



Joint and Individual Component Regression

Peiyao Wang^a, Haodong Wang^a, Quefeng Li^b, Dinggang Shen^{c,d,e}, and Yufeng Liu^{a,b,f} 

^aDepartment of Statistics and Operations Research, University of North Carolina at Chapel Hill, Chapel Hill, NC; ^bDepartment of Biostatistics, University of North Carolina at Chapel Hill, Chapel Hill, NC; ^cSchool of Biomedical Engineering & State Key Laboratory of Advanced Medical Materials and Devices, ShanghaiTech University, Shanghai, China; ^dShanghai United Imaging Intelligence Co., Ltd., Shanghai, China; ^eShanghai Clinical Research and Trial Center, Shanghai, China; ^fDepartment of Genetics, University of North Carolina at Chapel Hill, Chapel Hill, NC

ABSTRACT

Multi-group data, which include the same set of variables on separate groups of samples, are commonly seen in practice. Such data structure consists of data from multiple groups and can be challenging to analyze due to data heterogeneity. We propose a novel Joint and Individual Component Regression (JICO) model to analyze multi-group data. Our proposed model decomposes the response into shared and group-specific components, which are driven by low-rank approximations of joint and individual structures from the predictors respectively. The joint structure has the same regression coefficients across multiple groups, whereas individual structures have group-specific regression coefficients. We formulate this framework under the representation of latent components and propose an iterative algorithm to solve for the joint and individual scores. We use the Continuum Regression (CR) to estimate the latent scores, which provides a unified framework that covers the Ordinary Least Squares (OLS), the Partial Least Squares (PLS), and the Principal Component Regression (PCR) as its special cases. We show that JICO attains a good balance between global and group-specific models and remains flexible due to the usage of CR. We conduct simulation studies and analysis of an Alzheimer's disease dataset to further demonstrate the effectiveness of JICO. R package of JICO is available online at <https://CRAN.R-project.org/package=JICO>. Supplementary materials for this article are available online.

ARTICLE HISTORY

Received September 2022
Accepted October 2023

KEYWORDS

Continuum regression;
Heterogeneity; Latent
component regression;
Multi-group data

1. Introduction

Many fields of scientific research involve the analysis of heterogeneous data. In particular, data may appear in the form of multiple matrices, with data heterogeneity arising from either variables or samples. One example is the multi-view/source data, which include different sets of variables on the same set of samples. The sets of variables may come from different platforms/sources/modalities. For instance, in genomics studies, measurements are collected as different biomarkers, such as mRNA and miRNA (Muniategui et al. 2013). Another example is the multi-group data, which include the same set of variables on disparate sets of samples, which leads to heterogeneous subpopulations/subgroups in the entire population. For instance, in the Alzheimer's Disease (AD) study, subjects can have different subtypes, such as Normal Control (NC), Mild Cognitive Impairment (MCI), and AD.

We study the classical regression problem with one continuous response for multi-group data. Although there are many well-established regression techniques for homogeneous data (Hoerl and Kennard 1970; Tibshirani 1996), they may not be suitable for multi-group data. One naive approach is to ignore data heterogeneity and fit a global model using these techniques. However, a single global model can be too restrictive because

the diverse information from different subgroups may not be identified. On the other hand, one can train separate group-specific models. Despite its flexibility, the information that is jointly shared across different groups is not sufficiently captured. Therefore, it is desirable to build a flexible statistical model that can simultaneously quantify the jointly shared global information and individual group-specific information for heterogeneous data.

There are several existing methods in the literature under the context of regression for multi-group data. Meinshausen and Bühlmann (2015) took a conservative approach and proposed a maxmin effect method that is reliable for all possible subsets of the data. Zhao, Cheng, and Liu (2016) proposed a partially linear regression framework for massive heterogeneous data, and the goal is to extract common features across all subpopulations while exploring heterogeneity of each subpopulation. Tang and Song (2016), Ma and Huang (2017), and Chen et al. (2021) proposed fused penalties to estimate regression coefficients that capture subgroup structures in a linear regression framework. Wang, Liu, and Shen (2018) studied a locally-weighted penalized model to perform subject-wise variable selection. Wang et al. (2023) proposed a factor regression model for heterogeneous subpopulations under

the high-dimensional factor decomposition. However, these models either are not specifically designed to identify the globally-shared and group-specific structures, or impose strong theoretical assumptions on the covariates. On the other hand, there exist some related works for multi-source data. Lock et al. (2013) proposed JIVE to learn joint and individual structures from multiple data matrices by low-rank approximations. Some extensions of JIVE can be found in Feng et al. (2018) and Gaynanova and Li (2019). All of these decomposition methods are fully unsupervised. Recently, Li and Li (2021) proposed a supervised integrative factor regression model for multi-source data and studied its statistical properties with hypothesis tests. Palzer et al. (2022) proposed sJIVE that extends JIVE with supervision from the response. These methods are supervised, but focused on regressions for multi-source data.

In this article, we consider the supervised learning problem of predicting a response with multi-group data. We propose a Joint and Individual COmponent Regression (JICO), a novel latent component regression model that covers JIVE as a special case. Our proposed model decomposes the response into jointly shared and group-specific components, which are driven by low-rank approximations of joint and individual structures from the predictors respectively. The joint structure shares the same coefficients across all groups, whereas individual structures have group-specific coefficients. Moreover, by choosing different ranks of joint and individual structures, our model covers global and group-specific models as special cases. To estimate JICO, we propose an iterative algorithm to solve for joint and individual scores using latent component representation. To construct the latent scores, we use the Continuum Regression (CR) (Stone and Brooks 1990), which provides a unified framework that covers OLS, PLS, and PCR as special cases. Some follow-up studies and modern extensions of CR can be found in Björkström and Sundberg (1996) and Lee and Liu (2013). Embracing this flexibility and generalizability from CR, our proposed JICO model extends to the heterogeneous data setup and is able to achieve different model configurations on the spectrum of CR under this more complicated setting. It attains a good balance between global and group-specific models, and further achieves its flexibility by extending CR.

A diagram illustration of JICO is shown in Figure 1. The left side shows the decomposition of the data matrices into the joint, individual and error matrices that share the same column dimension. The joint and individual matrices are further decomposed into the score and loading matrices. The right side shows the decomposition of the response vector into the joint, individual and error components. The joint and individual components are further regressed on the score matrices to obtain the joint and individual coefficients.

As noted by the reviewers of our paper, the success of JICO relies on a good choice of tuning parameters, which is at the cost of extra computational time. For that reason, we give more details of tuning parameter selections in Appendix D of the supplementary materials. We further study the impact of different initial values on JICO in Appendix E, and its convergence and computational time in Appendix F. We conclude that JICO is robust to different choices of initial values and its extra computational time is affordable compared with traditional methods, such as PCR and PLS. Given its superior performance than these

methods, we believe the extra computational cost is worthwhile. More details on the tradeoffs between performance improvement and computational cost can also be found in Appendix F.

The rest of this article is organized as follows. In Section 2, we briefly review JIVE and introduce our proposed JICO model. We further provide sufficient conditions to make JICO identifiable. In Section 3, after two motivating special cases, we introduce our iterative algorithm. In Sections 4 and 5, we evaluate the performance of JICO by simulation studies and real data analysis on the Alzheimer's disease dataset, respectively. In Section 6, we conclude this article with some discussion and possible extensions. Proofs and additional details are provided in the supplementary materials.

2. Motivation and Model Framework

Suppose we observe data pairs $(\mathbf{X}_g, \mathbf{Y}_g)_{g=1}^G$ from G groups, where $\mathbf{X}_g \in \mathbb{R}^{n_g \times p}$ and $\mathbf{Y}_g \in \mathbb{R}^{n_g}$ are the data matrix and the response vector for the g th group, respectively. Each data matrix has the same set of p explanatory variables, whereas the samples vary across groups. We let $\mathbf{X} = [\mathbf{X}'_1, \dots, \mathbf{X}'_G]' \in \mathbb{R}^{n \times p}$ and $\mathbf{Y} = [\mathbf{Y}'_1, \dots, \mathbf{Y}'_G]' \in \mathbb{R}^n$, where $n = \sum_{g=1}^G n_g$.

Our model is closely related to JIVE, which provides a general formulation to decompose multiple data matrices into joint and individual structures. The JIVE decomposes \mathbf{X}_g as

$$\mathbf{X}_g = \mathbf{J}_g + \mathbf{A}_g + \mathbf{E}_g, \quad (2.1)$$

where $\mathbf{J}_g \in \mathbb{R}^{n_g \times p}$ represents the joint structure of \mathbf{X}_g , $\mathbf{A}_g \in \mathbb{R}^{n_g \times p}$ represents the individual structure of \mathbf{X}_g , and $\mathbf{E}_g \in \mathbb{R}^{n_g \times p}$ is the error matrix. We consider that \mathbf{Y}_g has a similar decomposition into joint and individual signals

$$\mathbf{Y}_g = \mathbf{J}_g^Y + \mathbf{A}_g^Y + \mathbf{e}_g, \quad (2.2)$$

where $\mathbf{e}_g \in \mathbb{R}^{n_g}$ is the noise from the g th group. Let $\tilde{\mathbf{X}}_g = \mathbf{J}_g + \mathbf{A}_g$ and $\tilde{\mathbf{Y}}_g = \mathbf{J}_g^Y + \mathbf{A}_g^Y$ be the noiseless counterparts of \mathbf{X}_g and \mathbf{Y}_g . Lemma 1 gives conditions to ensure that \mathbf{J}_g , \mathbf{A}_g , \mathbf{J}_g^Y , and \mathbf{A}_g^Y are identifiable.

Lemma 1. Given $\{\tilde{\mathbf{X}}_g, \tilde{\mathbf{Y}}_g\}_{g=1}^G$, where $\tilde{\mathbf{Y}}_g \in \text{col}(\tilde{\mathbf{X}}_g)$. There exist unique \mathbf{J}_g and \mathbf{A}_g such that:

- (i) $\tilde{\mathbf{X}}_g = \mathbf{J}_g + \mathbf{A}_g$;
- (ii) $\text{row}(\mathbf{J}_1) = \dots = \text{row}(\mathbf{J}_G) \subset \text{row}(\tilde{\mathbf{X}}_g)$;
- (iii) $\text{row}(\mathbf{J}_g) \perp \text{row}(\mathbf{A}_g)$, for $g = 1, \dots, G$;
- (iv) $\bigcap_{g=1}^G \text{row}(\mathbf{A}_g) = \{\mathbf{0}_{p \times 1}\}$.

Moreover, if $\text{col}(\mathbf{J}_g) \perp \text{col}(\mathbf{A}_g)$, then there exist unique \mathbf{J}_g^Y and \mathbf{A}_g^Y such that $\tilde{\mathbf{Y}}_g = \mathbf{J}_g^Y + \mathbf{A}_g^Y$ and they satisfy $\mathbf{J}_g^Y \in \text{col}(\mathbf{J}_g)$ and $\mathbf{A}_g^Y \in \text{col}(\mathbf{A}_g)$.

Lemma 1 shows that $\tilde{\mathbf{X}}_g$ can be uniquely decomposed into the sum of \mathbf{J}_g and \mathbf{A}_g if we require them to satisfy conditions (ii) – (iv), following similar statements as in Feng et al. (2018). To ensure the unique decomposition of $\tilde{\mathbf{Y}}_g$, we need to further require $\text{col}(\mathbf{J}_g) \perp \text{col}(\mathbf{A}_g)$, which is different from Palzer et al. (2022), that requires $\text{row}(\mathbf{J}_g) \perp \text{row}(\mathbf{A}_g)$.

In practice, only \mathbf{X}_g and \mathbf{Y}_g are observable. In Lemma 2, we show in (a) that the identifiable conditions in Lemma 1 can still

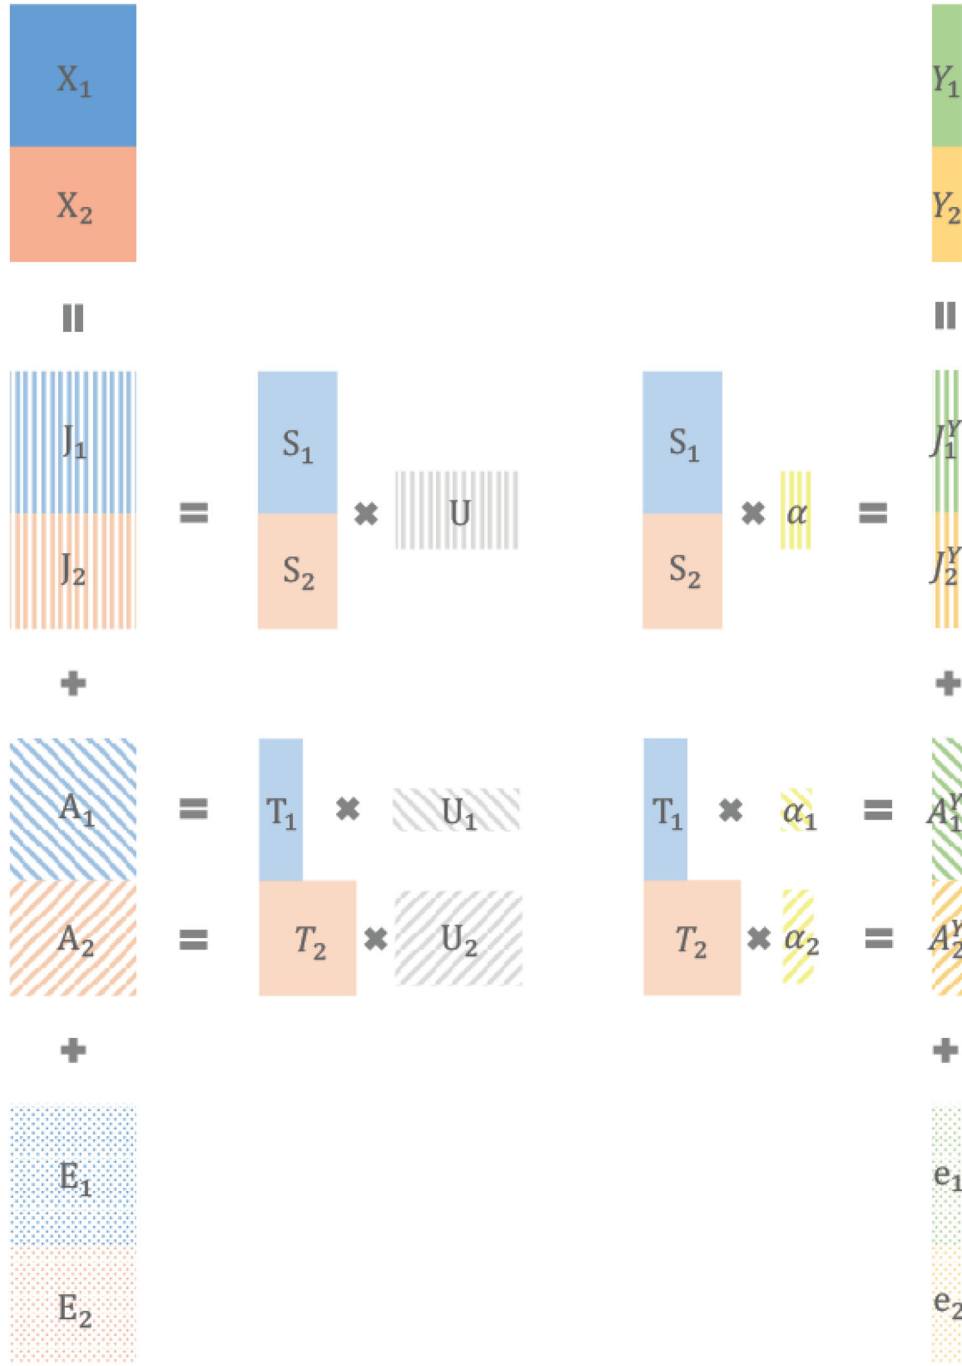


Figure 1. Diagram of JICO for two groups. $X_g \in \mathbb{R}^{n_g \times p}$: data matrix from group g . $J_g \in \mathbb{R}^{n_g \times p}$: submatrix of X_g corresponding to the joint structure. $A_g \in \mathbb{R}^{n_g \times p}$: submatrix of X_g corresponding to the individual structures. $E_g \in \mathbb{R}^{n_g \times p}$: error matrix. $Y_g \in \mathbb{R}^{n_g}$: response from group g . $J_g^Y \in \mathbb{R}^{n_g}$: subvector of Y_g corresponding to the joint component. $A_g^Y \in \mathbb{R}^{n_g}$: subvector of Y_g corresponding to the individual components. $e_g \in \mathbb{R}^{n_g}$: noise vector. $S_g \in \mathbb{R}^{n_g \times K}$: joint score matrix. $U \in \mathbb{R}^{K \times p}$: joint loading matrix. $T_g \in \mathbb{R}^{n_g \times K_g}$: individual score matrix. $U_g \in \mathbb{R}^{K_g \times p}$: individual loading matrix. $\alpha \in \mathbb{R}^K$: coefficients for the joint component. $\alpha_g \in \mathbb{R}^{K_g}$: coefficients for the individual components.

be achieved given the observed $\{X_g, Y_g\}_{g=1}^G$, when we construct \tilde{X}_g and \tilde{Y}_g under certain identifiability constraints. We achieve this by finding an arbitrary set of bases $w_1, \dots, w_K \in \mathbb{R}^p$ that span $\bigcap_{g=1}^G \text{row}(X_g)$ and $w_{g,1}, \dots, w_{g,K_g}$ that spans the space of solutions to (A.1) as shown in the proof of Lemma 2 in Appendix A. Using these bases, we can construct J_g and A_g such that they satisfy other identifiable conditions in Lemma 1, and construct \tilde{X}_g as $\tilde{X}_g = J_g + A_g$. Moreover, in Lemma 2(b), we show that if J_g and A_g are assumed to have low ranks, they can

be further decomposed as $J_g = S_g U$, where S_g is a $n_g \times K$ score matrix, U is a $K \times p$ loading matrix, and $K = \text{rank}(J_g)$; and $A_g = T_g U_g$, where T_g is a $n_g \times K_g$ score matrix and U_g is a $K_g \times p$ loading matrix, and $K_g = \text{rank}(A_g)$. Under this formulation, if $S_g' T_g = \mathbf{0}_{K \times K_g}$, then (2.1) and (2.2) can be expressed as

$$X_g = S_g U + T_g U_g + E_g, \quad (2.3)$$

$$Y_g = S_g \alpha + T_g \alpha_g + e_g, \quad (2.4)$$

where $\alpha \in \mathbb{R}^K$ and $\alpha_g \in \mathbb{R}^{K_g}$ are the coefficients of the joint and individual components respectively. Model (2.4) gives a unified framework to model multi-group data. When $K = 0$, the joint term $\mathbf{S}_g \alpha$ vanishes and (2.4) reduces to a group-specific model of $\mathbf{Y}_g = \mathbf{T}_g \alpha_g + \mathbf{e}_g$. On the other hand, when $K_1 = \dots = K_g = 0$, the individual term $\mathbf{T}_g \alpha_g$ vanishes and (2.4) reduces to a global model of $\mathbf{Y}_g = \mathbf{S}_g \alpha + \mathbf{e}_g$. When $K \neq 0$ and $K_g \neq 0$, (2.4) has both global and group-specific components, thus lies between the above two extreme cases.

Lemma 2. Given $\{\mathbf{X}_g, \mathbf{Y}_g\}_{g=1}^G$, we can construct $\tilde{\mathbf{X}}_g$ and $\tilde{\mathbf{Y}}_g$ such that

- There exist unique matrices \mathbf{J}_g and \mathbf{A}_g such that $\tilde{\mathbf{X}}_g = \mathbf{J}_g + \mathbf{A}_g$, conditions (ii) – (iv) in Lemma 1 and $\text{col}(\mathbf{J}_g) \perp \text{col}(\mathbf{A}_g)$ still hold.
- There exist matrices $\mathbf{U}, \mathbf{U}_g, \mathbf{S}_g, \mathbf{T}_g$ such that \mathbf{J}_g and \mathbf{A}_g can be expressed as $\mathbf{J}_g = \mathbf{S}_g \mathbf{U}$ and $\mathbf{A}_g = \mathbf{T}_g \mathbf{U}_g$, where $\mathbf{S}'_g \mathbf{T}'_g = \mathbf{0}_{K \times K_g}$. Moreover, there exist \mathbf{J}'_g and \mathbf{A}'_g that gives $\tilde{\mathbf{Y}}_g = \mathbf{J}'_g + \mathbf{A}'_g$ and can be expressed as $\mathbf{J}'_g = \mathbf{S}_g \alpha$ and $\mathbf{A}'_g = \mathbf{T}_g \alpha_g$, where $\alpha = (\mathbf{S}'\mathbf{S})^{-1} \mathbf{S}'\mathbf{Y}$ and $\alpha_g = (\mathbf{T}'_g \mathbf{T}_g)^{-1} \mathbf{T}'_g \mathbf{Y}_g$, with $\mathbf{S} = [\mathbf{S}'_1, \dots, \mathbf{S}'_G]'$.

Corollary 1. There exist matrices $\mathbf{W} \in \mathbb{R}^{p \times K}$ and $\mathbf{W}_g \in \mathbb{R}^{p \times K_g}$ such that $\mathbf{J}_g = \mathbf{S}_g \mathbf{U}$ and $\mathbf{A}_g = \mathbf{T}_g \mathbf{U}_g$ defined by $\mathbf{S}_g = \mathbf{X}_g \mathbf{W}$, $\mathbf{T}_g = \mathbf{X}_g \mathbf{W}_g$, $\mathbf{U} = (\mathbf{W}'\mathbf{W})^{-1} \mathbf{W}'$ and $\mathbf{U}_g = (\mathbf{W}'_g \mathbf{W}_g)^{-1} \mathbf{W}'_g$ as in Lemma 2(b) satisfy conditions (ii) – (iv) in Lemma 1 and $\text{col}(\mathbf{J}_g) \perp \text{col}(\mathbf{A}_g)$, if $\mathbf{W}'\mathbf{W}_g = \mathbf{0}$ and $\mathbf{W}'\mathbf{X}'_g \mathbf{X}_g \mathbf{W}_g = \mathbf{0}$, for all g .

Corollary 1 follows directly from the proof of Lemma 2(b). As a remark, the columns of \mathbf{W} and \mathbf{W}_g form the sets of bases that span the row spaces of \mathbf{J}_g and \mathbf{A}_g , respectively. Hence, $\mathbf{W}'\mathbf{W}_g = \mathbf{0}$ is a sufficient and necessary condition for $\text{row}(\mathbf{J}_g) \perp \text{row}(\mathbf{A}_g)$. Moreover, note that in Lemma 2(b), $\mathbf{S}'_g \mathbf{T}'_g = \mathbf{0}$ directly implies that $\mathbf{A}'_g \mathbf{J}_g = \mathbf{0}$, the latter being a sufficient condition for $\text{col}(\mathbf{J}_g) \perp \text{col}(\mathbf{A}_g)$. Therefore, in Corollary 1, $\mathbf{W}'\mathbf{X}'_g \mathbf{X}_g \mathbf{W}_g = \mathbf{0}$ provides a sufficient condition for $\text{col}(\mathbf{J}_g) \perp \text{col}(\mathbf{A}_g)$, which satisfies one of the identifiability constraints for the unique decomposition of \mathbf{X}_g in Lemma 1. In Section 3, we describe the algorithm to solve for \mathbf{W} and \mathbf{W}_g , respectively.

3. Model Estimation

The key to estimate (2.3) and (2.4) is the constructions of score matrices \mathbf{S}_g and \mathbf{T}_g . To motivate our estimation procedure, in Sections 3.1, we discuss the joint and individual score estimation under two special cases, respectively. In Section 3.2, we introduce an iterative algorithm for the general case.

3.1. Joint and Individual Score Estimation

We first consider a special case that $K_g = 0, g = 1, \dots, G$. Under this setup, the individual components vanish and (2.3) and (2.4) reduce to the following model:

$$\mathbf{X} = \mathbf{S}\mathbf{U} + \mathbf{E}, \quad \mathbf{Y} = \mathbf{S}\alpha + \mathbf{e}, \quad (3.1)$$

where $\mathbf{S} = [\mathbf{S}'_1, \dots, \mathbf{S}'_G]'$, $\mathbf{E} = [\mathbf{E}'_1, \dots, \mathbf{E}'_G]'$, and $\mathbf{e} = [\mathbf{e}'_1, \dots, \mathbf{e}'_G]'$.

The formulation of (3.1) covers many existing classic methods. For example, in PCR, \mathbf{S} is chosen to be the score matrix of the first K principal components of $\mathbf{X}'\mathbf{X}$. However, the principal components are inherently unsupervised and ignore the information from \mathbf{Y} . Among the other supervised methods, PLS regression is a popular approach that incorporates regression on the latent scores. When $K = 1$ and $p < n$, the standard OLS can also be cast under the above setup.

According to the proof of our Lemma 2, \mathbf{S} can be constructed with the basis matrix \mathbf{W} . For the estimation of \mathbf{W} , we use the continuum regression (CR) (Stone and Brooks 1990) algorithm, the result of which covers OLS, PLS, and PCR as special cases. For $k = 1, \dots, K$, CR sequentially solves \mathbf{w}_k from the following optimization problem:

$$\begin{aligned} \max_{\mathbf{w}} \quad & \text{cov}(\mathbf{X}\mathbf{w}, \mathbf{Y})^2 \text{var}(\mathbf{X}\mathbf{w})^{\gamma-1} \\ \text{s.t.} \quad & \mathbf{w}'\mathbf{w} = 1, \mathbf{w}'\mathbf{X}'\mathbf{X}\mathbf{w}_j = 0; \quad j = 1, \dots, k-1 \text{ if } k \geq 2, \end{aligned} \quad (3.2)$$

where $\text{cov}(\mathbf{X}\mathbf{w}, \mathbf{Y}) = \mathbf{w}'\mathbf{X}'\mathbf{Y}$ and $\text{var}(\mathbf{X}\mathbf{w}) = \mathbf{w}'\mathbf{X}'\mathbf{X}\mathbf{w}$, once columns of \mathbf{X} and \mathbf{Y} are centralized to have mean zero. Here, $\gamma \geq 0$ is a tuning parameter that controls how much variability of \mathbf{X} is taken into account for the construction of \mathbf{S} . When $\gamma \rightarrow \infty$, the objective function in (3.2) is dominated by $\text{var}(\mathbf{X}\mathbf{w})^{\gamma-1}$ and \mathbf{Y} does not play a role. The CR solution of \mathbf{W} then seeks to find the principal component directions that maximize the variation of \mathbf{X} . It can be shown that (3.2) coincides with OLS and PLS solutions when $\gamma = 0$ and 1, respectively.

Let $\hat{\mathbf{W}}$ denote the solution to (3.2) and $\hat{\mathbf{S}} = \mathbf{X}\hat{\mathbf{W}}$. Then α can be estimated by $\hat{\alpha} = (\hat{\mathbf{S}}'\hat{\mathbf{S}})^{-1} \hat{\mathbf{S}}'\mathbf{Y}$. As illustrated in Lemma 2, \mathbf{J}_g is the projection of \mathbf{X}_g onto the column space spanned by \mathbf{W} . Hence, we have $\hat{\mathbf{J}}_g = \mathbf{X}_g \hat{\mathbf{W}} (\hat{\mathbf{W}}'\hat{\mathbf{W}})^{-1} \hat{\mathbf{W}}'$, which further gives $\hat{\mathbf{U}} = (\hat{\mathbf{W}}'\hat{\mathbf{W}})^{-1} \hat{\mathbf{W}}'$.

Next we consider our model estimation under the special case that $K = 0$. In this case, the joint component vanishes, and (2.3) and (2.4) reduce to the following individual model:

$$\mathbf{X}_g = \mathbf{T}_g \mathbf{U}_g + \mathbf{E}_g, \quad \mathbf{Y}_g = \mathbf{T}_g \alpha_g + \mathbf{e}_g. \quad (3.3)$$

Same as the above discussion for joint score estimation, we use CR to construct $\mathbf{T}_g = \mathbf{X}_g \mathbf{W}_g$ as linear transformation of \mathbf{X}_g , where \mathbf{W}_g is a $p \times K_g$ basis matrix, whose columns span $\text{row}(\mathbf{A}_g)$. Let $\mathbf{W}_g = [\mathbf{w}_{g1}, \dots, \mathbf{w}_{gK_g}]$. Given group g , for $k = 1, \dots, K_g$, CR sequentially solves \mathbf{w}_{gk} from the following optimization problem:

$$\begin{aligned} \max_{\mathbf{w}} \quad & \text{cov}(\mathbf{X}_g \mathbf{w}, \mathbf{Y}_g)^2 \text{var}(\mathbf{X}_g \mathbf{w})^{\gamma-1} \\ \text{s.t.} \quad & \mathbf{w}'\mathbf{w} = 1, \\ & \mathbf{w}'\mathbf{X}'_g \mathbf{X}_g \mathbf{w}_j = 0; \quad j = 1, \dots, k-1 \text{ if } k \geq 2. \end{aligned} \quad (3.4)$$

Denote $\hat{\mathbf{W}}_g$ the solution to (3.4). Similar to the joint estimation, once $\hat{\mathbf{T}}_g = \mathbf{X}_g \hat{\mathbf{W}}_g$ is constructed, α_g can be obtained as the least square solution: $\hat{\alpha}_g = (\hat{\mathbf{T}}'_g \hat{\mathbf{T}}_g)^{-1} \hat{\mathbf{T}}'_g \mathbf{Y}_g$. Afterwards, we can have $\hat{\mathbf{A}}_g = \mathbf{X}_g \hat{\mathbf{W}}_g (\hat{\mathbf{W}}'_g \hat{\mathbf{W}}_g)^{-1} \hat{\mathbf{W}}'_g$ and $\hat{\mathbf{U}}_g = (\hat{\mathbf{W}}'_g \hat{\mathbf{W}}_g)^{-1} \hat{\mathbf{W}}'_g$.

3.2. JICO Algorithm

In this section, we consider the general case where K or K_g can be both nonzero. Since solving (3.2) and (3.4) simultaneously

can be hard with both joint and individual structures specified in the full model (2.3) and (2.4), we propose to iteratively solve one of them while fixing the other. This leads to the following iterative procedure.

- Given $\hat{\mathbf{W}}_g$, solve the following constrained problem sequentially for $\mathbf{w}_1, \dots, \mathbf{w}_K$:

$$\begin{aligned} \max_{\mathbf{w}} \quad & \text{cov}(\mathbf{X}^{\text{Joint}}\mathbf{w}, \mathbf{Y}^{\text{Joint}})^2 \text{var}(\mathbf{X}^{\text{Joint}}\mathbf{w})^{\gamma-1} \\ \text{s.t.} \quad & \mathbf{w}'\mathbf{w} = 1, \\ & \mathbf{w}'\mathbf{X}^{\text{Joint}'}\mathbf{X}^{\text{Joint}}\mathbf{w}_j = 0; \quad j = 1, \dots, k-1 \text{ if } k \geq 2, \\ & \hat{\mathbf{W}}_g'\mathbf{w} = \mathbf{0}_{K_g \times 1}; \quad g = 1, \dots, G, \\ & \hat{\mathbf{W}}_g'\mathbf{X}_g^{\text{Indiv}'}\mathbf{X}_g^{\text{Joint}}\mathbf{w} = \mathbf{0}_{K_g \times 1}; \quad g = 1, \dots, G. \end{aligned} \quad (3.5)$$

- Given $\hat{\mathbf{W}}$, for any $1 \leq g \leq G$, solve the following constrained problem sequentially for $\mathbf{w}_{g,1}, \dots, \mathbf{w}_{g,K_g}$:

$$\begin{aligned} \max_{\mathbf{w}} \quad & \text{cov}(\mathbf{X}_g^{\text{Indiv}}\mathbf{w}, \mathbf{Y}_g^{\text{Indiv}})^2 \text{var}(\mathbf{X}_g^{\text{Indiv}}\mathbf{w})^{\gamma-1} \\ \text{s.t.} \quad & \mathbf{w}'\mathbf{w} = 1, \\ & \mathbf{w}'\mathbf{X}_g^{\text{Indiv}'}\mathbf{X}_g^{\text{Indiv}}\mathbf{w}_{gj} = 0; \quad j = 1, \dots, k-1 \text{ if } k \geq 2, \\ & \hat{\mathbf{W}}'\mathbf{w} = \mathbf{0}_{K \times 1}; \\ & \hat{\mathbf{W}}'\mathbf{X}_g^{\text{Joint}'}\mathbf{X}_g^{\text{Indiv}}\mathbf{w} = \mathbf{0}. \end{aligned} \quad (3.6)$$

- Repeat the above two procedures until convergence.

Note that in (3.5) and (3.6), we denote

$$\mathbf{X}^{\text{Joint}} = \begin{bmatrix} \mathbf{X}_1 - \mathbf{T}_1\mathbf{U}_G \\ \vdots \\ \mathbf{X}_G - \mathbf{T}_G\mathbf{U}_G \end{bmatrix}, \quad \mathbf{Y}^{\text{Joint}} = \begin{bmatrix} \mathbf{Y}_1 - \mathbf{T}_1\boldsymbol{\alpha}_1 \\ \vdots \\ \mathbf{Y}_G - \mathbf{T}_G\boldsymbol{\alpha}_G \end{bmatrix},$$

and $\mathbf{X}_g^{\text{Indiv}} = \mathbf{X}_g - \mathbf{S}_g\mathbf{U}$, $\mathbf{Y}_g^{\text{Indiv}} = \mathbf{Y}_g - \mathbf{S}_g\boldsymbol{\alpha}$; $g = 1, \dots, G$. Moreover, the last two constraints in (3.5) and (3.6) correspond to the two sufficient conditions in Corollary 1 to satisfy the identifiability conditions $\text{row}(\mathbf{J}_g) \perp \text{row}(\mathbf{A}_g)$ and $\text{col}(\mathbf{J}_g) \perp \text{col}(\mathbf{A}_g)$ needed in Lemma 1.

We formulate (3.5) and (3.6) into a generic CR problem, and derive an algorithm to solve it in Appendix B of the supplementary materials. Furthermore, we describe the convergence criterion for the iterative procedure and give its pseudo code in Appendix C. Empirically, the algorithm always meets our convergence criteria, albeit there are no theoretical guarantees. In practice, we recommend starting the algorithm with multiple initial values and choose the one with the smallest cross-validated mean squared error. To predict the response using JICO estimates, we let $\mathbf{S}_{g,\text{test}} = \mathbf{X}_{g,\text{test}}\hat{\mathbf{W}}$ and $\mathbf{T}_{g,\text{test}} = \mathbf{X}_{g,\text{test}}\hat{\mathbf{W}}_g$, where $\mathbf{X}_{g,\text{test}}$ is the test set. Then the prediction of response is given by $\mathbf{S}_{g,\text{test}}\hat{\boldsymbol{\alpha}} + \mathbf{T}_{g,\text{test}}\hat{\boldsymbol{\alpha}}_g$.

In practice, we need to select tuning parameters K , K_g , and γ . As a rule of thumb, we propose to select the optimal γ by fine-tuning it in a wide range with a coarse grid search of ranks K and K_g . For rank selection, we recommend tuning ranks by performing an exhaustive search on $K \in \{0, 1, \dots, D_1\}$ and $K_1 = \dots = K_G \in \{0, 1, \dots, D_2\}$, where D_1 and D_2 are two user-defined integers. We describe how to select the tuning parameters in more details and perform a sensitivity study on how the

selection affects JICO's numerical performance in Appendix D of the supplementary materials.

Finally, we point out that our method includes JIVE-predict (Kaplan and Lock 2017) as a special case. JIVE-predict is a two-stage method that implements JIVE on \mathbf{X} first and then regresses the responses on the loading matrix. When we let $\gamma \rightarrow \infty$ in (3.2) and (3.4), JICO is equivalent as performing JIVE on \mathbf{X} . For that reason, our method in that case is equivalent to JIVE-predict.

4. Simulation Studies

One significant advantage of our proposed model is its flexibility of lying in between global and group-specific models. Moreover, the choice of parameter γ in CR allows it to identify the model that best fits the data. In this section, we conduct multiple simulation studies to further demonstrate the advantage of our proposed model.

We consider three simulation settings in this section. In the first two settings, we generate data according to models that contain both global and group-specific components. The data are generated in a way that PCR and PLS solutions are favored, respectively. In the last setting, we simulate data from two special cases: a global model and a group-specific model. The data are simulated so that the OLS is favored for both cases. For all three settings, JICO can adaptively choose the correct model parameter γ so that it has the optimal performance. Moreover, we further illustrate how the rank selection impacts the performance of JICO by examining the results using mis-specified ranks.

We fix $G = 2, p = 200, n_1 = n_2 = 50$. In each replication, we generate 100 training samples to train the models and evaluate the corresponding Mean Squared Error (MSE) in an independent test set of 100 samples. We repeat simulations for 50 times.

For $g = 1, \dots, G$, we generate \mathbf{X}_g as iid samples from $\mathcal{N}(\mathbf{0}, \mathbf{I}_{p \times p})$. For the sake of simplicity, we generate \mathbf{Y}_g by the following model with two latent components:

$$\mathbf{Y}_g = \boldsymbol{\alpha}\mathbf{S}_g + \boldsymbol{\alpha}_g\mathbf{T}_g + \mathbf{e}_g, \quad (4.1)$$

where $\mathbf{S}_g = \mathbf{X}_g\mathbf{w} \in \mathbb{R}^{n_g}$ is the joint latent score vector with an coefficient $\boldsymbol{\alpha}$, $\mathbf{T}_g = \mathbf{X}_g\mathbf{w}_g \in \mathbb{R}^{n_g}$ is the individual latent score vector with an coefficient $\boldsymbol{\alpha}_g$, and \mathbf{e}_g is generated as iid samples from $\mathcal{N}(0, 0.04)$. Here, \mathbf{w} and \mathbf{w}_g are all $p \times 1$ vectors, and are constructed such that $\mathbf{w}'\mathbf{w}_g = 0$. We vary the choices of \mathbf{w} , \mathbf{w}_g , $\boldsymbol{\alpha}$, and $\boldsymbol{\alpha}_g$, which will be discussed later.

4.1. PCR Setting

In this section, we simulate the model which favors $\gamma = \infty$. In this case, CR solutions to (3.2) and (3.4) coincide with PCR, which are essentially the top eigenvectors of the corresponding covariance matrices.

To simulate this setup, given training data $\mathbf{X} = [\mathbf{X}'_1, \mathbf{X}'_2]'$, we let \mathbf{w} be the top eigenvector corresponding to $\mathbf{X}'\mathbf{X}$. We further set \mathbf{w}_g as the top eigenvector of $\tilde{\mathbf{X}}_g'\tilde{\mathbf{X}}_g$, where $\tilde{\mathbf{X}}_g = \mathbf{X}_g(\mathbf{I} - \mathbf{w}\mathbf{w}')$ is the data matrix after projecting \mathbf{X}_g into the linear subspace that is orthogonal to \mathbf{w} . This projection ensures that the construction of \mathbf{w} and \mathbf{w}_g satisfies $\mathbf{w}'\mathbf{w}_g = 0$. To generate \mathbf{Y}_g , we let $\boldsymbol{\alpha} = 1, \boldsymbol{\alpha}_g = 1, g = 1, 2$.

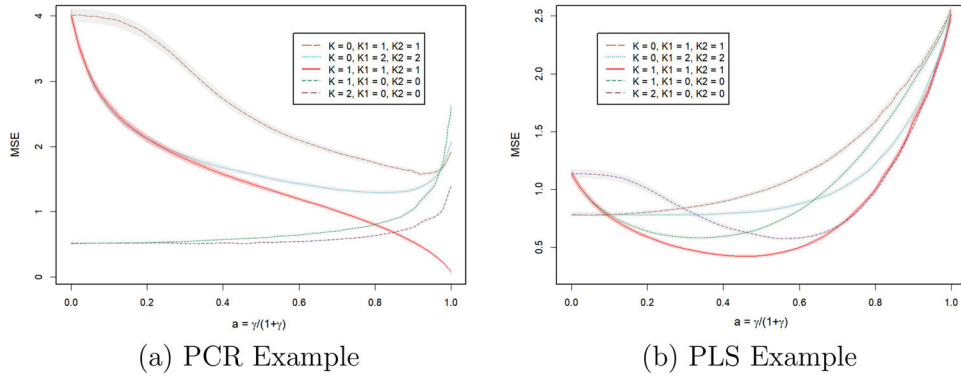


Figure 2. MSE curves for JICO models with different ranks under the PCR setting (a) and PLS setting (b).

We train JICO on a wide range of $\gamma \in [0, \infty)$, using different combinations of K, K_1, K_2 , with 300 iterations. Figure 2(a) demonstrates the MSEs evaluated on the test data over 50 repetitions. For better illustration, we plot MSE curves as a function of a , with $a = \gamma/(\gamma + 1)$, which is a one-to-one monotone map from $\gamma \in [0, \infty)$ to $a \in [0, 1]$. In particular, when $a = 0, 0.5$ and 1 , we have $\gamma = 0, 1$ and ∞ , which correspond to the cases of OLS, PLS and PCR respectively. The solid curve demonstrates the model performance given true ranks $K = K_1 = K_2 = 1$, whereas the gray curves show the performance of models with mis-specified ranks. In particular, we consider four mis-specified rank combinations. Among them, two rank combinations ($K = 1, K_1 = K_2 = 0$; $K = 2, K_1 = K_2 = 0$) correspond to joint models. The other two combinations ($K = 0, K_1 = K_2 = 1$; $K = 0, K_1 = K_2 = 2$) correspond to group-specific models. We can see from Figure 2(a) that the absolute minimum is given by the model with true ranks and $a = 1$, which refers to the underlying true model. When we look at the curves on the spectrum of a as a whole, the joint models with $K = 1$ or $2, K_1 = K_2 = 0$ always perform worse than the model with $K = K_1 = K_2 = 1$, because they are unable to capture the group-specific information from the underlying model. The model with true ranks performs better than the individual models with $K = 0, K_1 = K_2 = 1$ or 2 for larger values of a , because the latter models cannot capture as much global information as the former. However, the model with $K = K_1 = K_2 = 1$ performs worse than the individual models for smaller values of a , where the latter achieves much more acceptable performances. This means that the choice of optimal ranks for our model can be sensitive to the choice of γ . For smaller γ values, individual models tend to be more reliable under the PCR setting. We notice that the end of the curve is not very smooth when $K = 2, K_1 = K_2 = 0$. One possible reason is that the solution path of CR can sometimes be discontinuous with respect to γ (Björkström and Sundberg 1996), consequently the CR algorithm may be numerically unstable for certain γ values.

We further illustrate the performance of JICO by comparing it with several existing methods. In particular, we include ridge regression (Ridge), partial least squares (PLS) and principal component regression (PCR). For JICO, we select the models trained under true ranks $K = K_1 = K_2 = 1$ (performance as illustrated by the solid curve in Figure 2(a)), with $\gamma = 0, 1, \infty$, which correspond to the cases of OLS, PLS, and PCR,

respectively. For a fair comparison, for PLS and PCR methods, we fix the number of components to be 2 for both a global fit and a group-specific fit. Table 1(a) summarizes the MSEs of these methods. The numbers provided in the brackets represent the standard error. The first two columns summarize the performance for each group ($g = 1, 2$), and the last column summarizes the overall performance. The JICO model with $\gamma = \infty$ performs significantly better than the rest, because it agrees with the underlying true model. Among other mis-specified methods, group-specific PLS is relatively more robust to model mis-specification.

4.2. PLS Setting

In this section, we consider the model setup that is more favorable to $\gamma = 1$. In this scenario, the CR solutions to (3.2) and (3.4) coincide with the PLS solutions. Same as in Section 4.1, we still consider the construction of weights as linear transformations of the eigenvectors.

Given training data $\mathbf{X} = [\mathbf{X}'_1, \mathbf{X}'_2]'$, denote $\mathbf{V}_{p \times q}$ as the matrix of top q eigenvectors of $\mathbf{X}'\mathbf{X}$. We let $\mathbf{w} = \mathbf{V}\mathbf{1}_q/\sqrt{q}$, where $\mathbf{1}_q$ denotes a $q \times 1$ vector with elements all equal to 1. In this way, the q top eigenvectors contribute equally to the construction of \mathbf{S}_g . Similarly, we let $\tilde{\mathbf{X}}_g = \mathbf{X}_g(\mathbf{I} - \mathbf{w}\mathbf{w}')$ and \mathbf{V}_g be the $p \times q_g$ matrix of top q_g eigenvectors of $\tilde{\mathbf{X}}'_g\tilde{\mathbf{X}}_g$. Then we let $\mathbf{w}_g = \mathbf{V}_g\mathbf{1}_{q_g}/\sqrt{q_g}$. To construct a model more favorable to PLS, in this section, we let $q = n/2$ and $q_g = n_g/2$. We generate \mathbf{Y}_g from (4.1) by letting $\alpha = 1$ and $\alpha_g = 0.5$.

Similar to the PCR setting, in Figure 2(b), we illustrate the MSE curves of JICO models with different rank combinations on a spectrum of a , where $\gamma = a/(1 - a)$. Again, the solid curve represents the model with true ranks, while the gray curves represent models with mis-specified ranks. The absolute minimum is given by the solid curve at a around 0.5, which corresponds to the underlying true model. Moreover, the solid curve gives almost uniformly the best performance on the spectrum of a compared with the gray curves, except on a small range of a close to 0. Hence, under the PLS setting, the optimal ranks can be less sensitive to the choice of γ . At initial values of a , the solid curve almost overlaps with the gray curve that represents the joint model with $K = 1, K_1 = K_2 = 0$. This means that when γ is close to 0, the individual signals identified by the full model with $K = K_1 = K_2 = 1$ can be ignored. Therefore, the two

Table 1. Groupwise and overall MSEs under the PCR, PLS, and OLS settings.

		Method	$g = 1$	$g = 2$	Overall
(a) PCR example	JICO	$\gamma = 0$	1.994(0.063)	2.012(0.068)	2.003(0.053)
		$\gamma = 1$	0.679(0.018)	0.701(0.026)	0.69(0.017)
		$\gamma = \infty$	0.04 (0.001)	0.04 (0.001)	0.04 (0.001)
	Global	Ridge	1.734(0.056)	1.78(0.065)	1.757(0.052)
		PLS	1.163(0.033)	1.194(0.045)	1.178(0.031)
		PCR	0.946(0.04)	0.977(0.044)	0.961(0.022)
	Group-specific	Ridge	0.252(0.009)	0.27(0.009)	0.261(0.005)
		PLS	0.254(0.009)	0.272(0.009)	0.263(0.005)
		PCR	0.68(0.042)	0.71(0.05)	0.695(0.032)
(b) PLS example	JICO	$\gamma = 0$	0.57(0.023)	0.567(0.021)	0.569(0.018)
		$\gamma = 1$	0.211 (0.008)	0.218 (0.008)	0.215 (0.006)
		$\gamma = \infty$	1.236(0.038)	1.277(0.037)	1.256(0.025)
	Global	Ridge	1.698(0.064)	1.742(0.06)	1.72(0.055)
		PLS	0.3(0.011)	0.297(0.01)	0.299(0.008)
		PCR	1.229(0.036)	1.298(0.041)	1.263(0.025)
	Group-specific	Ridge	0.375(0.013)	0.425(0.016)	0.4(0.01)
		PLS	0.412(0.014)	0.406(0.016)	0.409(0.008)
		PCR	1.234(0.037)	1.25(0.036)	1.242(0.024)
(c) OLS example (a)	JICO	$\gamma = 0$	0.082 (0.002)	0.083 (0.003)	0.082 (0.002)
		$\gamma = 1$	0.403(0.011)	0.419(0.011)	0.411(0.007)
		$\gamma = \infty$	1.006(0.031)	1.07(0.03)	1.038(0.02)
	Global	Ridge	0.084 (0.004)	0.084 (0.003)	0.084 (0.003)
		PLS	0.221(0.007)	0.226(0.006)	0.223(0.005)
		PCR	0.991(0.032)	1.069(0.030)	1.030(0.020)
	Group-specific	Ridge	0.574(0.017)	0.599(0.024)	0.586(0.013)
		PLS	0.572(0.016)	0.599(0.024)	0.585(0.013)
		PCR	0.996(0.032)	1.061(0.030)	1.028(0.021)
(d) OLS example (b)	JICO	$\gamma = 0$	0.063 (0.002)	0.066 (0.004)	0.064 (0.002)
		$\gamma = 1$	0.257(0.009)	0.27(0.009)	0.264(0.006)
		$\gamma = \infty$	1.004(0.031)	1.002(0.03)	1.003(0.023)
	Global	Ridge	0.646(0.021)	0.673(0.024)	0.66(0.019)
		PLS	0.957(0.027)	0.971(0.032)	0.964(0.022)
		PCR	1.023(0.031)	1.016(0.031)	1.019(0.023)
	Group-specific	Ridge	0.076(0.004)	0.072(0.005)	0.074(0.003)
		PLS	0.113(0.003)	0.116(0.005)	0.115(0.003)
		PCR	0.978(0.03)	0.987(0.029)	0.983(0.022)

NOTE: Bolded values highlight the methods with the best performance. Numbers in brackets are standard errors.

group-specific models that capture more individual information give the best performance in this case. For a values closer to 1, the gray curve that represents the joint model with $K = 2, K_1 = K_2 = 0$ is very close to the solid curve. This means that the effects of individual components estimated by JICO tend to become more similar across groups for larger γ .

In Table 1(b), we summarize the MSEs of JICO models trained with true ranks $K = K_1 = K_2 = 1$ and $\gamma = 0, 1, \infty$, along with other methods as described in Section 4.1. JICO with $\gamma = 1$ shows the best performance among all methods, followed by the global PLS method, since the true model favors PLS and the coefficient $\alpha_g = 0.5$ for the group-specific component is relatively small.

4.3. OLS Setting

In this section, we simulate the setting that favors $\gamma = 0$, which corresponds to the case of OLS in CR. It is shown in Stone and Brooks (1990) that when $\gamma = 0$, there is only one nondegenerate direction that can be constructed from the CR algorithm. Hence, under the JICO framework, the model that favors $\gamma = 0$ embraces two special cases: a global model with $K = 1, K_g = 0$ and a group-specific model with $K = 0, K_g = 1$.

For the two cases, we simulate \mathbf{Y}_g with (a) $\alpha = 1, \alpha_g = 0$ and (b) $\alpha = 0, \alpha_g = 1$, respectively. The construction of \mathbf{w} and \mathbf{w}_g is the same as that in Section 4.2 with $q = n$ and $q_g = n_g$.

Figure 3 illustrates MSE curves of the two cases, where (a) represents the case of the global model and (b) represents the case of the group-specific model. In both cases, the absolute minimum can be found on the solid curves at $a = 0$, which represents the MSE curves from the models with true ranks $K = 1, K_g = 0$ and $K = 0, K_g = 1$, respectively. In (a), there are two competitive models against the model with true ranks: another global model with $K = 2, K_g = 0$ and the model with $K = K_1 = K_2 = 1$. They both achieve the same performance with the solid curve at $a = 0$, and stay low at larger values of a . This is because, when γ is mis-specified, additional model ranks help capture more information from data. The $K = 2, K_g = 0$ model performs better because the underlying model is a global model. This is also true for (b). The global minimum can be found at $a = 0$ on the solid curve, while the $K = 0, K_g = 2$ model performs better when a gets larger. Again, this is because larger K_g helps capture more information from data. The $K = K_1 = K_2 = 1$ model does not perform as well, because the estimated joint information dominates, which does not agree with the true model. We observe some discontinuities

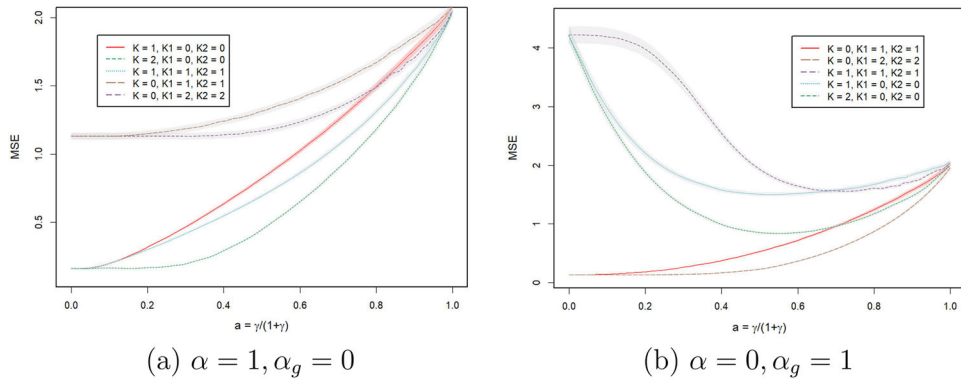


Figure 3. MSE curves for JICO models with different ranks under OLS settings (a) and (b). (a) is generated under a global model and (b) is generated under a group-specific model.

on the $K = 2, K_g = 0$ curve, since the CR solution path can sometimes be discontinuous with respect to γ as discussed in the PCR setting in Section 4.1.

In Table 1(c) and (d), we summarize the MSEs of JICO models trained with the true ranks with $\gamma = 0, 1, \infty$ and other methods described in Section 4.1. For a fair comparison, the number of components for PCR and PLS is chosen to be 1 for both global and group-specific fits. The JICO model with $\gamma = 0$, along with Ridge always achieve better performance than all other methods. It is interesting to notice that in (c), the JICO models with $\gamma = 1$ and ∞ coincide with global PLS and PCR, respectively, and hence they achieve the same performances. Similarly, in (d), JICO models with $\gamma = 1$ and ∞ coincide with group-specific PLS and PCR respectively, and they achieve the same performance correspondingly. In addition, when $K = 1, K_g = 0$, the solution of CR algorithm coincides with the global OLS model. Thus, the JICO model with $\gamma = 0$ and the global Ridge have similar performance in (c). Similarly, when $K = 0, K_g = 1$, the JICO model with $\gamma = 0$ and group-specific Ridge have similar performance in (d).

5. Applications to ADNI Data Analysis

We apply our proposed method to analyze data from the Alzheimer’s Disease (AD) Neuroimaging Initiative (ADNI). It is well known that AD accounts for most forms of dementia characterized by progressive cognitive and memory deficits. The increasing incidence of AD makes it a very important health issue and has attracted a lot of scientific attentions. To predict the AD progression, it is very important and useful to develop efficient methods for the prediction of disease status and clinical scores (e.g., the Mini Mental State Examination (MMSE) score and the AD Assessment Scale-Cognitive Subscale (ADAS-Cog) score). In this analysis, we are interested in predicting the ADAS-Cog score by features extracted from 93 brain regions scanned from structural magnetic resonance imaging (MRI). All subjects in our analysis are from the ADNI2 phase of the study. There are 494 subjects in total in our analysis and three subgroups: NC (178), eMCI (178), and AD (145), where the numbers in parentheses indicate the sample sizes for each subgroup. As a reminder, NC stands for the Normal Control, and eMCI stands for the early stage of Mild Cognitive Impairment in AD progression.

For each group, we randomly partition the data into two parts: 80% for training the model and the rest for testing the performance. We repeat the random split for 50 times. The testing MSEs and the corresponding standard errors are reported in Table 2. Both groupwise and overall performance are summarized. We compare our proposed JICO model with ridge regression (Ridge), PLS, and PCR. We perform both a global and a group-specific fit for Ridge, PLS, and PCR, where the regularization parameter in Ridge and the number of components in PCR or PLS are tuned by 10-fold cross validation (CV). For our proposed JICO model, we demonstrate the result by fitting the model with fixed $\gamma = 0, 0.25, 1, \infty$, or tuned γ , respectively. In practice, using exhaustive search to select the optimal values for K and K_g can be computationally cumbersome, because the number of combinations grows exponentially with the number of candidates for each parameter. Based on our numerical experience, we find that choosing K_g to be the same does not affect the performance on prediction too much. Details are discussed in Appendix D of the supplementary materials. Therefore, in all these cases, the optimal ranks for JICO are chosen by an exhaustive search in $K \in \{0, 1\}$ and $K_1 = K_2 = K_3 \in \{0, 1\}$ to see which combination gives the best MSE. We choose K and K_g to be small to improve our model interpretations. The optimal value of γ is chosen by 10-fold CV.

As shown in Table 2, JICO performs the best among all competitors. Fitting JICO with $\gamma = 0.25$ yields the smallest overall MSE. JICO with parameters chosen by CV performs slightly worse, but is still better than the other global or group-specific methods. The results of JICO with $\gamma = 0, 1$ and ∞ , which correspond to OLS, PLS, and PCR, are also provided in Table 2. Even though their prediction is not the best, an interesting observation is that they always have better performance than their global or group-specific counterparts. For example, when $\gamma = 1$, JICO has much better overall prediction than the group-specific PLS. This indicates that it is beneficial to capture global and individual structures for regression when subpopulations exist in the data.

In Table 2, global models perform the worst, because they do not take group heterogeneity into consideration. The group-specific Ridge appears to be the most competitive one among group-specific methods. Note that for the AD group, our JICO model with $\gamma = 0.25$ or tuned γ outperforms the group-specific Ridge method by a great margin.

Table 2. Groupwise and overall MSEs on the ADNI data.

Method		NC	eMCI	AD	Overall
JICO	$\gamma = 0$	6.671(0.137)	11.319(0.309)	55.798(1.556)	22.821(0.466)
	$\gamma = .25$	6.316 (0.121)	10.394(0.279)	40.853 (1.294)	17.951 (0.394)
	$\gamma = 1$	6.443(0.124)	10.353(0.291)	44.054(1.449)	18.929(0.441)
	$\gamma = \infty$	6.608(0.138)	11.121(0.308)	49.997(1.832)	21.013(0.558)
	CV	6.414(0.129)	10.333 (0.289)	41.297(1.348)	18.096(0.401)
Global	Ridge	23.450(0.751)	21.276(0.796)	63.989(2.657)	34.692(0.840)
	PLS	26.310(0.787)	22.672(0.915)	68.193(3.183)	37.442(0.982)
	PCR	25.228(0.771)	21.966(0.802)	69.541(2.969)	37.209(0.907)
Group-specific	Ridge	6.336(0.116)	10.353(0.278)	42.271(1.315)	18.364(0.392)
	PLS	6.656(0.136)	11.298(0.306)	48.434(1.725)	20.629(0.524)
	PCR	6.656(0.136)	11.346(0.304)	47.357(1.484)	20.327(0.454)

NOTE: Bolded values highlight the methods with the best performance. Numbers in brackets are standard errors.

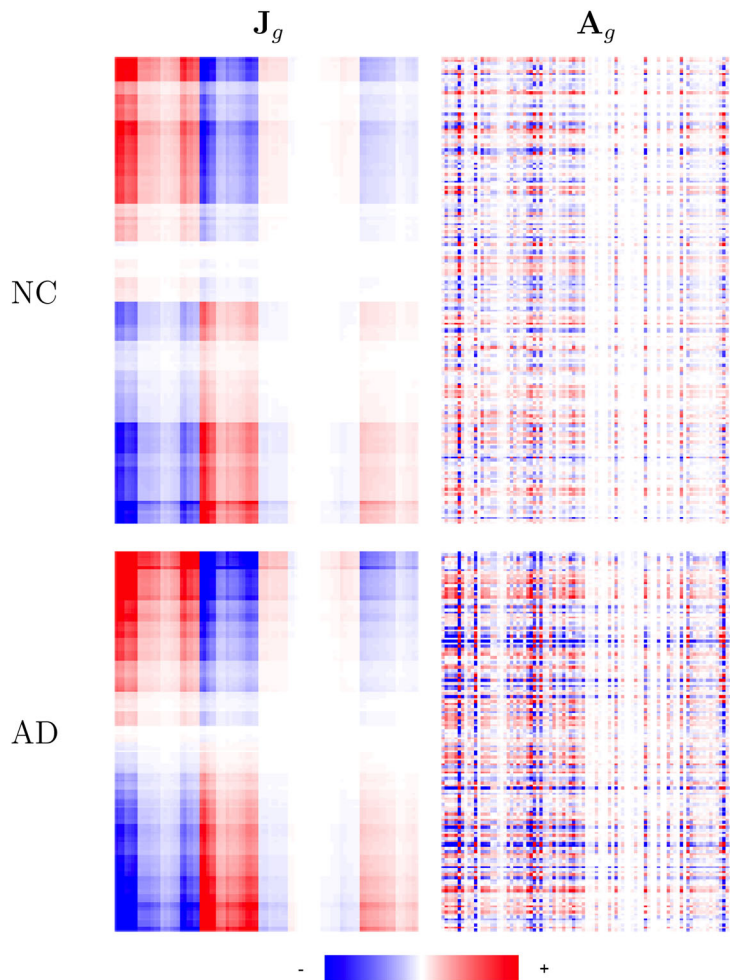


Figure 4. Heatmaps of joint and individual structures from NC and AD estimated from JICO.

To get our results more interpretable, we further apply the JICO model to NC and AD groups. We run 50 replications of 10-fold CV to see which combination of tuning parameters gives the smallest overall MSE. The best choice is $\gamma = \infty, K = 1, K_{NC} = K_{AD} = 3$. Then, we apply JICO using this choice and tuning parameters and display the heatmaps of the estimated \hat{J}_g (left column) and \hat{A}_g (right column) in Figure 4. Rows of heatmaps represent samples and columns represent MRI features. We use the Ward's linkage to perform hierarchical clustering on the rows of \hat{J}_g , and arrange the rows of \hat{J}_g and \hat{A}_g in the same order for each group. Moreover, we apply the same clustering

algorithm to the columns of \hat{J}_g to arrange the columns in the same order across the two disease groups for both joint and individual structures. Figure 4 shows that JICO separates joint and individual structures effectively. The joint structures across different disease groups share a very similar pattern, whereas the individual structures appear to be very distinct. We further magnify the right column of Figure 4 in Figure 5 with the brain region names listed. We find that the variation in \hat{A}_g for the AD group is much larger than the counterpart for the NC group. We highlight the brain regions that differ the most between the two groups. The highlighted regions play crucial roles in humans




Figure 5. Heatmaps of individual structures from NC and AD estimated from JICO with MRI feature names. The highlighted regions are frontal lobe WM, temporal lobe WM, occipital lobe WM, parietal lobe WM, superior parietal lobule, middle temporal gyrus. The top row is from NC and the bottom row is from AD.

cognition, thus, are important in AD early diagnosis (Killiany et al. 1993; Michon et al. 1994). For example, Michon et al. (1994) suggested that anosognosia in AD results in part from frontal dysfunction. Killiany et al. (1993) showed that the temporal horn of the lateral ventricles can be used as antemortem markers of AD.

6. Discussion

In this article, we propose JICO, a latent component regression model for multi-group heterogeneous data. Our proposed model decomposes the response into jointly shared and group-specific components, which are driven by low-rank approximations of joint and individual structures from the predictors, respectively. For model estimation, we propose an iterative procedure to solve for model components, and use CR algorithm that covers OLS, PLS, and PCR as special cases. As a result, the proposed procedure is able to extend many regression algorithms covered by CR to the setting of heterogeneous data. Extensive simulation studies and a real data analysis on ADNI data further demonstrate the competitive performance of JICO.

JICO is designed to be very flexible for multi-group data. It is able to choose the optimal parameter to determine the regression algorithm that suits the data the best, so that the prediction power is guaranteed. At the same time, it is also able to select the optimal joint and individual ranks

that best describe the degree of heterogeneity residing in each subgroup. The JICO application to ADNI data has effectively illustrated that our proposed model can provide nice visualization on identifying joint and individual components from the entire dataset without losing much of the prediction power.

Supplementary Materials

Appendix: Proofs of all theoretical results, Algorithm details, and additional numerical results. (Supplement.pdf, pdf file)

Acknowledgments

The authors are indebted to the editor, the associate editor, and two reviewers, whose helpful suggestions led to a much improved presentation.

Disclosure Statement

The authors report there are no competing interests to declare.

Funding

The authors were supported in part by NSF grants DMS-2100729 (Liu), SES-2217440 (Liu) and NIH grants R01AG073259 (Li) and R01GM126550 (Liu).

ORCID

Yufeng Liu  <https://orcid.org/0000-0002-1686-0545>

References

- Björkström, A., and Sundberg, R. (1996), “Continuum Regression Is Not Always Continuous,” *Journal of the Royal Statistical Society, Series B*, 58, 703–710. [2,6]
- Chen, J., Tran-Dinh, Q., Kosorok, M. R., and Liu, Y. (2021), “Identifying Heterogeneous Effect Using Latent Supervised Clustering with Adaptive Fusion,” *Journal of Computational and Graphical Statistics*, 30, 43–54. [1]
- Feng, Q., Jiang, M., Hannig, J., and Marron, J. (2018), “Angle-based Joint and Individual Variation Explained,” *Journal of Multivariate Analysis*, 166, 241–265. [2]
- Gaynanova, I., and Li, G. (2019), “Structural Learning and Integrative Decomposition of Multi-View Data,” *Biometrics*, 75, 1121–1132. [2]
- Hoerl, A. E., and Kennard, R. W. (1970), “Ridge Regression: Biased Estimation for Nonorthogonal Problems,” *Technometrics*, 12, 55–67. [1]
- Kaplan, A., and Lock, E. F. (2017), “Prediction with Dimension Reduction of Multiple Molecular Data Sources for Patient Survival,” *Cancer Informatics*, 16, 1176935117718517. [5]
- Killiany, R. J., Moss, M. B., Albert, M. S., Sandor, T., Tieman, J., and Jolesz, F. (1993), “Temporal Lobe Regions on Magnetic Resonance Imaging Identify Patients with Early Alzheimer’s Disease,” *Archives of Neurology*, 50, 949–954. [10]
- Lee, M. H., and Liu, Y. (2013), “Kernel Continuum Regression,” *Computational Statistics and Data Analysis*, 68, 190–201. [2]
- Li, Q., and Li, L. (2021), “Integrative Factor Regression and its Inference for Multimodal Data Analysis,” *Journal of the American Statistical Association*, 117, 2207–2221. [2]
- Lock, E. F., Hoadley, K. A., Marron, J. S., and Nobel, A. B. (2013), “Joint and Individual Variation Explained (JIVE) for Integrated Analysis of Multiple Data Types,” *The Annals of Applied Statistics*, 7, 523–542. [2]
- Ma, S., and Huang, J. (2017), “A Concave Pairwise Fusion Approach to Subgroup Analysis,” *Journal of the American Statistical Association*, 112, 410–423. [1]
- Meinshausen, N., and Bühlmann, P. (2015), “Maximin Effects in Inhomogeneous Large-Scale Data,” *The Annals of Statistics*, 43, 1801–1830. [1]
- Michon, A., Deweer, B., Pillon, B., Agid, Y., and Dubois, B. (1994), “Relation of Anosognosia to Frontal Lobe Dysfunction in Alzheimer’s Disease,” *Journal of Neurology, Neurosurgery and Psychiatry*, 57, 805–809. [10]
- Muniatégui, A., Pey, J., Planes, F. J., and Rubio, A. (2013), “Joint Analysis of miRNA and mRNA Expression Data,” *Briefings in Bioinformatics*, 14, 263–278. [1]
- Palzer, E. F., Wendt, C. H., Bowler, R. P., Hersh, C. P., Safo, S. E., and Lock, E. F. (2022), “sjive: Supervised Joint and Individual Variation Explained,” *Computational Statistics & Data Analysis*, 175, 107547. [2]
- Stone, M., and Brooks, R. J. (1990), “Continuum Regression: Cross-Validated Sequentially Constructed Prediction Embracing Ordinary Least Squares, Partial Least Squares and Principal Components Regression,” *Journal of the Royal Statistical Society, Series B*, 52, 237–258. [2,4,7]
- Tang, L., and Song, P. X. (2016), “Fused Lasso Approach in Regression Coefficients Clustering: Learning Parameter Heterogeneity in Data Integration,” *Journal of Machine Learning Research*, 17, 3915–3937. [1]
- Tibshirani, R. (1996), “Regression Shrinkage and Selection via the Lasso,” *Journal of the Royal Statistical Society, Series B*, 58, 267–288. [1]
- Wang, P., Li, Q., Shen, D., and Liu, Y. (2023), “High-Dimensional Factor Regression for Heterogeneous Subpopulations,” *Statistica Sinica*, 33, 1–27. [1]
- Wang, P., Liu, Y., and Shen, D. (2018), “Flexible Locally Weighted Penalized Regression with Applications on Prediction of Alzheimer’s Disease Neuroimaging Initiative’s Clinical Scores,” *IEEE Transactions on Medical Imaging*, 38, 1398–1408. [1]
- Zhao, T., Cheng, G., and Liu, H. (2016), “A Partially Linear Framework for Massive Heterogeneous Data,” *Annals of Statistics*, 44, 1400–1437. [1]