Z}$ and the estimated reduction $R(\mathbf{X})$. We considered three choices for $R(\mathbf{X})$: (i) the reduction is given by $\widehat{\boldsymbol{\alpha}}_{\text{PFC}}^T\mathbf{X}$, with $\widehat{\boldsymbol{\alpha}}_{\text{PFC}}$ computed as in standard PFC for continuous variables; (ii) the reduction is $\widehat{\boldsymbol{\alpha}}_{\text{ORD}}^T\mathbf{X}$, with $\widehat{\boldsymbol{\alpha}}_{\text{ORD}}$ computed as proposed here but applied on the observed ordinal data $\mathbf{X}$; (iii) $R(\mathbf{X}) = E(\widehat{\boldsymbol{\alpha}}_{\text{ORD}}^T\mathbf{Z}|\mathbf{X})$ as proposed in Section 2.

Figure 3-(a) shows boxplots of the obtained results for 100 runs of the experiment and a sample of size $n = 500$. The angle is measured in degrees. It can be seen that the mean value of the angle is significantly smaller for the proposed method for ordinal predictors, compared to using choices (i) or (ii). It can be seen that the variance is somewhat increased, but the gain in accuracy clearly worths the price. It is less obvious from the figure that $\widehat{\boldsymbol{\alpha}}_{\text{ORD}}$ provides a better estimation of the true subspace spanned by $\boldsymbol{\alpha}$ than the standard $\widehat{\boldsymbol{\alpha}}_{\text{PFC}}$ estimator. The 95% normal confidence interval for the difference $[\text{angle}(\boldsymbol{\alpha}, \widehat{\boldsymbol{\alpha}}_{\text{PFC}}) - \text{angle}(\boldsymbol{\alpha}, \widehat{\boldsymbol{\alpha}}_{\text{ORD}})]$ is $(9.59, 11.02)$ degrees. These results show that for ordinal data, estimation of the subspace spanned by $\boldsymbol{\alpha}$ using the proposed method for ordered predictors clearly outperforms standard PFC as derived for continuous predictors.

It seems fair to ask whether the gain in performance discussed in this example still holds when the normality assumption for the underlying latent variables does not hold. For standard PFC, Cook and Forzani [11, Theorem 3.5] showed that the estimator is still consistent when $\mathbf{Z}|Y$ deviates from multivariate normality. To assess the performance of the proposed method in this scenario, we generated data similarly as before, but with $\boldsymbol{\epsilon}$ non-normally distributed. In particular, we assumed that $\boldsymbol{\epsilon}$ had chi-squared distributed coordinates $(\boldsymbol{\epsilon})_j \sim \chi^2(5)$. The rest of the simulation parameters remained fixed as before. The results obtained are shown in Figure 3-(b). It can be seen that the angles obtained with both methods are very close to those obtained for conditionally normal data. This confirms the superiority of the proposed method and algorithm to estimate $\text{span}(\boldsymbol{\alpha})$ when the predictors are ordered categories.

5.3. Inference about d

In this section, we carry out a simulation study to evaluate methods to infer the dimension of the reduction subspace from the data. In particular, we compare the accuracy of permutation testing against 10-fold cross validation (CV) and well-known information criteria like Akaike's

Figure 3: Angle between span($\widehat{\boldsymbol{\alpha}}$) and span($\boldsymbol{\alpha}$), for $\widehat{\boldsymbol{\alpha}}$ obtained using `PFC` and the proposed approach for ordinal predictors. (a) normal conditional model for $\mathbf{Z}|(Y = y)$; (b) non-normal conditional model for $\mathbf{Z}|(Y = y)$.

information criterion (AIC) and Bayes' information criterion (BIC). For this study, we generate data as described in Section 5.2, with $p = 10$, $d = 2$ and we use a polynomial basis with degree $r = 4$ for $\mathbf{f}_Y$. Since $d \leq \min(r, p)$, we search for the true value of d within the set $\{0, 1, 2, 3, 4\}$. For permutation testing, we build the permutation distribution of $\hat{\Lambda}_m$ resampling the data 500 times as discussed in Section 4 and use a significance level of 0.01. For CV we used the averaged mean-squared prediction error over the test partition as the driving measure of performance, taking simple k-NN regression as the prediction rule. For information criteria, we take

$$\hat{d}_{IC} = \arg\min_{d_o}\{2\mathcal{Q}_{d_o}(\boldsymbol{\alpha}) + c_{IC}h_\Theta(d_o)\},$$

where $c_{AIC} = 2$, $c_{BIC} = \log(n)$, $h_\Theta(d_o) = rd_o + d_o(p - d_o) + p(p + 3)/2 + n_\Theta$ and n_Θ is the number of thresholds estimated during the computation of $\widehat{\boldsymbol{\alpha}}$. The experiment was repeated using two different sample sizes, $n = 200$ and $n = 300$, and it was run 500 times for each sample size.

Table 1 shows the obtained results. For $n = 300$, permutation testing finds the true dimension 83% of the runs, while CV finds it 66% of the time. In addition, for this sample size, both methods hardly ever pick less than two directions, and therefore no information is lost. The fraction of the runs that at most one extra direction is chosen, that is, $\hat{d} = 2$ or $\hat{d} = 3$, is 0.898 for permutation testing and 0.85 using cross-validation. On the other hand, information criteria show a poorer performance. AIC picks the right dimension around half of the runs. Nevertheless, the rest of the runs it picks $\hat{d} = 1$, thus losing information. Underestimation of the true dimension is even more severe with BIC, since it picks only one direction almost always.

For the smaller sample size $n = 200$, permutation testing is less accurate. It finds the true dimension 67% of the runs in this scenario, but in 24% of the runs it picked only one direction for projection instead of two. CV finds the true dimension 59% of the runs, but tends to overestimate the required dimension d, leading to a potential loss in efficiency but preserving information. A test for the difference in the proportion of choices $\hat{d} \geq 2$ between permutation testing and CV

13

Table 1: Fraction of times a given value of d was chosen

		PERMUTATION	CV	AIC	BIC
	$d = 1$	0.240	0.000	0.794	0.998
$n = 200$	$d = 2$	0.670	0.591	0.206	0.002
	$d = 3$	0.067	0.214	0	0
	$d = 4$	0.023	0.195	0	0
	$d = 1$	0.077	0.000	0.488	0.986
$n = 300$	$d = 2$	0.832	0.657	0.512	0.014
	$d = 3$	0.066	0.191	0	0
	$d = 4$	0.025	0.152	0	0

gives a p-value of $\approx 10^{-15}$, evidencing a statistical significant advantage of the cross-validation in this scenario of small samples in order to avoid information loss. On the other hand, in this setting AIC and BIC underestimate the true dimension even more frequently than for the larger sample size.

Summarizing, both permutation testing and CV provide more accurate results compared to AIC and BIC, with information criteria typically underestimating the true dimension of the reduction. Permutation testing seems to be a better procedure to infer d when the sample size is large enough. Nevertheless, cross-validation can provide a safer solution regarding information loss when the available data is limited.

5.4. Performance of the proposed method including regularization

Finally, we conducted a simulation study to assess the performance of the regularized version of the proposed method. Unlike the previous setting, the reduction depends now on a subset of the predictors only. We chose that the first four predictors conveyed information about the response, that is, $\boldsymbol{\alpha} = (\mathbf{A} \ \mathbf{0}_{2 \times p-4})^T$, with

$$\mathbf{A} = \begin{pmatrix} 1/2 & 1/2 & 1/2 & 1/2 \\ -1/2 & 1/2 & 1/2 & -1/2 \end{pmatrix}.$$

We set the values for the rest of the parameters as described in Section 5.2. The reduction was estimated using PFC for continuous predictors, the non-regularized method introduced in Section 4, and the regularized version proposed in Section 5. In all cases we used a polynomial basis with degree $r = 2$ for $\mathbf{f}_Y$. For each estimator, we computed the angle between the subspaces $\mathcal{S}_\alpha$ and $\mathcal{S}_{\hat{\alpha}}$ in each of 100 runs of the experiment. Figure 4 shows the obtained results for $p = 10$. It can be seen that the obtained angles are smaller than in the case where all the predictors are relevant. Since all the methods are applied to identical data, it is clear from the boxplot that the estimators specifically tailored to ordinal predictors still perform significantly better than the standard PFC approach. Moreover, for this situation, the regularized estimator (from now on, reg-PFCord) clearly proves to be superior to the non-regularized version.

To further study the performance of the proposed regularized estimator, another set of experiments was carried out in order to evaluate the stability of the subset of variables chosen by the algorithm. Denote by $S_0 \subset \{1, 2, \ldots, p\}$ the index set for the subset of variables that are truly relevant for describing the response Y and let S_0^c be its complement. Similarly, let $\widehat{S}$ be the subset of

14

Figure 4: Angle between span($\widehat{\boldsymbol{\alpha}}$) and span($\boldsymbol{\alpha}$), for $\widehat{\boldsymbol{\alpha}}$ obtained using PFC, the proposed approach for ordinal predictors without regularization, and the regularized approach introduced in Section 5.

variables chosen by the regularized estimator, in the sense that $\boldsymbol{\alpha}_j = \mathbf{0}$ for $j \notin \widehat{S}$, and let $\mathbf{Z}_{S_0}$ be the random vector with entries Z_j for $j \in S_0$, with a similar definition for $\mathbf{Z}_{S_0^c}$. We are interested in assessing: (i) $\Pr(S_0 \subset \widehat{S})$, as an indicator of the relevant variables that are indeed retained; (ii) the average cardinality of the set $\widehat{S}$ ($\#\widehat{S}$) as a measure of the amount of non-relevant variables that are preserved. Moreover, we are interested in evaluating how these performance measures vary for different amounts of correlation between $\mathbf{Z}_{S_0}$ and $\mathbf{Z}_{S_0^c}$. To measure this dependence we use the distance correlation measure as defined in [47, 46], denoted here by $dCor_n(\mathbf{Z}_{S_0}, \mathbf{Z}_{S_0^c})$ when it is computed from a sample of size n. This is a generalized nonparametric measure of correlation that is suitable for random vectors of different size and it does not require tuning any parameter for its computation. To control this quantity, we adjust the value of $\boldsymbol{\Delta}$ used to generate the data. In particular, we set $\boldsymbol{\Delta} = 4\mathbf{I}_p + \rho\boldsymbol{\alpha}\mathbf{B}\boldsymbol{\alpha}^T$ for different values of ρ. We obtained $dCor_n(\mathbf{Z}_{S_0}, \mathbf{Z}_{S_0^c}) \approx 0.3$ for $\rho = 0.2$, 0.5 for $\rho = 0.3$ and approximately 0.7 for $\rho = 0.5$. During all the experiments, the number of relevant variables in S_0 was held fixed at 4, with $p = 20$.

Table 2 shows the results obtained using 100 replicates of the experiment for each assessed condition. It can be seen that for a large enough sample ($n = 500$), the penalized versions of both PFCord and PFC achieve perfect accuracy in selecting the true active set when the predictors are not correlated each other. For moderate levels of correlation between the predictors ($\rho = 0.3$), using reg-PFCord the true active set S_0 is contained in the solution 96% of the time, with a very low fraction of false insertions. On the other hand, using the regularized version of the standard PFC (from now on, reg-PFC) in the same scenario allows picking the true active set of variables only 63% of the time. This difference is statistically significant at the 0.001 level. Moreover, the average number of variables picked by the reg-PFC is around 3.50, meaning that some information is typically lost in the procedure. For very high levels of correlation between the pre-

Table 2: Performance of variable selection algorithms when using `reg-PFCord` or `reg-PFC`.

n	ρ	reg-PFCord		reg-PFC	
		$\Pr(S_0 \subseteq \widehat{S})$	$\#\widehat{S}$	$\Pr(S_0 \subseteq \widehat{S})$	$\#\widehat{S}$
	0.0	0.98	3.87 (± 0.15)	0.92	3.88 (± 0.63)
200	0.2	0.90	3.91 (± 0.29)	0.13	2.36 (± 0.76)
	0.3	0.53	3.49 (± 0.61)	0.04	2.23 (± 0.46)
	0.5	0.21	2.86 (± 0.63)	0.01	2.07 (± 0.38)
	0.0	1.00	4.11 (± 0.00)	1.00	4.00 (± 0.00)
500	0.2	0.98	4.05 (± 0.10)	0.96	4.52 (± 0.73)
	0.3	0.96	3.92 (± 0.38)	0.63	3.52 (± 0.98)
	0.5	0.54	4.05 (± 1.08)	0.17	2.53 (± 0.83)

dictors ($\rho = 0.5$), the true set of active variables is picked only 17% of the time using `reg-PFC`, whereas `reg-PFCord` still finds it half the time. The loss of accuracy of `reg-PFCord` usually involved replacing one of the predictors in the true active set by a highly correlated alternative, mantaining the average cardinality of the estimated set $\widehat{S}$ close to 4. When the sample size is smaller ($n = 200$), the performance of `reg-PFC` for variable selection degrades much faster with the level of correlation than the ordinal counterpart. When the predictors are uncorrelated, `reg-PFC` picks the true active set S_0 92% of the time, whereas `reg-PFCord` does it 98% of the time. But for low levels of correlation between the predictors ($\rho = 0.2$), the performance of `reg-PFC` decreases quickly to 13%, while for the ordinal counterpart it is still greater than 90%. For very high levels of correlation, both procedures tend to underestimate the number of relevant variables, with this trend being stronger for the `reg-PFC`. These results show that using the proposed method especially tailored to ordinal data provides significantly higher accuracy when variable selection is needed.

6. Real data analysis: SES index construction

For many social protection and welfare programs carried out by governments and NGOs, a classification of households or individuals into different socio-economic groups is required. For example, over the last decades, it has become common in developing countries that governments establish economic aid programs focused on the most deprived households. This tailoring of the aid is achieved via a focalization index, which basically mounts to a Socio-Economic Status (*SES*) index, as it is commonly known in the related literature. In particular, in many Latin American countries, this focalization index (called *Índice de Focalización de Pobreza*) has been used to implement several programs to reduce poverty (e.g., the CAS in Chile, Sisben in Colombia, SISFOH in Perú, Tekoporá in Paraguay, SIERP in Honduras, and PANES in Uruguay, among others).

Income or consumption expenditures constitute a traditional focus of poverty analysis, and some countries take an income-based poverty line from a household survey to infer the socioeconomic situation of the population [38, 42]. However, the collection of income data presents many problems in terms of unavailability or unreliability [48, 18]. For this reason, asset-based indexes are often constructed as a proxy of income, taking into account some housing and household

16

variables that are easier to observe. This proxy, usually called the SES index, is most of the time computed using principal component analysis (`PCA`) [37, 24]. Since the observable variables used to construct these indexes are ordered categorical variables, [28] proposed a variant of `PCA` adapted for ordinal data using polychoric correlations between predictors instead of the standard covariance matrix.

In this example, we provide a different approach to the construction of an SES index, based on the proposed SDR methodology for ordinal data. The main idea is to obtain a single index to predict a unidimensional measure of some socioeconomic aspect, such as household income or the poverty condition. Therefore, in this case we fix the dimension of the reduction to be $d = 1$ and derive the index as a normalized version of the supervised dimension reduction $\widehat{\boldsymbol{\alpha}}^T \mathrm{E}(\mathbf{Z}|\mathbf{X}) \in \mathbb{R}$. Unlike `PCA`-based indexes, this new approach uses information about the response under analysis.

The data comes from the microdata of the *Encuesta Permanente de Hogares* (EPH) of Argentina, taking the fourth trimester of 2013. The EPH is the main household survey in Argentina and is carried out by the *Instituto Nacional de Estadísticas y Censos* (INDEC). We consider nine ordinal variables about household living conditions, and two socio-economic variables of heads of households (educational attainment and work situation). More details about these variables can be found in Appendix C. To take into account regional heterogeneity, we estimate separate SES indexes for the following five regions: the metropolitan area of Buenos Aires ($n = 2351$ households), Humid Pampas ($n = 5003$), the Argentine Northwest ($n = 2852$), the Northeast ($n = 1594$), and Patagonia ($n = 2398$). Two cases with different types of response are considered: a continuous one, household income per capita (`ipcf`), and a binary one based on income (`poverty`) that indicates whether a household is poor or not. We are interested in demonstrating that the proposed `reg-PFCord` provides a superior alternative to a `PCA`-based method for constructing an SES index, while retaining a predictive power comparable to the full set of predictors. To do this, the predictive performance of the proposed response-driven index is compared to the following strategies:

- The full set of predictors is included without dimension reduction and they are treated as continuous (metric) predictors. We will refer to this apporach as `FULL`.

- The full set of predictors is included without dimension reduction and they are treated through dummy variables. We will refer to this approach as `FULL-I`.

- The full set of predictors is considered but using a group-lasso-type procedure for ordinal predictors that induces variable selection [21]. We will refer to this approach as `LASSOord`.

- A variant of `PCA` tailored to ordered categorical predictors using polychoric correlations [28]. We will refer to this method as `PCApoly`.

- A nonlinear variant of `PCA` that uses special scaling to take into account the ordered categories [34]. We will refer to this approach as `NLPCA`.

- Standard moments-based sufficient dimension reduction methods `SIR`, `SAVE` and `DR`.

The first three strategies are included in order to provide a reference for the performance achievable using the full set of predictors, but it should be clear that they do not provide an index. Actually, only the last three approaches in the list are competing methods for SES-index extraction. Among them, `PCApoly` and `NLPCA` can deal explicitly with ordinal predictors and they will be further compared later.

For every strategy, we fit a logistic regression for the `poverty` response and a linear regression for the `ipcf` response. When computing the reduction, we use a different choice of $\mathbf{f}_Y$ for each type of outcome. For the continuous response, we use a polynomial basis with degree $r = 2$. For the binary response, $\mathbf{f}_Y$ is simply a centered indicator variable. The data was partitioned into ten disjoint sets to allow for ten replications of the experiment. In each run, one of the subsets was used as the test set, while the rest of them formed the training sample. Averaged 10-fold cross-validation MSEs obtained with each method are shown in Table 3, along with the corresponding standard deviations.

From the table it can be seen that for the continuous response, using dummy variables for the full set of predictors, as in `FULL-I` and `LASSOord`, is more effective than considering the full set of predictors as continuous variables (`FULL`). Among the SES indexes, scores show that `reg-PFCord` is superior to both `PCApoly` and `NLPCA`, with the latter being slightly superior to the former. In addition, all these methods specifically targeted to ordinal predictors perform better than standard sufficient dimension reduction methods represented here by `SIR`, `SAVE` and `DR`, which were originally aimed to continuous predictors only. Moreover, prediction errors obtained with `reg-PFCord` are very close to those attained with `FULL` across all the regions.

On the other hand, it is interesting to see that for the binary outcome, unlike the continuous case, `FULL` performs better than `FULL-I` and `LASSOord`. Among the SES indexes, the predictive performance of `reg-PFCord` is again very close to or identical with that of `FULL`, and it outperforms the PCA-based methods and standard moment-based sufficient dimension reduction methods in this scenario too. Moreover, `reg-PFCord` outperforms `LASSOord` for three of the regions when the discrete response is considered. It should also be remarked that, unlike `LASSOord`, obtaining the indexes from the SDR-based techniques allows us to use any predictive method.

As an illustration of the obtained fit, Figure 5 shows marginal model plots for the regression of `ipcf` on the SES index obtained for the whole database. A quadratic term SES^2 was added to correct for curvature in the estimated regression function and the response was transformed by $ipcf \leftarrow ipcf^{1/3}$ following a Box–Cox transformation analysis. It can be seen that for SES modeled using `PCA` on polychoric correlations, the index values are concentrated mainly in the interval $[0.5; 1.0]$ whereas, for SES modeled using `reg-PFCord`, the spread of the index values is more regular over the whole interval $[0; 1.0]$. This allows for a better fit of the linear model, as shown by an R^2 value of 0.302 compared to 0.231 obtained with the SES index based on `PCA`.

Tables 4 and 5 show the estimated coefficient vectors that define the SES index using `ipcf` and `poverty` as response variables, respectively. To keep the analysis clear, standard methods not targeted to ordinal predictors like `SIR`, `SAVE` and `DR` were not included, since they showed a clearly inferior performance for prediction in Table 3. Note that for the proposed method, some of the elements of $\widehat{\alpha}$ have been pushed to zero in the regularized estimation, whereas for `PCApoly` and `NLPCA` only *working hours* seems not to be relevant for constructing the index. In addition, several differences can be appreciated between the `reg-PFCord` and `PCA`-based approaches

18

Figure 5: Marginal model plots showing the fit of linear model of income as a function of the obtained SES index. `PCA`-based SES index is shown on the left and the proposed method on the right.

(i.e., `PCApoly` and `NLPCA`) from the reported results. First, the relative importance of each predictor in the SES index obtained is different for the two methods. For instance, *overcrowding* obtains the highest weight with `reg-PFCord` across all the regions for both responses, whereas *toilet facility* and *water location* appear as the most important in index construction based on `PCApoly`, and *toilet facility* and *toilet drainage* for `NLPCA`. Second, SES indexes constructed using both `PCA` methods give similar weights to the predictors across the different regions. On the other hand, SES indexes based on `reg-PFCord` capture the regional economic divergence explained by different factor endowments, productivity, activity levels and regional economic growth patterns. Moreover, in the richest Argentinian urban regions (specifically, Buenos Aires and Humid Pampas) the regularized estimation of `reg-PFCord` often sets to zero the variables with more weight in `PCA`-derived SES index. This difference is appealing, since these regions have in general better services and public infrastructure, so that variables related to drainage, source of water and toilet facility are less important for measuring socio-economic status. In this way, other variables, such as overcrowding or schooling, are needed in order to have a better SES index to predict household income. In the same line, for regions with higher levels of poverty (Northwest and Northeast) the `reg-PFCord`-based SES index shows that other variables, such as *housing location*, *source of drinking water* or *water location*, become important for determining socio-economic status.

Comparing both `PCA` methods, it can be noted that the `NLPCA` is more sensitive to the regional heterogeneity than is the `PCApoly`, but differences in the index weights compared to those of the `reg-PFCord` remain substantial. Additionally, it can be appreciated that SES indexes obtained using `reg-PFCord` are sensitive to the response variable used to characterize a social phenomenon of interest. For example, in Buenos Aires, Humid Pampas and Patagonia, *schooling* has a considerable weight in the SES index to explain per capita income but not to predict poverty. This makes sense, since for these richest regions it is easier for all the population to get access to basic levels of schooling. On the other hand, the decision to pursue higher levels of education is often driven by income. Moreover, for these richest regions, some variables, such as *toilet drainage*, *toilet facility* or *toilet sharing* become relevant to explaining whether a household is poor or not (following poverty line criteria). Such differences cannot be captured by an SES index based on `PCA` methods.

19

Table 3: 10-fold MSE for SES index

Response	Method	Buenos Aires	Humid Pampas	Northwest	Northeast	Patagonia
		\multicolumn{5}{c}{Prediction Errors -MSE}				
Per capita Income	REG-PFCORD	7.29	4.72	4.73	3.34	12.80
(continuous)		(2.91)	(1.73)	(2.69)	(1.52)	(3.72)
	PCApoly	7.60	5.10	5.07	3.68	14.7
		(2.45)	(0.90)	(1.77)	(0.90)	(4.01)
	NLPCA	7.38	4.95	4.89	3.52	13.67
		(2.29)	(0.61)	(1.48)	(0.65)	(3.71)
	SIR	7.36	6.21	5.61	6.15	14.41
		(4.38)	(4.73)	(3.72)	(0.86)	(1.16)
	SAVE	9.06	6.09	5.87	4.12	16.21
		(4.04)	(0.99)	(2.73)	(1.24)	(3.74)
	DR	8.96	5.76	5.85	4.10	15.95
		(3.96)	(1.02)	(2.74)	(1.24)	(3.70)
	FULL	7.22	4.69	4.68	3.32	13.14
		(3.50)	(0.88)	(2.47)	(0.93)	(3.34)
	FULL-I	7.01	4.52	4.48	3.08	12.92
		(2.46)	(0.83)	(1.74)	(0.76)	(3.80)
	LASSOord	7.00	4.51	4.42	3.05	12.88
		(2.46)	(0.83)	(1.77)	(0.76)	(3.84)
Poverty	reg-PFCord	0.204	0.169	0.278	0.288	0.126
(discrete)		(0.017)	(0.012)	(0.031)	(0.029)	(0.025)
	PCApoly	0.213	0.188	0.324	0.357	0.133
		(0.024)	(0.020)	(0.026)	(0.053)	(0.027)
	NLPCA	0.212	0.188	0.325	0.358	0.134
		(0.023)	(0.019)	(0.025)	(0.055)	(0.027)
	SIR	0.229	0.204	0.366	0.392	0.1314
		(0.018)	(0.010)	(0.032)	(0.047)	(0.022)
	SAVE	0.229	0.204	0.362	0.378	0.133
		(0.019)	(0.009)	(0.031)	(0.042)	(0.022)
	DR	0.230	0199	0.357	0.364	0.133
		(0.019)	(0.009)	(0.035)	(0.040)	(0.023)
	FULL	0.202	0.162	0.274	0.287	0.129
		(0.021)	(0.008)	(0.026)	(0.036)	(0.020)
	FULL-I	0.206	0.171	0.286	0.298	0.126
		(0.023)	(0.021)	(0.024)	(0.065)	(0.028)
	LASSOord	0.206	0.171	0.286	0.298	0.129
		(0.023)	(0.021)	(0.024)	(0.065)	(0.027)

Note: standard deviations in parentheses. Database: EPH (2013)

Table 4: Comparison of SES index results for `ordinal PCF` and `PCApoly` to predict household per capita income.

Variables	Buenos Aires			Humid Pampas			Northwest		
	reg-PFCord	PCApoly	NLPCA	reg-PFCord	PCApoly	NLPCA	reg-PFCord	PCApoly	NLPCA
housing location	0	-0.1690	-0.0943	0	-0.1903	-0.0976	-0.1314	-0.1068	-0.0835
housing quality	-0.1985	-0.3768	-0.2199	0.2591	-0.3557	-0.1985	0	-0.3278	-0.1849
sources of cooking fuel	-0.4646	-0.3788	-0.2080	0.3627	-0.3609	-0.1678	-0.1070	-0.3287	-0.1582
overcrowding	-0.7272	-0.2888	-0.1788	0.8300	-0.2351	-0.1329	-0.8798	-0.1991	-0.1194
schooling	-0.2676	-0.2275	-0.1474	0.2668	-0.2075	-0.1135	-0.3614	-0.2197	-0.1201
toilet drainage	-0.0873	-0.3381	-0.2047	0	-0.3519	-0.2333	-0.0901	-0.3623	-0.2462
toilet facility	0	-0.4061	-0.2246	0	-0.4105	-0.2411	0	-0.4217	-0.2545
toilet sharing	0	-0.2759	-0.1186	0	-0.3176	-0.1699	0	-0.2579	-0.1334
water location	0.1700	-0.3918	-0.1790	0	-0.3933	-0.1941	-0.2054	-0.4202	-0.2309
water source	0	-0.2023	-0.1033	0	-0.2461	-0.1129	0	-0.3646	-0.1374
working hours	-0.3283	0	0	0.2029	0	0	-0.1277	0	0

Variables	Northeast			Patagonia					
	reg-PFCord	PCApoly	NLPCA	reg-PFCord	PCApoly	NLPCA			
housing location	-0.1509	-0.1809	-0.0978	-0.1149	-0.1437	-0.0981			
housing quality	0	-0.3727	-0.2130	-0.3046	-0.3258	-0.1844			
sources of cooking fuel	-0.0742	-0.1648	-0.0646	-0.1797	-0.4026	-0.1810			
overcrowding	-0.8496	-0.2052	-0.1040	-0.7263	-0.2207	-0.0984			
schooling	-0.3507	-0.1869	-0.1009	-0.3670	-0.1284	-0.0516			
toilet drainage	0	-0.3572	-0.2573	-0.1383	-0.4122	-0.2566			
toilet facility	0	-0.4383	-0.2735	0	-0.4376	-0.2622			
toilet sharing	-0.2284	-0.2921	-0.1344	-0.1204	-0.2937	-0.1734			
water location	0	-0.4227	-0.2377	0.1877	-0.4169	-0.2196			
water source	-0.2574	-0.3733	-0.1384	0.2473	-0.1525	-0.0522			
working hours	-0.0922	0	0	-0.2637	0	0			

Table 5: Comparison of SES index results for `ordinal PCF` and `PCApoly` for a discrete response (poverty).

Variables	Buenos Aires			Humid Pampas			Northwest		
	reg-PFCord	PCApoly	NLPCA	reg-PFCord	PCApoly	NLPCA	reg-PFCord	PCApoly	NLPCA
housing location	0	-0.1690	-0.0943	0	-0.1903	-0.0976	-0.2434	-0.1068	-0.0835
housing quality	-0.4033	-0.3768	-0.2199	0.3347	-0.3557	-0.1985	0	-0.3278	-0.1849
sources of cooking fuel	-0.5240	-0.3788	-0.2080	0.3579	-0.3609	-0.1678	0	-0.3287	-0.1582
overcrowding	-0.7076	-0.2888	-0.1788	0.7216	-0.2351	-0.1329	-0.7939	-0.1991	-0.1194
schooling	0	-0.2275	-0.1474	0	-0.2075	-0.1135	-0.2094	-0.2197	-0.1201
toilet drainage	0	-0.3381	-0.2047	0	-0.3519	-0.2333	0	-0.3623	-0.2462
toilet facility	0	-0.4061	-0.2246	0.3990	-0.4105	-0.2411	-0.1528	-0.4217	-0.2545
toilet sharing	-0.1836	-0.2759	-0.1186	0	-0.3176	-0.1699	0	-0.2579	-0.1334
water location	-0.1208	-0.3918	-0.1790	0.2647	-0.3933	-0.1941	-0.4933	-0.4202	-0.2309
water source	0	-0.2023	-0.1033	0	-0.2461	-0.1129	0	-0.3646	-0.1374
working hours	-0.1173	0	0	0.0990	0	0	0	0	0
	Northeast			Patagonia					
	reg-PFCord	PCApoly	NLPCA	reg-PFCord	PCApoly	NLPCA			
housing location	-0.1982	-0.1809	-0.0978	-0.1187	-0.1437	-0.0981			
housing quality	0	-0.3727	-0.2130	-0.3693	-0.3258	-0.1844			
sources of cooking fuel	-0.2509	-0.1648	-0.0646	-0.2788	-0.4026	-0.1810			
overcrowding	-0.7063	-0.2052	-0.1040	-0.3987	-0.2207	-0.0984			
schooling	-0.1442	-0.1869	-0.1009	-0.0766	-0.1284	-0.0516			
toilet drainage	0	-0.3572	-0.2573	-0.1887	-0.4122	-0.2566			
toilet facility	0	-0.4383	-0.2735	-0.1313	-0.4376	-0.2622			
toilet sharing	-0.3477	-0.2921	-0.1344	-0.1289	-0.2937	-0.1734			
water location	0	-0.4227	-0.2377	0.2785	-0.4169	-0.2196			
water source	-0.5071	-0.3733	-0.1384	0.6585	-0.1525	-0.0522			
working hours	0	0	0	-0.1626	0	0			

7. Conclusions

The approximate expectation-maximization (EM) algorithm presented here for dimension reduction in regression problems with ordinal predictors is proved to outperform the standard inverse regression methods derived for continuous predictors, both in simulation settings and with real data sets involving ordered categorical predictors. Experiments showed that this advantage is emphasized in variable selection applications, where the proposed method clearly outperforms its counterpart for continuous data when the counterpart is naively applied to ordinal predictors. This is not a minor issue since many analyses in the applied sciences usually treat them as continuous variables, not taking into account their discrete nature. Moreover, it has better computing efficiency due to the proposed approximate EM algorithm's rendering the method feasible for a much larger set of problems compared to using the exact computation of the truncated moments. This savings also allows permutation testing and cross validation procedures for inferring the dimension of the eduction, which proved reasonably accurate in simulations. Finally, the application of the proposed methodology to socio-economic status (SES) index construction showed many advantages over common PCA-based indexes. In particular, the method not only helps get better predictions but also allows understanding the relations between the predictors and the response. More precisely, for the SES index, it gives varying weights capturing regional, historical and/or cultural differences, as well as various social measurement criteria (such as household per capita income or the poverty line), which it is not possible with PCA-based methods. This property of the proposed method has relevant implications for the applied social analysis.

Considering that many applications involve predictors of different natures (such as ordinal, continuous, and binary variables), further developments in SDR-based methods could be in this direction. The Principal Fitted Components (PFC) method for ordinal variables here proposed constitutes the first step to this extension. In particular, the combination of ordinal and continuous predictors could be treated by taking all of them as continuous variables, where some of them are latent and the others are observable. Then, from the results here found, the reduction is identified, and the parameters can be estimated via maximum likelihood using the EM method on the latent variables and the PFC conventional method on the observed continuous variables. Nevertheless, the combination with binary variables requires a more exhaustive treatment, taking into account that the assumption of the existence of a latent normal variable on a binary variable may be naive and not make sense when the binary variable does not have a natural order (e.g., gender). Therefore, it is necessary to find a proper representation for the binary predictors, and search for a way to combine them with the other types of variables.

Appendix A. The EM algorithm

In order to simplify the notation we will always omit the conditioning on $\boldsymbol{\Omega}^{(k-1)}$ when taking expectations. We will also omit the conditioning on some variables in the subscript. For instance, for any function g, we will call

$$E_{\mathbf{z}_i}\big(g(\mathbf{z}_i)|y_i\big) \doteq E_{\mathbf{z}_i|y_i,\boldsymbol{\Omega}^{(k-1)}}(g(\mathbf{z}_i)|y_i;\boldsymbol{\Omega}^{(k-1)}),$$
$$E_{\mathbf{z}_i}(g(\mathbf{z}_i)|\mathbf{x}_i,y_i) \doteq E_{\mathbf{z}_i|\mathbf{x}_i,y_i;\boldsymbol{\Omega}^{(k-1)}}(g(\mathbf{z}_i)|\mathbf{x}_i,y_i;\boldsymbol{\Omega}^{(k-1)}). \tag{A.1}$$

In order to obtain an explicit form of Q we compute the conditional expectation of the joint log-likelihood. Following (A.1), we will write

$$Q(\boldsymbol{\Omega}|\boldsymbol{\Omega}^{(k-1)}) = \sum_{i=1}^{n} E_{\mathbf{z}_i|y_i;\boldsymbol{\Omega}^{(k-1)}}\Big[\log f_{\mathbf{x}_i,\mathbf{z}_i}(\mathbf{x}_i,\mathbf{z}_i|y_i;\boldsymbol{\Omega})\big|y_i;\boldsymbol{\Omega}^{(k-1)}\Big] = \sum_{i=1}^{n} E_{\mathbf{z}_i}\Big[\log f_{\mathbf{x}_i,\mathbf{z}_i}(\mathbf{x}_i,\mathbf{z}_i|y_i;\boldsymbol{\Omega})\big|y_i\Big].$$

Therefore,

$$\begin{aligned}
Q(\boldsymbol{\Omega}|\boldsymbol{\Omega}^{(k-1)}) &= \sum_{i=1}^{n} E_{\mathbf{z}_i}\Big[\log f_{\mathbf{x}_i,\mathbf{z}_i}(\mathbf{x}_i,\mathbf{z}_i|y_i;\boldsymbol{\Omega})\big|y_i\Big] \\
&= \sum_{i=1}^{n} E_{\mathbf{z}_i}\Big[\log\Big((2\pi)^{-p/2}|\boldsymbol{\Delta}|^{-1/2}e^{-\frac{1}{2}\operatorname{tr}(\boldsymbol{\Delta}^{-1}(\mathbf{z}_i-\boldsymbol{\Delta}\boldsymbol{\alpha}\boldsymbol{\xi}\bar{\mathbf{f}}_{y_i})(\mathbf{z}_i-\boldsymbol{\Delta}\boldsymbol{\alpha}\boldsymbol{\xi}\bar{\mathbf{f}}_{y_i})^T)}I_{\{\mathbf{z}_i\in C(\mathbf{x}_i,\boldsymbol{\Theta})\}}\Big)\big|y_i\Big] \\
&= -\frac{pn}{2}\log(2\pi) - \frac{n}{2}\log|\boldsymbol{\Delta}| \\
&\quad -\frac{n}{2}\operatorname{tr}\Big[\boldsymbol{\Delta}^{-1}\Big(\frac{1}{n}\sum_{i=1}^{n}E_{\mathbf{z}_i}(\mathbf{z}_i\mathbf{z}_i^T|\mathbf{x}_i,y_i) - \frac{2}{n}\boldsymbol{\Delta}\boldsymbol{\alpha}\boldsymbol{\xi}\sum_{i=1}^{n}\bar{\mathbf{f}}_{y_i}E_{\mathbf{z}_i}(\mathbf{z}_i^T|\mathbf{x}_i,y_i) + \frac{1}{n}\boldsymbol{\Delta}\boldsymbol{\alpha}\boldsymbol{\xi}\sum_{i=1}^{n}\bar{\mathbf{f}}_{y_i}f_{y_i}^T\boldsymbol{\xi}^T\boldsymbol{\alpha}^T\boldsymbol{\Delta}\Big)\Big] \\
&= -\frac{pn}{2}\log(2\pi) - \frac{n}{2}\log|\boldsymbol{\Delta}| - \frac{n}{2}\operatorname{tr}\Big[\boldsymbol{\Delta}^{-1}(\mathbf{S} - \frac{2\boldsymbol{\Delta}\boldsymbol{\alpha}\boldsymbol{\xi}\mathbf{F}^T\mathbf{M}}{n} + \frac{\boldsymbol{\Delta}\boldsymbol{\alpha}\boldsymbol{\xi}\mathbf{F}^T\mathbf{F}\boldsymbol{\xi}^T\boldsymbol{\alpha}^T\boldsymbol{\Delta}}{n})\Big], \tag{A.2}
\end{aligned}$$

where $\mathbf{S}\in\mathbb{R}^{p\times p}$, $\mathbf{F}\in\mathbb{R}^{n\times r}$ and $\mathbf{M}\in\mathbb{R}^{n\times p}$ are given by $\mathbf{S} = \frac{1}{n}\sum_{i=1}^{n}E_{\mathbf{z}_i}(\mathbf{z}_i\mathbf{z}_i^T|\mathbf{x}_i,y_i)$, $\mathbf{F}^T = [\bar{\mathbf{f}}_{y_1},\dots,\bar{\mathbf{f}}_{y_n}]$ and $\mathbf{M}^T = [E_{\mathbf{z}_1}(\mathbf{z}_1|\mathbf{x}_1,y_1),\dots,E_{\mathbf{z}_n}(\mathbf{z}_n|\mathbf{x}_n,y_n))]$, respectively.

Appendix A.1. Maximizing the Q-function (A.2).

From (A.2), we have

$$Q(\boldsymbol{\Delta},\boldsymbol{\alpha},\boldsymbol{\xi}) = -\frac{pn}{2}\log(2\pi) - \frac{n}{2}\log|\boldsymbol{\Delta}| - \frac{n}{2}\operatorname{tr}(\boldsymbol{\Delta}^{-1}\mathbf{S}) + \operatorname{tr}(\boldsymbol{\alpha}\boldsymbol{\xi}\mathbf{F}^T\mathbf{M}) - \frac{1}{2}\operatorname{tr}(\boldsymbol{\alpha}\boldsymbol{\xi}\mathbf{F}^T\mathbf{F}\boldsymbol{\xi}^T\boldsymbol{\alpha}^T\boldsymbol{\Delta}).$$

Since Q is a quadratic form in $\boldsymbol{\xi}$, the maximum will be attained at $\boldsymbol{\xi}^{(k)} = (\boldsymbol{\alpha}^T\boldsymbol{\Delta}\boldsymbol{\alpha})^{-1}\boldsymbol{\alpha}^T\mathbf{M}^T\mathbf{F}(\mathbf{F}^T\mathbf{F})^{-1}$. Replacing $\boldsymbol{\xi}^{(k)}$ in the Q-function, we obtain the partial log-likelihood

$$\begin{aligned}
Q(\boldsymbol{\Delta}^{-1},\boldsymbol{\alpha}) &= -\frac{pn}{2}\log(2\pi) - \frac{n}{2}\log|\boldsymbol{\Delta}| - \frac{n}{2}\operatorname{tr}(\boldsymbol{\Delta}^{-1}\mathbf{S}) + \frac{1}{2}\operatorname{tr}\big[(\boldsymbol{\alpha}^T\boldsymbol{\Delta}\boldsymbol{\alpha})^{-1}\boldsymbol{\alpha}^T\mathbf{M}^T\mathbf{F}(\mathbf{F}^T\mathbf{F})^{-1}\mathbf{F}^T\mathbf{M}\boldsymbol{\alpha}\big] \\
&= -\frac{pn}{2}\log(2\pi) - \frac{n}{2}\log|\boldsymbol{\Delta}| - \frac{n}{2}\operatorname{tr}(\boldsymbol{\Delta}^{-1}\mathbf{S}) + \frac{n}{2}\operatorname{tr}\big[(\boldsymbol{\alpha}^T\boldsymbol{\Delta}\boldsymbol{\alpha})^{-1}\boldsymbol{\alpha}^T\mathbf{S}_{\text{fit}}\boldsymbol{\alpha}\big], \tag{A.3}
\end{aligned}$$

where $\mathbf{S}_{\text{fit}} = \frac{1}{n}\mathbf{M}^T\mathbf{F}(\mathbf{F}^T\mathbf{F})^{-1}\mathbf{F}^T\mathbf{M}$. In order to maximize Q with respect to $\boldsymbol{\Delta}^{-1}$ observe that, by Proposition 5.14 in [19], if $\boldsymbol{\alpha}\in\mathbb{R}^{p\times d}$ is fixed and $\boldsymbol{\alpha}_0\in\mathbb{R}^{p\times(p-d)}$ is the semi-orthogonal complement of

$\boldsymbol{\alpha}$, we have a one to one correspondence between $\boldsymbol{\Delta}^{-1}$ and $(\mathbf{H}_1, \mathbf{H}_2, \mathbf{H}_3)$, with $\mathbf{H}_1 = \boldsymbol{\alpha}^T \boldsymbol{\Delta} \boldsymbol{\alpha}$; $\mathbf{H}_2 = (\boldsymbol{\alpha}_0^T \boldsymbol{\Delta}^{-1} \boldsymbol{\alpha}_0)^{-1}$ and $\mathbf{H}_3 = (\boldsymbol{\alpha}^T \boldsymbol{\Delta} \boldsymbol{\alpha})^{-1} \boldsymbol{\alpha}^T \boldsymbol{\Delta} \boldsymbol{\alpha}_0$. From [40] and [12] we have

$$\boldsymbol{\Delta}^{-1} = \boldsymbol{\alpha}(\boldsymbol{\alpha}^T \boldsymbol{\Delta} \boldsymbol{\alpha})^{-1} \boldsymbol{\alpha}^T + \boldsymbol{\Delta}^{-1} \boldsymbol{\alpha}_0 (\boldsymbol{\alpha}_0^T \boldsymbol{\Delta}^{-1} \boldsymbol{\alpha}_0)^{-1} \boldsymbol{\alpha}_0^T \boldsymbol{\Delta}^{-1}, \tag{A.4}$$

$$|\boldsymbol{\Delta}| = |\boldsymbol{\alpha}_0^T \boldsymbol{\Delta}^{-1} \boldsymbol{\alpha}_0|^{-1} |\boldsymbol{\alpha}^T \boldsymbol{\Delta} \boldsymbol{\alpha}|. \tag{A.5}$$

Now the identity $(\boldsymbol{\alpha}^T \boldsymbol{\Delta} \boldsymbol{\alpha})^{-1} \boldsymbol{\alpha}^T \boldsymbol{\Delta} \boldsymbol{\alpha}_0 = -\boldsymbol{\alpha}^T \boldsymbol{\Delta}^{-1} \boldsymbol{\alpha}_0 (\boldsymbol{\alpha}_0^T \boldsymbol{\Delta}^{-1} \boldsymbol{\alpha}_0)^{-1}$ implies that

$$\boldsymbol{\alpha}^T \boldsymbol{\Delta}^{-1} \boldsymbol{\alpha}_0 = -(\boldsymbol{\alpha}^T \boldsymbol{\Delta} \boldsymbol{\alpha})^{-1} \boldsymbol{\alpha}^T \boldsymbol{\Delta} \boldsymbol{\alpha}_0 \boldsymbol{\alpha}_0^T \boldsymbol{\Delta}^{-1} \boldsymbol{\alpha}_0 = -\mathbf{H}_3 \mathbf{H}_2^{-1},$$

which, together with $\boldsymbol{\alpha}\boldsymbol{\alpha}^T + \boldsymbol{\alpha}_0 \boldsymbol{\alpha}_0^T = \mathbf{I}_d$,

$$\boldsymbol{\Delta}^{-1} \boldsymbol{\alpha}_0 = \boldsymbol{\alpha}\boldsymbol{\alpha}^T \boldsymbol{\Delta}^{-1} \boldsymbol{\alpha}_0 + \boldsymbol{\alpha}_0 \boldsymbol{\alpha}_0^T \boldsymbol{\Delta}^{-1} \boldsymbol{\alpha}_0 = -\boldsymbol{\alpha}\mathbf{H}_3 \mathbf{H}_2^{-1} + \boldsymbol{\alpha}_0 \mathbf{H}_2^{-1} = (-\boldsymbol{\alpha}\mathbf{H}_3 + \boldsymbol{\alpha}_0)\mathbf{H}_2^{-1}.$$

With all this together in (A.4) we get

$$\boldsymbol{\Delta}^{-1} = \boldsymbol{\alpha}\mathbf{H}_1^{-1} \boldsymbol{\alpha}^T + (-\boldsymbol{\alpha}\mathbf{H}_3 + \boldsymbol{\alpha}_0)\mathbf{H}_2^{-1}(-\mathbf{H}_3^T \boldsymbol{\alpha}^T + \boldsymbol{\alpha}_0^T), \tag{A.6}$$

Therefore, finding $\widehat{\mathbf{H}}_i$, $i = 1, 2, 3$ is equivalent to finding $(\boldsymbol{\Delta}^{-1})^{(k)}$. In order to write the Q-function in terms of $\mathbf{H}_1$, $\mathbf{H}_2$ and $\mathbf{H}_3$, let us write $\log|\boldsymbol{\Delta}|$ and $\mathrm{tr}(\boldsymbol{\Delta}^{-1}\mathbf{S})$ in terms of them. And therefore, using (A.6) and (A.5), the Q-function is then written in terms of $\mathbf{H}_1$, $\mathbf{H}_2$, and $\mathbf{H}_3$ as

$$Q(\mathbf{H}_1, \mathbf{H}_2, \mathbf{H}_3, \boldsymbol{\alpha}) = -\frac{pn}{2}\log(2\pi) - \frac{n}{2}\log|\mathbf{H}_1| - \frac{n}{2}\log|\mathbf{H}_2| - \frac{n}{2}\mathrm{tr}\left[\boldsymbol{\alpha}\mathbf{H}_1^{-1}\boldsymbol{\alpha}^T(\mathbf{S} - \mathbf{S}_{\text{fit}})\right]$$
$$- \frac{n}{2}\mathrm{tr}\left[(\boldsymbol{\alpha}_0 - \boldsymbol{\alpha}\mathbf{H}_3)\mathbf{H}_2^{-1}(\boldsymbol{\alpha}_0^T - \mathbf{H}_3^T\boldsymbol{\alpha}^T)\mathbf{S}\right]. \tag{A.7}$$

Now, since Q is quadratic in $\mathbf{H}_3$, the maximum of Q for $\mathbf{H}_3$ is attained at

$$\widehat{\mathbf{H}}_3 = (\boldsymbol{\alpha}^T\mathbf{S}\boldsymbol{\alpha})^{-1}(\boldsymbol{\alpha}^T\mathbf{S}\boldsymbol{\alpha}_0). \tag{A.8}$$

Replacing (A.8) in (A.7) and calling $\mathbf{S}_{\text{res}} = \mathbf{S} - \mathbf{S}_{\text{fit}}$ (which is semidefinite positive), we have the partial log-likelihood function

$$Q(\mathbf{H}_1, \mathbf{H}_2, \boldsymbol{\alpha}) = -\frac{pn}{2}\log(2\pi) - \frac{n}{2}\log|\mathbf{H}_1| - \frac{n}{2}\log|\mathbf{H}_2| - \frac{n}{2}\mathrm{tr}\left[\mathbf{H}_1^{-1}\boldsymbol{\alpha}^T\mathbf{S}_{\text{res}}\boldsymbol{\alpha}\right]$$
$$- \frac{n}{2}\mathrm{tr}\left[\mathbf{H}_2^{-1}(\boldsymbol{\alpha}_0^T - \boldsymbol{\alpha}_0^T\mathbf{S}\boldsymbol{\alpha}(\boldsymbol{\alpha}^T\mathbf{S}\boldsymbol{\alpha})^{-1}\boldsymbol{\alpha}^T)\mathbf{S}(\boldsymbol{\alpha}_0 - \boldsymbol{\alpha}(\boldsymbol{\alpha}^T\mathbf{S}\boldsymbol{\alpha})^{-1}\boldsymbol{\alpha}^T\mathbf{S}\boldsymbol{\alpha}_0)\right].$$

The maximum of Q over $\mathbf{H}_1$ and $\mathbf{H}_2$ is attained at

$$\widehat{\mathbf{H}}_1 = \boldsymbol{\alpha}^T\mathbf{S}_{\text{res}}\boldsymbol{\alpha};$$
$$\widehat{\mathbf{H}}_2 = \boldsymbol{\alpha}_0^T\mathbf{S}\boldsymbol{\alpha}_0 - \boldsymbol{\alpha}_0^T\mathbf{S}\boldsymbol{\alpha}(\boldsymbol{\alpha}^T\mathbf{S}\boldsymbol{\alpha})^{-1}\boldsymbol{\alpha}^T\mathbf{S}\boldsymbol{\alpha}_0 = (\boldsymbol{\alpha}_0^T\mathbf{S}^{-1}\boldsymbol{\alpha}_0)^{-1}.$$

After substitution of the maximum for $\mathbf{H}_1$, $\mathbf{H}_2$, and $\mathbf{H}_3$ into (A.6), we get that the maximum for $\boldsymbol{\Delta}^{-1}$ is attained at

$$(\boldsymbol{\Delta}^{-1})^{(k)} = \mathbf{S}^{-1} + \boldsymbol{\alpha}(\boldsymbol{\alpha}^T\mathbf{S}_{\text{res}}\boldsymbol{\alpha})^{-1}\boldsymbol{\alpha}^T - \boldsymbol{\alpha}(\boldsymbol{\alpha}^T\mathbf{S}\boldsymbol{\alpha})^{-1}\boldsymbol{\alpha}^T.$$

Since this estimated matrix could not have unit elements in its diagonal, we scale it in order to have an unit-diagonal one. With this estimator of $\boldsymbol{\Delta}^{-1}$, the partially maximized log-likelihood reads

$$Q(\boldsymbol{\alpha}) = -\frac{pn}{2}\log(2\pi) - \frac{n}{2}\log|\boldsymbol{\alpha}^T\mathbf{S}_{\text{res}}\boldsymbol{\alpha}| - \frac{n}{2}\log|(\boldsymbol{\alpha}_0^T\mathbf{S}^{-1}\boldsymbol{\alpha}_0)^{-1}| - \frac{nd}{2} - \frac{n(p-d)}{2}$$
$$= -\frac{pn}{2}[\log(2\pi) + 1] - \frac{n}{2}\log|\boldsymbol{\alpha}^T\mathbf{S}_{\text{res}}\boldsymbol{\alpha}| - \frac{n}{2}\log|\mathbf{S}| + \frac{n}{2}\log|\boldsymbol{\alpha}^T\mathbf{S}\boldsymbol{\alpha}|.$$

where in the last equality we have used (A.5). Finnally, the maximum in $\boldsymbol{\alpha}$ is attended at

$$\boldsymbol{\alpha}^{(k)} = \mathbf{S}^{-1/2}\hat{\boldsymbol{\zeta}}\mathbf{N},$$

where $\hat{\boldsymbol{\zeta}}$ are the first d eigenvectors of $\mathbf{S}^{-1/2}\mathbf{S}_{\text{fit}}\mathbf{S}^{-1/2}$ and $\mathbf{N}$ a matrix such that $\widehat{\boldsymbol{\alpha}}^T\widehat{\boldsymbol{\alpha}} = \mathbf{I}_d$.

Appendix B. Approximating $\mathbf{S}$ and $\mathbf{M}$

Given $(\mathbf{\Omega}^{(k-1)}, \mathbf{\Theta}^{(k)})$ and y_i fixed, we need to estimate $\mathbf{S}$ and $\mathbf{M}$ in order to compute the Q-function. Each entry of matrix $\mathbf{S}$ can be written as $s_{jk} = \sum_{i=1}^{n} E_{\mathbf{z}_i}(z_{i,j} z_{i,k}|\mathbf{x}_i)$ with $j, k = 1, \ldots, p$. So, for $j = k$ we have the conditional second moment $E_{\mathbf{z}_i}(z_{i,j}^2|\mathbf{x}_i)$. Following [23], when $j \neq k$ the terms $E_{\mathbf{z}_i}(z_{i,j} z_{i,k}|\mathbf{x}_i)$ can be approximated by $E_{\mathbf{z}_i}(z_{i,j} z_{i,k}|\mathbf{x}_i) \approx E_{\mathbf{z}_i}(z_{i,j}|\mathbf{x}_i) E_{\mathbf{z}_i}(z_{i,k}|\mathbf{x}_i)$. With this, we can obtain an estimator of $\mathbf{S}$ through the estimation of first and second moments. The following is a modification of the procedure to compute these moments developed by [23], adapted to the case of conditional distributions. We can write $\mathbf{x}_i$ as $\mathbf{x}_i = (x_{i,j}, \mathbf{x}_{i,-j})$ and $\mathbf{z}_i$ as $\mathbf{z}_i = (z_{i,j}, \mathbf{z}_{i,-j})$ where $\mathbf{x}_{i,-j} = (x_{i,1}, \ldots, x_{i,j-1}, x_{i,j+1}, \ldots, x_{i,p})$ and $\mathbf{z}_{i,-j} = (z_{i,1}, \ldots, z_{i,j-1}, z_{i,j+1}, \ldots, z_{i,p})$. So, the first moment is

$$
\begin{aligned}
E_{\mathbf{z}_i}(z_{i,j}|\mathbf{x}_i) &= \int_{\mathbb{R}^p} z_{i,j} f_{\mathbf{z}_i}(\mathbf{z}_i|\mathbf{x}_i) d\mathbf{z}_i \\
&= \int_{\mathbb{R}^p} z_{i,j} f_{z_{i,j}}(z_{i,j}|\mathbf{z}_{i,-j}, x_{i,j}, \mathbf{x}_{i,j}) f_{\mathbf{z}_{i,-j}}(\mathbf{z}_{i,-j}|\mathbf{x}_i) d\mathbf{z}_i \\
&= \int_{\mathbb{R}^{p-1}} \left[\int_{\mathbb{R}} z_{i,j} f_{z_{i,j}}(z_{i,j}|z_{i,-j}, x_{i,j}) dz_{i,j} \right] f_{\mathbf{z}_{i,-j}}(\mathbf{z}_{i,-j}|\mathbf{x}_i) d\mathbf{z}_{i,-j} \\
&= E_{\mathbf{z}_{i,-j}} \left\{ E_{z_{i,j}}(z_{i,j}|\mathbf{z}_{i,-j}, x_{i,j})|\mathbf{x}_i \right\}.
\end{aligned}
\tag{B.1}
$$

In the same way, the second moment can be written as

$$
E_{\mathbf{z}_i}(z_{i,j}^2|\mathbf{x}_i) = E_{\mathbf{z}_{i,-j}} \left\{ E_{z_{i,j}}(z_{i,j}^2|\mathbf{z}_{i,-j}, x_{i,j})|\mathbf{x}_i \right\}.
\tag{B.2}
$$

Given y_i and $\mathbf{\Omega}^{(k-1)}$, $(z_{i,1}, z_{i,2}, \ldots, z_{i,p})$ has a multivariate normal distribution with mean $\mu_i = \mathbf{\Psi}^{(k-1)} \bar{\mathbf{f}}_{y_i} = \mathbf{\Delta}^{(k-1)} \boldsymbol{\alpha}^{(k-1)} \boldsymbol{\xi}^{(k-1)}$, and covariance matrix $\mathbf{\Delta}^{(k-1)}$. Taking $\mathbf{\Delta}_{j,j}^{(k-1)} = 1$, for each $j = 1, \ldots, p$ we can write

$$
\mathbf{\Delta}^{(k-1)} = \begin{pmatrix} 1 & \mathbf{\Delta}_{j,-j}^{(k-1)} \\ \mathbf{\Delta}_{-j,j}^{(k-1)} & \mathbf{\Delta}_{-j,-j}^{(k-1)} \end{pmatrix} \qquad \text{and} \qquad \mu_i = \begin{pmatrix} (\mathbf{\Psi}^{(k-1)} \bar{\mathbf{f}}_{y_i})_j \\ (\mathbf{\Psi}^{(k-1)} \bar{\mathbf{f}}_{y_i})_{-j} \end{pmatrix},
$$

and therefore the conditional distribution of $z_{i,j}$ given $\mathbf{z}_{i,-j}$ is

$$
z_{i,j}|\mathbf{z}_{i,-j} \sim N(\widetilde{\mu}_{i,j}, \widetilde{\delta}_{i,j}),
$$

where the mean is $\widetilde{\mu}_{i,j} = (\mathbf{\Psi}^{(k-1)} \bar{\mathbf{f}}_{y_i})_j + \mathbf{\Delta}_{j,-j}^{(k-1)} (\mathbf{\Delta}_{-j,-j}^{(k-1)})^{-1} \left(\mathbf{z}_{i,-j} - (\mathbf{\Psi}^{(k-1)} \bar{\mathbf{f}}_{y_i})_{-j} \right)^T$ and the variance $\widetilde{\delta}_{i,j}^2 = 1 - \mathbf{\Delta}_{j,-j}^{(k-1)} (\mathbf{\Delta}_{-j,-j}^{(k-1)})^{-1} \mathbf{\Delta}_{-j,j}^{(k-1)}$. In addition, the conditional distribution of $z_{i,j}$ on observed data $x_{i,j}$ is equivalent to conditioning on $z_{i,j} \in C(x_{i,j}, \mathbf{\Theta}) = [\theta_{x_{i,j}-1}^{(j)}, \theta_{x_{i,j}}^{(j)})$, which follows a truncated normal distribution with density

$$
f(z_{i,j}|C(x_{i,j}, \mathbf{\Theta}), \mathbf{z}_{i,-j}) = \frac{\frac{1}{\widetilde{\delta}_{i,j}} \phi(\widetilde{z}_{i,j})}{\Phi(\widetilde{\theta}_{x_{i,j}}^{(j)}) - \Phi(\widetilde{\theta}_{x_{i,j}-1}^{(j)})} I_{\{z_{i,j} \in C(x_{i,j}, \mathbf{\Theta})\}}.
$$

Here, $\widetilde{z}_{i,j} = (z_{i,j} - \widetilde{\mu}_{i,j})/\widetilde{\delta}_{i,j}$, $\widetilde{\theta}_{x_{i,j}}^{(j)} = (\theta_{x_{i,j}}^{(j)} - \widetilde{\mu}_{i,j})/\widetilde{\delta}_{i,j}$ and $\widetilde{\theta}_{x_{i,j}-1}^{(j)} = (\theta_{x_{i,j}-1}^{(j)} - \widetilde{\mu}_{i,j})/\widetilde{\delta}_{i,j}$. From the moment generating function of the truncated normal distribution, the first and second moment of $z_{i,j}$ are given by

$$
E(z_{i,j}|\mathbf{z}_{i,-j}, x_{i,j}) = \widetilde{\mu}_{i,j} + \widetilde{\delta}_{i,j} a_{i,j},
\tag{B.3}
$$
$$
E(z_{i,j}^2|\mathbf{z}_{i,-j}, x_{i,j}) = \widetilde{\mu}_{i,j}^2 + \widetilde{\delta}_{i,j}^2 + 2 a_{i,j} \widetilde{\mu}_{i,j} \widetilde{\delta}_{i,j} + b_{i,j} \widetilde{\delta}_{i,j}^2,
\tag{B.4}
$$

where

26

$$a_{i,j} = \frac{\phi(\widetilde{\theta}^{(j)}_{x_{i,j}-1}) - \phi(\widetilde{\theta}^{(j)}_{x_{i,j}})}{\Phi(\widetilde{\theta}^{(j)}_{x_{i,j}}) - \Phi(\widetilde{\theta}^{(j)}_{x_{i,j}-1})}, \qquad b_{i,j} = \frac{\widetilde{\theta}^{(j)}_{x_{i,j}-1}\phi(\widetilde{\theta}^{(j)}_{x_{i,j}-1}) - \widetilde{\theta}^{(j)}_{x_{i,j}}\phi(\widetilde{\theta}^{(j)}_{x_{i,j}})}{\Phi(\widetilde{\theta}^{(j)}_{x_{i,j}}) - \Phi(\widetilde{\theta}^{(j)}_{x_{i,j}-1})}.$$

Using (B.3) and (B.4) in (B.1) and (B.2), respectively, the first and second moments read

$$E_{\mathbf{z}_i}(z_{i,j}|\mathbf{x}_i) = E_{\mathbf{z}_{i,-j}}\left(\widetilde{\mu}_{i,j}|\mathbf{x}_i\right) + \widetilde{\delta}_{i,j}E_{\mathbf{z}_{i,-j}}\left(a_{i,j}|\mathbf{x}_i\right). \tag{B.5}$$

$$E_{\mathbf{z}_i}(z^2_{i,j}|\mathbf{x}_i) = E_{\mathbf{z}_{i,-j}}\left(\widetilde{\mu}^2_{i,j}|\mathbf{x}_i\right) + \widetilde{\delta}^2_{i,j} + 2\widetilde{\delta}_{i,j}E_{\mathbf{z}_{i,-j}}\left(a_{i,j}\widetilde{\mu}_{i,j}|\mathbf{x}_i\right) + \widetilde{\delta}^2_{i,j}E_{\mathbf{z}_{i,-j}}\left(b_{i,j}|\mathbf{x}_i\right). \tag{B.6}$$

Here $\widetilde{\mu}_{i,j}$ is linear in $\mathbf{z}_{i,-j}$ then, for $\mathbf{\Delta}^{(k-1)}$, $\widehat{\alpha}$ and y_i fixed, we have

$$E_{\mathbf{z}_{i,-j}}\left(\widetilde{\mu}_{i,j}|\mathbf{x}_i\right) = (\mathbf{\Psi}^{(k-1)}\bar{\mathbf{f}}_{y_i})_j + \mathbf{\Delta}^{(k-1)}_{j,-j}(\mathbf{\Delta}^{(k-1)}_{-j,-j})^{-1}E_{\mathbf{z}_{i,-j}}\left(\mathbf{z}^T_{i,-j}|\mathbf{x}_i\right) - \mathbf{\Delta}^{(k-1)}_{j,-j}(\mathbf{\Delta}^{(k-1)}_{-j,-j})^{-1}(\mathbf{\Psi}^{(k-1)}\bar{\mathbf{f}}_{y_i})^T_{-j} \tag{B.7}$$

$$E_{\mathbf{z}_{i,-j}}\left(\widetilde{\mu}^2_{i,j}|\mathbf{x}_i\right) = \left(\mathbf{\Psi}^{(k-1)}\bar{\mathbf{f}}_{y_i}\right)^2_j + 2\left(\mathbf{\Psi}^{(k-1)}\bar{\mathbf{f}}_{y_i}\right)_j \mathbf{\Delta}^{(k-1)}_{j,-j}(\mathbf{\Delta}^{(k-1)}_{-j,-j})^{-1}\left[E_{\mathbf{z}_{i,-j}}\left(\mathbf{z}_{i,-j}|\mathbf{x}_i\right) - \left(\mathbf{\Psi}^{(k-1)}\bar{\mathbf{f}}_{y_i}\right)_{-j}\right]^T$$

$$+\mathbf{\Delta}^{(k-1)}_{j,-j}(\mathbf{\Delta}^{(k-1)}_{-j,-j})^{-1}\left[E_{\mathbf{z}_{i,-j}}\left(\mathbf{z}^T_{i,-j}\mathbf{z}_{i,-j}|\mathbf{x}_i\right) - E_{\mathbf{z}_{i,-j}}\left(\mathbf{z}^T_{i,-j}|\mathbf{x}_i\right)\left(\mathbf{\Psi}^{(k-1)}\bar{\mathbf{f}}_{y_i}\right)_{-j}\right.$$

$$\left. - \left(\mathbf{\Psi}^{(k-1)}\bar{\mathbf{f}}_{y_i}\right)^T_{-j}E_{\mathbf{z}_{i,-j}}\left(\mathbf{z}_{i,-j}|x_i\right) + \left(\mathbf{\Psi}^{(k-1)}\bar{\mathbf{f}}_{y_i}\right)^T_{-j}\left(\mathbf{\Psi}^{(k-1)}\bar{\mathbf{f}}_{y_i}\right)_{-j}\right](\mathbf{\Delta}^{(k-1)}_{-j,-j})^{-1}(\mathbf{\Delta}^{(k-1)}_{-j,-j})^T \tag{B.8}$$

On the other hand, we have that the functions $a_{i,j}$ and $b_{i,j}$ are nonlinear in $\widetilde{\theta}^j_{x_{i,j}}$ and $\widetilde{\theta}^j_{x_{i,j}-1}$ who are linear functions of $\widetilde{\mu}_{i,j}$ and thus of $\mathbf{z}_{i,-j}$. So, we can write $a_{i,j}$ and $b_{i,j}$ as $a_{i,j}(\mathbf{z}_{i,-j})$ and $b_{i,j}(\mathbf{z}_{i,-j})$. Conditioning on $\mathbf{x}_i$, $\mathbf{z}_{i,-j}$ has a truncated normal distribution with mean $\widetilde{\mathbf{v}}_{i,-j} = E_{\mathbf{z}_{i,-j}}(\mathbf{z}_{i,-j}|\mathbf{x}_i)$ and covariance matrix $\widetilde{\mathbf{V}} = E_{\mathbf{z}_{i,-j}}\left((\mathbf{z}_{i,-j} - \widetilde{\mathbf{v}}_{i,-j})(\mathbf{z}_{i,-j} - \widetilde{\mathbf{v}}_{i,-j})^T|\mathbf{x}_i\right)$. If we assume that $a_{i,j}$ and $b_{i,j}$ have continuous first partial derivatives, by the first order delta method we have that

$$n^{1/2}\left\{a_{i,j}(\mathbf{z}_{i,-j}|\mathbf{x}_i) - a_{i,j}(\widetilde{\mathbf{v}}_{i,-j})\right\} \xrightarrow{\mathcal{D}} N\left(0, \nabla a_{i,j}(\widetilde{\mathbf{v}}_{i,-j})\widetilde{\mathbf{V}}\nabla^T a_{i,j}(\widetilde{\mathbf{v}}_{i,-j})\right) \text{ and}$$

$$n^{1/2}\left\{b_{i,j}(\mathbf{z}_{i,-j}|\mathbf{x}_i) - b_{i,j}(\widetilde{\mathbf{v}}_{i,-j})\right\} \xrightarrow{\mathcal{D}} N\left(0, \nabla b_{i,j}(\widetilde{\mathbf{v}}_{i,-j})\widetilde{\mathbf{V}}\nabla^T b_{i,j}(\widetilde{\mathbf{v}}_{i,-j})\right),$$

so we can approximate the expectation $E_{\mathbf{z}_{i,-j}}\left(a_{i,j}|\mathbf{x}_i\right)$ with $a_{i,j}(\widetilde{\mathbf{v}}_{i,-j})$ and $E_{\mathbf{z}_{i,-j}}\left(b_{i,j}|\mathbf{x}_i\right)$ with $b_{i,j}(\widetilde{\mathbf{v}}_{i,-j})$, i.e.

$$E_{\mathbf{z}_{i,-j}}\left(a_{i,j}|\mathbf{x}_i\right) \approx a_{i,j}(\widetilde{\mathbf{v}}_{i,-j}) = \frac{\phi(\widetilde{\widetilde{\theta}}^{(j)}_{x_{i,j}-1}) - \phi(\widetilde{\widetilde{\theta}}^{(j)}_{x_{i,j}})}{\Phi(\widetilde{\widetilde{\theta}}^{(j)}_{x_{i,j}}) - \Phi(\widetilde{\widetilde{\theta}}^{(j)}_{x_{i,j}-1})} \tag{B.9}$$

$$E_{\mathbf{z}_{i,-j}}\left(b_{i,j}|\mathbf{x}_i\right) \approx b_{i,j}(\widetilde{\mathbf{v}}_{i,-j}) = \frac{\widetilde{\widetilde{\theta}}^{(j)}_{x_{i,j}-1}\phi(\widetilde{\widetilde{\theta}}^{(j)}_{x_{i,j}-1}) - \widetilde{\widetilde{\theta}}^{(j)}_{x_{i,j}}\phi(\widetilde{\widetilde{\theta}}^{(j)}_{x_{i,j}})}{\Phi(\widetilde{\widetilde{\theta}}^{(j)}_{x_{i,j}}) - \Phi(\widetilde{\widetilde{\theta}}^{(j)}_{x_{i,j}-1})} \tag{B.10}$$

with

$$\widetilde{\widetilde{\theta}}^{(j)}_{x_{i,j}-k} = \frac{\widetilde{\theta}^{(j)}_{x_{i,j}-k} - E_{\mathbf{z}_{i,-j}}(\widetilde{\mu}_{i,j}|\mathbf{x}_i)}{\widetilde{\delta}_{i,j}} = \frac{\widetilde{\theta}^{(j)}_{x_{i,j}-k} - \left(\mathbf{\Psi}^{(k-1)}\bar{\mathbf{f}}_{y_i}\right)_j + \mathbf{\Delta}^{(k-1)}_{j,-j}(\mathbf{\Delta}^{(k-1)}_{-j,-j})^{-1}\left[\widetilde{\mathbf{v}}_{i,-j} - \left(\mathbf{\Psi}^{(k-1)}\bar{\mathbf{f}}_{y_i}\right)_{-j}\right]^T}{\widetilde{\delta}_{i,j}}$$

Using (B.7), (B.8), (B.9), (B.10) and the approximation $E_{\mathbf{z}_{i,-j}}\left(a_{i,j}\widetilde{\mu}_{i,j}|\mathbf{x}_i\right) \approx E_{\mathbf{z}_{i,-j}}\left(a_{i,j}|\mathbf{x}_i\right)E_{\mathbf{z}_{i,-j}}\left(\widetilde{\mu}_{i,j}|\mathbf{x}_i\right)$, the conditional expectation in (B.5) can be approximated by

$$
E_{\mathbf{z}_i}\left(z_{i,j}|\mathbf{x}_i\right) \approx (\boldsymbol{\Psi}^{(k-1)}\bar{\mathbf{f}}_{y_i})_j + \boldsymbol{\Delta}_{j,-j}^{(k-1)}(\boldsymbol{\Delta}_{-j,-j}^{(k-1)})^{-1}\left[E_{\mathbf{z}_{i,-j}}\left(\mathbf{z}_{i,-j}^T|\mathbf{x}_i\right) - (\boldsymbol{\Psi}^{(k-1)}\bar{\mathbf{f}}_{y_i})_{-j}\right]^T
$$
$$
+ \widetilde{\delta}_{i,j}\frac{\phi(\widetilde{\overline{\theta}}_{x_{i,j}-1}^{(j)}) - \phi(\widetilde{\overline{\theta}}_{x_{i,j}}^{(j)})}{\Phi(\widetilde{\overline{\theta}}_{x_{i,j}}^{(j)}) - \Phi(\widetilde{\overline{\theta}}_{x_{i,j}-1}^{(j)})}, \tag{B.11}
$$

and the second moment in (B.6) by

$$
E_{\mathbf{z}_i}\left(z_{i,j}^2|\mathbf{x}_i\right) \approx \left(\boldsymbol{\Psi}^{(k-1)}\bar{\mathbf{f}}_{y_i}\right)_j^2 + 2\left(\boldsymbol{\Psi}^{(k-1)}\bar{\mathbf{f}}_{y_i}\right)_j\boldsymbol{\Delta}_{j,-j}^{(k-1)}(\boldsymbol{\Delta}_{-j,-j}^{(k-1)})^{-1}\left[E_{\mathbf{z}_{i,-j}}\left(\mathbf{z}_{i,-j}|\mathbf{x}_i\right) - \left(\boldsymbol{\Psi}^{(k-1)}\bar{\mathbf{f}}_{y_i}\right)_{-j}\right]^T
$$
$$
+ \boldsymbol{\Delta}_{j,-j}^{(k-1)}(\boldsymbol{\Delta}_{-j,-j}^{(k-1)})^{-1}\left[E_{\mathbf{z}_{i,-j}}\left(\mathbf{z}_{i,-j}^T\mathbf{z}_{i,-j}|\mathbf{x}_i\right) - E_{\mathbf{z}_{i,-j}}\left(\mathbf{z}_{i,-j}^T|\mathbf{x}_i\right)\left(\boldsymbol{\Psi}^{(k-1)}\bar{\mathbf{f}}_{y_i}\right)_{-j}\right.
$$
$$
\left. - \left(\boldsymbol{\Psi}^{(k-1)}\bar{\mathbf{f}}_{y_i}\right)_{-j}^T E_{\mathbf{z}_{i,-j}}\left(\mathbf{z}_{i,-j}|x_i\right) + \left(\boldsymbol{\Psi}^{(k-1)}\bar{\mathbf{f}}_{y_i}\right)_{-j}^T\left(\boldsymbol{\Psi}^{(k-1)}\bar{\mathbf{f}}_{y_i}\right)_{-j}\right](\boldsymbol{\Delta}_{-j,-j}^{(k-1)})^{-1}(\boldsymbol{\Delta}_{-j,-j}^{(k-1)})^T + \widetilde{\delta}_{i,j}
$$
$$
\tag{B.12}
$$
$$
+ 2\widetilde{\delta}_{i,j}\frac{\phi(\widetilde{\overline{\theta}}_{x_{i,j}-1}^{(j)}) - \phi(\widetilde{\overline{\theta}}_{x_{i,j}}^{(j)})}{\Phi(\widetilde{\overline{\theta}}_{x_{i,j}}^{(j)}) - \Phi(\widetilde{\overline{\theta}}_{x_{i,j}-1}^{(j)})}\left[(\boldsymbol{\Psi}^{(k-1)}\bar{\mathbf{f}}_{y_i})_j + \boldsymbol{\Delta}_{j,-j}^{(k-1)}(\boldsymbol{\Delta}_{-j,-j}^{(k-1)})^{-1}\left[E_{\mathbf{z}_{i,-j}}\left(\mathbf{z}_{i,-j}^T|\mathbf{x}_i\right) - (\boldsymbol{\Psi}^{(k-1)}\bar{\mathbf{f}}_{y_i})_{-j}\right]^T\right]
$$
$$
+ \widetilde{\delta}_{i,j}^2\frac{\widetilde{\overline{\theta}}_{x_{i,j}-1}^{(j)}\phi(\widetilde{\overline{\theta}}_{x_{i,j}-1}^{(j)}) - \widetilde{\overline{\theta}}_{x_{i,j}}^{(j)}\phi(\widetilde{\overline{\theta}}_{x_{i,j}}^{(j)})}{\Phi(\widetilde{\overline{\theta}}_{x_{i,j}}^{(j)}) - \Phi(\widetilde{\overline{\theta}}_{x_{i,j}-1}^{(j)})}.
$$

Equations (B.11) and (B.12) give recursive expressions for computing (iteratively) $\mathbf{S}$ and $\mathbf{M}$, respectively.

Appendix C. Description of Variables for SES index construction

The following variables are used to construct the SES indices:

- *Housing location*: indicates if the housing is located in a disadvantaged or vulnerable area. More precisely, it considers if housing: (i) is located in a shanty town, (ii) or/and near to landfill sites, (iii) or/and in a floodplain. It has 4 categories: 1 for houses that jointly present the (i)-(iii) characteristics, 2 for housing presenting two of (i)-(iii), 3 if housing has only one of them, and 4 if the house has none of these characteristics.

- *Housing quality*: jointly contemplates the quality of roof, walls and floor based on the CALMAT's methodology [25] used in the population censuses of Argentina. It has 4 categories in increasing order in terms of housing quality.

- *Sources of cooking fuel*: indicates the kind of fuel used for cooking in the housing. It has 3 categories: 1 if the main source of cooking fuel in the housing is kerosene, wood or charcoal, 2 for bottled gas, and 3 for natural gas by pipeline.

- *Overcrowding*: characterizes the overcrowding by computing the ratio between rooms and number of household members. It has 4 categories: 1 if this ratio is less or equal than 1, 2 if the ratio is in the interval $(1, 2]$, 3 if it is in $(2, 3]$, and 4 if this ratio is greater than 3.

- *Schooling*: indicates the formal education attained by the head of household. It has 7 categories: 1 if the head of household has no formal education, 2 in the case of incomplete elementary level, 3 for complete elementary level, 4 for incomplete secondary school, 5 for a complete level of secondary school, 6 for an incomplete higher education and 7 if the head of household achieved a university or tertiary degree.

- *Working hours*: describes the labor situation of head of household. It has 4 categories: 1 for unemployment or inactive cases, 2 when the head of household works less than 40 hours per week, 3 for 40-45 per week working hours, and 4 when the head of household is employed for more than 45 hours per week.

- *Toilet drainage*: indicates the type of drainage of the housing. It has 4 categories: 1 if drainage is a hole, 2 if drainage is only in a cesspool, 3 for cesspool and septic tank, and 4 for drain pipes in a public network.

- *Toilet facility*: indicates the toilet facility available in the housing. It has 3 categories: 1 for latrines, 2 for toilets without flush water, and 3 for flushing toilets.

- *Toilet sharing*: indicates if the toilet is shared or not. It has 3 categories: 1 if the toilet is shared with other housing, 2 if the toilet is shared with other households into the same housing, and 3 if the toilet is used exclusively by the household.

- *Water location*: indicates the nearest location of drinking water. It has 3 categories: 1 if drinking water is gotten outside the plot of land of housing, 2 if water is into plot of land but outside of housing, and 3 of drinking water is obtained inside housing by pipe.

- *Water source*: indicates the source of the water in the housing. It has 3 categories: 1 if drinking water comes from a hand pump or from a public tap shared with neighbours, 2 if drinking water is obtained by an automated drilling pump, and 3 for housing with piped drinking water.

Founding

References

[1] Adragni, K. and Cook, R. [2009], 'Sufficient dimension reduction and prediction in regression', *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* **367**(1906), 4385–4405.

[2] Bobadilla, J., Ortega, F., Hernando, A. and Bernal, J. [2012], 'Generalization of recommender systems: Collaborative filtering extended to groups of users and restricted to groups of items', *Expert Systems with Applications* **39**(1), 172–186.

[3] Bollen, K., Glanville, J. and Stecklov, G. [2001], 'Socioeconomic status and class in studies on fertility and health in developing countries', *Annual Review of Sociology* **27**, 153–185.

[4] Bura, E. and Cook, R. [2001], 'Estimating the structural dimension of regressions via parametric inverse regression', *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* **63**(2), 393–410.

[5] Bura, E., Duarte, S. and Forzani, L. [2015], 'Sufficient reductions in regressions with exponential family inverse predictors.', *To appear in Journal of the American Statistical Association* .

[6] Chen, X., Zou, C. and Cook, R. [2010], 'Coordinate-independent sparse sufficient dimension reduction and variable selection', *The Annals of Statistics* **38**, 3696–3723.

[7] Chiaromonte, F., Cook, R. and Li, B. [2002], 'Sufficient dimensions reduction in regressions with categorical predictors', *The Annals of Statistics.* **30**(2), 475–497.

[8] Cook, R. [1996], 'Graphics for regressions with a binary response', *Journal of the American Statistical Association* **91**, 983–992.

[9] Cook, R. [1998], *Regression Graphics*, Wiley, New York.

[10] Cook, R. [2007], 'Fisher lecture: Dimension reduction in regression (with discussion)', *Statistical Science* **22**, 1–26.

[11] Cook, R. and Forzani, L. [2008], 'Principal fitted components for dimension reduction in regression', *Statistical Science* **23**, 485–501.

[12] Cook, R. and Forzani, L. [2009], 'Likelihood-Based sufficient dimension reduction', *Journal of the American Statistical Association* **104**(485), 197–208.

[13] Cook, R. and Lee, H. [1999], 'Dimension reduction in binary response regression', *Journal of the American Statistical Association* **94**(448), 1187–1200.

[14] Cook, R. and Li, B. [2002], 'Dimension reduction for conditional mean in regression', *The Annals of Statistics* **30**(2), 455–474.

[15] Cook, R. and Ni, L. [2005], 'Sufficient dimension reduction via inverse regression: A minimum discrepancy approach', *Journal of the American Statistical Association* **100**(470), 410–428.

[16] Cook, R. and Weisberg, S. [1991], 'Discussion of sliced inverse regression for dimension reduction', *Journal of the American Statistical Association* **86**, 328–332.

[17] Cook, R. and Yin, X. [2001], 'Dimension reduction and visualization in discriminant analysis (invited, with discussion)', *Australia & New Zeland Journal of Statistics* **43**, 147–200.

[18] Doocy, S. and Burnham, G. [2006], 'Assessment of socio-economic status in the context of food insecurity: Implications for field research', *World Health and Population* **8**(3), 32–42.

[19] Eaton, M. [1983], *Multivariate Statistics*, Wiley, New York.

[20] Feeny, S., Mcdonald, L. and Posso, A. [2014], 'Are poor people less happy? findings from melanesia', *World Development* **64**, 448–459.

[21] Gertheiss, J. and Tutz, G. [2010], 'Sparse modelling of categorical explanatory variables', *The Annals of Applied Statistics* **4**, 2150–2180.

[22] Greene, W. and Hensher, D. [2010], *Modeling Ordered Choices: A Primer*, Cambridge University Press.

[23] Guo, J., Levina, E., Michailidis, G. and Zhu, J. [2015], 'Graphical models for ordinal data', *Journal of Computational and Graphical Statistics* **24**(1), 183–204.

[24] Hoque, S. [2014], 'Asset-based poverty analysis in rural bangladesh: A comparison of principal component analysis and fuzzy set theory', *SRI Papers, Sustainability Research Institute* **59**.

[25] INDEC [2003], *Calidad de los materiales de la vivienda -CALMAT-Hábitat y vivienda por medio de datos censales. Dirección Nacional de Estadísticas Sociales y de Población, Dirección de Estadísticas Poblacionales, área de Información Derivada.*, DNESyP/DEP/P5/PID Serie Hábitat y Vivienda DT No. 13, Buenos Aires.

[26] Jackman, S. [2009], *Bayesian Analysis for the Social Sciences*, Wiley Series in Probability and Statistics, Wiley.

[27] Kamakura, W. and Mazzon, J. [2013], 'Socioeconomic status and consumption in an emerging economy', *International Journal of Research in Marketing* **30**, 4–18.

[28] Kolenikov, S. and Angeles, G. [2009], 'Socioeconomic status measurement with discrete proxy variables: Is principal component analysis a reliable answer?', *The Review of Income and Wealth* **55**(1), 128–165.

[29] Lee, G. and Scott, C. [2012], 'Em algorithms for multivariate gaussian mixture models with truncated and censored data', *Computational Statistics and Data Analysis* **56**(9), 2816–2829.

[30] Li, B. and Wang, S. [2007], 'On directional regression for dimension reduction', *Journal of the American Statistical Association* **102**(479), 997–1008.

[31] Li, B., Zha, H. and Chiaromonte, C. [2005], 'Contour regression: a general approach to dimension reduction', *The Annals of Statistics* **33**(4), 1580–1616.

[32] Li, K. [1992], 'On principal hessian directions for data visualization and dimension reduction: Another application of steinś lemma', *Journal of the American Statistical Association* **87**(420), 1025–1039.

[33] Li, K. C. [1991], 'Sliced inverse regression for dimension reduction (with discussion)', *Journal of the American Statistical Association* **86**, 316–342.

[34] Linting, M. and van der Kooij, A. J. [2009], 'Nonlinear principal components analysis with catpca: A tutorial.', *Journal of Personality Assessment* **94**(1), 12–25.

[35] Long, J. [1997], *Regression models for categorical and limited dependent variables*, 2 edn, SAGE Publications.

[36] Mazzonna, F. [2014], 'The long-lasting effects of family background: A european cross-country comparison', *Economics of Education Review* **40**, 25–42.

[37] Merola, G. and Baulch, B. [2014], Using sparse categorical principal components to estimate asset indices new methods with an application to rural south east asia, Rome, Italy.

[38] Mokomane, Z. [2012], 'Social protection as a mechanism for family protection in sub-saharan africa', *International Journal of Social Welfare* **22**(3), 248–259.

[39] Murasko, J. [2007], 'Socioeconomic status, height and obesity in children', *Economics and Human Biology* **7**(3), 376–386.

[40] Rao, C. [1973], *Linear statistical inference and its applications, 2nd. Edition.*, Wiley, New York.

[41] Reddi, S. and Poczos, B. [2013], 'Scale Invariant Conditional Dependence Measures', *JMLR Workshop and Conference Proceedings* **3**(28), 1355-1363.

[42] Richardson, D. and Bradshaw, J. [2012], *Family-oriented anti-poverty policies in developed countries*, Department of Economic and Social Affairs, Division for Social Policy and Development, United Nations, New York, New York.

[43] Roberts, W. [2014], 'Factor analysis parameter estimation from incomplete data', *Computational Statistics & Data Analysis* **70**(0), 61–66.

[44] Roy, K. and Chaudhuri, A. [2009], 'Influence of socioeconomic status, wealth and financial empowerment on gender differences in health and healthcare utilization in later life: Evidence from india', *Social Science and Medicine* **66**, 1951–1962.

[45] Skrondal, A. and Rabe-Hesketh, S. [2004], *Generalized Latent Variable Modeling: Multilevel, Longitudinal, and Structural Equation Models*, Chapman & Hall/CRC.

[46] Székely, G. and Rizzo, M. [2009], 'Brownian distance covariance', *The Annals of Applied Statistics* **3**(4), 1236–1265.

[47] Székely, G., Rizzo, M. and Bakirov, N. [2007], 'Measuring and testing dependence by correlation of distances', *The Annals of Statistics* **35**(6), 2769–2794.

[48] Vyas, S. and Kumaranayake, L. [2006], 'Constructing socio-economic status indices: How to use principal components analysis', *Health Policy and Planning* **21**(6), 459–468.

[49] Xia, Y., Tong, H., Li, W. and Zhu, L. [2002], 'An adaptative estimation of dimension reduction space', *Journal of the Royal Statistical Society, Series B* **64**, 363–410.

[50] Yin, X., Li, B. and Cook, R. D. [2008], 'Successive direction extraction for estimating the central subspace in a multiple-index regression', *Journal of Multivariate Analysis* **99**, 1733–1757.

[51] Zhu, Y. and Zeng, P. [2006], 'Fourier methods for estimating the central subspace and the central mean subspace in regression', *Journal of the American Statistical Association* **101**(476), 1638–1651.