



Integrative Factor Regression and Its Inference for Multimodal Data Analysis

Quefeng Li^{a,b} and Lexin Li^{a,b}

^aUniversity of North Carolina, Chapel Hill, NC; ^bDepartment of Biostatistics and Epidemiology, University of California, Berkeley, CA

ABSTRACT

Multimodal data, where different types of data are collected from the same subjects, are fast emerging in a large variety of scientific applications. Factor analysis is commonly used in integrative analysis of multimodal data, and is particularly useful to overcome the curse of high dimensionality and high correlations. However, there is little work on statistical inference for factor analysis-based supervised modeling of multimodal data. In this article, we consider an integrative linear regression model that is built upon the latent factors extracted from multimodal data. We address three important questions: how to infer the significance of one data modality given the other modalities in the model; how to infer the significance of a combination of variables from one modality or across different modalities; and how to quantify the contribution, measured by the goodness of fit, of one data modality given the others. When answering each question, we explicitly characterize both the benefit and the extra cost of factor analysis. Those questions, to our knowledge, have not yet been addressed despite wide use of factor analysis in integrative multimodal analysis, and our proposal bridges an important gap. We study the empirical performance of our methods through simulations, and further illustrate with a multimodal neuroimaging analysis.

KEYWORDS

Data integration; Dimension reduction; Factor analysis; High-dimensional inference; Multimodal neuroimaging; Principal components analysis

1. Introduction

Thanks to rapid technological advances, multiple types of data are now frequently collected for a common set of experimental subjects. Such a new data structure, often referred as multi-view, multi-source or multimodal data, is fast emerging in a wide range of scientific fields. Examples include multi-omics data in genomics, multimodal neuroimaging data in neuroscience, multimodal electronic health records data in health care administration, among others. Numerous empirical studies have found that, by combining diverse but usually complementary information from different types of data, an integrative analysis of multimodal data is often beneficial; see Uludag and Roebroeck (2014), Li, Wu, and Ngom (2016), Richardson, Tseng, and Sun (2016) for reviews and the references therein.

In view of the promise of multimodal data, a number of statistical methods have recently been developed for integrative analysis. An important class of such solutions is matrix or tensor factorization, which decomposes multimodal data into the components that capture joint variations shared across modalities, and the components that characterize modality-specific variations (Lock et al. 2013; Yang and Michailidis 2015; Li and Jung 2017; Lock and Li 2018; Gaynanova and Li 2019). Another class is canonical correlation analysis, which seeks maximum correlations between different data modalities through decomposition of the between-modality dependency structure (Li and Gaynanova 2018; Shu, Wang, and Zhu 2020). However, all these methods are unsupervised, in the sense that there is not a response variable involved. Li, Liu, and Chen (2018) recently proposed an integrative reduced-rank regression to

model multivariate responses given multi-view data as predictors. Xue and Qu (2020) developed an estimating equations approach to accommodate block missing patterns in multimodal data. Their methods are supervised, but both focused on parameter estimation and variable selection instead of statistical inference.

It is of ubiquitous interest to study the predictive associations between responses and multimodal predictors. However, there are some unique characteristics of multimodal data that make the problem challenging. First, multimodal data are often high-dimensional. In plenty of applications, even a single modality contains more variables than the sample size. Second, some variables in multimodal data can be highly correlated. This can happen for the variables within a single modality, or for the related variables across multiple modalities as they often measure-related features of the same subject. This phenomenon has been constantly observed, and is actually the base upon which those matrix or tensor factorization solutions are built (Lock et al. 2013). Such high correlations pose challenges when directly applying many standard high-dimensional methods, such as LASSO (Tibshirani 1996), to multimodal data, as they usually require the predictors not to be highly correlated in order to achieve some desired statistical properties, for example, the variable selection consistency. Finally, multimodal data pose new questions; for instance, how to quantify the contribution and statistical significance of one data modality conditioning on the other data modalities in the regression model.

Factor analysis is a well-known approach to both reduce high dimensionality and high correlations among the variables.

For data with a single modality, Fan, Liao, and Mincheva (2013) employed a factor model to estimate a non-sparse covariance matrix. Kneip and Sarda (2011) proposed to include the latent factors as additional explanatory variables in a high-dimensional linear regression, and established the model selection consistency. Fan, Ke, and Wang (2016) proposed a factor-adjusted model selection method for a general high-dimensional M -estimation problem. They separated the latent factors from the idiosyncratic components to reduce correlations among the covariates, and showed that their method can reach the variable selection consistency under milder conditions than standard selection methods. Li et al. (2018) studied estimation of a covariance matrix of variables. They showed that leveraging on additional auxiliary variables can improve the estimation, when the auxiliary variables share some common latent factors with the variables of interest. For data with multiple modalities, Shen, Wang, and Mo (2013) proposed an integrative clustering method based on identifying common latent factors from multi-omics data. Zhang, Tang, and Qu (2020) developed an imputed factor regression model for dimension reduction and prediction of multimodal data with missing blocks. Despite these efforts, however, there is little work on statistical inference for supervised modeling of multimodal data. Moreover, there is no explicit quantification of the benefit of factor analysis in a multimodal regression setting, and many important inference-related questions remain unanswered.

In this article, we aim to bridge this gap. We consider an integrative linear regression model built upon the latent factors extracted from multimodal data. We show that this model alleviates high dimensionality and high correlations of multimodal data. Based on this model, we address three important questions: how to infer the significance of one data modality given the other modalities in the model; how to infer the significance of a combination of variables from one or more modalities; and how to quantify the contribution, measured by the goodness-of-fit, of one data modality given the others. When answering each question, we explicitly characterize both the benefit and the extra cost of factor analysis. First, by resorting to a relatively small number of latent factors, it effectively reduces the dimensionality and turns a high-dimensional test to a low-dimensional one when testing the significance of a whole modality. As a result, it enables us to derive a closed form for the limiting distribution of the test statistic; see [Theorem 1](#). Second, our method can consistently estimate the support and nonzero components of the covariate coefficients in the regression model. More importantly, by using the decorrelated idiosyncratic components from factor analysis as the pseudo predictors, instead of the original highly correlated covariates, it requires much weaker conditions to reach the variable selection and estimation consistency; see [Theorem 3](#). In addition, we show that, when there are enough variables in each modality so that the latent factors can be well estimated, the resulting estimation error can reach the minimax optimal rate, and under weaker conditions. Third, such an improvement in selection and estimation in turn benefits the inference of the significance of a linear combination of predictors, by requiring less stringent conditions to establish the limiting distribution of the test statistic; see [Theorem 4](#). Finally, by leveraging on the latent factors shared across modalities, it enables us to obtain a

closed-form measure of the variance of the response explained by one modality in addition to the others. Such a measure facilitates the quantification of the contribution of an individual modality.

Our proposal contributes on several fronts. Even though factor analysis has been widely used in multimodal data analysis, there has been no formal test developed to explicitly quantify the contribution and significance of an individual modality or a related set of variables across different modalities. Our proposal provides the first inferential tools to address those important questions. Moreover, our work is built on careful examination of the benefit and tradeoff of factor analysis in regression. Compared to the existing literature, the proof techniques are much more involved than those of the standard setting when the design matrix is observed and fixed with a single data modality. The technical tools we develop here are not limited to our setting alone, but are applicable to general supervised high-dimensional factor models. We also remark that, although we focus on a linear factor regression model, most of the inference-related results we obtain can be extended to more general M -estimation problems such as a generalized linear model.

We employ the following notation throughout this article. For a vector $\mathbf{a} \in \mathcal{R}^d$, let $\|\mathbf{a}\|_\infty = \max_j |a_j|$, $\|\mathbf{a}\|_1 = \sum_{j=1}^d |a_j|$, $\|\mathbf{a}\|_2 = (\sum_{j=1}^d a_j^2)^{1/2}$ denote its sup-norm, L_1 -norm and Euclidean norm, respectively. For an index set S , let \mathbf{a}_S denote the subvector of \mathbf{a} with indices in S . In particular, let the subscript m denote the index set of the m th modality, and the subscript $-m$ denote the index set of all other modalities. Let $\mathbf{a}^{\otimes 2} = \mathbf{a}\mathbf{a}'$ denote its outer product. Let $\text{supp}(\mathbf{a}) = \{j : a_j \neq 0\}$ denote the support of \mathbf{a} . For a square matrix $\mathbf{A} = (a_{ij}) \in \mathcal{R}^{d \times d}$, let $\lambda_{\min}(\mathbf{A})$ and $\lambda_{\max}(\mathbf{A})$ denote its minimum and maximum eigenvalues. Let $\|\mathbf{A}\|_\infty = \sup_{ij} |a_{ij}|$, $\|\mathbf{A}\|_1 = \max_{1 \leq j \leq d} \sum_{i=1}^d |a_{ij}|$, $\|\mathbf{A}\|_2 = \lambda_{\max}(\mathbf{A})$, $\|\mathbf{A}\|_F = (\sum_{i,j} a_{ij}^2)^{1/2}$ denote its element-wise sup-norm, L_1 -norm, L_2 -norm, and Frobenius norm, respectively. Let \mathbf{A}_S denote the submatrix of \mathbf{A} with row and column indices in S . For a rectangular matrix $\mathbf{B} = (b_{ij}) \in \mathcal{R}^{m \times n}$, let $\|\mathbf{B}\|_{L_\infty} = \max_{1 \leq i \leq m} \sum_{j=1}^n |b_{ij}|$. For two sequences a_n and b_n , write $a_n = o(b_n)$ if $a_n/b_n \rightarrow 0$, and $a_n \gg b_n$ if $b_n/a_n \rightarrow 0$. For an integer M , let $[M] = \{1, \dots, M\}$. For a set S , let $|S|$ denote the number of elements in S .

The rest of the article is organized as follows. We introduce the integrative factor regression model in [Section 2](#), and describe the parameter estimation in [Section 3](#). These results are mostly built upon the existing literature on factor analysis. Then we address the three questions, which to our knowledge have not been answered before. That is, we develop a test to evaluate the significance of an individual modality given the other modalities in [Section 4](#), develop a test for a linear combination of predictors in [Section 5](#), and derive a measure to quantify the contribution of an individual modality in [Section 6](#). We present the simulations and a multimodal neuroimaging data example in [Section 7](#). We relegate all technical proofs and some additional lemmas to the supplementary materials.

2. Integrative Factor Regression Model

Suppose there are M modalities of variables. Let $\mathbf{x}_m \in \mathcal{R}^{p_m}$ denote the vector of p_m random variables from the

m th modality, and y denote the response variable. Let $\mathbf{x} = (\mathbf{x}'_1, \dots, \mathbf{x}'_M) \in \mathcal{R}^p$, and $p = \sum_{m=1}^M p_m$. We assume \mathbf{x}_m is driven by some latent factors in that \mathbf{x}_m can be decomposed as

$$\mathbf{x}_m = \mathbf{\Lambda}_m \mathbf{f}_m + \mathbf{u}_m, \tag{1}$$

where $\mathbf{f}_m \in \mathcal{R}^{K_m}$ is the vector of K_m random latent factors, $\mathbf{u}_m \in \mathcal{R}^{p_m}$ is the vector of random idiosyncratic errors of variables in the m th modality that are uncorrelated with \mathbf{f}_m , and $\mathbf{\Lambda}_m \in \mathcal{R}^{p_m \times K_m}$ is the loading matrix of \mathbf{x}_m on the latent factors \mathbf{f}_m . To avoid the identifiability issue on $\mathbf{\Lambda}_m$ and \mathbf{f}_m , we adopt the usual assumption in the factor analysis literature by assuming that

$$\begin{aligned} \text{var}(\mathbf{f}_m) &= \mathbf{I}_{K_m}, \quad \text{and} \\ \mathbf{\Lambda}'_m \mathbf{\Lambda}_m &= \mathbf{D}_m \in \mathcal{R}^{K_m \times K_m} \text{ is a diagonal matrix.} \end{aligned}$$

We also assume that $\mathbf{f} = (\mathbf{f}'_1, \dots, \mathbf{f}'_M) \in \mathcal{R}^K$ is uncorrelated with $\mathbf{u} = (\mathbf{u}'_1, \dots, \mathbf{u}'_M) \in \mathcal{R}^p$, where $K = \sum_{m=1}^M K_m$. Let $\mathbf{\Lambda} = \text{diag}(\mathbf{\Lambda}_1, \dots, \mathbf{\Lambda}_M) \in \mathcal{R}^{p \times K}$ be the block diagonal matrix of the loading matrices from all modalities, and $\mathbf{\Sigma}_u = \text{var}(\mathbf{u})$ be the covariance matrix of the idiosyncratic errors. In the factor analysis literature, for example, Fan, Liao, and Mincheva (2013), it is often assumed that $\mathbf{\Sigma}_u$ is sparse, that is, after removing the variations contributed by the latent factors, the correlations among the idiosyncratic components are weak. Therefore, the idiosyncratic \mathbf{u} can be viewed as a decorrelated version of the original variables \mathbf{x} .

We next employ a linear model to connect \mathbf{x}_m with y , in that,

$$y = \sum_{m=1}^M \mathbf{x}'_m \boldsymbol{\beta}^*_m + \epsilon, \tag{2}$$

where $\boldsymbol{\beta}^*_m \in \mathcal{R}^{p_m}$ is the true effect of \mathbf{x}_m on the response y , and ϵ is an error uncorrelated of the covariates, with $E(\epsilon) = 0$ and $\text{var}(\epsilon) = \sigma_\epsilon^2$.

Suppose we have n iid realizations of the data, $\mathbf{Y} = (y_1, \dots, y_n)' \in \mathcal{R}^n$, $\mathbf{X}_m = (\mathbf{x}_{1,m}, \dots, \mathbf{x}_{n,m})' \in \mathcal{R}^{n \times p_m}$, $\mathbf{X} = (\mathbf{X}_1, \dots, \mathbf{X}_M) \in \mathcal{R}^{n \times p}$, $\boldsymbol{\beta}^* = (\boldsymbol{\beta}^{*1}, \dots, \boldsymbol{\beta}^{*M})' \in \mathcal{R}^p$, and $\epsilon = (\epsilon_1, \dots, \epsilon_n)' \in \mathcal{R}^n$. Then model (2) can be written as

$$\mathbf{Y} = \sum_{m=1}^M \mathbf{X}_m \boldsymbol{\beta}^*_m + \epsilon = \mathbf{X} \boldsymbol{\beta}^* + \epsilon.$$

By the factor model (1), we have $\mathbf{X}_m = \mathbf{F}_m \mathbf{\Lambda}'_m + \mathbf{U}_m$, where $\mathbf{F}_m = (\mathbf{f}_{1,m}, \dots, \mathbf{f}_{n,m})' \in \mathcal{R}^{n \times K_m}$ is the matrix of the K_m factors in the m th modality pertaining to the n subjects, and $\mathbf{U}_m = (\mathbf{u}_{1,m}, \dots, \mathbf{u}_{n,m})' \in \mathcal{R}^{n \times p_m}$ is the matrix of idiosyncratic errors. Then, we have,

$$\mathbf{Y} = \mathbf{F} \boldsymbol{\gamma}^* + \mathbf{U} \boldsymbol{\beta}^* + \epsilon, \tag{3}$$

where $\mathbf{F} = (\mathbf{F}_1, \dots, \mathbf{F}_M) \in \mathcal{R}^{n \times K}$, $\boldsymbol{\gamma}^* = (\boldsymbol{\beta}^{*1} \mathbf{\Lambda}_1, \dots, \boldsymbol{\beta}^{*M} \mathbf{\Lambda}_M)' \in \mathcal{R}^K$, and $\mathbf{U} = (\mathbf{U}_1, \dots, \mathbf{U}_M) \in \mathcal{R}^{n \times p}$. We call model (3) an integrative factor regression model. In the remainder of this article, we aim to show that model (3) can benefit estimation, selection and inference about $\boldsymbol{\beta}^*$. The intuition is that, after the latent factors \mathbf{F} are separated, the idiosyncratic error \mathbf{U} can be treated as the pseudo predictors. Such a decorrelation eases selection of $\boldsymbol{\beta}^*$, which in turn benefits the inference on $\boldsymbol{\beta}^*$. In addition, the factor decomposition also

serves as a dimension reduction tool, which enables us to derive some closed-form results in inference. We remark that, in model (3), the coefficient $\boldsymbol{\beta}^*$ associated with \mathbf{U} is the same as that associated with the original predictor \mathbf{X} . It is the main object of interest in our inference, as its component $\boldsymbol{\beta}^*_m$ reflects the effect of the m th modality \mathbf{x}_m on the response y . Meanwhile, we treat $\boldsymbol{\gamma}^*$ as a nuisance parameter. We also remark that, one does *not* necessarily have to perform factor decomposition for all data modalities. In practice, we first estimate the number of factors for each modality. For a particular modality that does not admit a factor structure, we can set the corresponding $\mathbf{\Lambda}_m = \mathbf{0}$ and $\mathbf{x}_m = \mathbf{u}_m$ in (1).

3. Estimation

To fit model (3), we first estimate the latent variables \mathbf{F} and \mathbf{U} , along with the number of latent factors K_m , using some well established methods in the factor analysis literature. We then estimate $\boldsymbol{\beta}^*_m$ through a penalized least-square approach.

First, we estimate the latent variables \mathbf{F} and \mathbf{U} , we adopt the method in Bai and Li (2012) and Fan, Liao, and Mincheva (2013), by running principal components analysis (PCA) on each individual modality \mathbf{X}_m . We then estimate \mathbf{F}_m by \sqrt{n} times eigenvectors corresponding to the largest K_m eigenvalues of $\mathbf{X}_m \mathbf{X}'_m$. Denote this estimator by $\hat{\mathbf{F}}_m$. We next estimate $\mathbf{\Lambda}_m$ by $\hat{\mathbf{\Lambda}}_m = (1/n) \mathbf{X}'_m \hat{\mathbf{F}}_m$, and estimate \mathbf{U}_m by $\hat{\mathbf{U}}_m = \mathbf{X}_m - \hat{\mathbf{F}}_m \hat{\mathbf{\Lambda}}'_m$, accordingly. Fan, Liao, and Mincheva (2013) showed that $\hat{\mathbf{F}}_m$ is a consistent estimator, up to a rotation, of \mathbf{F}_m , under a pervasive condition that the latent factors should affect many variables; see Condition 2 and its discussion in Section 4. We remark that, there are alternative ways to estimate the factors. For instance, methods such as Ma (2013); Cai, Ma, and Wu (2013), and Lock et al. (2013) may also be applicable, as long as the resulting factor estimates are consistent.

Next, we determine the number of the latent factors K_m in each modality, which is usually unknown in practice. We use the method of Bai and Ng (2002) to estimate K_m by

$$\begin{aligned} \hat{K}_m &= \text{argmin}_{0 \leq k \leq \tilde{M}} \log \left\{ \frac{1}{np_m} \left\| \mathbf{X}_m - n^{-1} \hat{\mathbf{F}}_{m_k} \hat{\mathbf{F}}'_{m_k} \mathbf{X}_m \right\|_F^2 \right\} \\ &\quad + kg(n, p_m), \end{aligned} \tag{4}$$

where \tilde{M} is a predefined upper bound on K_m , $\hat{\mathbf{F}}_{m_k}$ is \sqrt{n} times eigenvectors corresponding to the largest m_k eigenvalues of $\mathbf{X}_m \mathbf{X}'_m$, and $g(n, p_m)$ is a penalty function that,

$$\begin{aligned} g(n, p_m) &= \frac{n + p_m}{np_m} \log \left(\frac{np_m}{n + p_m} \right), \quad \text{or} \\ g(n, p_m) &= \frac{n + p_m}{np_m} \log (\min\{n, p_m\}). \end{aligned}$$

For both choices, Bai and Ng (2002) showed that \hat{K}_m is a consistent estimator of K_m under some regularity conditions. We make two additional remarks about K_m . First, in this article, we treat K_m as fixed, which is reasonable in numerous scientific applications. As K_m is related to the number of spiked eigenvalues of $\mathbf{X}_m \mathbf{X}'_m$, it is usually small. Second, following a common practice of the factor analysis literature, in our subsequent theoretical

analysis, we treat K_m as known. All the theoretical results remain valid conditioning on that there is a consistent estimator \hat{K}_m of K_m .

Next, we estimate β^* and γ^* . We replace F and U with the corresponding estimators $\hat{F} = (\hat{F}_1, \dots, \hat{F}_M)$ and $\hat{U} = (\hat{U}_1, \dots, \hat{U}_M)$, and solve a penalized least-square problem,

$$(\hat{\gamma}, \hat{\beta}) = \operatorname{argmin}_{(\gamma, \beta)} \frac{1}{2n} \sum_{i=1}^n (y_i - \hat{F}'_i \gamma - \hat{U}'_i \beta)^2 + \lambda \sum_{j=1}^p p(|\beta_j|), \quad (5)$$

where $p(\cdot)$ is some general folded-concave penalty function, and λ is a tuning parameter. This class of penalty functions includes SCAD (Fan and Lv 2011) and MCP (Zhang 2010). It assumes that $p(t)$ is increasing and concave in $t \geq 0$, and has a continuous first derivative $\dot{p}(t)$ with $\dot{p}(0+) > 0$. This optimization problem can be solved by standard proximal gradient descent algorithms (Parikh and Boyd 2014). We also briefly comment that, in our optimization (5), we do not impose the linear constraint that $\gamma = \Lambda' \beta$, mainly because both F and Λ are unknown and unidentifiable in our setting. Besides, as we later show in Theorem 3 that, even without this constraint, the estimator $\hat{\beta}$ from Equation (5) achieves the minimax rate, and can consistently select the support and estimate the nonzero components of β^* , as long as there are enough variables to estimate the latent factors well. We tune λ in Equation (5) using the standard cross-validation method, following Fan and Lv (2011), while alternative criteria, for example, the extended BIC (Chen and Chen 2008), can also be used to tune λ .

Finally, we estimate σ_ϵ^2 by $\hat{\sigma}_\epsilon^2 = n^{-1} \sum_{i=1}^n (y_i - \mathbf{x}'_i \tilde{\beta})^2$, where $\tilde{\beta}$ is the Lasso estimator of β^* obtained by $\tilde{\beta} = \operatorname{argmin}_\beta (2n)^{-1} \sum_{i=1}^n (y_i - \mathbf{x}'_i \beta)^2 + \lambda_\epsilon \|\beta\|_1$, where λ_ϵ is the tuning parameter. We show in the Supplementary Materials that $\hat{\sigma}_\epsilon^2$ is consistent to σ_ϵ^2 . Actually, any consistent estimator of σ_ϵ^2 would suffice for the subsequent hypothesis testing procedures. Alternative methods such as scaled LASSO (Sun and Zhang 2013), refitted cross-validation (Fan, Guo, and Hao 2012), or directly using the residuals from Equation (5) can all be applied to estimate σ_ϵ^2 as well.

4. Hypothesis Test of a Whole Modality

A crucial question in multimodal data analysis is to evaluate if a whole modality is significantly associated with the outcome, given other modalities in the model. For instance, in multi-omics analysis, it is of interest to test if DNA methylation correlates with the phenotypic traits related to genetic disorders given gene expression level (Richardson, Tseng, and Sun 2016). In multimodal neuroimaging analysis, it is of interest to evaluate if functional imaging quantification for hypometabolism associates with the diagnosis of Alzheimer's disease, given structural magnetic resonance imaging of brain atrophy measurement (Zhang et al. 2011). The challenge here is that even a single modality often contains many more variables than the sample size.

This is essentially a problem of testing a high-dimensional subvector of β^* in a high-dimensional regression model. Related testing problems have been extensively studied for a single modality data. For example, Zhang and Zhang (2014),

van de Geer et al. (2014), and Javanmard and Montanari (2014) developed bias-corrected or de-sparsifying methods to test if a fixed-dimensional subvector of β^* in a high-dimensional linear or generalized linear model equals zero. In particular, under a general M -estimation framework, Ning and Liu (2017) proposed a decorrelated score test for the same problem, that is, to test if a subvector $\beta_S^* = \mathbf{0}$. They first showed that their score test statistic has a closed-form limiting distribution when the dimension of the subset $|S|$ is fixed. They then extended to the case where β_S^* can be any arbitrary subvector of β^* with $|S|$ diverging and even when $|S| > n$. Built on a pioneering work by Chernozhukov et al. (2013), they showed that the distribution of the supremum of the decorrelated score functions can be approximated by a multiplier bootstrap approach. Consequently, they employed bootstrap simulations to obtain the critical values of the limiting distribution to form the rejection region. Our test differs from Ning and Liu (2017). When $|S|$ diverges, the test of Ning and Liu (2017) no longer has a closed-form limiting distribution, and they had to resort to bootstrap for critical values. By contrast, we are able to obtain a closed-form limiting distribution for our test when $|S|$ diverges. This is due to that, instead of using the observed likelihood, we perform factor decomposition on X_m first, then use the factor model as a dimension reduction tool to reduce a high-dimensional test to a fixed-dimensional one. Our method does pay the extra price that we need to estimate the latent factors to plug into the likelihood function. However, as we show later, this extra cost can be well controlled. We also numerically compare with Ning and Liu (2017) in Section 7.1. We show that our test is as powerful, and often more powerful than the test of Ning and Liu (2017).

Formally, for our multimodal analysis, we aim at testing the following pair of hypotheses:

$$H_0 : \beta_m^* = \mathbf{0} \quad \text{versus} \quad H_a : \beta_m^* \neq \mathbf{0}. \quad (6)$$

We perform factor decomposition on the m th modality following Equation (1). Then,

$$\mathbf{x}' \beta = \mathbf{x}'_{-m} \beta_{-m} + \mathbf{f}'_m \gamma_m + \mathbf{u}'_m \beta_m,$$

where $\gamma_m = \Lambda'_m \beta_m$. The null hypothesis $\beta_m^* = \mathbf{0}$ implies that $\gamma_m^* = \mathbf{0}$, where $\gamma_m^* = \Lambda'_m \beta_m^*$. Therefore, under the null hypothesis, testing γ_m^* is the same as testing β_m^* . The difference is that $\gamma_m^* \in \mathcal{R}^{K_m}$ is a low-dimensional vector, while $\beta_m^* \in \mathcal{R}^{p_m}$ is high-dimensional. As such, the factor model plays the role of dimension reduction for our testing problem. Actually, directly testing β_m^* is challenging, since the dimension of β_m^* diverges with the sample size, and there is not a closed form for the limiting distribution of $\hat{\beta}_m - \beta_m^*$, where $\hat{\beta}_m$ is an estimator of β_m^* . On the other hand, we note that, under the alternative hypothesis, the magnitude of γ_m^* can be different from that of β_m^* . As such, the power of the test that is built on γ_m^* can be different from the one that is built on β_m^* . We later study the local power property of the test based on γ_m^* in detail.

Next, we develop a factor-adjusted decorrelated score test, and show that it is asymptotically efficient when the latent factors can be well estimated. Following Ning and Liu (2017), and

based on the Gaussian quasi-likelihood, we define the decorrelated score function as

$$S(\boldsymbol{\beta}, \boldsymbol{\gamma}_m) = \frac{1}{n\hat{\sigma}_\epsilon^2} \sum_{i=1}^n (y_i - \mathbf{f}'_{i,m} \boldsymbol{\gamma}_m - \mathbf{z}'_i \boldsymbol{\beta}) (\mathbf{f}_{i,m} - \mathbf{W}^{*'} \mathbf{z}_i),$$

where $\mathbf{z}_i = (\mathbf{x}'_{i,-m}, \mathbf{u}'_{i,m})' \in \mathcal{R}^p$, $p_{-m} = p - p_m$, and $\mathbf{W}^* = \mathbf{E}(\mathbf{x}_{i,-m}^{\otimes 2})^{-1} \mathbf{E}(\mathbf{x}_{i,-m} \mathbf{f}'_{i,m}) \in \mathcal{R}^{p_{-m} \times K_m}$, which is essentially the projection of the latent factors onto the linear space spanned by \mathbf{x}_{-m} . Such a projection is needed to control the variability of the high-order terms in establishing the central limit theorem in Theorem 1 (Ning and Liu 2017). In the high-dimensional setting, we need some sparsity condition on \mathbf{W}^* ; see Condition 4, and solve a regularized problem to obtain its consistent estimator. We treat the score function as a function of $\boldsymbol{\gamma}_m$. Under the null hypothesis, we propose to estimate $S(\boldsymbol{\beta}, \boldsymbol{\gamma}_m)$ by,

$$\hat{S}(\hat{\boldsymbol{\beta}}_{-m}, \mathbf{0}) = \frac{1}{n\hat{\sigma}_\epsilon^2} \sum_{i=1}^n (y_i - \mathbf{x}'_{i,-m} \hat{\boldsymbol{\beta}}_{-m}) (\hat{\mathbf{f}}_{i,m} - \hat{\mathbf{W}}' \mathbf{x}_{i,-m}),$$

where $\hat{\mathbf{f}}_{i,m}$ is the i th row of the estimated latent factor matrix $\hat{\mathbf{F}}_m$, $\hat{\sigma}_\epsilon^2$ is any consistent estimator of σ_ϵ^2 that satisfies Condition 5 below, and $\hat{\boldsymbol{\beta}}_{-m} \in \mathcal{R}^{p_{-m}}$ and $\hat{\mathbf{W}} \in \mathcal{R}^{p_{-m} \times K_m}$ are obtained by solving the following optimization problems,

$$(\hat{\boldsymbol{\beta}}_{-m}, \hat{\boldsymbol{\gamma}}_m) = \operatorname{argmin}_{(\boldsymbol{\beta}_{-m}, \boldsymbol{\gamma}_m)} \frac{1}{2n} \sum_{i=1}^n (y_i - \mathbf{x}'_{i,-m} \boldsymbol{\beta}_{-m} - \hat{\mathbf{f}}'_{i,m} \boldsymbol{\gamma}_m)^2 + \lambda_1 \|\boldsymbol{\beta}_{-m}\|_1, \tag{7}$$

$$\hat{\mathbf{W}} = \operatorname{argmin} \|\mathbf{W}\|_1, \quad \text{such that}$$

$$\left\| \frac{1}{n} \sum_{i=1}^n \mathbf{x}_{i,-m} (\hat{\mathbf{f}}_{i,m} - \mathbf{x}'_{i,-m} \mathbf{W}) \right\|_\infty \leq \lambda_2. \tag{8}$$

We make a few remarks regarding our score function and compare it to Ning and Liu (2017). First, we need to estimate the latent factors in the decorrelated score function, while in the score function of Ning and Liu (2017), the covariates are fully observed. This introduces an additional layer of complexity when analyzing the statistical property of the score function. Later in Theorems 1 and 4, we carefully evaluate the extra cost of estimating the latent factors. Second, we need to involve $\boldsymbol{\gamma}_m$ in (7), even under the null hypothesis $\boldsymbol{\gamma}_m = \mathbf{0}$. This is because, if $\boldsymbol{\gamma}_m$ is removed from Equation (7), $\hat{\boldsymbol{\beta}}_{-m}$ is no longer consistent to $\boldsymbol{\beta}^*_{-m}$ under the alternative, which would in turn impact the power of the test. On the other hand, $\hat{\boldsymbol{\gamma}}_m$ is not used in constructing the test statistic, but only $\hat{\boldsymbol{\beta}}_{-m}$ is. Finally, in order to consistently estimate \mathbf{W}^* , we need to solve a non-typical Dantzig problem (8), where the latent factors are replaced by their corresponding estimators.

Next, we compute the variance of the score function by using the Fisher information. By the sandwich formula, the information matrix is

$$\mathbf{I}^*_{\boldsymbol{\gamma}_m | \boldsymbol{\beta}_{-m}} = \sigma_\epsilon^{-2} \{ \mathbf{E}(\mathbf{f}_{i,m}^{\otimes 2}) - \mathbf{E}(\mathbf{f}_{i,m} \mathbf{x}'_{i,-m}) \mathbf{E}(\mathbf{x}_{i,-m}^{\otimes 2})^{-1} \mathbf{E}(\mathbf{x}_{i,-m} \mathbf{f}'_{i,m}) \} \in \mathcal{R}^{K_m \times K_m},$$

which can be estimated by

$$\hat{\mathbf{I}}_{\boldsymbol{\gamma}_m | \boldsymbol{\beta}_{-m}} = \hat{\sigma}_\epsilon^{-2} \left\{ \frac{1}{n} \sum_{i=1}^n \hat{\mathbf{f}}_{i,m}^{\otimes 2} - \hat{\mathbf{W}}' \left(\frac{1}{n} \sum_{i=1}^n \mathbf{x}_{i,-m} \hat{\mathbf{f}}'_{i,m} \right) \right\}.$$

Then, our test statistic is given by

$$T_n = \sqrt{n} \hat{\mathbf{I}}_{\boldsymbol{\gamma}_m | \boldsymbol{\beta}_{-m}}^{-1/2} \hat{S}(\hat{\boldsymbol{\beta}}_{-m}, \mathbf{0}).$$

We next show that, under the null hypothesis, the asymptotic distribution of T_n is $N(\mathbf{0}, \mathbf{I}_{K_m})$. In other words, T_n is asymptotically efficient. We first begin with a set of conditions.

Condition 1. For $m \in [M]$, suppose $\{(\mathbf{f}'_{i,m}, \mathbf{u}'_{i,m})'\}_{i=1}^n$ are iid uncorrelated sub-Gaussian random vectors with zero mean. That is, $\mathbf{E}(\mathbf{f}_{i,m}) = \mathbf{0}$, $\mathbf{E}(\mathbf{u}_{i,m}) = \mathbf{0}$, and $\mathbf{E}(\mathbf{f}_{i,m} \mathbf{u}'_{i,m}) = \mathbf{0}$. Moreover, $\mathbf{E}\{\exp(\mathbf{t}\boldsymbol{\alpha}' \mathbf{f}_{i,m})\} \leq \exp(C \|\boldsymbol{\alpha}\|_2^2 t^2/2)$, and $\mathbf{E}\{\exp(\mathbf{t}\boldsymbol{\alpha}' \mathbf{u}_{i,m})\} \leq \exp(C \|\boldsymbol{\alpha}\|_2^2 t^2/2)$, for some constant C . In addition, for all $k \in [K_m]$, $\mathbf{x}'_{i,-m} \mathbf{w}_k^*$ are iid sub-Gaussian such that $\mathbf{E}\{\exp(\mathbf{t}\mathbf{x}'_{i,-m} \mathbf{w}_k^*)\} \leq \exp(Ct^2/2)$, where \mathbf{w}_k^* is the k th column of \mathbf{W}^* . Additionally, $\{\epsilon_i\}_{i=1}^n$ are iid sub-Gaussian with zero mean, and ϵ_i is uncorrelated with $(\mathbf{f}'_{i,m}, \mathbf{u}'_{i,m})'$ for all $m \in [M]$.

Condition 2. For $m \in [M]$, suppose $0 < c \leq \lambda_{\min}(\boldsymbol{\Lambda}'_m \boldsymbol{\Lambda}_m / p_m) \leq \lambda_{\max}(\boldsymbol{\Lambda}'_m \boldsymbol{\Lambda}_m / p_m) \leq C < \infty$, for some positive constants c and C .

Condition 3. For $m \in [M]$, $s, t \in [p_m]$, $i, j \in [n]$, suppose $\mathbf{E}[p_m^{-1/2} \{\mathbf{u}'_{i,m} \mathbf{u}_{j,m} - \mathbf{E}(\mathbf{u}'_{i,m} \mathbf{u}_{j,m})\}^4] \leq C$, and $\mathbf{E} \left\| p_m^{-1/2} \boldsymbol{\Lambda}'_m \mathbf{u}_{i,m} \right\|_2^4 \leq C$. Moreover, $\|\boldsymbol{\Lambda}_m\|_\infty \leq C$, $\lambda_{\min}(\boldsymbol{\Sigma}_{u_m}) > c$, $\|\boldsymbol{\Sigma}_{u_m}\|_1 \leq C$, where $\boldsymbol{\Sigma}_{u_m} = \operatorname{var}(\mathbf{u}_m)$, and $\min_{s,t \in [p_m]} \operatorname{var}(u_{i,m_s} u_{i,m_t}) > c$.

Condition 4. Let $s_w^* = \max_{k \in [K_m]} |\operatorname{supp}(\mathbf{w}_k^*)|$. Suppose $s_w^* \log(p_{-m}) \{1 \vee (n^{1/4}/\sqrt{p_m})\} = o(n^{1/2})$, and $[s_{-m}^* \{\sqrt{(\log p_{-m})/n + 1/\sqrt{p_m}}\}] \cdot \sqrt{\log(p_{-m})} \{1 \vee (n^{1/4}/\sqrt{p_m})\} = o(1)$.

Condition 5. Suppose $\hat{\sigma}_\epsilon^2 = \sigma_\epsilon^2 + o_P(1)$.

Condition 6. Suppose $0 < c \leq \lambda_{\min}(\mathbf{I}^*_{\boldsymbol{\gamma}_m | \boldsymbol{\beta}_{-m}})$.

Condition 1 is a typical sub-Gaussian assumption for high-dimensional problems. Condition 2 is the pervasive condition, and is common in factor analysis (Fan, Liao, and Mincheva 2013). It requires that the latent factors affect a large number of variables. This is reasonable for a variety of multimodal data. For instance, in multi-omics data, some genetic factors are believed to impact both gene expression and DNA methylation, and in multimodal neuroimaging, some neurological factors affect both brain structures and functions. Condition 3 imposes some technical requirements on the loading matrix and idiosyncratic component. Together, Conditions 2 and 3 ensure that $\mathbf{f}_{i,m}$ and $\mathbf{u}_{i,m}$ can be consistently estimated by the PCA method (Fan, Liao, and Mincheva 2013). Condition 4 is a sparsity condition on \mathbf{W}^* and $\boldsymbol{\beta}^*_{-m}$, which requires s_w^* and s_{-m}^* to be much smaller than n . Under such a condition, \mathbf{W}^* and $\boldsymbol{\beta}^*_{-m}$ can both be consistently estimated, even if the latent factors are unknown; see Lemmas 6 and 8 in the supplementary materials for more details. We remark that this sparsity assumption on \mathbf{W}^* is weaker and more flexible than requiring both $\mathbf{E}(\mathbf{x}_{i,-m}^{\otimes 2})^{-1}$ and $\mathbf{E}(\mathbf{x}_{i,-m} \mathbf{f}'_{i,m})$ are sparse. Condition 5 ensures the estimator of σ_ϵ^2 is consistent. Condition 6 ensures the information matrix is

invertible. We also remark that, if there is no factor in the m th modality, [Condition 1](#) reduces to the sub-Gaussian assumption on $\mathbf{x}_{i,m}$, and [Conditions 2](#) and [3](#) are no longer needed for that modality.

We next obtain a closed-form limiting distribution for the test statistic T_s .

Theorem 1. Suppose [Conditions 1–6](#) hold. Suppose $\lambda_1 \asymp \sqrt{(\log p_{-m})/n} + 1/\sqrt{p_m}$, and $\lambda_2 \asymp \sqrt{(\log p_{-m})/n} \vee (n^{1/4}/\sqrt{p_m})$. Then, under $H_0 : \boldsymbol{\beta}_m^* = \mathbf{0}$, it holds that

$$T_n \xrightarrow{D} N(\mathbf{0}, \mathbf{I}_{K_m}).$$

By [Theorem 1](#), we reject the null hypothesis if $n\{\hat{\mathbf{S}}(\hat{\boldsymbol{\beta}}_{-m}, \mathbf{0})\}' \hat{\mathbf{T}}_{\mathbf{y}_m | \boldsymbol{\beta}_{-m}}^{-1} \hat{\mathbf{S}}(\hat{\boldsymbol{\beta}}_{-m}, \mathbf{0}) > \chi_{\alpha}^2(K_m, 0)$, where $\chi_{\alpha}^2(K_m, 0)$ is the α -upper quantile of the χ^2 -distribution with K_m degrees of freedom.

We next explicitly discuss the benefit and the extra cost of our factor-based test when compared with Ning and Liu (2017). The main difference is that, through latent factors, we obtain a closed-form limiting distribution and do not have to resort to bootstrap. The price we pay mainly lies in [Condition 4](#) and the choices of λ_1 and λ_2 . Actually, the extra term $1/\sqrt{p_m}$ appearing in both [Condition 4](#) and λ_1, λ_2 reflects the estimation error caused by using $\hat{\mathbf{f}}_{i,m}$ to estimate $\boldsymbol{\beta}_{-m}^*$. The term $n^{1/4}/\sqrt{p_m}$ is due to the same reason for estimating \mathbf{w}_k^* . Therefore, the choices of the tuning parameters λ_1 and λ_2 need to be adjusted accordingly, by taking into account such extra estimation errors.

We further consider three scenarios. First, when $p_m \gg n$, both $1/\sqrt{p_m}$ and $n^{1/4}/\sqrt{p_m}$ are dominated by $\sqrt{(\log p_{-m})/n}$. Therefore, the estimation errors of $\hat{\boldsymbol{\beta}}_{-m}$ and $\hat{\mathbf{W}}$ reach the optimal oracle rate, that is, the best rate as if the latent factors were known; see Lemmas 6 and 8 in the supplementary materials. In this case, using the factor estimates actually does not incur any extra cost. The reason is that many variables are used to estimate the latent factors, and its estimation error is so small that it would not affect the inference on $\boldsymbol{\beta}_m^*$. Second, when $p_m = o(n)$, the estimation errors of $\hat{\boldsymbol{\beta}}_{-m}$ and $\hat{\mathbf{W}}$ would be greater than the optimal rate. However, the central limit theorem still holds, given proper choices of λ_1 and λ_2 , and more stringent sparsity conditions on $\boldsymbol{\beta}_{-m}^*$ and \mathbf{W}^* in [Condition 4](#). Third, in a special case where variables in all modalities are driven by exactly the same latent factors, even we perform the hypothesis test on the m th modality, we could use variables from all different modalities to estimate the latent factors. Then, the terms $1/\sqrt{p_m}$ and $n^{1/4}/\sqrt{p_m}$ become $1/\sqrt{p}$ and $n^{1/4}/\sqrt{p}$, respectively, which are naturally dominated by $\sqrt{(\log p_{-m})/n}$. In this case, the optimal oracle rate is again attained. Such a result can be viewed as a blessing of the dimensionality for the factor model. In summary, our method is most suitable for testing the significance of a modality containing many variables, or for multimodal data with a large number of variables driven by some common latent factors.

Next, we study the power of the proposed test under the local alternative $H_{a_n} : \boldsymbol{\beta}_m^* = \mathbf{b}_{m_n}$, where \mathbf{b}_{m_n} is a sequence converging to $\mathbf{0}$ as $n \rightarrow \infty$. Since we use the latent factors to transform the test on $\boldsymbol{\beta}_m^*$ to the one on $\boldsymbol{\gamma}_m^*$, we show that the local power of the test depends on $\mathbf{c}_{m_n} = \boldsymbol{\Lambda}_m' \mathbf{b}_{m_n}$. We consider the following

parameter space under the local alternative, $\mathcal{N} = \{\boldsymbol{\beta}_m^* : \boldsymbol{\beta}_m^* = \mathbf{b}_{m_n}, |\text{supp}(\boldsymbol{\beta}_{-m}^*)| = s_{-m}^*$, where $s_{-m}^* \ll n\}$. The next theorem gives the limiting distribution of $Q_n = \mathbf{T}_n' \mathbf{T}_n$ uniformly for all $\boldsymbol{\beta}_m^* \in \mathcal{N}$ under the local alternative.

Theorem 2. Suppose the conditions of [Theorem 1](#) hold. Suppose $\lambda_1 \asymp \sqrt{(\log p_{-m})/n} + 1/\sqrt{p_m}$, $\lambda_2 \asymp \sqrt{(\log p_{-m})/n} \vee (n^{1/4}/\sqrt{p_m})$, $\|\mathbf{b}_{m_n}\|_2 = o(1/\sqrt{\log n})$, and $\|\mathbf{c}_{m_n}\|_2 = o(1/\sqrt{\log n})$. Then, under the H_{a_n} , it holds uniformly for all $\boldsymbol{\beta}_m^* \in \mathcal{N}$ that

$$\sup_{x>0} |\Pr(Q_n \leq x) - \Pr\{\chi^2(K_m, h_{m_n}) \leq x\}| \rightarrow 0,$$

where $h_{m_n} = n\mathbf{c}_{m_n}' \mathbf{T}_{\mathbf{y}_m | \boldsymbol{\beta}_{-m}}^* \mathbf{c}_{m_n}$.

Since $h_{m_n} \asymp n\|\mathbf{c}_{m_n}\|_2^2$, [Theorem 2](#) implies that the local power of our test essentially depends on $\|\mathbf{c}_{m_n}\|_2$. If we let $\|\mathbf{c}_{m_n}\|_2 = Cn^{-\phi_{\gamma_m}}$, the local power is to exhibit some transition behavior depending on the value of ϕ_{γ_m} , which is summarized in the next corollary.

Corollary 1. Suppose the conditions of [Theorem 2](#) hold. Then,

- $\lim_{n \rightarrow \infty} \sup_{\boldsymbol{\beta}_m^* \in \mathcal{N}} \sup_{x>0} |\Pr(Q_n \leq x) - \Pr(\chi^2(K_m, 0) \leq x)| \rightarrow 0$, if $\phi_{\gamma_m} > 1/2$;
- $\lim_{n \rightarrow \infty} \sup_{\boldsymbol{\beta}_m^* \in \mathcal{N}} \sup_{x>0} |\Pr(Q_n \leq x) - \Pr(\chi^2(K_m, h) \leq x)| \rightarrow 0$, if $\phi_{\gamma_m} = 1/2$;
- $\liminf_{n \rightarrow \infty} \sup_{\boldsymbol{\beta}_m^* \in \mathcal{N}} \Pr(Q_n > x) = 1$, if $\phi_{\gamma_m} < 1/2$;

where $h = \lim_{n \rightarrow \infty} n\mathbf{c}_{m_n}' \mathbf{T}_{\mathbf{y}_m | \boldsymbol{\beta}_{-m}}^* \mathbf{c}_{m_n}$ in (b), and (c) holds for any $x > 0$.

We make some remarks. First, [Corollary 1](#) shows that, the local power is to converge to the Type I error if $\phi_{\gamma_m} > 1/2$; to a non-central χ^2 -distribution if $\phi_{\gamma_m} = 1/2$; and to 1 if $\phi_{\gamma_m} < 1/2$. Such a transition behavior is analogous to the classical local power results, which showed that the root- n local alternative is the transition point of the local power (van der Vaart 2000). Second, the existing debiased method (van de Geer et al. 2014) and the decorrelated method (Ning and Liu 2017) only established the root- n local power results when the dimension of the parameters being tested is fixed. Moreover, even though Ning and Liu (2017) used a multiplier bootstrap method to extend their test from a single parameter to arbitrarily many parameters, they only studied the local power when testing a single parameter. Our local power result differs in that we allow the dimension of $\boldsymbol{\beta}_m^*$ to grow with n , whereas we fix the dimension of $\boldsymbol{\gamma}_m^*$. Finally, [Theorem 2](#) shows that the local power depends on the magnitude of $\boldsymbol{\gamma}_m^*$, or $\mathbf{c}_{m_n} = \boldsymbol{\Lambda}_m' \mathbf{b}_{m_n}$. This is again due to that we transform the test of $\boldsymbol{\beta}_m^*$ to that of $\boldsymbol{\gamma}_m^*$. Therefore, the power of our test depends on the relation between the loadings and where the alternative hypothesis occurs.

Next, we give some specific examples to further illustrate the power behavior of our proposed test. To simplify the discussion, we set $K_m = 1$.

Example 1. Let $\boldsymbol{\Lambda}_m = (D_L, 1, \dots, 1)'$ and $\mathbf{b}_{m_n} = (D_S n^{-1/2}, 0, \dots, 0)'$, where D_L and D_S are two constants. In this case, $\|\mathbf{c}_{m_n}\|_2 = D_S D_L n^{-1/2}$, and thus the power of our test depends on the product $D_S D_L$. On the contrary, even one

had known apriori that the alternative only occurs at the first coordinate, and performs a debiased or decorrelated test on that coordinate, its local power depends on D_S . Therefore, when D_L is large and D_S is small, our testing method gains power. On the other hand, when D_L is small and D_S is large, the alternative methods may be more powerful.

Example 2. Let $\mathbf{\Lambda}_m = (1, 1, \dots, 1)'$, and $\mathbf{b}_{m_n} = (C_1 n^{-2/3}, C_2 n^{-2/3}, \dots, C_L n^{-2/3}, 0, \dots, 0)'$. In this case, if $L \gg n^{1/6}$, $\|\mathbf{c}_{m_n}\|_2 \gg n^{-1/2}$, then the power of our test converges to 1. On the contrary, if one performs a debiased or decorrelated test on each element of $\boldsymbol{\beta}_m^*$, there is no power. Our testing method gains power in this example too.

Example 3. Let $\mathbf{\Lambda}_m = (0, 1, \dots, 1)'$, and $\mathbf{b}_{m_n} = (c_n, 0, \dots, 0)'$. In this case, no matter how large c_n is, our testing method has no power to detect the alternative, because $\mathbf{c}_{m_n} = \mathbf{0}$.

As we have seen in these examples, when transforming the test from $\boldsymbol{\beta}_m^*$ to $\boldsymbol{\gamma}_m^*$, our method does not necessarily lose power, but can gain power in some situations. For instance, in [Example 1](#), the variables with large loadings on the latent factors have nonzero coefficients, while in [Example 2](#), many variables with nonzero loadings have nonzero coefficients. In such cases, our test gains power. On the other hand, in [Example 3](#), the product of the loadings and the nonzero coefficients is small, then our test has little power.

5. Hypothesis Test of Linear Combinations of Predictors of One or More Modalities

Another important question in multimodal data analysis is to test if some linear combinations of predictors, within the same modality or across different modalities, is significantly correlated with the response. This is because multimodal data often measures different aspects of related quantities. For instance, in multi-omics studies, expression data measures how genes are expressed, methylation data measures how DNA molecules are methylated, and both data may be related to the same set of genes. In multimodal neuroimaging analysis, brain structures, functions, and chemical constituents of the same brain regions are often measured simultaneously. As such, it is of great scientific interest to test if various measurements on a particular gene or brain region are associated with the outcome.

Shi et al. (2019) considered a similar testing problem in a high-dimensional generalized linear model for a single modality data. They derived the corresponding partially penalized likelihood ratio test, score test and Wald test, and showed that the three tests are asymptotically equivalent. They allowed the dimension of the model to grow with the sample size, as long as the dimension of the subvector being tested and the number of linear combinations are smaller than the sample size. Our method differs from Shi et al. (2019) in several ways. Shi et al. (2019) treated the design matrix \mathbf{X} as fixed, while we treat \mathbf{X} as iid random realizations from some distributions. More importantly, we do not directly use the observed \mathbf{X} , but instead perform a factor decomposition and use the decorrelated idiosyncratic components as the pseudo design matrix. We explicitly show in [Theorem 3](#) that such a factor-adjusted step

leads to less stringent conditions to reach the variable selection and estimation consistency. Moreover, since variable selection consistency is needed to correctly calculate the variance of the test statistic, as shown in [Theorem 4](#), our method also requires less stringent conditions to establish the limiting distribution of the test statistic. Moreover, our model concerns with data with multiple modalities, instead of a single modality as in Shi et al. (2019). High correlations are commonly observed in multimodal data, and as such the factor-adjusted decorrelation step becomes essential. Relatedly, Zhu and Bradic (2018) proposed a test for a linear combination of predictors under a unimodal linear regression model. Even though they did not restrain the size or the sparsity of the model, they only considered a single linear combination, and required the eigenvalues of the covariate covariance matrix $\text{var}(\mathbf{x})$ to be bounded. By contrast, both Shi et al. (2019) and we consider jointly testing multiple linear combinations of predictors, and we do not require the eigenvalues of $\text{var}(\mathbf{x})$ to be bounded. We further numerically compare with Shi et al. (2019) and Zhu and Bradic (2018) in [Section 7.2](#).

Formally, we consider testing the following pair of hypotheses:

$$H_0 : \mathbf{A}\boldsymbol{\beta}_T^* = \mathbf{b} \quad \text{versus} \quad H_a : \mathbf{A}\boldsymbol{\beta}_T^* \neq \mathbf{b}, \tag{9}$$

where $\mathbf{A} \in \mathcal{R}^{r \times t}$, $\mathbf{b} \in \mathcal{R}^r$, $\boldsymbol{\beta}_T^* \in \mathcal{R}^t$ is a subvector of $\boldsymbol{\beta}^*$, and $T \subset [p]$ is a low-dimensional index set with $|T| = t < n$. This simultaneously tests r linear combinations of $\boldsymbol{\beta}_T^*$, with $r < n$. We next develop a factor-adjusted Wald test.

To construct the test statistic, we first consider a penalized least-square problem,

$$(\hat{\boldsymbol{\gamma}}_a, \hat{\boldsymbol{\beta}}_a) = \underset{(\boldsymbol{\gamma}, \boldsymbol{\beta})}{\text{argmin}} \frac{1}{2n} \sum_{i=1}^n (y_i - \hat{\mathbf{F}}_i' \boldsymbol{\gamma} - \hat{\mathbf{U}}_i' \boldsymbol{\beta})^2 + \lambda_a \sum_{j \notin T} p(|\beta_j|). \tag{10}$$

This is essentially the same as Equation (5), except that, instead of penalizing all variables in $\boldsymbol{\beta}$, we do not penalize β_j for $j \in T$. This is to avoid introducing bias when estimating β_j^* for $j \in T$, which is needed for [Theorem 4](#). A similar idea was also adopted in Shi et al. (2019).

Given $\hat{\boldsymbol{\beta}}_a$, our factor-adjusted Wald test statistic is given by

$$T_w = (\mathbf{A}\hat{\boldsymbol{\beta}}_{a,T} - \mathbf{b})' (\mathbf{A}\hat{\boldsymbol{\Omega}}_T \mathbf{A}')^{-1} (\mathbf{A}\hat{\boldsymbol{\beta}}_{a,T} - \mathbf{b}) / \hat{\sigma}_\epsilon^2,$$

where $\hat{\boldsymbol{\beta}}_{a,T}$ is the subvector of $\hat{\boldsymbol{\beta}}_a$ with indices in T , $\hat{\boldsymbol{\Omega}}_T$ is the first T rows and columns of

$$\hat{\boldsymbol{\Omega}}_{T \cup \hat{S}_a} = n \begin{pmatrix} \hat{\mathbf{U}}_T' \hat{\mathbf{U}}_T & \hat{\mathbf{U}}_T' \hat{\mathbf{U}}_{\hat{S}_a} \\ \hat{\mathbf{U}}_{\hat{S}_a}' \hat{\mathbf{U}}_T & \hat{\mathbf{U}}_{\hat{S}_a}' \hat{\mathbf{U}}_{\hat{S}_a} \end{pmatrix}^{-1},$$

$\hat{S}_a = \{j \in T^c : \hat{\beta}_{a,j} \neq 0\}$, and $\hat{\sigma}_\epsilon^2$ is any consistent estimator of σ_ϵ^2 . In this test statistic, \hat{S}_a plays a critical role in calculating the variance of $\mathbf{A}\hat{\boldsymbol{\beta}}_{a,T} - \mathbf{b}$. In fact, \hat{S}_a needs to be consistent to $S_a = \{j \in T^c : \beta_j^* \neq 0\}$ in order for the variance to be valid. Such a consistency is guaranteed by [Theorem 3](#).

We next present a set of regularity conditions.

Condition 7. Suppose $c \leq \lambda_{\min}\{E(\mathbf{u}^{\otimes 2})\} \leq \lambda_{\max}\{E(\mathbf{u}^{\otimes 2})\} \leq C$ for some positive constants c and C .

Condition 8. Suppose $\|E(\mathbf{u}_{TUS_a}^{\otimes 2})^{-1}\|_{L_\infty} \leq C$.

Condition 9. Suppose $\|E(\mathbf{u}_{(TUS_a)^c} \mathbf{u}'_{TUS_a})\{E(\mathbf{u}_{TUS_a}^{\otimes 2})^{-1}\}\|_{L_\infty} \leq C$.

Condition 10. Suppose $d_n = \min\{|\beta_j^*| : \beta_j^* \neq 0\}/2 \gg \lambda_a \gg \delta_n$, where $\delta_n = \sqrt{(\log p)/n} \vee (n^{1/4}/\sqrt{p_{\min}})$, $p_{\min} = \min_{m \in [M]} p_m$, and $\lambda_a \dot{p}(d_n) = o(\delta_n)$, where \dot{p} is the first derivative.

We first note that **Conditions 7–9** are imposed on \mathbf{u} , instead of on \mathbf{x} . Since \mathbf{u} can be viewed as the residual of \mathbf{x} after the latent factors are removed, the correlations among the variables in \mathbf{u} are much weaker than those in \mathbf{x} . In particular, **Conditions 7 and 8** are needed to avoid singularity of $E(\mathbf{u}^{\otimes 2})$ and $E(\mathbf{u}_{TUS_a}^{\otimes 2})$. **Condition 9** is the well-known irrepresentable condition, which is necessary for establishing the variable selection consistency. Shi et al. (2019) required such a condition to hold for the Gram matrix $X'X$, which essentially requires the correlations among X must be small. This condition hardly holds for multimodal data. By contrast, we only impose such a condition on $E(\mathbf{u}^{\otimes 2})$, which requires the idiosyncratic components not to be highly correlated. This condition is well accepted in the factor model literature. Indeed, when an exact factor model is assumed, $E(\mathbf{u}^{\otimes 2})$ is a diagonal matrix, then **Condition 7** naturally holds.

We now establish the variable selection and estimation consistency of the estimator $\hat{\beta}_a$ in (10), which is essential for deriving the asymptotic distribution of T_w .

Theorem 3. Suppose **Conditions 1–3** and **7–10** hold. Then there exists a solution $(\hat{\gamma}_a, \hat{\beta}_a)$ of (10) such that, with probability tending to 1, the following results hold:

- (a) (sign consistency) $\text{sign}(\hat{\beta}_a) = \text{sign}(\beta^*)$;
- (b) (L_∞ consistency) $\|\hat{\beta}_{a,TUS_a} - \beta_{TUS_a}^*\|_\infty = O_P(\delta_n)$;
- (c) (L_2 consistency) $\|\hat{\beta}_{a,TUS_a} - \beta_{TUS_a}^*\|_2 = O_P(\sqrt{t + s_a} \delta_n)$, where $s_a = |S_a|$;
- (d) (asymptotic expansion) $\sqrt{n}(\hat{\beta}_{a,TUS_a} - \beta_{TUS_a}^*) = n^{-1/2} \mathbf{K}_n^{-1} \mathbf{U}'_{TUS_a} \boldsymbol{\epsilon} + o_P(1)$, where $\mathbf{K}_n = (1/n) \mathbf{U}'_{TUS_a} \mathbf{U}_{TUS_a}$, if we further have that $p_{\min} \gg n^{3/2}$, and $\sqrt{n} \lambda_a \dot{p}(d_n) = o(1)$.

We again explicitly examine the benefit and the extra cost of our factor-based test compared with Shi et al. (2019). The main difference is that we obtain the variable selection and estimation consistency under much weaker conditions than Shi et al. (2019). The price we pay lies in δ_n , which reflects the convergence rates in (b) and (c) of **Theorem 3**. Particularly, the component $n^{1/4}/\sqrt{p_{\min}}$ in δ_n is due to the factor estimation. We consider two scenarios. First, when all data modalities have a large number of variables, i.e. $p_{\min} \gg n^{1/2}$, then $\delta_n = \sqrt{(\log p)/n}$, which makes the convergence rates in (b) and (c) to be minimax optimal. This is because when there are enough variables to estimate the latent factors well, the extra factor estimation error becomes so small that it would not affect the estimation error on $\hat{\beta}_a$. Second, when one modality has only

a small number of variables, i.e. $p_m = o(n^{1/4})$ for some $m \in [M]$, estimating the latent factors in that modality becomes challenging, and the resulting estimation error would slow the convergence of $\hat{\beta}_a$. In this case, one possible alternative solution is to skip factor decomposition for that particular modality, but directly use X_m in Equation (10) and solve

$$(\hat{\gamma}_a, \hat{\beta}_a) = \underset{(\gamma, \beta)}{\text{argmin}} \frac{1}{2n} \sum_{i=1}^n \left(y_i - X'_{i,m} \beta_m - \hat{F}'_{i,-m} \gamma - \hat{U}'_{i,-m} \beta_{-m} \right)^2 + \lambda_a \sum_{j \notin T} p(|\beta_j|),$$

where $X'_{i,m}$, $\hat{F}'_{i,-m}$ and $\hat{U}'_{i,-m}$ denote the i th row of X_{-m} , \hat{F}_{-m} and \hat{U}_{-m} , respectively. Finally, we note that the variable selection and estimation consistency of $\hat{\beta}$ in Equation (5) is directly implied by **Theorem 3** if we treat T as the empty set.

Next, we study the asymptotic distribution of our test statistic T_w , and show that it can be uniformly approximated by a χ^2 -distribution under both H_0 and H_a . We need two more regularity conditions.

Condition 11. Suppose $\|\mathbf{h}_n\|_2 = O(\sqrt{r/n})$, and $\lambda_{\max}\{(\mathbf{A}\mathbf{A}')^{-1}\} \leq C$ for some constant C , where $\mathbf{h}_n = \mathbf{A}\boldsymbol{\beta}_T^* - \mathbf{b}$.

Condition 12. Suppose $r^{1/4} n^{-1/2} E|\mathbf{u}'_{TUS_a} \boldsymbol{\Sigma}_{u,TUS_a}^{-1} \mathbf{u}_{TUS_a}|^{3/2} \rightarrow 0$, where $\boldsymbol{\Sigma}_{u,TUS_a}^{-1}$ is the inverse of the submatrix of $\boldsymbol{\Sigma}_u$ with rows and columns in $T \cup S_a$.

Condition 11 regulates the local alternative \mathbf{h}_n and avoids singularity of $\mathbf{A}\mathbf{A}'$. **Condition 12** is a Lyapunov condition to ensure the asymptotic normality of $\hat{\beta}_{a,TUS_a}$, which is the key to establish the χ^2 -approximation.

Theorem 4. Suppose the conditions of **Theorem 3** and **Conditions 11** and **12** hold, $p_{\min} \gg n^{3/2}$, $\sqrt{n} \lambda_a \dot{p}(d_n) = o(1)$, and $t + s_a = o(n^{1/3})$. Then it holds that

$$\sup_x |\Pr(T_w \leq x) - \Pr\{\chi^2(r, \nu_n) \leq x\}| \rightarrow 0,$$

where $\nu_n = n \mathbf{h}'_n (\mathbf{A}\boldsymbol{\Omega}_T \mathbf{A}')^{-1} \mathbf{h}_n / \sigma_\epsilon^2$, $\boldsymbol{\Omega}_T$ is the submatrix of $\boldsymbol{\Sigma}_{u,TUS_a}^{-1}$ with rows and columns in T .

By **Theorem 4**, we reject $H_0 : \mathbf{A}\boldsymbol{\beta}_T^* = \mathbf{b}$ if $T_w > \chi_\alpha^2(r, 0)$, where $\chi_\alpha^2(r, 0)$ is the α -upper quantile of the χ^2 -distribution with r degrees of freedom. The limiting distribution we establish in **Theorem 4** is the same as the classical Wald test result for a low-dimensional linear regression model (Shi et al. 2019).

We also remark that, the requirement of $p_{\min} \gg n^{3/2}$ in **Theorem 4** ensures that the latent factors in each modality can be well estimated. Therefore, the extra factor estimation error would not affect the limiting distribution of T_w . This condition is more stringent than that of $p_{\min} \gg n^{1/2}$, which guarantees the minimax optimal rate of estimation in **Theorem 3**. This is because hypothesis testing is a more challenging task than estimation.

Finally, write $\|(\mathbf{A}\boldsymbol{\Omega}_T \mathbf{A}')^{-1/2} \mathbf{h}_n\|_2 = Cn^{-\phi_\nu}$ for some constant $C > 0$, and let $\tilde{\mathcal{N}} = \{\boldsymbol{\beta}^* : \|\mathbf{A}\boldsymbol{\beta}_T^* - \mathbf{b}\|_2 = O(\sqrt{r/n}), t + s_a = o(n^{1/3})\}$. **Theorem 4** implies the following corollary

regarding the local power of the proposed test. Its proof is similar to that for Corollary 1 and is omitted.

Corollary 2. Suppose the conditions of Theorem 4 hold. Then,

- (a) $\lim_{n \rightarrow \infty} \sup_{\beta^* \in \tilde{\mathcal{N}}} \sup_{x > 0} |\Pr(T_w \leq x) - \Pr(\chi^2(r, 0) \leq x)| \rightarrow 0$, if $\phi_v > 1/2$;
- (b) $\lim_{n \rightarrow \infty} \sup_{\beta^* \in \tilde{\mathcal{N}}} \sup_{x > 0} |\Pr(T_w \leq x) - \Pr(\chi^2(r, \nu) \leq x)| \rightarrow 0$, if $\phi_v = 1/2$;
- (c) $\liminf_{n \rightarrow \infty} \sup_{\beta^* \in \tilde{\mathcal{N}}} \Pr(T_w > x) = 1$, if $\phi_v < 1/2$;

where $\nu = \lim_{n \rightarrow \infty} \nu_n$ in (b), and (c) holds for any $x > 0$.

6. Quantification of Contribution of a Single Modality

In addition to testing the significance of a whole data modality, it is of equal interest to quantify the amount of contribution of a modality conditioning on other data modalities in the regression model. As an example, in heritability analysis, the goal is to evaluate the contribution of genetic effects to the phenotype in addition to the environmental effects (Lynch and Walsh 1998). Motivated by the proportion of the response variance explained in the classical linear regression, we propose a measure of the contribution of a single data modality in our integrative factor regression model.

Let $\mathbf{x}_{-m} \in \mathcal{R}^{p-p_m}$ denote the subvector of $\mathbf{x} \in \mathcal{R}^p$ excluding $\mathbf{x}_m \in \mathcal{R}^{p_m}$, and $\mathbf{X}_{-m} \in \mathcal{R}^{n \times (p-p_m)}$ denote the submatrix of $\mathbf{X} \in \mathcal{R}^{n \times p}$ excluding $\mathbf{X}_m \in \mathcal{R}^{n \times p_m}$. To evaluate the contribution of \mathbf{x}_m , our key idea is to compare the goodness of fit of regressing y on \mathbf{x} to that of regressing y on \mathbf{x}_{-m} . Toward that end, under model (2), we define

$$\sigma_{m|-m}^2 = \text{var}(\mathbf{x}'_m \beta_m^* | \mathbf{x}_{-m}).$$

Next, we present a proposition regarding $\sigma_{m|-m}^2$, where statements (a) and (b) show in two different ways that $\sigma_{m|-m}^2$ can be interpreted as the improvement of the goodness-of-fit, or equivalently, additional variance of the response explained, given by the m th modality in addition to all other modalities. This justifies why $\sigma_{m|-m}^2$ can be used to quantify the contribution of a single modality. To simplify the presentation, we only consider the case where $p < n$. We then discuss that such an interpretation of $\sigma_{m|-m}^2$ holds for $p > n$ as well. Next, statement (c) shows that, if \mathbf{x}_m and \mathbf{x}_{-m} share some common factors, in that $\mathbf{x}_m = \mathbf{\Lambda}_m \mathbf{f} + \mathbf{u}_m$, and $\mathbf{x}_{-m} = \mathbf{\Lambda}_{-m} \mathbf{f} + \mathbf{u}_{-m}$, where $\mathbf{f} \in \mathcal{R}^K$, we then have a closed-form expression for $\sigma_{m|-m}^2$. This expression holds true regardless of $p < n$ or $p > n$, and thus provides a unified way of computing $\sigma_{m|-m}^2$ in practice.

Proposition 1. Suppose \mathbf{x} follows a multivariate normal distribution and $p < n$. Let $\hat{\mathbf{Y}}$ and $\hat{\mathbf{Y}}_{-m}$ denote the predicted response by regressing y on \mathbf{x} , and regressing y on \mathbf{x}_{-m} , respectively, via least squares. Let $\sigma_y^2 = \text{var}(y)$. Then the following results hold:

- (a) $\sigma_{m|-m}^2 = E \left\| \mathbf{Y} - \hat{\mathbf{Y}}_{-m} \right\|_2^2 / (n - p_{-m}) - \sigma_\epsilon^2$;
- (b) $\sigma_{m|-m}^2 = \sigma_y^2 (r^2 - r_{-m}^2)$, where $r^2 = 1 - E \left\| \mathbf{Y} - \hat{\mathbf{Y}} \right\|_2^2 / \{(n - p)\sigma_y^2\}$, and $r_{-m}^2 = 1 - E \left\| \mathbf{Y} - \hat{\mathbf{Y}}_{-m} \right\|_2^2 / \{(n - p_{-m})\sigma_y^2\}$;
- (c) $\sigma_{m|-m}^2 = \beta_m^{*'} \left\{ \mathbf{\Lambda}_m (\mathbf{I}_K + \mathbf{\Lambda}'_{-m} \mathbf{\Sigma}_{u_{-m}}^{-1} \mathbf{\Lambda}_{-m})^{-1} \mathbf{\Lambda}'_m + \mathbf{\Sigma}_{u_m} \right\} \beta_m^*$.

By Proposition 1(a), when regressing y using all but the m th modality, we have $E \left\| \mathbf{Y} - \hat{\mathbf{Y}}_{-m} \right\|_2^2 = (n - p_{-m})(\sigma_\epsilon^2 + \sigma_{m|-m}^2)$. On the other hand, when regressing y on all data modalities, we have $E \left\| \mathbf{Y} - \hat{\mathbf{Y}} \right\|_2^2 = (n - p)\sigma_\epsilon^2$. Therefore, from a goodness-of-fit perspective, ignoring \mathbf{x}_m leads to a “worsened” prediction by an amount of $\sigma_{m|-m}^2$.

For Proposition 1(b), recall in the classical linear regression model, the adjusted R^2 measures the percentage of total variation in the response that has been explained by the predictors, and is defined as $R^2 = 1 - \{\text{RSS}/(n - p)\} / \{\text{TSS}/(n - 1)\}$, where RSS and TSS are the residual sum of squares and total sum of squares, respectively. Then, r^2 in Proposition 1(b) can be viewed as an “expected” percentage of total variation in the response explained, in that,

$$r^2 = 1 - \frac{E(\text{RSS})/(n - p)}{E(\text{TSS})/(n - 1)} = 1 - \frac{E \left\| \mathbf{Y} - \hat{\mathbf{Y}} \right\|_2^2}{(n - p)\sigma_y^2}.$$

As we show in the proof of Proposition 1, when using all but the m th modality, the “expected” percentage of total variation in the response explained is $r_{-m}^2 = 1 - (\sigma_\epsilon^2 + \sigma_{m|-m}^2) / \sigma_y^2$, where $\sigma_\epsilon^2 = E \left\| \mathbf{Y} - \hat{\mathbf{Y}} \right\|_2^2 / (n - p)$. On the other hand, when using all data modalities, the “expected” percentage of total variation in the response explained is $r^2 = 1 - \sigma_\epsilon^2 / \sigma_y^2$. Therefore, using the m th modality improves the “expected” percentage of total variation in the response explained by an amount of $\sigma_{m|-m}^2 / \sigma_y^2$.

We have so far justified $\sigma_{m|-m}^2$ in the setting where $p < n$. In the setting, where $p > n$ and the true model is sparse, in that there are only s variables associate with y with $s < n$, we can still use $\sigma_{m|-m}^2$ to quantify the contribution of an individual modality. This is because Proposition 1(a) and (b) continue to hold if we replace $\hat{\mathbf{Y}}$ and $\hat{\mathbf{Y}}_{-m}$ with $\tilde{\mathbf{Y}}$ and $\tilde{\mathbf{Y}}_{-m}$, and replace $(n - p)$ with $(n - s)$, where $\tilde{\mathbf{Y}}$ denotes the predicted response by regressing y on the s true variables via least squares, and $\tilde{\mathbf{Y}}_{-m}$ is defined similarly but excluding the variables in the m th modality. In practice, of course, which subset are the true variables is unknown. However, Proposition 1(a) and (b) only provide conceptual justifications of $\sigma_{m|-m}^2$. We always resort to Proposition 1(c) to compute $\sigma_{m|-m}^2$, which holds regardless of $p < n$ or $p > n$. Besides, it does not require the knowledge of the true variables, nor any extra variable selection step to identify them.

Next, we present a plug-in estimator of $\sigma_{m|-m}^2$ given the data. We first apply Bai and Ng (2002) in Equation (4) to the concatenated data matrix $\mathbf{X} = (\mathbf{X}_1, \dots, \mathbf{X}_M) \in \mathcal{R}^{n \times p}$ to estimate the number of shared factors. We then apply PCA to obtain $\hat{\mathbf{F}}$. We next estimate $\mathbf{\Lambda}_m$ by $\hat{\mathbf{\Lambda}}_m = (1/n)\mathbf{X}'_m \hat{\mathbf{F}}$, and obtain $\hat{\mathbf{U}}_m = \mathbf{X}_m - \hat{\mathbf{F}} \hat{\mathbf{\Lambda}}'_m$. We solve for $\hat{\beta}$ following Equation (5). We then apply the thresholding method of Fan, Liao, and Mincheva (2013) to estimate $\mathbf{\Sigma}_u$ by $\hat{\mathbf{\Sigma}}_u$, whose (i, j) th element is $\hat{\sigma}_{u,ij}^2 = s(n^{-1} \sum_{\ell=1}^n \hat{U}_{\ell i} \hat{U}_{\ell j}, \omega)$, $s(x, \omega)$ is a thresholding function, ω is the threshold, and $\hat{\mathbf{U}} = (\hat{\mathbf{U}}_1, \dots, \hat{\mathbf{U}}_M)$. Let $\hat{\beta}_m$ denote the subvector of $\hat{\beta}$ with indices in the m th modality. By Theorem 3, $\hat{\beta}_m$ is a consistent estimator of β_m^* . By Theorem 3.3 of Fan, Liao, and Mincheva (2013), $\hat{\mathbf{\Lambda}}_m$ and $\hat{\mathbf{\Lambda}}_{-m}$ are two consistent

estimators of the loading matrices. In addition, by Theorem 3.1 of Fan, Liao, and Mincheva (2013), $\hat{\Sigma}_u$ is a consistent estimator of Σ_u . Plugging all these estimators into Proposition 1(c) gives a consistent estimator of $\sigma_{m|-m}^2$.

We make two additional remarks about $\sigma_{m|-m}^2$. First, the closed-form expression of $\sigma_{m|-m}^2$ uses the factors commonly shared by \mathbf{x}_m and \mathbf{x}_{-m} . Indeed, such factors determine the correlations between \mathbf{x}_m and \mathbf{x}_{-m} . When no such common factors exist, \mathbf{x}_m and \mathbf{x}_{-m} are uncorrelated. In that case, $\text{var}(\mathbf{x}'_m \boldsymbol{\beta}_m^* | \mathbf{x}_{-m}) = \text{var}(\mathbf{x}'_m \boldsymbol{\beta}_m^*) = \boldsymbol{\beta}_m^{*'} \Sigma_{x_m} \boldsymbol{\beta}_m^* = \boldsymbol{\beta}_m^{*'} \Sigma_{u_m} \boldsymbol{\beta}_m^*$. Therefore, the closed-form expression in Proposition 1(c) can be viewed as a more general form of this special case by taking the correlations between \mathbf{x}_m and \mathbf{x}_{-m} into account. Second, the computation of $\sigma_{m|-m}^2$ only requires to invert a sparse high-dimensional matrix $\Sigma_{u_{-m}}$ and a low-dimensional matrix $\mathbf{I}_K + \Lambda'_{-m} \Sigma_{u_{-m}}^{-1} \Lambda_{-m}$. If an exact factor model is further adopted such that Σ_u becomes a diagonal matrix, $\sigma_{m|-m}^2$ can be easily computed, then as one only needs to invert a low-dimensional matrix. On the contrary, if one does not employ a factor model, then $\text{var}(\mathbf{x}'_m \boldsymbol{\beta}_m^* | \mathbf{x}_{-m}) = \boldsymbol{\beta}_m^{*'} (\Sigma_{x_m} - \Sigma_{x_m, x_{-m}} \Sigma_{x_{-m}}^{-1} \Sigma_{x_{-m}, x_m}) \boldsymbol{\beta}_m^*$, where $\Sigma_{x_m} = E(\mathbf{x}_m \mathbf{x}_m^{\otimes 2})$, $\Sigma_{x_m, x_{-m}} = E(\mathbf{x}_m \mathbf{x}'_{-m})$, $\Sigma_{x_{-m}} = E(\mathbf{x}_{-m} \mathbf{x}_{-m}^{\otimes 2})$, and $\Sigma_{x_{-m}, x_m} = E(\mathbf{x}_{-m} \mathbf{x}'_m)$. Consequently, a large dense matrix $\Sigma_{x_{-m}}$ has to be inverted.

7. Numerical Analysis

7.1. Test of a Whole Modality

We evaluate the empirical performance of the factor-adjusted score test of a whole modality proposed in Section 4. We generate $M = 3$ modalities, and consider two cases of \mathbf{x} . Specifically, for each modality $m = 1, 2, 3$, \mathbf{x}_m are n iid random samples generated from $N_{p_m}(0, \Sigma_m)$. For Case 1, $\Sigma_m = \Lambda_m \Lambda_m' + 0.5 \mathbf{I}_{p_m}$, where each column of $\Lambda_m \in \mathcal{R}^{p_m \times K_m}$ is generated from $N_{p_m}(0, 2)$, and the number of factors $K_m = 2$. For Case 2, the diagonal elements of Σ_m equal 1 and the off-diagonal elements equal 0.4. Accordingly, in Case 1, \mathbf{x}_m indeed follows a factor model setup, and in Case 2, although \mathbf{x}_m does not strictly follow a factor model, its covariance matrix has spiked eigenvalues. In both cases, we aim to test if the first modality \mathbf{x}_1 is significantly associated with the response, i.e. $H_0 : \beta_{11}^* = \dots = \beta_{1p_1}^* = 0$. We then consider two types of alternatives. The first alternative is $H_{A1} : \beta_{11}^* = \dots = \beta_{1p_1}^* = \delta/p$, where δ is a sequence approaching zero. As such, there is a weak signal in each variable of \mathbf{x}_1 and the overall signal is dense. The second alternative is $H_{A2} : \beta_{11}^* = \dots = \beta_{15}^* = \delta/5, \beta_{16}^* = \dots = \beta_{1p_1}^* = 0$. As such, the overall signal is sparse, as all signals come only from the first 5 variables, whereas the rest do not associate with the response. For the other two modalities \mathbf{x}_2 and \mathbf{x}_3 , we set $\beta_{21}^* = 1, \beta_{22}^* = 2, \beta_{23}^* = \dots = \beta_{2p_2}^* = 0$, and $\beta_{31}^* = -1, \beta_{32}^* = -1, \beta_{33}^* = \dots = \beta_{3p_3}^* = 0$. We generate the error ϵ as n iid samples from $N(0, 0.5)$, and generate \mathbf{y} based on model (2). We set $p_1 = p_2 = p_3 = p/3$. We consider two combinations $(n, p) = (100, 600)$, and $(200, 900)$. We compare our test with the score test of Ning and Liu (2017), where the critical values are obtained by bootstrap.

We report the proportion of rejections of H_0 by both tests out of 600 data replications as we vary the value of δ . When

$\delta = 0$, this gives the empirical size, and when $\delta > 0$, it gives the empirical power of the two tests. Figures 1 and 2 report the results for Cases 1 and 2, respectively. In both cases, we see that our proposed test controls the Type I error at the nominal level of $\alpha = 0.05$ when $\delta = 0$. However, the test of Ning and Liu (2017) often yields an inflated size. This may be due to that their multiplier bootstrap method rejects the null hypothesis if the maximum of the decorrelated score functions of variables in that modality is greater than a threshold, and as such, it is easier to reject the null hypothesis. Moreover, our test achieves an as good or often a better power than the test of Ning and Liu (2017) as δ increases.


Next, we simulate data from the three examples as we discussed in Section 4 to further examine the performance of the proposed test. For all three examples, we generate $M = 2$ modalities with $K_m = 1$ factor in each modality, and set $n = 100, p_m = 200$, and $\alpha = 0.05$. We aim to test the significance of the first modality. We generate \mathbf{x}_m as n iid random samples from $N_{p_m}(0, \Sigma_m)$, where $\Sigma_m = \Lambda_m \Lambda_m' + 0.5 \mathbf{I}_{p_m}$ for $m = 1, 2$. For the second modality, we always choose $\Lambda_2 = (1, 1, 1, \dots, 1)'$, and set its coefficients as $\beta_{21}^* = 1, \beta_{22}^* = 2, \beta_{23}^* = \dots = \beta_{2p_2}^* = 0$. For the first modality, we choose different loadings and test different local alternatives. For Example 1, we choose $\Lambda_1 = (20, 1, 1, \dots, 1)'$, and the alternative $H_A : \beta_{11}^* = 0.08, \beta_{12}^* = \dots = \beta_{1p_1}^* = 0$. For Example 2, we choose $\Lambda_1 = (1, 1, 1, \dots, 1)'$, and the alternative $H_A : \beta_{11}^* = \dots = \beta_{18}^* = 0.01, \beta_{19}^* = \dots = \beta_{1p_1}^* = 0$. For Example 3, we choose $\Lambda_1 = (0, 1, 1, \dots, 1)'$, and the alternative $H_A : \beta_{11}^* = 0.2, \beta_{12}^* = \dots = \beta_{1p_1}^* = 0$. Table 1 reports the empirical size and power of the proposed factor-adjusted score test and the score test of Ning and Liu (2017) based on 600 data replications. For Example 1, thanks to the large loading of the first variable, the proposed test achieves a better power than the test of Ning and Liu (2017). For Example 2, the nonzero coefficients are spread out in eight covariates, and they all have loadings on the latent factor. Therefore, thanks to the cumulative effects, our method again achieves a better power to detect the alternative. For Example 3, while the nonzero coefficient β_{11}^* is large, the corresponding loading is zero. Our test thus has no power in detecting such an alternative. These numerical results agree with our discussion in Section 4 on the local power of the proposed test.

7.2. Test of a Linear Combination of Predictors


We next evaluate the empirical performance of the factor-adjusted Wald test of a linear combination of predictors proposed in Section 5. We again generate $M = 3$ modalities and two cases of \mathbf{x} similarly as in Section 7.1, except that in

Table 1. Empirical size and power of testing a whole modality in Examples 1 to 3 for the factor-adjusted score test, and the decorrelated score test of Ning and Liu (2017).


	Factor-adjusted test		Test of Ning and Liu	
	Size	Power	Size	Power
Example 1	0.06	0.50	0.05	0.40
Example 2	0.06	0.60	0.07	0.45
Example 3	0.06	0.07	0.05	0.35




$H_{A1}: (n, p) = (100, 600)$



$H_{A1}: (n, p) = (200, 900)$



$H_{A2}: (n, p) = (100, 600)$




$H_{A2}: (n, p) = (200, 900)$

Figure 1. Empirical size and power of testing a whole modality in Case 1 for the factor-adjusted score test (solid line), and the score test of Ning and Liu (2017) (dashed line).


the second case we increase the off-diagonal elements of Σ_m to 0.8 for $m = 1, 2, 3$. We set $\beta^* = (\beta_1^*, \beta_2^*, \beta_3^*)'$, $\beta_1^* = (-2 + \delta, -1, 0, \dots, 0)'$, $\beta_2^* = (1 + \delta, 2, 0, \dots, 0)'$, $\beta_3^* = (1 + \delta, 1, 0, \dots, 0)'$, and aim to test the linear combination of the first variable in each modality that $H_0 : \beta_{11}^* + \beta_{21}^* + \beta_{31}^* = 0$ versus $H_A : \beta_{11}^* + \beta_{21}^* + \beta_{31}^* \neq 0$. The rest of the simulation setup is the same as that in Section 7.1. Since we only consider testing a single linear combination in this study, we compare our test with both the partially penalized Wald test proposed

in Shi et al. (2019), and the test proposed in Zhu and Bradic (2018).


We again report the proportion of rejections of H_0 out of 600 data replications as we vary the value of δ . Figure 3 reports the results for both Cases 1 and 2. In both cases, we see that our factor-adjusted Wald test and the test of Shi et al. (2019) achieve a good control of the Type I error at the nominal level $\alpha = 0.05$ when $\delta = 0$. But our test achieves a much improved power as δ increases. The main reason is that, in this




$$H_{A1}: (n, p) = (100, 600)$$



$$H_{A1}: (n, p) = (200, 900)$$



$$H_{A2}: (n, p) = (100, 600)$$




$$H_{A2}: (n, p) = (200, 900)$$


Figure 2. Empirical size and power of testing a whole modality in Case 2 for the factor-adjusted score test (solid line), and the score test of Ning and Liu (2017) (dashed line).

example, the variables are highly correlated with each other. The factor adjustment alleviates such high correlations, and yields a better variable selection and estimation of the true regression coefficients, which in turn benefits the inference. Meanwhile, in Case 1, the test of Zhu and Bradic (2018) yields a Type I error that is much larger than the nominal level. This is because, in this case, the variables are driven by the latent factors, which leads to a wide spectrum of the eigenvalues of the covariate covariance matrix. However, Zhu and Bradic (2018) required


the eigenvalues to be bounded. In Case 2, the variables are not generated from a factor model. In this case, the test of Zhu and Bradic (2018) enjoys the best power. Nevertheless, it still suffers from an inflated Type I error, which is around 0.09 and is about twice as large as the nominal level. Moreover, we recall that both our test and the test of Shi et al. (2019) can jointly test multiple linear combinations, while the test of Zhu and Bradic (2018) was designed to test a single linear combination.




Case 1: $(n, p) = (100, 600)$



Case 1: $(n, p) = (200, 900)$



Case 2: $(n, p) = (100, 600)$



Case 2: $(n, p) = (200, 900)$

Figure 3. Empirical size and power of testing a linear combination of predictors for the factor-adjusted Wald test (solid line), the Wald test of Shi et al. (2019) (dashed line), and the test of Zhu and Bradic (2018) (dotted line).

7.3. Multimodal Neuroimaging Analysis

We illustrate our methods with a multimodal neuroimaging analysis to study Alzheimer’s disease (AD). AD is an irreversible neurodegenerative disorder characterized by progressive impairment of cognitive and memory functions. It is the leading form of dementia in elderly subjects, and is the sixth leading cause of death in the United States. In 2018, AD affects over 5.5 million Americans, and without any effective treatment and prevention, this number is projected to almost triple by 2050

(Alzheimer’s Association 2018). Tau is a hallmark pathological protein of AD, and is believed to be part of the driving mechanism of the disorder. It is present in the brains of both AD subjects and the elderly absent of dementia. Brain atrophy is another well known characteristic that differentiates between AD and normal aging. We study a dataset with $n = 125$ subjects. Each subject receives a positron emission tomography (PET) scan with AV-1451 tracer that measures accumulation of tau protein, as well as an anatomical magnetic resonance imaging (MRI)

Table 2. The identified brain regions with the coefficient estimates and the corresponding p -values of the factor-adjusted Wald test for the significance of the brain regions.

Coefficient	$\hat{\beta}_1$	R.rostancng	0.07	L.superiorparietal	0.12	L.inferiorparietal	0	L.middletemporal	0
p -value	$\hat{\beta}_2$		0		0		-0.18		-0.01
			0.05		0.03		0.001		0.1
Coefficient	$\hat{\beta}_1$	L.parahippocampal	0	L.rostancng	0	R.parstriangularis	0	R.superiortemporal	0
p -value	$\hat{\beta}_2$		0.07		0.06		0.01		0.12
			0.02		0.03		0.14		0.03
Coefficient	$\hat{\beta}_1$	R.supramarginal	0	R.temppole	0				
p -value	$\hat{\beta}_2$		0.11		0.04				
			0.02		0.05				

scan that measures brain gray matter cortical thickness. We map both types of images to a common brain atlas from Free Surfer, then summarize each PET and MRI image by a 58-dimensional vector, with each entry measuring the tau accumulation and cortical thickness of a particular brain region-of-interest (ROI), respectively. We remove some regions with quality issues for the PET images, which result in $p_1 = 51$ ROIs for PET, and $p_2 = 58$ ROIs for MRI, for each subject. In our integrative analysis, the tau and cortical thickness measurements form the two modalities \mathbf{x}_1 and \mathbf{x}_2 . Memory score is a critical measure of cognitive decline for AD, and in our analysis, the memory score after removing potential age and sex effects is the response variable y .

We estimate the number of latent factors using the method of Bai and Li (2012), which concludes that there are $\hat{K}_1 = 3$ factors in the tau modality and $\hat{K}_2 = 1$ factor in the cortical thickness modality. We then estimate β^* and γ^* using Equation (5) with a SCAD penalty, where the tuning parameter was chosen by cross-validation. We then apply the three methods we develop in this article. We first test the significance of the entire modality using the factor-adjusted score test in Section 4. The p -values are 4.9×10^{-7} and 1.2×10^{-3} , for testing the significance of tau and cortical thickness modality, respectively. As such, both modalities are clearly significantly associated with the memory outcome. We then report the estimated nonzero coefficients from our integrative factor model and their corresponding brain regions in Table 2. We further carry out the factor-adjusted Wald test in Section 5 to evaluate if the identified regions are significantly correlated with the outcome in either modality. We report the corresponding p -values in Table 2 as well. Our findings agree with the AD literature. For instance, the ROI with the smallest p -value we found is inferior parietal lobe, which is one of brain regions that is known to be associated with progression from healthy aging to AD (Greene and Killiany 2010). Another significant ROI is parahippocampal gyrus, and cortical thinning of this region has been identified as an early biomarker of AD (Echavarri et al. 2011; Krumm et al. 2016). Finally, we evaluate the contribution of each individual modality. If we include the tau modality \mathbf{x}_1 in the model first, and add the cortical thickness modality \mathbf{x}_2 next, then we have $\hat{\sigma}_1^2 = \hat{\text{var}}(\mathbf{x}'_1 \beta_1^*) = 0.11$, and $\hat{\sigma}_{2|1}^2 = \hat{\text{var}}(\mathbf{x}'_2 \beta_2^* | \mathbf{x}_1) = 0.18$. Correspondingly, $\hat{\sigma}_1^2 / \hat{\sigma}_y^2 = 14\%$, and $\hat{\sigma}_{2|1}^2 / \hat{\sigma}_y^2 = 24\%$. In other words, the tau modality explains 14% total variation in the response, and adding the cortical thickness modality explains an additional 24% total variation. On the other hand, if we include the cortical thickness modality \mathbf{x}_2 in the model first, and add the tau modality \mathbf{x}_1 next, we

have $\hat{\sigma}_2^2 = \hat{\text{var}}(\mathbf{x}'_2 \beta_2^*) = 0.19$, and $\hat{\sigma}_{1|2}^2 = \hat{\text{var}}(\mathbf{x}'_1 \beta_1^* | \mathbf{x}_2) = 0.08$. Correspondingly, $\hat{\sigma}_2^2 / \hat{\sigma}_y^2 = 25\%$, and $\hat{\sigma}_{1|2}^2 / \hat{\sigma}_y^2 = 11\%$. In other words, the cortical thickness modality explains 25% total variation in the response, and adding the tau modality explains an additional 11% total variation. We also note that, the explained variation depends on which modality is already included in the model, and the total explained variations do not necessarily match if the two modalities enter the model in different orders.

8. Discussion

In recent years, high-dimensional inference has seen many fruitful results. Particularly, Javanmard and Lee (2020), Cai and Guo (2017), and Zhu and Bradic (2017) have developed a family of flexible debiased methods to test the hypothesis of the form $\beta^* \in \mathcal{C}$, where \mathcal{C} is a general set. By choosing different \mathcal{C} , these methods can handle a wide range of inference problems, and can potentially address the inference questions we target in this paper too. We view the proposed solution and the debiased method as two legitimate and complementary inferential approaches, each with its own strength and limitation. Specifically, the debiased method does not require the beta-min condition that is difficult to check in practice, but instead requires the eigenvalues of $\text{var}(\mathbf{x})$ to be bounded. Besides, it can test the linear combination of the whole p -dimensional parameters β^* , while the existing work so far focuses on a single linear combination. By contrast, our proposal assumes the variables are driven by latent factors that satisfy a pervasive condition, and thus allows the largest eigenvalue of $\text{var}(\mathbf{x})$ to diverge. This may be desirable in the context of multimodal data analysis, as the predictor covariance matrix can have spiked eigenvalues. Nevertheless, our proposal requires the beta-min condition. In addition, we aim at testing multiple linear combinations, but require the number of linear combinations r and the number of involved parameters t in Equation (9) to be smaller than the sample size. In conclusion, we believe our proposal offers a useful solution to address some important inference questions in multimodal data analysis. Meanwhile, developing the debiased counterpart provides an important alternative, and is warranted for future research.

Acknowledgments

The authors thank the editor, the associate editor, and two referees for their constructive comments and suggestions.

Funding

Quefeng Li's research was partially supported by NIH grant R01HL149683. Lexin Li's research was partially supported by NIH grants R01AG061303, R01AG062542, and R01AG034570.

References

- Alzheimer's Association. (2018), "Alzheimer's Disease Facts and Figures," *Alzheimer's & Dementia*, 14, 367–429. [2219]
- Bai, J., and Li, K. (2012), "Statistical Analysis of Factor Models of High Dimension," *The Annals of Statistics*, 40, 436–465. [2209,2220]
- Bai, J., and Ng, S. (2002), "Determining the Number of Factors in Approximate Factor Models," *Econometrica*, 70, 191–221. [2209,2215]
- Chen, J., and Chen, Z. (2008), "Extended Bayesian Information Criteria for Model Selection With Large Model Spaces," *Biometrika*, 95, 759–771. [2210]
- Chernozhukov, V., Chetverikov, D., and Kato, K. (2013), "Gaussian Approximations and Multiplier Bootstrap for Maxima of Sums of High-Dimensional Random Vectors," *The Annals of Statistics*, 41, 2786–2819. [2210]
- Cai, T. T., and Guo, Z. (2017), "Confidence Intervals for High-dimensional Linear Regression: Minimax Rates and Adaptivity," *The Annals of Statistics*, 45, 615–646. [2220]
- Cai, T. T., Ma, Z., and Wu, Y. (2013), "Sparse PCA: Optimal Rates and Adaptive Estimation," *Annals of Statistics*, 41, 3074–3110. [2209]
- Echávარი, C., Aalten, P., Uylings, H. B., Jacobs, H. I., Visser, P. J., Groenenschild, E. H., Verhey, F. R., Burgmans, S. (2011), "Atrophy in the Parahippocampal Gyrus as an Early Biomarker of Alzheimer's Disease," *Brain Structure and Function*, 215, 265–271. [2220]
- Fan, J., Guo, S., and Hao, N. (2012), "Variance Estimation Using Refitted Cross-validation in Ultrahigh Dimensional Regression," *Journal of the Royal Statistical Society, Series B*, 74, 37–65. [2210]
- Fan, J., Ke, Y., and Wang, K. (2016), "Factor-adjusted Regularized Model Selection," arXiv:1612.08490. [2208]
- Fan, J., Liao, Y., and Mincheva, M. (2013), "Large Covariance Estimation by Thresholding Principal Orthogonal Complements," *Journal of the Royal Statistical Society, Series B*, 75, 603–680. [2208,2209,2211,2215,2216]
- Fan, J., and Lv, J. (2011), "Nonconcave Penalized Likelihood With NP-dimensionality," *Information Theory, IEEE Transactions*, 57, 5467–5484. [2210]
- Gaynanova, I., and Li, G. (2019), "Structural Learning and Integrative Decomposition of Multi-view Data," *Biometrics*, 75, 1121–1132. [2207]
- Greene, S. J., and Killiany, R. J. (2011), "Subregions of the Inferior Parietal Lobule Are Affected in the Progression to Alzheimer's Disease," *Neurobiology of Aging*, 31, 1304–1311. [2220]
- Javanmard, A., and Lee, J. D. (2020), "A Flexible Framework for Hypothesis Testing in High Dimensions," *Journal of the Royal Statistical Society, Series B*, 82, 685–718. [2220]
- Javanmard, A., and Montanari, A. (2014), "Confidence Intervals and Hypothesis Testing for High-dimensional Regression," *The Journal of Machine Learning Research*, 15, 2869–2909. [2210]
- Kneip, A., and Sarda, P. (2011), "Factor Models and Variable Selection in High-dimensional Regression Analysis," *The Annals of Statistics*, 39, 2410–2447. [2208]
- Krumm, S., Kivisaari, S.L., Probst, A., Monsch, A. U., Reinhardt, J., Ulmer S., Stippich, C., Kressig, R. W., Taylor, K. I. (2016), "Cortical Thinning of Parahippocampal Subregions in Very Early Alzheimer's Disease," *Neurobiology of Aging*, 38, 188–196. [2220]
- Li, G., and Gaynanova, I. (2018), "A General Framework for Association Analysis of Heterogeneous Data," *The Annals of Applied Statistics*, 12, 1700–1726. [2207]
- Li, G., and Jung, S. (2017), "Incorporating Covariates into Integrated Factor Analysis of Multi-view Data," *Biometrics*, 73, 1433–1442. [2207]
- Li, G., Liu, X., and Chen, K. (2018), "Integrative Multi-view Reduced-rank Regression: Bridging Group-sparse and Low-rank Models," *Biometrics*, 75, 593–602. [2207]
- Li, Q., Cheng, G., Fan, J., and Wang, Y. (2018), "Embracing the Blessing of Dimensionality in Factor Models," *Journal of the American Statistical Association*, 113, 380–389. [2208]
- Li, Y., Wu, F.-X., and Ngom, A. (2016), "A Review on Machine Learning Principles for Multi-view Biological Data Integration," *Briefings in Bioinformatics*, 19, 325–340. [2207]
- 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. [2207,2209]
- Lock, E. F., and Li, G. (2018), "Supervised Multiway Factorization," *Electronic Journal of Statistics*, 12, 1150. [2207]
- Lynch, M., and Walsh, B. (1998), *Genetics and Analysis of Quantitative Traits*, Vol. 1. Sunderland, MA: Sinauer Associates, Inc. [2215]
- Ma, Z. (2013), "Sparse Principal Component Analysis and Iterative Thresholding," *The Annals of Statistics*, 41, 772–801. [2209]
- Ning, Y., and Liu, H. (2017), "A General Theory of Hypothesis Tests and Confidence Regions for Sparse High Dimensional Models," *The Annals of Statistics*, 45, 158–195. [2210,2211,2212,2216,2217,2218]
- Parikh, N., and Boyd, S. (2014), "Proximal Algorithms," *Foundations and Trends® in Optimization*, 1, 127–239. [2210]
- Raskutti, G., Wainwright, M. J., and Yu, B. (2011), "Minimax Rates of Estimation for High-dimensional Linear Regression Over L_q -balls," *IEEE Transactions on Information Theory*, 57, 6976–6994.
- Richardson, S., Tseng, G. C., and Sun, W. (2016), "Statistical Methods in Integrative Genomics," *Annual Review of Statistics and Its Applications*, 3, 181–209. [2207,2210]
- Shen, R., Wang, S., and Mo, Q. (2013), "Sparse Integrative Clustering of Multiple Omics Data Sets," *The Annals of Applied Statistics*, 7, 269–294. [2208]
- Shi, C., Song, R., Chen, Z., and Li, R. (2019), "Linear Hypothesis Testing for High Dimensional Generalized Linear Models," *The Annals of Statistics*, to appear. [2213,2214,2217,2218,2219]
- Shu, H., Wang, X., and Zhu, H. (2020), "D-CCA: A Decomposition-based Canonical Correlation Analysis for High-dimensional Datasets," *Journal of the American Statistical Association*, 115, 292–306. [2207]
- Sun, T., and Zhang, C.-H. (2013), "Sparse Matrix Inversion With Scaled Lasso," *The Journal of Machine Learning Research*, 14, 3385–3418. [2210]
- Tibshirani, R. (1996), "Regression Shrinkage and Selection Via the Lasso," *Journal of the Royal Statistical Society, Series B*, 58, 267–288. [2207]
- Uludag, K., and Roebroeck, A. (2014), "General Overview on the Merits of Multimodal Neuroimaging Data Fusion," *Neuroimage*, 102, 3–10. [2207]
- Van de Geer, S., Bühlmann, P., Ritov, Y., and Dezeure, R. (2014), "On Asymptotically Optimal Confidence Regions and Tests for High-dimensional Models," *The Annals of Statistics*, 42, 1166–1202. [2210,2212]
- Van der Vaart, A.W. (2000). *Asymptotic Statistics*, Vol. 3. Cambridge, UK: Cambridge University Press. [2212]
- Xue, F., and Qu, A. (2020), "Integrating Multi-source Block-wise Missing Data in Model Selection," *Journal of the American Statistical Association*, in press. [2207]
- Yang, Z., and Michailidis, G. (2015), "A Non-negative Matrix Factorization Method for Detecting Modules in Heterogeneous Omics Multi-modal Data," *Bioinformatics*, 32, 1–8. [2207]
- Zhang, C.-H. (2010), "Nearly Unbiased Variable Selection Under Minimax Concave Penalty," *The Annals of Statistics*, 38, 894–942. [2210]
- Zhang, C.-H., and Zhang, S. S. (2014), "Confidence Intervals for Low Dimensional Parameters in High Dimensional Linear Models," *Journal of the Royal Statistical Society, Series B*, 76, 217–242. [2210]
- Zhang, D., Wang, Y., Zhou, L., Yuan, H., Shen, D., and the Alzheimers Disease Neuroimaging Initiative. (2011), "Multimodal Classification of Alzheimer's Disease and Mild Cognitive Impairment," *Neuroimage*, 55, 856–867. [2210]
- Zhang, Y., Tang, N., and Qu, A. (2020), "Imputed Factor Regression for High-dimensional Block-wise Missing Data," *Statistica Sinica*, 30, 631–651. [2208]
- Zhu, Y., and Bradic, J. (2017), "A Projection Pursuit Framework for Testing General High-dimensional Hypothesis," *arXiv preprint 1705.01024*. [2220]
- (2018), "Linear Hypothesis Testing in Dense High-dimensional Linear Models," *Journal of the American Statistical Association*, 113, 1583–1600. [2213,2217,2218,2219]