



HHS Public Access

Author manuscript

Comput Stat Data Anal. Author manuscript; available in PMC 2023 November 01.

Published in final edited form as:

Comput Stat Data Anal. 2022 November ; 175: . doi:10.1016/j.csda.2022.107547.

sJIVE: Supervised Joint and Individual Variation Explained*

Elise F. Palzer^a, Christine H. Wendt^b, Russell P. Bowler^c, Craig P. Hersh^d, Sandra E. Safo^a, Eric F. Lock^{a,*}

^aDivision of Biostatistics, University of Minnesota, Minneapolis, 55455, USA

^bDivision of Pulmonary, Allergy and Critical Care, University of Minnesota, Minneapolis, 55455, USA

^cDivision of Pulmonary, Critical Care and Sleep Medicine, Department of Medicine, National Jewish Health, Denver, CO, USA

^dChanning Division of Network Medicine, Brigham and Women's Hospital, Harvard Medical School, Boston, MA, USA

Abstract

Analyzing multi-source data, which are multiple views of data on the same subjects, has become increasingly common in molecular biomedical research. Recent methods have sought to uncover underlying structure and relationships within and/or between the data sources, and other methods have sought to build a predictive model for an outcome using all sources. However, existing methods that do both are presently limited because they either (1) only consider data structure shared by all datasets while ignoring structures unique to each source, or (2) they extract underlying structures first without consideration to the outcome. The proposed method, supervised joint and individual variation explained (sJIVE), can simultaneously (1) identify shared (joint) and source-specific (individual) underlying structure and (2) build a linear prediction model for an outcome using these structures. These two components are weighted to compromise between explaining variation in the multi-source data and in the outcome. Simulations show sJIVE to outperform existing methods when large amounts of noise are present in the multi-source data. An application to data from the COPDGene study explores gene expression and proteomic patterns associated with lung function.

Keywords

Data integration; Dimension reduction; Genomic data; High-dimensional prediction; Multi-source data; Multi-view learning

*Additional proofs, simulations, and details can be found in the supplement. All code is available online at <http://github.com/lockEF/r.jive>.

^{*}Corresponding author at: Division of Biostatistics, University of Minnesota, A460 Mayo Building, MMC 303, 420 Delaware Street S.E., Minneapolis, MN 55455, USA. elock@umn.edu, tel: 612-625-0651.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

1. Introduction

Access to multiple sets of characteristics, or views, on the same group of individuals is becoming increasingly common. Each view is distinct but potentially related to one another. For example, Genetic Epidemiology of COPD (COPDGene) [1] is an ongoing observational study that is obtaining numerous views on the same group of people including clinical, RNA sequencing (RNAseq), and proteomic data that aims to uncover the pathobiology in COPD. Such datasets, referred to as multi-view or multi-source data, have fueled an active area of statistical research to develop methods that [1] seek to uncover the relationships between and within each dataset, and/or [2] use multi-source data to create prediction models.

Underlying signals and associations that are shared across all datasets are called joint structures. Numerous methods can identify joint structure, many of which are extensions of canonical correlation analysis (CCA) [2, 3, 4, 5]. CCA constructs canonical variates, linear combinations of the variables, for each dataset such that the correlation between two datasets' variates is maximized [6]. Extensions of CCA [7] allow for more than two datasets, incorporate penalties for sparsity and regularization, or may maximize other measures of association such as covariance. Further extensions of CCA have incorporated supervision into the method in order for an outcome to influence the construction of the canonical variates [7, 8, 9, 10, 11]. By choosing canonical variates associated with an outcome, these CCA-based methods can build a prediction model based on the joint structure in the data. Other methods have combined CCA with discriminant analysis to build similar joint prediction models. In particular, joint association and classification analysis (JACA) [12] and sparse integrative discriminant analysis (SIDA) [13] combine linear discriminant and canonical correlation analyses to identify latent vectors that explain the association in multi-source data and that optimally separate subjects into different groups.

A limitation of predictive methods based on CCA is that they seek signal that is shared across all data sources, while relevant signal can also be specific to a single data source. One way to capture all variation in a dataset, shared or not shared, is by principle components analysis (PCA) [14]. PCA uses singular value decomposition (SVD) for dimension reduction while maximizing the variance. The low-rank output from PCA can then act as the design matrix in a regression framework for prediction [15]. In the multi-source setting, PCA for prediction can be applied to each dataset individually, or to a concatenated matrix of all datasets. However, often multi-source datasets will contain both joint (i.e., shared) and individual (i.e., source-specific) signal, which PCA does not distinguish. Since 2010, several more flexible methods have been developed to capture both joint and individual structures within multi-source data [16, 17, 18, 19, 20]. In particular, joint and individual variation explained (JIVE) [21, 22] is an extension of PCA and SVD that decomposes the data into low-rank and orthogonal joint and individual components. JIVE and the other methods referenced are solely exploratory, in that they do not inherently involve prediction or supervision for an outcome.

Few methods follow a supervised approach using both joint and individual signals from the data. Supervised integrated factor analysis (SIFA) [23] uses the outcome to supervise the construction of the joint and individual components, but does not have a predictive

model for the outcome. Recently, a bayesian method for joint association and prediction that incorporates prior functional information was proposed [24]. JIVE-predict [25] uses the joint and individual scores from the JIVE output to formulate a predictive model, and this approach has successfully been applied to COPD [26] and brain imaging data [27]. However, this 2-step approach always determines the joint and individual components without consideration for the outcome, which may hinder the predictive accuracy of the method.

In this paper, we propose supervised joint and individual variance explained (sJIVE) to find joint and individual components while simultaneously predicting a continuous outcome. This 1-step approach allows for joint and individual components to be influenced by their association with the outcome, explaining variation in both the multi-source data and the outcome in a single step.

The rest of the article is organized as follows. In Section 2, we review JIVE and introduce sJIVE's methodology and estimation technique. In Section 3, we compare sJIVE to JIVE and JIVE-predict under a variety of conditions. In Section 4, we compare sJIVE to existing methods. In Section 5, we apply sJIVE and other competing methods to COPD data, and conduct a pathway analysis to interpret the results. In Section 6, we discuss limitations and future research. Additional methodology details and simulation results can be found in our supplemental material.

2. Method and Estimation

2.1. Framework and Notation

Throughout this article, bold uppercase letters (\mathbf{A}) will denote matrices, bold lowercase letters (\mathbf{a}) will denote vectors, and unbolded lowercase letters (a) will denote scalars. Define the squared Frobenius norm for an $m \times n$ matrix \mathbf{A} by $\|\mathbf{A}\|_F^2 = \sum_{i=1}^m \sum_{j=1}^n a_{ij}^2$. Define the row space of a matrix \mathbf{A} by $\text{row}(\mathbf{A})$.

For our context, there are k views of data on the same n subjects. Each dataset \mathbf{X}_i , $i = 1, \dots, k$ must contain complete data and are of size $p_i \times n$ where p_i is the number of variables in the i th dataset. By default each row, or variable, in \mathbf{X}_i , $i = 1, \dots, k$ is centered and scaled to have mean 0 and variance 1. Let \mathbf{X} (without a subscript) represent the concatenation of all views of data, $\mathbf{X} = [\mathbf{X}_1^T \dots \mathbf{X}_k^T]^T$. Let \mathbf{y} be a length n outcome vector, which by default we center and scale to have mean 0 and variance 1. Our goal is to identify underlying structures within and between $\{\mathbf{X}_i\}_{i=1}^k$ and simultaneously build a predictive model for \mathbf{y} from \mathbf{X} .

2.2. Review of JIVE and JIVE-predict

We first review JIVE and JIVE-predict before explaining our proposed method. When $k = 2$, JIVE decomposes \mathbf{X} as follows:

$$\begin{aligned} \mathbf{X}_1 &= \mathbf{U}_1 \mathbf{S}_J + \mathbf{W}_1 \mathbf{S}_1 + \mathbf{E}_1, \text{ and} \\ \mathbf{X}_2 &= \mathbf{U}_2 \mathbf{S}_J + \mathbf{W}_2 \mathbf{S}_2 + \mathbf{E}_2, \end{aligned}$$

where \mathbf{U}_i and \mathbf{S}_J make up the joint component and \mathbf{W}_i and \mathbf{S}_i are the individual components with some error matrix, \mathbf{E}_i , $i = 1, 2$. In order to uniquely define the decomposition and distinguish the joint and individual signals, the joint and individual components in the i th dataset are orthogonal to each other. These structures are represented by reduced-rank matrices with the ranks r_J , r_1 , and r_2 where the ranks are pre-determined by a permutation [21] or Bayesian information criterion [22] approach, and r_J is the rank of the joint structure while r_i , $i = 1, \dots, k$ are the ranks of the individual structures. The loadings of the i th dataset, $\mathbf{U}_i \{p_i \times r_J\}$ and $\mathbf{W}_i \{p_i \times r_i\}$, map the p_i predictors to the low-rank subspace. Conversely, the score matrices, \mathbf{S}_J , \mathbf{S}_1 , and \mathbf{S}_2 , map the data from the low-rank subspace to the n subjects. Note that \mathbf{S}_J , the joint scores, are the same for \mathbf{X}_1 and \mathbf{X}_2 and thus capture structure in the samples that is shared across the sources.

JIVE-predict uses the JIVE output for prediction of an outcome, \mathbf{y} . The scores, \mathbf{S}_J and $\mathbf{S}_i \forall i = 1, \dots, k$, can be used as a design matrix for a regression model. For example, when $k = 2$, we can model \mathbf{y} linearly by

$$\mathbf{y} = \theta_1 \mathbf{S}_J + \theta_2 \mathbf{S}_1 + \theta_3 \mathbf{S}_2 + \mathbf{e}_Y,$$

where θ_1 , θ_2 , and θ_3 are vectors of length r_J , r_1 , and r_2 respectively such that the total number of regression parameters is the sum of the ranks. Since JIVE calculates the joint and individual components without consideration to an outcome, the predictive accuracy of this two-stage approach may be limited. Thus, our proposed method, sJIVE, allows for us to simultaneously construct joint and individual components while building a linear regression model. This allows for the components to be influenced by their association with \mathbf{y} , which could increase predictive accuracy.

2.3. Proposed model and objective

Consider the simple case when $k = 2$. Then, \mathbf{X} and \mathbf{y} can be decomposed by

$$\begin{aligned} \mathbf{X}_1 &= \mathbf{U}_1 \mathbf{S}_J + \mathbf{W}_1 \mathbf{S}_1 + \mathbf{E}_1, \\ \mathbf{X}_2 &= \mathbf{U}_2 \mathbf{S}_J + \mathbf{W}_2 \mathbf{S}_2 + \mathbf{E}_2, \text{ and} \\ \mathbf{y} &= \theta_1 \mathbf{S}_J + \theta_2 \mathbf{S}_1 + \theta_3 \mathbf{S}_2 + \mathbf{e}_Y. \end{aligned}$$

Similar to JIVE-predict, the scores from the decomposition of \mathbf{X}_1 and \mathbf{X}_2 are used to predict \mathbf{y} . However, the calculation of \mathbf{S}_J and $\mathbf{S}_i \forall i = 1, \dots, k$ are influenced by their association with the outcome. This can be accomplished by minimizing the loss of both \mathbf{X} and \mathbf{y} through the following optimization problem:

$$\begin{aligned} \underset{\mathbf{S}_J, \theta_1, \theta_2, \mathbf{U}_i, \mathbf{W}_i, \mathbf{S}_i, i = 1, \dots, k}{\operatorname{argmin}} & \sum_{i=1}^k \eta \|\mathbf{X}_i - \mathbf{U}_i \mathbf{S}_J - \mathbf{W}_i \mathbf{S}_i\|_F^2 \\ & + (1 - \eta) \|\mathbf{y} - \theta_1 \mathbf{S}_J - \sum_{i=1}^k \theta_2 \mathbf{S}_i\|_2^2, \end{aligned} \quad (1)$$

where η is a weight parameter between 0 and 1 to signify the relative importance of \mathbf{X} and \mathbf{y} when calculating the loss. The joint structure across all k datasets is represented by \mathbf{U}_i

$i = i, \dots, k$ and \mathbf{S}_J . Unique structure from the i th dataset is accounted for by \mathbf{W}_i and \mathbf{S}_i . Low-rank approximations are used for both the joint and individual components such that \mathbf{U}_i is $p_i \times r_J$ and \mathbf{S}_J is $r_J \times n$ for the joint components, and \mathbf{W}_i is $p_i \times r_i$ and \mathbf{S}_i is $r_i \times n$ for the i th individual component. Section 2.7 will discuss how ranks are selected. The second expression in (1) allows for the joint and individual scores to linearly predict the outcome, \mathbf{y} . The coefficients, $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_{2i}$ $i = 1, \dots, k$, are vectors with lengths equal to the joint and individual ranks.

2.4. Identifiability

Consider the sJIVE approximation without error,

$$\mathbf{X}_i = \mathbf{J}_i + \mathbf{A}_i \text{ for } i = 1, \dots, k \text{ and } \mathbf{y} = \mathbf{j}_y + \mathbf{a}_y, \quad (2)$$

where $\mathbf{J}_i = \mathbf{U}_i \mathbf{S}_J$, $\mathbf{A}_i = \mathbf{W}_i \mathbf{S}_i$, $\mathbf{j}_y = \boldsymbol{\theta}_1 \mathbf{S}_J$ and $\mathbf{a}_y = \sum_{i=1}^k \boldsymbol{\theta}_{2i} \mathbf{S}_i$. Let $\mathbf{X} = [\mathbf{X}_1^T \dots \mathbf{X}_k^T]^T$, $\mathbf{J} = [\mathbf{J}_1^T \dots \mathbf{J}_k^T]^T$, and $\mathbf{A} = [\mathbf{A}_1^T \dots \mathbf{A}_k^T]^T$. Theorem 1 describes conditions for the identifiability of the terms in (2).

Theorem 1. Consider $\{\mathbf{X}_1, \dots, \mathbf{X}_k, \mathbf{y}\}$ where $\mathbf{y} \in \text{row}(\mathbf{X})$. There exists a uniquely defined decomposition (2) satisfying the following conditions:

1. $\text{row}(\mathbf{J}_i) = \text{row}(\mathbf{J}) \subset \text{row}(\mathbf{X}_i)$ for $i = 1, \dots, k$,
2. $\text{row}(\mathbf{J}) \perp \text{row}(\mathbf{A}_i)$ for $i = 1, \dots, k$,
3. $\cap_{i=1}^k \text{row}(\mathbf{A}_i) = \{0\}$,
4. $\mathbf{j}_y \in \text{row}(\mathbf{J})$ and $\mathbf{a}_y \in \text{row}(\mathbf{A})$.

A proof of Theorem 1 can be found in Section 1 of the supplemental material. The first three conditions are equivalent to those for the JIVE decomposition (see Lemma 1 for angle-based JIVE [28]), which similarly enforces orthogonality between the joint and individual structures (condition 2.). Condition (4.) implies that predictions for the outcome can be uniquely decomposed into the contributions of joint and individual structure. If the row spaces of each \mathbf{A}_i are linearly independent ($\text{rank}(\mathbf{A}) = \sum_{i=1}^k r_i$), then the contribution of individual structure for each data source, $\boldsymbol{\theta}_{2i} \mathbf{S}_i$, are also uniquely determined. We identify more granular terms by the SVD of joint or individual structures: $[\mathbf{U}_1^T \dots \mathbf{U}_k^T \boldsymbol{\theta}_1^T]^T$ are given by the normalized left singular vectors of $[\mathbf{J}^T \mathbf{j}_y^T]^T$ and $[\mathbf{W}_i^T \boldsymbol{\theta}_{2i}^T]^T$ are the normalized left singular vectors of $[\mathbf{A}_i^T \boldsymbol{\theta}_{2i} \mathbf{S}_i]$ for $i \in 1, \dots, k$.

2.5. Estimation

Our optimization function assumes the weighting parameter η is fixed, but generally its value is not predetermined for a given application. In practice, we recommend running sJIVE with 5-fold cross validation (CV) for a range of possible η values, choosing the one with the lowest mean squared error (MSE) when predicting the outcome for the test set. Our accompanying R software considers a grid of η values between 0.01 and 0.99 and can be

run in parallel to decrease computation time needed to tune η . When $\eta = 1$, our objective reduces to that of JIVE and will produce identical results. When η shifts away from 1 and closer to 0, the results from sJIVE will start to deviate from JIVE in order to minimize the squared error of the y expression. This allows the joint and individual components to be selected based on their association with the outcome. When $\eta = 0$ the model is not estimable, because without consideration of the X_j the joint and individual components can always fit y with no error.

By distributing η throughout the expression, let $\tilde{X}_i = \eta^{1/2}X_i$, $\tilde{U}_i = \eta^{1/2}U_i$, $\tilde{W}_i = \eta^{1/2}W_i$, $\tilde{y} = (1 - \eta)^{1/2}y$, $\tilde{\theta}_1 = (1 - \eta)^{1/2}\theta_1$ and $\tilde{\theta}_{2i} = (1 - \eta)^{1/2}\theta_{2i} \forall i \in 1, \dots, k$. Furthermore, let $\tilde{W}S$ denote $\left[(\tilde{W}_1 S_1)^T \dots (\tilde{W}_k S_k)^T \right]^T$. Then, we can minimize equation (1) as described in Algorithm 1.

Algorithm 1: Pseudocode for sJIVE estimation once η, r_j , and r_i is known $\forall i \in 1, \dots, k$

```

1. Initialize  $\tilde{U}_i, \tilde{W}_i, S_i, S_j, \tilde{\theta}_1, \tilde{\theta}_{2i} \forall i \in 1, \dots, k$ 
2. Loop until convergence
(a)  $S_j = \left[ \begin{matrix} \tilde{U}_j^T \\ \tilde{\theta}_1^T \end{matrix} \right] \left( \begin{bmatrix} \tilde{X} \\ \tilde{y} \end{bmatrix} - \begin{bmatrix} \tilde{W}S \\ \sum_{i=1}^k \tilde{\theta}_{2i} S_i \end{bmatrix} \right)$ 
(b)  $\left[ \tilde{U}^T \tilde{\theta}_1^T \right]^T$  = the first  $r_j$  left singular vectors from  $\text{SVD} \left( \begin{bmatrix} \tilde{X} \\ \tilde{y} \end{bmatrix} - \begin{bmatrix} \tilde{W}S \\ \sum_{i=1}^k \tilde{\theta}_{2i} S_i \end{bmatrix} \right)$ 
(c) for each dataset  $i = 1, \dots, k$ 
    i.  $S_i = \left[ \begin{matrix} \tilde{W}_i \\ \tilde{\theta}_{2i} S_i \end{matrix} \right] \left( \begin{bmatrix} \tilde{X} \\ \tilde{y} - \sum_{j \neq i} \tilde{\theta}_{2j} S_j \end{bmatrix} - \begin{bmatrix} \tilde{U}_i \\ \tilde{\theta}_1 \end{bmatrix} S_j \right) P_j^C$ 
    ii.  $\left[ \tilde{W}_i^T \tilde{\theta}_{2i}^T \right]^T$  = the first  $r_i$  left singular vectors from  $\text{SVD} \left( \begin{bmatrix} \tilde{X} \\ \tilde{y} - \sum_{j \neq i} \tilde{\theta}_{2j} S_j \end{bmatrix} - \begin{bmatrix} \tilde{U}_i \\ \tilde{\theta}_1 \end{bmatrix} S_j \right) P_j^C$ 
(d) Check for convergence of  $\left\| \begin{bmatrix} \tilde{X} \\ \tilde{y} \end{bmatrix} - \begin{bmatrix} \tilde{U} \\ \tilde{\theta}_1 \end{bmatrix} S_j - \begin{bmatrix} \tilde{W}S \\ \sum_{i=1}^k \tilde{\theta}_{2i} S_i \end{bmatrix} \right\|_F^2$ 

```

The algorithm will converge to a coordinate-wise optimum, and the solution is robust to initialization. By default, the joint parameters are initialized by taking the SVD of the concatenated data matrix, $\left[\tilde{X}^T \tilde{y}^T \right]^T$ and setting $\left[\tilde{U}^T \tilde{\theta}_1^T \right]^T$ to be equal to the first r_j left eigenvectors. S_j is set to the product of the first r_j eigenvalues and right eigenvectors. Similarly, the individual parameters for the i th dataset ($\tilde{\theta}_{2i}$, \tilde{W}_i , and S_i) are initialized by taking the SVD of the individual matrix $\left[\tilde{X}_i^T \tilde{y}_i^T \right]^T$ and then projecting that result onto the orthogonal complement of the joint space.

After initializing the parameters, our model iteratively solves for the joint and individual components until convergence of our optimization function. In step 2c of Algorithm 1, the individual components are projected onto the orthogonal complement of the joint subspace, P_j^C , to retain orthogonality. After the algorithm converges, additional scaling is needed for the results to be identifiable. The joint loadings and regression coefficients are scaled such that $\left[U_1^T \dots U_k^T \theta_1^T \right]^T$ has a squared Frobenius norm of 1. The joint scores, S_j absorb this scaling as to not change the overall joint effect, J , where $J = \left[U_1^T \dots U_k^T \theta_1^T \right]^T S_j$. The same scaling can be done for each individual effect such that $\left[W_i^T \theta_{2i}^T \right]^T \forall i \in 1, \dots, k$ have a squared Frobenius norm of 1 and S_i absorbs this scaling for the i th dataset $i = 1, \dots, k$.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

2.6. sJIVE-prediction for a Test Set

After running sJIVE, we may want to predict new outcomes with external data or with a test set. Let m be the number of out-of-sample observations, and let us have complete data in \mathbf{X}_i^* , $i = 1, \dots, k$, $i = 1, \dots, k$ for each of the m samples. By extracting $\widehat{\mathbf{U}}_i$, $\widehat{\mathbf{W}}_i$, $\widehat{\theta}_1$, and $\widehat{\theta}_{2i} \forall i = 1, \dots, k$ from the fitted sJIVE model and treating them as fixed, we solve the following optimization function to obtain joint and individual scores for the new data,

$$\underset{\mathbf{S}_J, \mathbf{S}_i, i=1, \dots, k}{\operatorname{argmin}} \sum_{i=1}^k \|\mathbf{X}_i^* - \widehat{\mathbf{U}}_i \mathbf{S}_J - \widehat{\mathbf{W}}_i \mathbf{S}_i\|_F^2, \quad (3)$$

where \mathbf{S}_J is $r_J \times m$, and \mathbf{S}_i is $r_i \times m$, $i = 1, \dots, k$. We can iteratively solve for \mathbf{S}_J and \mathbf{S}_i , $i = 1, \dots, k$ using the closed-form solutions,

$$\widehat{\mathbf{S}}_J = \begin{bmatrix} \widehat{\mathbf{U}}_1 \\ \vdots \\ \widehat{\mathbf{U}}_k \end{bmatrix}^T \begin{bmatrix} \mathbf{X}_1^* \\ \vdots \\ \mathbf{X}_k^* \end{bmatrix} - \begin{bmatrix} \widehat{\mathbf{W}}_1 \widehat{\mathbf{S}}_1 \\ \vdots \\ \widehat{\mathbf{W}}_k \widehat{\mathbf{S}}_k \end{bmatrix}, \text{ and} \\ \widehat{\mathbf{S}}_i = \widehat{\mathbf{W}}_i^T (\mathbf{X}_i^* - \widehat{\mathbf{U}}_i \widehat{\mathbf{S}}_J).$$

Using the newly-obtained scores, the fitted outcomes can be estimated by $\widehat{\mathbf{y}}^* = \widehat{\theta}_1 \widehat{\mathbf{S}}_J + \sum_{i=1}^k \widehat{\theta}_{2i} \widehat{\mathbf{S}}_i$ using $\widehat{\theta}_1$ and $\widehat{\theta}_{2i} \forall i = 1, \dots, k$ from the original sJIVE model.

2.7. Rank Selection

Choosing an appropriate reduced rank for the joint and individual components is necessary for optimal model performance. By default, sJIVE selects ranks via a forward selection 5-fold cross validation (CV) approach. The ranks are iteratively added in order to minimize the average test MSE for \mathbf{y} . Once adding an additional rank fails to lower the MSE, the function stops and the ranks are recorded. For more details, see Section 2 of the supplemental material.

In contrast to this approach, JIVE uses a permutation approach to select ranks. For both approaches, the joint rank must be $\min(n, p_1, \dots, p_k)$ and the individual rank for dataset i must be $\min(n, p_i) \forall i = 1, \dots, k$. In section 3.4, we will compare the accuracy of sJIVE's CV approach to JIVE's permutation approach.

3. Compare sJIVE to JIVE-predict

3.1. Simulation Set-up

Since our work is an extension of JIVE and JIVE-predict, we will first assess how these models compare to supervised JIVE. Datasets were simulated to reflect the sJIVE framework such that the true joint and individual components are known.

Data were simulated as follows: Predictors \mathbf{X}_1 and \mathbf{X}_2 , with $p_1 = p_2 = 200$ rows and $n = 200$ columns, were generated such that the true reduced joint and individual ranks are all equal to 1. Both the joint and individual components contribute equally to the total amount

of variation in \mathbf{X} with varying levels of noise. Similarly, both the joint and individual scores contain equal signal to the composition of \mathbf{y} . We assessed the performance of each model by the MSE of an independent test set of $n = 200$. Each scenario was simulated 100 times. For additional details about generating datasets, see Section 3 of the supplemental material.

For Sections 3.2 and 3.3 we use the true ranks, and we compare the accuracy of rank selection approaches for sJIVE and JIVE in Section 3.4. We varied the amount of noise in \mathbf{X} and in \mathbf{y} across simulations. For example, when \mathbf{X} error is 10%, 10% of the variation in \mathbf{X} is due to error, and the joint and individual components can account for 90% of the total variation in \mathbf{X} . Each test MSE is the average of the 100 simulations, and the final column records how often sJIVE outperformed JIVE-predict.

3.2. Comparing Test MSE

The results displayed in Table 1 show that when there is a relatively small amount of noise in \mathbf{X} or a relatively large noise in \mathbf{y} , sJIVE tends to perform similar to JIVE. However, when the amount of noise in \mathbf{X} is between 30% and 99% with low error in \mathbf{y} , sJIVE consistently outperforms JIVE-predict. In these cases, the large amount of noise in \mathbf{X} makes it difficult for JIVE to capture the signal. sJIVE incorporates \mathbf{y} when determining the joint and individual components, which helps uncover these signals when there is error in \mathbf{X} .

3.3. Recovering True Structures

Similarly, we tested how well sJIVE and JIVE were able to reconstruct the true joint and individual components. Accuracy of each component was summarized by the standardized squared Frobenius norm difference, e.g., $\|\hat{\mathbf{J}} - \mathbf{J}\|_F^2 / \|\mathbf{J}\|_F^2$ where $\hat{\mathbf{J}}$ is the estimated joint component and \mathbf{J} is the true joint component. A similar measure can be obtained for the accuracy of each individual component. When comparing the two methods to each other, sJIVE tended to identify the individual components more accurately than JIVE for all levels of error in \mathbf{X} and \mathbf{y} (Table 2).

The largest eigenvalue of the signal in \mathbf{X} drops below that of the noise when the error in \mathbf{X} is above 97%. When this occurs, statistical models tend to struggle at identifying the signal; rather, they capture a mix of the true signal and the noise. However, when \mathbf{X} error was 99%, sJIVE continued to perform mildly better in terms of test MSE and recovering the joint structure compared to JIVE-predict, especially when error in \mathbf{y} was low, suggesting that sJIVE can more effectively separate noise from signal in \mathbf{X} when large amounts of noise are present.

3.4. Comparing Rank Selection

In addition to the simulations with known rank, we also compared the rank selection techniques of sJIVE and JIVE. JIVE uses a permutation approach while sJIVE uses a forward-selection CV method. For sJIVE's CV approach, all ranks are initially set to zero. Ranks are added if the additional rank results in the largest reduction in test MSE after running 5-fold CV. We simulated all combinations of joint and individual ranks of size 1 to 4 with defaults $n = 100$, $k = 2$, $p_1 = p_2 = 100$ for each dataset, joint and individual components have equal weight, 50% error in \mathbf{X} , and 10% error in \mathbf{y} . Table 3 shows the percent of the time

when sJIVE and JIVE were able to accurately specify the ranks. When the true joint rank was equal to 1, JIVE's permutation approach tended to perform equal or better than sJIVE's CV approach. However, when the true joint rank was greater than 1, JIVE rarely identified the true joint rank while sJIVE selected the true joint rank about 20% of the time. In total, sJIVE correctly identified all 3 ranks 2% of the time, while JIVE had a 10% chance.

In Table 4, we further look at the probability of over- and under-estimating the ranks. JIVE's permutation method regularly underestimated the joint rank (80.5%) and overestimated the individual ranks (58.5% and 50.2% for datasets X_1 and X_2 respectively). Additionally, JIVE correctly specified the joint rank about 20% of the time across all simulations and 40% of the time for each individual rank. In contrast, the sJIVE's probability of overestimating, underestimating, or correctly specifying the joint rank was similar to that of the individual ranks. sJIVE was twice as likely to underestimate each of the ranks compared to overestimate, and had about a 25% chance of specifying any rank correctly. We also found that on average, sJIVE underestimated the joint rank by 0.6 (sd=1.7) and underestimated the individual rank by 0.5 (sd=1.8). In comparison, JIVE underestimated the joint rank by 1.6 (sd=1.2) and overestimated the individual rank by 0.9 (sd=1.2).

4. Compare sJIVE to Other Methods

Next, we compared sJIVE to not only JIVE-predict, but also to concatenated PCA (i.e., PCA of the concatenated data), individual PCA (i.e., PCA of each dataset individually), and CVR [9]. The data was generated in the same manner as earlier simulations. Our default parameters included $k = 2$ datasets, all ranks set to 1, and equal contribution of the joint and individual components. Additionally, we set X error to 90%, y error to 1%, $n = 200$, and $p_1 = p_2 = 200$. We tested the following 10 scenarios: (1.) the listed default parameters, (2.) increasing the number of predictors in each dataset, p , to 500, (3.) increasing the error in X to 99%, (4.) increasing the number of datasets, k , to 4, (5.) increasing the joint signal to be 20 times that of the individual signals, (6.) increasing the individual signals to be 20 times greater than the joint signal, (7.) increasing all ranks to be 10, and (8.) increasing all ranks to be 10, but only allowing the first rank to be predictive of y . For scenarios (9.) and (10.), we set the true ranks to 2 where where for each term the first component has 2 times the strength compared to the second component, where strength is defined by its singular value. Only the second, smaller signal predicts y . In (9.) we set all methods to underestimate the true ranks by setting ranks equal to 1, and (10.) used the correct model ranks by setting model ranks equal to 2. Though CVR can be extended to $k > 2$ datasets, its R package [29] only works for exactly 2 datasets. Thus, we could not obtain a test MSE for CVR when k was greater than 2. Each scenario was run 100 times and the percent of time when sJIVE performed the best was recorded.

The results can be found in Table 5. sJIVE outperformed all other methods at least 70% of the time for 8 of the 10 scenarios, with the exceptions of when a large joint weight or large individual weight are present. For a large joint weight, concatenated PCA marginally outperformed the other methods. Concatenated PCA only looks at the joint signal, ignoring the individual signals, which allows for better predictive accuracy in the presence of large joint effects. When there was a large individual signal, sJIVE and JIVE-predict performed

equally well. Setting the error in X to 99% resulted in better predictive accuracy in sJIVE compared each of the other methods, but sJIVE only saw a modest increase in predictive accuracy in the first two scenarios.

sJIVE performed the best compared to other methods in scenario (9.) when there existed a large signal in X that did not predict y and ranks were assumed to be 1. Unsurprisingly, the JIVE-predict and PCA tended to isolate the major component since these methods aim to find the largest source of variation regardless of its association with an outcome. This resulted in poor performance since the major component does not predict y (JIVE-predict MSE=0.830; concatenated PCA MSE=0.828). sJIVE performed better because its 1-step approach can find minor components with a strong association to the outcome (MSE=0.406). In scenario (10.) where sJIVE and JIVE-predict are run with ranks=2, sJIVE still performs consistently better than JIVE-predict, though the gain in MSE is smaller (sJIVE MSE=0.384; JIVE-predict MSE=0.464).

5. Application to COPDGene Data

Genetic Epidemiology of COPD (COPDGene) is a multi-center longitudinal observational study to identify genetic factors associated with chronic obstructive pulmonary disease (COPD). To attain this goal, multiple views of data were collected on the same group of individuals including proteomic data at baseline [30, 31] and RNA sequencing (RNAseq) data at their 5 year follow-up [32]. Additionally, spirometry was performed at baseline to calculate a percent predicted forced expiratory volume in 1 second (FEV1% predicted, or FEV1pp) value, a measure of pulmonary function which is significantly lower in COPD patients compared to their healthy counterparts.

RNAseq data was not available at baseline. However, it is reasonable to expect some joint variation between the baseline proteomic and follow-up RNA-Seq data, due to shared effects or associations that persist over time. Similarly, we would expect association between baseline FEV1pp and follow-up RNA-Seq data if the genomic determinants of lung function are relatively static. We have complete data on 359 participants for 21,669 RNAseq targets, 1,318 proteins, and a value for FEV1pp. We compare sJIVE to JIVE-predict, concatenated PCA, individual PCA, and CVR. The data were split into 2/3 training set and 1/3 test set with accuracy measured by test MSE, and ranks were determined by JIVE's permutation approach.

The results for each model can be found in Table 6. Of the 6 methods tested, sJIVE had the lowest test MSE at 0.6980. This implies that over 30% of the variation in FEV1pp can be explained by the gene expression and proteomic profile. This result is substantially better than each of the PCA methods and CVR, though JIVE-predict had a similar test MSE of 0.6991. In terms of computation time on a 2.4 GHz computer with 8 GB RAM, sJIVE took 13.1 minutes to run, which is almost a third of the time needed to run JIVE-predict, and substantially quicker than CVR. However, the run time for sJIVE does not include the time needed to estimate the joint and individual ranks. In high-dimensional settings, JIVE and sJIVE map their data to smaller dimensions to increase computational efficiency and the tuning parameter for sJIVE was estimated in parallel using 7 cores. This allows for both

methods to handle larger datasets. See Section 4 of the supplemental material for more details on this *a priori* data compression. Overall, sJIVE results in only a modest increase in accuracy compared to JIVE-predict for the COPDGene data, but a substantial increase in accuracy compared to CVR and PCA approaches.

In addition to testing the predictive accuracy of sJIVE, we further investigated the fit of the estimated model. The underlying data structures are graphically displayed as heatmaps in Figure 1, with strong positive and negative values in red and blue, respectively. The column ordering follows that of the FEV1pp values, and associations are apparent in the joint structure and individual proteomic structure. To further assess the predictive model, Figure 2 compares the true FEV1pp values to the estimated ones, showing that a linear fit is reasonable and that there is no systematic over- or underestimation of FEV1pp.

The results of the predicted model are further summarized in Table 7 to assess the effect of the joint and individual components on FEV1% predict, and we use an F-test to assess the significance of each component in the multivariate model. JIVE's permutation method selected a joint rank of 1, and individual ranks of 27 and 24 for the RNAseq and proteomic data respectively. The single joint rank accounts for 6.9% of the total variation in FEV1% predicted ($p = 0.001$). As for the individual effects, proteomics account for 33.5% of the variation ($p < 0.001$) and RNAseq accounts for 19.1%. However, the individual effect of RNAseq failed to reach statistical significance ($p = 0.217$). Our proteomic data exhibited a large, significant association with FEV1pp after removing the joint structure while RNAseq failed to attain significance beyond its contribution to the joint component. This suggests that some shared signal between proteomic data and RNAseq may be invariant to temporal changes. However, the lack of significance for the individual structure of RNAseq may be due to the separate timepoint.

To further investigate the proteomic results, we conducted a pathway analysis using the WEB-based GENE SeT AnaLysis Toolkit (WebGestalt) [33]. Using the results from sJIVE, we calculated the meta-loadings for each predictor in a similar manner to that in [26], by taking the sum across the joint and individual loadings weighted by their regression coefficients, i.e. $\theta_1 U_j + \theta_2 W_j$. The meta-loadings for the proteomic dataset can be found in Figure 3. The top 20% of absolute meta-loadings were used to perform an over-representation enrichment analysis with KEGG pathway database and the top 10 pathways are shown in Table 8. The following were found to be statistically significant pathways (all $p < 0.001$): Glucagon signaling pathway, dopaminergic synapse, cholinergic synapse, wnt signaling pathway, B cell receptor signaling pathway, chemokine signaling pathway, AGE-RAGE signaling pathways in diabetic complications, VEGF signaling pathway, circadian entrainment, and insulin signaling pathway. The false discovery rate (FDR) for these 10 pathways remained under 0.005. Three of these pathways have been mechanistically linked to COPD, specifically the emphysema phenotype. RAGE or receptor for advanced glycosylation end product receptor has been identified as both a biomarker and mediator of emphysema [34, 35, 36]. Similarly, VEGF has been mechanistically linked to emphysema and sputum VEGF levels are reciprocally related to the level of COPD [37, 38]. The Wnt signaling pathway is associated with aging and down-regulation of this pathway in human airway epithelium in smokers is associated with smoking and COPD [39, 40].

Though the effect of RNAseq was hindered by its time of collection, sJIVE was able to find significant associations between the proteomic individual signal and the joint signal. Importantly, this approach has the potential to identify and link novel and existing pathways associated with COPD, opening new avenues for research.

6. Summary and Discussion

We have proposed sJIVE as a one-step approach to identify joint and individual components in multi-source data that relate to an outcome, and use those components to create a prediction model for the outcome. This approach facilitates interpretation by identifying concordant and complementary effects among the different sources, and has competitive predictive performance. When comparing sJIVE to a similar two-step approach, JIVE-predict, sJIVE performed well in the presence of large amounts of error in the multi-source data X when the error in y was relatively small. sJIVE performed best when there was a large signal in X that did not predict y and the ranks were underestimated. Even in scenarios when the largest eigenvalue of the signal is slightly smaller than that of the noise, sJIVE was able to capture the signal better than JIVE-predict. When comparing sJIVE to principal components and canonical correlation based approaches, sJIVE tended to perform best in almost all scenarios tested.

When applying our method to the COPDGene data, sJIVE and JIVE-predict also outperformed the other methods. Though sJIVE resulted in the lowest test MSE, the gain in accuracy between the two methods was marginal in this case. sJIVE found significant associations between the proteomic data and FEV1% predicted, as well as in the joint effect of proteomic and RNAseq data. After conducting a pathway analysis of the proteomic results, we uncovered 3 pathways that had previously been mechanistically linked to COPD, as well as additional pathways that could benefit by future research. While we consider predicting FEV1pp at baseline, the ongoing COPDGene study has collected FEV1pp and other outcomes at follow-up timepoints, in addition to other “omics” data, and more comprehensive longitudinal analyses are worth pursuing.

Our method has some limitations and avenues for future work. While our simulations demonstrate its good performance when the correct ranks are selected, the correct ranks may not be selected in practice. Depending on the application, JIVE’s permutation method or sJIVE’s CV method for rank selection may be preferred, but neither are ideal. Future research can explore and compare alternative approaches to determine the ranks. Other methods of scaling, such as scaling each X_j to have a Frobenius norm of 1, have been used, which is different from our method that scales each predictor to have variance 1. Moreover, sJIVE treats all predictors and the outcome as continuous. A useful extension to our method would be to allow for a binary outcome, or other distributional forms. Additionally, our method does not allow for missing data, so missing values must be imputed or the entire observation must be removed. Lastly, sJIVE does not explicitly capture signal that is shared between some, but not all, data sources. Extensions that allow for partially-shared structure, such as in the SLIDE method [20], would allow for more flexibility.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

The views expressed in this article are those of the authors and do not reflect the views of the United States Government, the Department of Veterans Affairs, the funders, the sponsors, or any of the authors' affiliated academic institutions.

Funding

This work was partially supported by grants R01-GM130622 and 1R35GM142695-01 from the National Institutes of Health and by Award Number U01 HL089897 and Award Number U01 HL089856 from the National Heart, Lung, and Blood Institute. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Heart, Lung, and Blood Institute or the National Institutes of Health.

COPDGene is also supported by the COPD Foundation through contributions made to an Industry Advisory Board comprised of AstraZeneca, Boehringer-Ingelheim, Genentech, GlaxoSmithKline, Novartis, Pfizer, Siemens, and Sunovion.

References

- [1]. Regan EA, Hokanson JE, Murphy JR, Make B, Lynch DA, Beaty TH, Curran-Everett D, Silverman EK, Crapo JD, Genetic epidemiology of COPD (COPDGene) study design, COPD: Journal of Chronic Obstructive Pulmonary Disease 7 (2011) 32–43. doi: 10.3109/15412550903499522.
- [2]. Xu M, Zhu Z, Zhang X, Zhao Y, Li X, Canonical correlation analysis with $l_{2,1}$ -norm for multiview data representation., IEEE Transactions on Cybernetics (2019). doi:10.1109/TCYB.2019.2904753.
- [3]. Guo Y, Ding X, Liu C, Xue J-H, Sufficient canonical correlation analysis., IEEE Trans Image Process 25 (6) (2016) 2610–2619. doi: 10.1109/TIP.2016.2551374. [PubMed: 27071172]
- [4]. Gossman A, Zille P, Calhoun V, Wang Y-P, FDR-corrected sparse canonical correlation analysis with applications to imaging genomics., IEEE transactions on medical imaging 37 (8) (2018) 1761–1774. doi:10.1109/TMI.2018.2815583. [PubMed: 29993802]
- [5]. Wilms I, Croux C, Robust sparse canonical correlation analysis., BMC systems biology 10 (1) (2016) 72. doi:10.1186/s12918-016-0317-9. [PubMed: 27516087]
- [6]. Hotelling H, Relations between two sets of variates, Biometrika 28 (3–4) (1936) 321–377. doi:10.1093/biomet/28.3-4.321.
- [7]. Rodosthenous T, Shahrezaei V, Evangelou M, Integrating multi-OMICS data through sparse canonical correlation analysis for the prediction of complex traits: A comparison study, Bioinformatics (2020). doi:10.1093/bioinformatics/btaa530.
- [8]. Gross SM, Tibshirani R, Collaborative regression, Biostatistics 16 (2015) 326–338. arXiv:1401.5823, doi:10.1093/biostatistics/kxu047. [PubMed: 25406332]
- [9]. Luo C, Liu J, Dey DK, Chen K, Canonical variate regression., Biostatistics 17 (3) (2016) 468–483. doi:10.1093/biostatistics/kxw001. [PubMed: 26861909]
- [10]. Witten D, Tibshirani R, Extensions of sparse canonical correlation analysis with applications to genomic data, Statistical Applications in Genetics and Molecular Biology 8 (2009) Article28. doi:10.2202/1544-6115.1470.
- [11]. Wilms I, Croux C, Sparse canonical correlation analysis from a predictive point of view, Biometrical Journal 57 (2015) 834–851. doi:10.1002/bimj.201400226. [PubMed: 26147637]
- [12]. Zhang Y, Gaynanova I, Joint association and classification analysis of multi-view data, Biometric Methodology (2021). doi:10.1111/biom.13536.
- [13]. Safo SE, Min EJ, Haine L, Sparse linear discriminant analysis for multi-view structured data, Biometric Methodology (2021). doi: 10.1111/biom.13458.

- [14]. Hotelling H, Analysis of a complex of statistical variables into principal components, *Journal of Educational Psychology* 24 (1933) 417–441. doi:10.1037/h0071325.
- [15]. Bair E, Hastie T, Paul D, Tibshirani R, Prediction by supervised principal components, *Journal of the American Statistical Association* 101 (473) (2006) 119–137. doi:10.1198/016214505000000628.
- [16]. Schoutenden M, Van Deun K, Wilderjans TF, Iven VM, Performing DISCO-SCA to search for distinctive and common information in linked data, *Behavior Research Methods* 46 (2) (2014) 576–87. doi:10.3758/s13428-013-0374-6. [PubMed: 24178130]
- [17]. Argelaguet R, Velten B, Arnol D, Dietrich S, Zenz T, Marioni JC, Buettner F, Huber W, Stegle O, Multi-omics factor analysis - a framework for unsupervised integration of multi-omics data sets, *Molecular System of Biology* 14 (e8124) (5 2018). doi:10.15252/msb.20178124.
- [18]. Shou G, Cichocki A, Zhang Y, Mandic D, Group component analysis for multiblock data: Common and individual feature extraction, *IEEE Transactions on Neural Networks and Learning Systems* 27 (11) (2015) 2426–2439. doi:10.1109/TNNLS.2015.2487364. [PubMed: 26529787]
- [19]. Zhu H, Li G, Lock EF, Generalized integrative principal component analysis for multi-type data with block-wise missing structure, *Biostatistics* 21 (2) (2020) 302–318. doi:10.1093/biostatistics/kxy052. [PubMed: 30247540]
- [20]. Gaynanova I, Li G, Structural learning and integrative decomposition of multi-view data, *Biometrics* 75 (4) (2019) 1121–1132. doi: 10.1111/biom.13108. [PubMed: 31254385]
- [21]. Lock EF, Hoadley KA, Marron JS, Nobel AB, Joint and individual variation explained (JIVE) for integrated analysis of multiple data types., *The annals of applied statistics* 7 (1) (2013) 523–542. doi:10.1214/12-AOAS597. [PubMed: 23745156]
- [22]. O’Connell M, Lock E, R.JIVE for exploration of multi-source molecular data, *Bioinformatics* 32 (2016) 2877–2879. doi:10.1093/bioinformatics/btw324. [PubMed: 27273669]
- [23]. Li G, Jung S, Incorporating covariates into integrated factor analysis of multi-view data, *Biometrics* 73 (1) (2017) 1433–1442. doi: 10.1111/biom.12698. [PubMed: 28407218]
- [24]. Chekouo T, Safo SE, Bayesian integrative analysis and prediction with application to atherosclerosis cardiovascular disease, *arXiv preprint arXiv:2005.11586* (2020).
- [25]. Kaplan A, Lock EF, Prediction with dimension reduction of multiple molecular data sources for patient survival., *Cancer informatics* 16 (2017) 1–11. doi:10.1177/1176935117718517. [PubMed: 28096648]
- [26]. Sandri BJ, Kaplan A, Hodgson SW, Peterson M, Avdulov S, Higgins L, Markowski T, Yang P, Limper AH, Griffin TJ, Bitterman P, Lock EF, Wendt CH, Multi-omic molecular profiling of lung cancer in COPD., *The European respiratory journal* 52 (1) (2018). doi: 10.1183/13993003.02665-2017.
- [27]. Zhao Y, Klein A, Castellanos FX, Milham MP, Brain age prediction: Cortical and subcortical shape covariation in the developing human brain, *NeuroImage* 202 (2019) 116149. doi:10.1016/j.neuroimage.2019.116149.
- [28]. Feng Q, Jiang M, Hannig J, Marron J, Angle-based joint and individual variation explained, *Journal of multivariate analysis* 166 (2018) 241–265. doi:10.1016/j.jmva.2018.03.008.
- [29]. Luo C, Chen K, CVR: Canonical Variate Regression, *r* package version 0.1.1 (2017).
- [30]. Raffield LM, Dang H, Pratte KA, Jacobson S, Gillenwater LA, Ampleford E, Barjaktarevic I, Basta P, Clish CB, Comellas AP, Cornell E, Curtis JL, Doerschuk C, Durda P, Emson C, Freeman CM, Guo X, Hastie AT, Hawkins GA, Herrera J, Johnson WC, Labaki WW, Comparison of proteomic assessment methods in multiple cohort studies, *PROTEOMICS* 20 (12) (2020) 1900278. arXiv: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/pmic.201900278>, doi:10.1002/pmic.201900278.
- [31]. Mastej E, Gillenwater L, Zhuang Y, Pratte KA, Bowler RP, Kechris K, Identifying protein-metabolite networks associated with copd phenotypes, *Metabolites* 10 (4) (2020). doi:10.3390/metabo10040124.
- [32]. Parker MM, Chase RP, Lamb A, Reyes A, Saferali A, Yun JH, Himes BE, Silverman EK, Hersh CP, Castaldi PJ, RNA sequencing identifies novel non-coding RNA and exon-specific effects associated with cigarette smoking, *BMC Medical Genomics* 10 (6) (2017) 58. doi:10.1186/s12920-017-0295-9. [PubMed: 28985737]

- [33]. Wang J, Vasaikar S, Shi Z, Greer M, Zhang B, WebGestalt 2017: a more comprehensive, powerful, flexible and interactive gene set enrichment analysis toolkit, *Nucleic Acids Research* 45 (W1) (2017) W130–W137. arXiv:<https://academic.oup.com/nar/article-pdf/45/W1/W130/18137438/gkx356.pdf>, doi:10.1093/nar/gkx356. [PubMed: 28472511]
- [34]. Carolan BJ, Hughes G, Morrow J, Hersh CP, O'Neal WK, Rennard S, Pillai SG, Belloni P, Cockayne DA, Comellas A. P. e. a., The association of plasma biomarkers with computed tomography-assessed emphysema phenotypes, *Respiratory Research* 15 (1) (2014). doi:10.1186/s12931-014-0127-9.
- [35]. Cheng DT, Kim DK, Cockayne DA, Belousov A, Bitter H, Cho MH, Duvoix A, Edwards LD, Lomas DA, Miller B. E. e. a., Systemic soluble receptor for advanced glycation endproducts is a biomarker of emphysema and associated with AGER genetic variants in patients with chronic obstructive pulmonary disease, *American Journal of Respiratory and Critical Care Medicine* 188 (8) (2013) 948–57. doi:10.1164/rccm.201302-0247OC. [PubMed: 23947473]
- [36]. Kanazawa H, Yoshikawa J, Elevated oxidative stress and reciprocal reduction of vascular endothelial growth factor levels with severity of COPD, *Chest* 128 (5) (2005) 3191–3197. doi:10.1378/chest.128.5.3191. [PubMed: 16304261]
- [37]. Lehmann M, Baarsma HA, Königshoff M, Wnt signaling in lung aging and disease, *Annals of the American Thoracic Society* 13 (Supplement 5) (2016) S411–S416. doi:10.1513/annalsats.201608-586aw. [PubMed: 28005418]
- [38]. Sanders KA, Delker DA, Huecksteadt T, Beck E, Wuren T, Chen Y, Zhang Y, Hazel MW, Hoidal JR, RAGE is a critical mediator of pulmonary oxidative stress, alveolar macrophage activation and emphysema in response to cigarette smoke, *Scientific Reports* 9 (1) (2019) 231. doi:10.1038/s41598-018-36163-z. [PubMed: 30659203]
- [39]. Tang K, Rossiter HB, Wagner PD, Breen EC, Lung-targeted VEGF inactivation leads to an emphysema phenotype in mice, *Journal of Applied Physiology* 97 (4) (2004) 1559–1566. doi:10.1152/jappphysiol.00221.2004. [PubMed: 15208295]
- [40]. Wang R, Ahmed J, Wang G, Hassan I, Strulovici-Barel Y, Hackett NR, Crystal RG, Down-regulation of the canonical wnt β -catenin pathway in the airway epithelium of healthy smokers and smokers with COPD, *PLoS ONE* 6 (4) (2011) e14793. doi:10.1371/journal.pone.0014793.

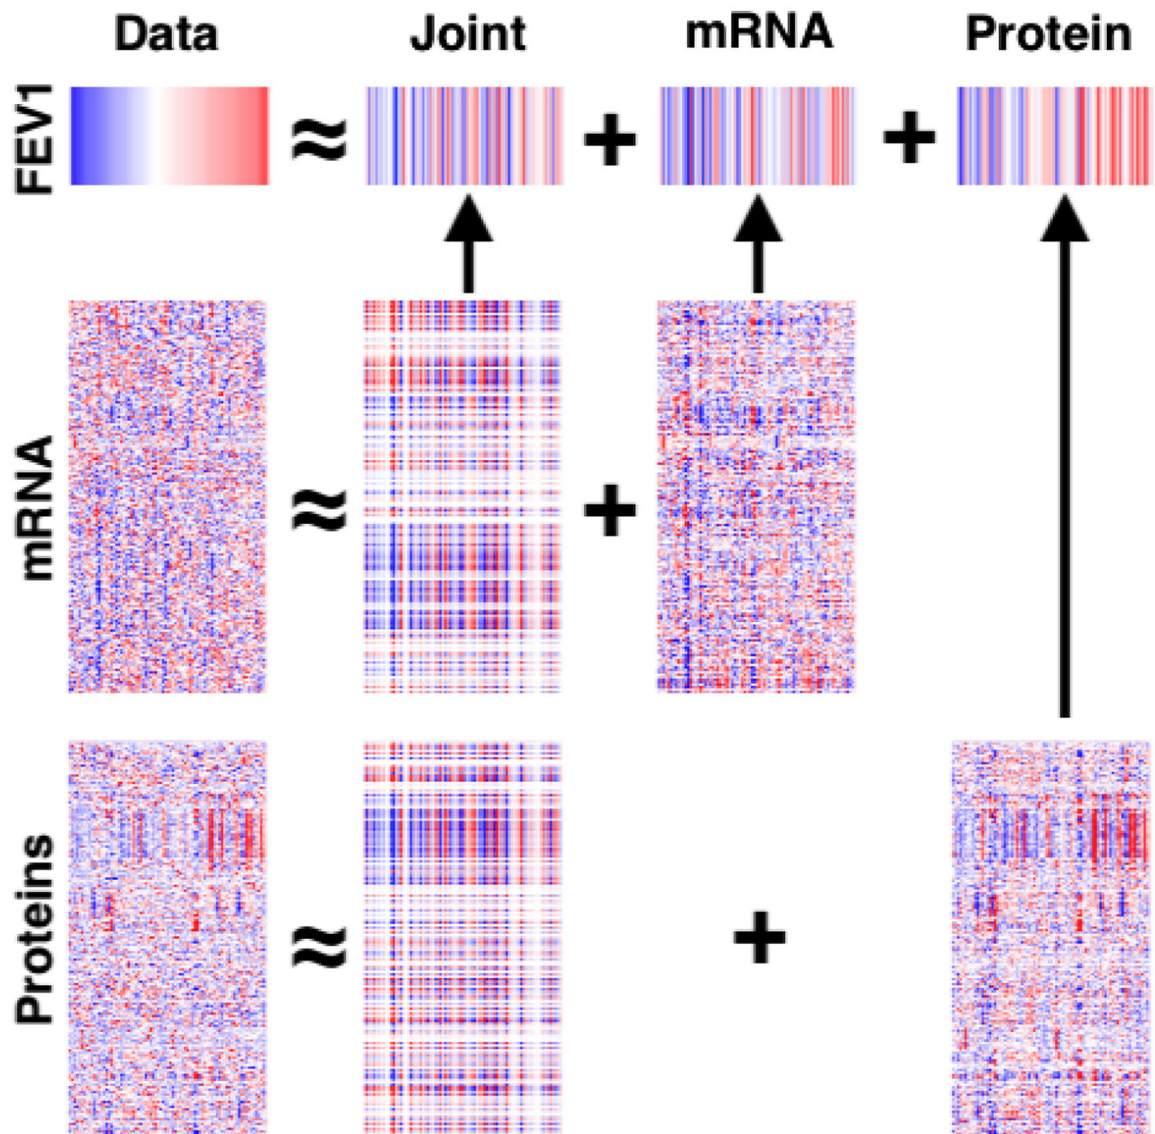


Figure 1: Heatmap of the joint and individual structures, and their contributions to the predictive model for FEV1pp, in the COPDGene training dataset. The same ordering of the sample set (column) is used for each heatmap; red corresponds to higher values and blue corresponds to lower values.




Figure 2:
Comparing true FEV1 percent predicted values to estimated values from sJIVE output.




Figure 3:
Meta-loadings from the sJIVE result for the proteomic dataset.

Table 1:

Comparing Test MSE (SE) of sJIVE, and JIVE-predict. Defaults: $n=200$ for training set, $n = 200$ for test set, $k = 2$, $p = 200$ for each dataset, joint and individual components have equal weight, and all ranks are 1. Each row is the average of 100 simulations

X Error	Y Error	sJIVE MSE	JIVE-predict MSE	% of Time sJIVE wins	Eigenvalue of X signal	Eigenvalue of X error
0.10	0.10	0.112 (0.002)	0.113 (0.002)	67	154.55	8.84
0.10	0.30	0.311 (0.005)	0.311 (0.005)	50	154.09	8.84
0.10	0.50	0.492 (0.006)	0.492 (0.006)	52	154.29	8.83
0.10	0.70	0.718 (0.008)	0.718 (0.008)	53	154.23	8.84
0.10	0.90	0.939 (0.007)	0.933 (0.007)	48	153.94	8.84
0.30	0.10	0.12 (0.002)	0.121 (0.002)	69	136.10	15.32
0.30	0.30	0.331 (0.005)	0.332 (0.005)	59	136.59	15.32
0.30	0.50	0.531 (0.008)	0.531 (0.008)	54	136.06	15.32
0.30	0.70	0.719 (0.008)	0.717 (0.008)	51	136.21	15.31
0.30	0.90	0.933 (0.007)	0.917 (0.006)	48	136.19	15.32
0.50	0.10	0.149 (0.003)	0.154 (0.004)	88	115.12	19.78
0.50	0.30	0.346 (0.005)	0.348 (0.005)	59	115.68	19.77
0.50	0.50	0.618 (0.088)	0.533 (0.007)	68	115.34	19.78
0.50	0.70	0.739 (0.009)	0.731 (0.008)	55	115.19	19.78
0.50	0.90	0.965 (0.028)	0.919 (0.007)	39	115.64	19.73
0.70	0.10	0.206 (0.007)	0.209 (0.006)	75	89.40	23.34
0.70	0.30	0.382 (0.006)	0.383 (0.007)	60	89.48	23.32
0.70	0.50	0.558 (0.008)	0.553 (0.008)	42	89.06	23.35
0.70	0.70	0.751 (0.009)	0.735 (0.007)	47	89.41	23.37
0.70	0.90	0.956 (0.01)	0.925 (0.007)	35	89.04	23.34
0.90	0.10	0.247 (0.007)	0.246 (0.007)	61	51.43	26.42
0.90	0.30	0.424 (0.007)	0.419 (0.006)	51	51.51	26.38
0.90	0.50	0.604 (0.009)	0.597 (0.007)	38	51.44	26.43
0.90	0.70	0.759 (0.007)	0.754 (0.007)	35	51.59	26.38
0.90	0.90	0.95 (0.008)	0.932 (0.006)	42	51.65	26.36
0.95	0.10	0.329 (0.006)	0.335 (0.006)	73	36.49	27.13
0.95	0.30	0.478 (0.007)	0.477 (0.006)	59	36.56	27.15
0.95	0.50	0.63 (0.008)	0.624 (0.007)	52	36.41	27.10
0.95	0.70	0.796 (0.008)	0.781 (0.007)	50	36.60	27.10
0.95	0.90	0.953 (0.008)	0.933 (0.005)	34	36.43	27.14
0.99	0.10	0.765 (0.008)	0.884 (0.006)	95	16.32	27.68
0.99	0.30	0.871 (0.01)	0.91 (0.005)	81	16.33	27.72
0.99	0.50	0.942 (0.007)	0.937 (0.005)	56	16.31	27.62
0.99	0.70	0.989 (0.008)	0.955 (0.003)	49	16.38	27.65
0.99	0.90	1.025 (0.009)	0.985 (0.002)	42	16.31	27.68

Table 2:

Comparing standardized squared Frobenius norm difference between true and estimated components. Defaults: $n=200$ for training set, $n = 200$ for test set, $k = 2$, $p = 200$ for each dataset, joint and individual components have equal weight, and all ranks are 1. Each row is the average of 100 simulations

X Error	Y Error	J	sJIVE A	J	JIVE A	% sJIVE J	beats A1	JIVE A2
0.10	0.10	0.042 (0.002)	0.035 (0.002)	0.041 (0.001)	0.034 (0.002)	43	39	49
0.10	0.30	0.042 (0.002)	0.03 (0.002)	0.042 (0.002)	0.029 (0.002)	49	43	51
0.10	0.50	0.039 (0.002)	0.03 (0.002)	0.039 (0.002)	0.03 (0.002)	58	58	52
0.10	0.70	0.04 (0.002)	0.031 (0.002)	0.04 (0.002)	0.031 (0.002)	64	63	59
0.10	0.90	0.114 (0.026)	0.066 (0.014)	0.042 (0.002)	0.029 (0.002)	56	56	52
0.30	0.10	0.106 (0.003)	0.065 (0.002)	0.12 (0.003)	0.073 (0.002)	97	97	93
0.30	0.30	0.12 (0.003)	0.073 (0.002)	0.124 (0.003)	0.075 (0.002)	80	77	67
0.30	0.50	0.122 (0.003)	0.072 (0.002)	0.125 (0.003)	0.074 (0.002)	86	88	80
0.30	0.70	0.128 (0.003)	0.082 (0.002)	0.129 (0.003)	0.082 (0.003)	67	67	65
0.30	0.90	0.216 (0.029)	0.119 (0.015)	0.124 (0.003)	0.072 (0.002)	72	71	60
0.50	0.10	0.278 (0.008)	0.155 (0.005)	0.329 (0.008)	0.187 (0.005)	100	100	98
0.50	0.30	0.308 (0.009)	0.174 (0.005)	0.333 (0.009)	0.189 (0.005)	93	94	89
0.50	0.50	0.319 (0.012)	0.177 (0.005)	0.319 (0.008)	0.181 (0.005)	91	87	86
0.50	0.70	0.364 (0.02)	0.2 (0.012)	0.341 (0.012)	0.187 (0.007)	85	81	87
0.50	0.90	0.433 (0.029)	0.242 (0.017)	0.322 (0.008)	0.179 (0.005)	75	70	70
0.70	0.10	0.79 (0.03)	0.418 (0.021)	0.962 (0.02)	0.52 (0.014)	92	95	95
0.70	0.30	0.888 (0.025)	0.465 (0.014)	0.974 (0.022)	0.516 (0.013)	91	91	96
0.70	0.50	0.926 (0.023)	0.474 (0.014)	0.952 (0.022)	0.492 (0.014)	86	90	86
0.70	0.70	0.914 (0.024)	0.467 (0.013)	0.924 (0.022)	0.468 (0.012)	85	84	83
0.70	0.90	0.963 (0.021)	0.516 (0.016)	0.957 (0.02)	0.505 (0.014)	76	72	72
0.90	0.10	1.696 (0.017)	0.749 (0.02)	1.719 (0.013)	0.84 (0.014)	59	96	96
0.90	0.30	1.729 (0.014)	0.769 (0.014)	1.747 (0.013)	0.81 (0.014)	62	97	97
0.90	0.50	1.706 (0.018)	0.817 (0.017)	1.711 (0.014)	0.822 (0.015)	58	88	86
0.90	0.70	1.681 (0.013)	0.818 (0.016)	1.695 (0.013)	0.821 (0.015)	63	70	82
0.90	0.90	1.662 (0.016)	0.835 (0.015)	1.685 (0.014)	0.835 (0.015)	70	61	58
0.95	0.10	2.491 (0.025)	1.091 (0.035)	2.523 (0.022)	1.235 (0.02)	63	91	94
0.95	0.30	2.468 (0.024)	1.127 (0.022)	2.5 (0.024)	1.248 (0.021)	66	96	94
0.95	0.50	2.481 (0.027)	1.22 (0.026)	2.504 (0.025)	1.259 (0.025)	69	91	96
0.95	0.70	2.497 (0.027)	1.237 (0.023)	2.544 (0.024)	1.24 (0.022)	67	80	76
0.95	0.90	2.421 (0.032)	1.294 (0.024)	2.491 (0.023)	1.267 (0.021)	70	46	59
0.99	0.10	5.754 (0.208)	6.081 (0.119)	8.752 (0.121)	6.315 (0.086)	88	36	44
0.99	0.30	6.23 (0.232)	6.276 (0.111)	8.755 (0.106)	6.446 (0.085)	87	46	47
0.99	0.50	6.419 (0.246)	6.263 (0.089)	8.78 (0.105)	6.248 (0.078)	91	46	54
0.99	0.70	7.195 (0.226)	6.349 (0.073)	8.643 (0.094)	6.341 (0.076)	84	61	58
0.99	0.90	7.891 (0.195)	6.274 (0.084)	8.74 (0.106)	6.302 (0.079)	62	57	62

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 3:

Comparing rank selection techniques of sJIVE and JIVE. Each row is the average of 10 simulations. Defaults: $n = 100$ $k = 2$, $p = 100$ for each dataset, joint and individual components have equal weight, 50% of variation in X is noise, 10% of variation in Y is noise, and all ranks are 1. Each row is the average of 10 simulations

Rank J	Rank A1	Rank A2	% sJIVE J	% JIVE J	% sJIVE A	% JIVE A
1	1	1	100	90	20	100
1	1	2	70	70	20	70
1	1	3	50	40	0	100
1	1	4	70	80	0	40
1	2	2	50	90	10	70
1	2	3	20	90	0	80
1	2	4	60	50	0	30
1	3	3	20	70	0	60
1	3	4	80	70	0	40
1	4	4	30	50	10	10
2	1	1	50	0	60	70
2	1	2	20	0	10	40
2	1	3	40	0	0	40
2	1	4	0	10	10	0
2	2	2	10	0	10	70
2	2	3	50	0	10	40
2	2	4	0	10	0	0
2	3	3	10	0	0	20
2	3	4	20	0	0	10
2	4	4	30	0	0	10
3	1	1	0	0	20	0
3	1	2	60	0	10	0
3	1	3	30	0	10	0
3	1	4	20	0	0	0
3	2	2	20	0	10	0
3	2	3	30	10	10	0
3	2	4	20	10	0	0
3	3	3	30	0	10	10
3	3	4	20	0	0	30
3	4	4	10	0	0	0
4	1	1	0	0	30	0
4	1	2	20	0	10	0
4	1	3	0	0	10	0
4	1	4	0	0	0	0
4	2	2	0	0	0	0
4	2	3	30	10	0	0
4	2	4	0	0	10	0

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Rank J	Rank A1	Rank A2	% sJIVE J	% JIVE J	% sJIVE A	% JIVE A
4	3	3	10	0	0	0
4	3	4	10	0	0	0
4	4	4	0	0	0	0

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 4:

Comparing rank selection techniques of sJIVE and JIVE. 400 total simulations were run with true ranks ranging from 1 to 4.

	Times Correctly Specified	Times Underestimated	Times Overestimated
sJIVE CV rank Selection			
- J rank	109 (27.3%)	208 (52.0%)	83 (20.8%)
- A1 rank	113 (28.3%)	180 (45.0%)	107 (26.8%)
- A2 rank	105 (26.3%)	197 (49.3%)	98 (24.5%)
JIVE Permutation rank Selection			
- J rank	75 (18.8%)	322 (80.5%)	3 (0.8%)
- A1 rank	156 (39.0%)	10 (2.5%)	234 (58.5%)
- A2 rank	170 (42.5%)	29 (7.2%)	201 (50.2%)

Table 5:

Comparing mean (SE) test MSE of sJIVE to other existing methods. Defaults: $n = 200$ for training and test set, $k = 2$, $p = 200$ for each dataset, joint and individual components have equal weight, 90% of variation in X is noise, 1% of variation in y is noise, and all ranks are 1. Each row is the average of 100 simulations

Scenario	sJIVE	JIVE Predict	Conc. PCA	Indiv. PCA 1	Indiv. PCA 2	CVR	%sJIVE Wins
1. Default	0.231 (0.056)	0.231 (0.053)	0.236 (0.054)	0.487 (0.101)	0.497 (0.114)	0.304 (0.054)	70
2. High Dimensional (p=500)	0.2 (0.046)	0.208 (0.051)	0.214 (0.053)	0.468 (0.109)	0.477 (0.114)	0.25 (0.05)	87
3. Large X error (0.99)	0.534 (0.058)	0.607 (0.052)	0.659 (0.054)	0.766 (0.061)	0.763 (0.06)	0.781 (0.056)	99
4. K=4	0.368 (0.055)	0.373 (0.056)	0.388 (0.055)	0.754 (0.086)	0.743 (0.088)	—	70
5. Large Joint weight	0.126 (0.016)	0.125 (0.016)	0.125 (0.016)	0.147 (0.019)	0.147 (0.016)	0.174 (0.021)	24
6. Large Individual weight	0.148 (0.017)	0.147 (0.018)	0.153 (0.02)	0.573 (0.13)	0.573 (0.122)	0.249 (0.029)	58
7. All ranks=10	0.728 (0.068)	0.763 (0.057)	0.793 (0.052)	0.862 (0.055)	0.867 (0.057)	1.027 (0.086)	76
8. Ranks=10, but only first rank predicts Y	0.765 (0.068)	0.803 (0.055)	0.828 (0.046)	0.893 (0.052)	0.885 (0.054)	1.045 (0.077)	76
9. Ranks=1, major component* does not predict Y	0.406 (0.050)	0.830 (0.144)	0.828 (0.145)	0.766 (0.190)	0.820 (0.165)	0.790 (0.130)	99
10. Ranks=2, major component* does not predict Y	0.384 (0.048)	0.464 (0.059)	0.568 (0.061)	0.695 (0.163)	0.745 (0.159)	0.674 (0.068)	99

* True ranks=2 where the first rank, or major component, has 2 times the signal compared to the second rank. Only the second rank predicts Y.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 6:

Comparing test MSE of sJIVE to other existing methods on the COPDGene data.

Model	Time	Test MSE
sJIVE	13.1 min	0.6980
JIVE-predict	34.6 min	0.6991
Concatenated PCA	4.5 sec	0.7477
Individual PCA 1	4.2 sec	0.7832
Individual PCA 2	0.4 sec	0.7610
CVR	110.9 hrs	0.9805

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 7:

Results of sJIVE model on COPDGene data.

Predictor	Rank	Partial R^2	P-value
Joint Component	1	0.069	0.001
RNAseq	27	0.191	0.217
Proteomics	24	0.335	<0.001

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 8:

Results to pathway analysis of Proteomic dataset.

	Pathway	Number of Genes	Ratio	P-value	FDR
1	Glucagon signaling pathway	18	3.58	4.70×10^{-7}	1.05×10^{-4}
2	Dopaminergic synapse	20	3.23	3.23×10^{-6}	4.36×10^{-4}
3	Cholinergic synapse	20	3.00	3.06×10^{-5}	2.29×10^{-3}
4	Wnt signaling pathway	32	2.45	7.15×10^{-5}	3.27×10^{-3}
5	B cell receptor signaling pathway	24	2.69	8.43×10^{-5}	3.27×10^{-3}
6	Chemokine signaling pathway	74	1.87	1.02×10^{-4}	3.27×10^{-3}
7	AGE-RAGE signaling pathways in diabetic complications	51	2.08	1.05×10^{-4}	3.27×10^{-3}
8	VEGF signaling pathway	33	2.37	1.19×10^{-4}	3.27×10^{-3}
9	Circadian entrainment	10	3.69	1.31×10^{-4}	3.27×10^{-3}
10	Insulin signaling pathway	31	2.38	1.86×10^{-4}	4.16×10^{-3}