

Generalized Liquid Association Analysis for Multimodal Data Integration

Lexin Li[†], Jing Zeng[‡], and Xin Zhang[‡]

[†]*University of California at Berkeley* and [‡]*Florida State University*

Abstract

Multimodal data are now prevailing in scientific research. One of the central questions in multimodal integrative analysis is to understand how two data modalities associate and interact with each other given another modality or demographic variables. The problem can be formulated as studying the associations among three sets of random variables, a question that has received relatively less attention in the literature. In this article, we propose a novel generalized liquid association analysis method, which offers a new and unique angle to this important class of problems of studying three-way associations. We extend the notion of liquid association of Li (2002) from the univariate setting to the sparse, multivariate, and high-dimensional setting. We establish a population dimension reduction model, transform the problem to sparse Tucker decomposition of a three-way tensor, and develop a higher-order orthogonal iteration algorithm for parameter estimation. We derive the non-asymptotic error bound and asymptotic consistency of the proposed estimator, while allowing the variable dimensions to be larger than and diverge with the sample size. We demonstrate the efficacy of the method through both simulations and a multimodal neuroimaging application for Alzheimer’s disease research.

Key Words: Liquid association; Multimodal neuroimaging; Sufficient dimension reduction; Tensor analysis; Tucker tensor decomposition.

1 Introduction

1.1 Motivation and problem formulation

Multimodal data are now prevailing in scientific research, where different types of data are acquired for a common set of experimental subjects. One example is multi-omics, where different genetic information such as gene expressions, copy number alternations, and methylation changes are jointly collected for the same biological samples (Richardson et al., 2016). Another example is multimodal neuroimaging, where distinct brain characteristics including brain structure, function, and chemical constituents are simultaneously measured for the same study subjects (Uludag and Roebroek, 2014). Integrative analysis of multimodal data aggregates such diverse and often complementary information, and consolidates knowledge across multiple data modalities.

In this article, we aim to address one of the questions of central interest in multimodal integrative analysis, i.e., to understand how different data modalities associate and interact with each other given other modalities or covariates. This problem is of a broad scientific interest; for instance, it is useful to understand how gene expressions and microRNA levels interact given the severity of ovarian cancer and other demographical variables (Cai et al., 2016), or how gene expressions and comparative genomic hybridization measures interact given the progression of breast cancer and demographics (Mai and Zhang, 2019). Our motivation is a multimodal positron emission tomography (PET) study for Alzheimer’s disease (AD). Amyloid-beta and tau are two hallmark proteins of AD, both of which can be measured *in vivo* by PET imaging using different nuclear tracers. The two proteins are closely associated in terms of spatial patterns of their accumulations, and such association patterns are believed to be affected by the subject’s age (Braak and Braak, 1991). Nevertheless, their specific age-dependent regional associations remain unclear. The data we study involve 81 elderly subjects, each receiving two PET scans that measure the depositions of amyloid-beta and tau, respectively. Each PET modality is represented by a vector of protein

measurements at a set of brain regions of interest, with 60 regions for amyloid-beta, and 26 for tau. Our goal is to find how and where in the brain the associations of the two proteins are the most contrastive as the subject’s age varies.

This problem can be formulated statistically as studying the associations of two sets of random variables $\mathbf{X} \in \mathbb{R}^{p_1}$ and $\mathbf{Y} \in \mathbb{R}^{p_2}$ conditional on the third set of random variables $\mathbf{Z} \in \mathbb{R}^{p_3}$. In our motivation example, \mathbf{X} denotes the amyloid-beta PET imaging with $p_1 = 60$, \mathbf{Y} denotes the tau PET imaging with $p_2 = 26$, and \mathbf{Z} denotes the subject’s age with $p_3 = 1$. Meanwhile, in plenty of multimodal applications, $\mathbf{X}, \mathbf{Y}, \mathbf{Z}$ can all be high-dimensional, and their dimensions can be even larger than the sample size. For instance, in imaging genetics (Nathoo et al., 2017), \mathbf{X}, \mathbf{Y} can represent different imaging modalities, whose dimensions can be in hundreds, and \mathbf{Z} can denote the genetic information, whose dimension can be in tens of thousands or more. In high-dimensional data analysis, it is common to postulate that the data information can be sufficiently captured by some low-dimensional representations, and most often, some linear combinations of the originally high-dimensional variables (Cook, 2007). Adopting this view, our question can be formulated as seeking linear combinations of \mathbf{X} and linear combinations of \mathbf{Y} whose conditional associations given \mathbf{Z} are the most contrastive. In other words, we seek linear combinations of \mathbf{X} and \mathbf{Y} that change the most as \mathbf{Z} changes.

1.2 Related work

There has been a rich statistical literature studying the associations between two sets of multivariate variables \mathbf{X} and \mathbf{Y} . A well studied and commonly used family of methods are canonical correlation analysis (CCA) and its variants (Witten et al., 2009; Gao et al., 2015; Li and Gaynanova, 2018; Shu et al., 2019; Mai and Zhang, 2019, among others). CCA explores the symmetric relations between \mathbf{X} and \mathbf{Y} , and looks for pairs of linear combinations that are most correlated. This goal, however, is different from ours, as the highly correlated linear combinations of \mathbf{X} and \mathbf{Y} are not necessarily the ones that are the

most contrastive. For instance, a pair of linear combinations of \mathbf{X} and \mathbf{Y} can be highly correlated, while this correlation remains a constant as the value of Z varies, and as such they are not the target of our problem. We later numerically compare our method with CCA to further demonstrate their differences. Another popular family of methods are sufficient dimension reduction (SDR), which looks for linear combinations of \mathbf{X} that capture full regression information of \mathbf{Y} given \mathbf{X} ; see Li (2018) for a review of this topic. Later we show that our proposed method is connected to several SDR methods, including principal Hessian directions (Li, 1992; Cook, 1998), and partial and groupwise sufficient dimension reduction (Chiaromonte et al., 2002; Li et al., 2010). However, the goals of the two are utterly different. Whereas SDR studies asymmetric relations of \mathbf{Y} conditioning on \mathbf{X} , we seek symmetric relations between \mathbf{X} and \mathbf{Y} conditioning on the third set of variables Z , in that the roles of \mathbf{X} and \mathbf{Y} are interchangeable, but not with the role of Z .

Compared to the setting of two sets of variables, there have been much fewer statistical methods studying the associations among three sets of multivariate variables in the form of \mathbf{X} and \mathbf{Y} given Z . In his groundbreaking work, Li (2002) proposed a novel three-way interaction metric, termed *liquid association*, that measures the extent to which the association of a pair of random variables depends on the value of a third variable. He showed that this metric is particularly useful in discovering co-expressed gene pairs that are regulated by another gene. However, Li (2002) only considered the scenario where all three variables X, Y, Z are one-dimensional. Li et al. (2004) extended the notion of liquid association to the scenario of a multivariate \mathbf{X} and a scalar Z , and sought two linear combinations $\gamma_1^\top \mathbf{X}$ and $\gamma_2^\top \mathbf{X}$ such that $\text{corr}(\gamma_1^\top \mathbf{X}, \gamma_2^\top \mathbf{X} | Z)$ varies the most with Z . Ho et al. (2011) and Yu (2018) developed some modified versions of liquid association, but still focused on the one-dimensional X, Y, Z scenario. Relatedly, Chen et al. (2011) proposed a bivariate conditional normal model to identify the variables that regulate the co-expressions between two genes. That corresponds to the scenario with a scalar X , a scalar Y and a multivariate Z . Abid et al. (2018) proposed contrastive principal component analysis for a multivariate \mathbf{X}

and a binary scalar Z , which sought linear combinations of \mathbf{X} that have the largest changes in the conditional variance given $Z = 0$ versus $Z = 1$. Moreover, Lock et al. (2013); Li and Jung (2017) developed a class of matrix and tensor factorization methods, which aimed to decompose the multimodal data into the components that capture joint variation shared across modalities, and the components that characterize modality-specific variation. Their goal is again different from ours, as their methods did not target the conditional distribution of \mathbf{X}, \mathbf{Y} given Z . Finally, Xia et al. (2019) analyzed a similar dataset as our motivation example, but tackled a totally different problem. They studied hypothesis testing of covariance between the two multivariate PET measurements, and worked on the residuals after regressing out the age effect, which involves no conditioning of any third set of variables.

1.3 Proposal and contributions

In this article, we study the three-way association among multivariate $\mathbf{X}, \mathbf{Y}, Z$, and seek a set of linear combinations of \mathbf{X} and \mathbf{Y} that has a varying association as Z varies. We generalize the notion of liquid association of Li (2002) from the univariate case to the multivariate case, and develop a population dimension reduction framework for three-way association analysis. Our extension is far from trivial though, resulting in a completely different estimation method and the associated asymptotic theory. For the estimation, we transform the liquid association analysis to the problem of sparse Tucker decomposition of a three-way tensor, and introduce sparsity for the linear combinations to improve the interpretability. We then develop a higher-order orthogonal iteration algorithm for parameter estimation, and establish its algorithmic convergence. For the theory, we establish a population model that is essential for the study of statistical properties. We derive the error bound and consistency, while allowing the variable dimensions p_1, p_2, p_3 to be larger than and to diverge to infinity along with the sample size n . As a result, our proposal makes some useful contributions from both the scientific and statistical perspectives.

Scientifically, characterizing the associations between different modalities given other

modalities or covariates is of crucial importance for multimodal integrative analysis. However, there is almost no existing statistical solution available for this type of problem, especially when the modalities involved are high-dimensional. Our proposal offers a unique angle for this important problem. As an illustration, for our multimodal PET study, understanding the patterns between amyloid-beta and tau with respect to age would offer insight about how pathological proteins of Alzheimer’s disease interact in the aging brains.

Statistically, our proposal of generalized liquid association analysis makes a useful addition to the toolbox of association analysis of more than two sets of variables. Moreover, our method involves sparse tensor decomposition, which is itself of independent interest. Tensor data analysis is gaining increasing attention in recent years (Kolda and Bader, 2009; Zhou et al., 2013; Sun et al., 2017; Bi et al., 2018; Tang et al., 2019; Zhang and Han, 2019; Hao et al., 2020, among others); see also Bi et al. (2021) for a review of tensor analysis in statistics. Nevertheless, our proposal differs in several ways. In particular, our sparse tensor decomposition algorithm is related to some recent singular value decomposition (SVD) type solutions for matrix denoising (Yang et al., 2016) and tensor denoising (Zhang and Han, 2019), in that they share a similar iterative hard thresholding SVD scheme. However, our algorithm is tailored to the tensor parameter estimation with more flexible initialization and tuning. As a result, our theoretical analysis differs considerably from the denoising problems. Whereas both Yang et al. (2016) and Zhang and Han (2019) achieved the minimax optimal estimation for their denoising problems, we establish the dimension reduction subspace recovery consistency, variable selection consistency, and tensor parameter estimation consistency. Our rate of convergence matches the optimal rate in previous works, and all the consistency results are established in the ultrahigh dimensional setting of $s \log(p) = o(n)$, where $s = s_1 s_2 s_3$ and $p = p_1 p_2 p_3$ are the products of the number of nonzero entries and dimensions, respectively, of \mathbf{X} , \mathbf{Y} and \mathbf{Z} . Our theoretical development is highly non-trivial, and may be of independent interest for future research involving tensor parameter estimation in a statistical model with i.i.d. data. In a sense, our work

further broadens the scope of higher-order sparse SVD and tensor analysis.

The rest of the article is organized as follows. Section 2 develops the concept of generalized liquid association, and the corresponding population model of generalized liquid association analysis. Section 3 introduces the estimation algorithm, and Section 4 establishes the theoretical guarantees. Section 5 presents the simulations, and Section 6 revisits the multimodal PET study. Section 7 concludes the paper with a discussion, and the Supplementary Appendix collects all technical proofs and additional results.

2 Generalized Liquid Association Analysis

We first generalize the concept of liquid association from the univariate case to the multivariate case. The conceptual generalization itself is straightforward. Nevertheless, it motivates us to develop a dimension reduction model, along with an optimization formulation, that connects to the problem of tensor decomposition. We show that the solution to this optimization problem is closely related to the low-dimensional representation in the generalized liquid association analysis that we seek. Our method, in a sense, provides a new dimension reduction framework for three-way association analysis.

2.1 Generalized liquid association

We begin with a brief review of the concept of liquid association (LA) proposed by Li (2002) for the univariate case. We then extend this notion to the multivariate case.

Suppose X, Y, Z are random variables with mean zero and variance one. Define $g(z) = E(XY|Z = z) : \mathbb{R} \mapsto \mathbb{R}$. Li (2002) defined the liquid association of X and Y given Z as,

$$\text{LA}(X, Y | Z) = E \left\{ \frac{dg(Z)}{dZ} \right\} = E \left\{ \frac{d}{dZ} E(XY | Z) \right\},$$

When Z follows a standard normal distribution, by Stein's Lemma (Stein, 1981), we have,

$$\text{LA}(X, Y | Z) = E \{g(Z)Z\} = E(XYZ).$$

Intuitively, $\text{LA}(X, Y \mid Z)$ characterizes the change of the association of X and Y conditioning on Z through $g(z)$, and the normality condition connects this quantity with the simple unconditional expectation $\text{E}(XYZ)$. In practice, the univariate Z is transformed to standard normal using the normal score transformation, and the LA measure is estimated by the sample mean $\text{E}(XYZ)$. Li et al. (2004) considered an extension of LA to a multivariate $\mathbf{X} \in \mathbb{R}^{p_1}$ and a scalar Z , by looking for two linear combinations, such that $\text{LA}(\gamma_1^\top \mathbf{X}, \gamma_2^\top \mathbf{X} \mid Z) = \gamma_1^\top \text{E}(\mathbf{X} \mathbf{X}^\top Z) \gamma_2$ is maximized. It has a close-form solution that $\gamma_1 = (\mathbf{v}_1 + \mathbf{v}_p)/\sqrt{2}$, $\gamma_2 = (\mathbf{v}_1 - \mathbf{v}_p)/\sqrt{2}$, where \mathbf{v}_1 and \mathbf{v}_p are the eigenvectors of the matrix $\text{E}(\mathbf{X} \mathbf{X}^\top Z) \in \mathbb{R}^{p \times p}$ with the largest and smallest eigenvalues.

We next extend the concept of liquid association to the multivariate case, where $\mathbf{X} \in \mathbb{R}^{p_1}$, $\mathbf{Y} \in \mathbb{R}^{p_2}$, and $\mathbf{Z} \in \mathbb{R}^{p_3}$. Without loss of generality, suppose each variable entry in \mathbf{X} , \mathbf{Y} , and \mathbf{Z} are standardized with mean zero and variance one. Define

$$\mathbf{g}(z) = \text{E}(\mathbf{X} \mathbf{Y}^\top \mid \mathbf{Z} = z) : \mathbb{R}^{p_3} \mapsto \mathbb{R}^{p_1 \times p_2}.$$

We introduce the generalized liquid association measure, which is a three-way tensor.

Definition 1. *The generalized liquid association (GLA) of \mathbf{X} and \mathbf{Y} with respect to \mathbf{Z} is,*

$$\Phi = \text{GLA}(\mathbf{X}, \mathbf{Y} \mid \mathbf{Z}) = \text{E} \left\{ \frac{d}{d\mathbf{Z}} \mathbf{g}(\mathbf{Z}) \right\} \in \mathbb{R}^{p_1 \times p_2 \times p_3}.$$

When \mathbf{Z} follows a multivariate normal distribution, by the multivariate version of Stein's Lemma (Liu, 1994, Lemma 1), we have the following property regarding Φ .

Proposition 1. *If $\mathbf{Z} \sim \text{Normal}(\mathbf{0}, \Sigma_{\mathbf{Z}})$, then $\Phi = \text{E}(\mathbf{X} \circ \mathbf{Y} \circ \mathbf{Z}) \times_3 \Sigma_{\mathbf{Z}}^{-1}$, where \circ denotes the outer product, and \times_3 denotes the mode-3 product between a tensor and a matrix.*

This conceptual extension from univariate to multivariate variables is straightforward. However, we recognize that all \mathbf{X} , \mathbf{Y} and \mathbf{Z} can be high-dimensional in that $p_1, p_2, p_3 > n$, and the dimension of Φ is $p_1 p_2 p_3$, which can be ultrahigh-dimensional. Besides, it involves

the inversion of a potentially high-dimensional covariance matrix $\Sigma_{\mathbf{Z}}$, which makes a direct calculation or any operation on Φ difficult, if not completely infeasible. Finally, the normality assumption can be restrictive, and there may be no simple way to transform a multivariate \mathbf{Z} to follow an approximate normal distribution. Next, we develop a dimension reduction model for Φ , which reduces the dimensionality, avoids $\Sigma_{\mathbf{Z}}^{-1}$, and improves the interpretability of the result. We also examine the normality assumption carefully, and show that it is *not* absolutely necessary for our generalized liquid association analysis.

2.2 Dimension reduction model for three-way association

We next propose a dimension reduction model for the three-way association analysis. Our goal is to seek the linear combinations of \mathbf{X} and \mathbf{Y} that change the most as \mathbf{Z} or its linear combinations change. Specifically, we first postulate that the matrix $\mathbf{g}(\mathbf{z}) = \mathbb{E}(\mathbf{X}\mathbf{Y}^\top | \mathbf{Z} = \mathbf{z}) \in \mathbb{R}^{p_1 \times p_2}$ varies within a low-dimensional subspace for all values of \mathbf{z} , in that

$$\mathbf{g}(\mathbf{z}) = \mathbb{E}(\mathbf{X}\mathbf{Y}^\top | \mathbf{Z} = \mathbf{z}) = \Gamma_1 \mathbf{f}_z(\mathbf{z}) \Gamma_2^\top,$$

for some semi-orthogonal basis matrices $\Gamma_k \in \mathbb{R}^{p_k \times r_k}$, $k = 1, 2$, and some latent function $\mathbf{f}_z : \mathbb{R}^{p_3} \mapsto \mathbb{R}^{r_1 \times r_2}$. This implies that the linear combinations $\Gamma_1^\top \mathbf{X}$ and $\Gamma_2^\top \mathbf{Y}$ capture all the variations in the first two modes of the generalized liquid association tensor Φ . Next, we further assume that $\mathbf{g}(\mathbf{z})$ depends on \mathbf{Z} only through a few linear combinations $\Gamma_3^\top \mathbf{Z}$ of \mathbf{Z} , for some semi-orthogonal basis matrix $\Gamma_3 \in \mathbb{R}^{p_3 \times r_3}$. Putting these two dimension reduction structures together, we obtain our dimension reduction model for the general three-way association analysis:

$$\mathbf{g}(\mathbf{z}) = \mathbb{E}(\mathbf{X}\mathbf{Y}^\top | \mathbf{Z} = \mathbf{z}) = \Gamma_1 \mathbf{f}(\Gamma_3^\top \mathbf{z}) \Gamma_2^\top, \quad (1)$$

for some latent function $\mathbf{f} : \mathbb{R}^{r_3} \mapsto \mathbb{R}^{r_1 \times r_2}$. This model is to serve as the basis for our subsequent generalized liquid association analysis. Later, we further introduce sparsity to $(\Gamma_1, \Gamma_2, \Gamma_3)$ to improve the interpretability of the linear combinations in model (1).

2.3 Generalized liquid association analysis via tensor decomposition

We propose to estimate the linear combination coefficient Γ_k in the dimension reduction model (1), or more accurately, the subspace $\text{span}(\Gamma_k)$ spanned by the columns of Γ_k , $k = 1, 2, 3$, by solving the following optimization problem,

$$\text{minimize}_{\mathbf{G}_1, \mathbf{G}_2, \mathbf{G}_3} \|\Delta - \Delta \times_1 \mathbf{P}_{\mathbf{G}_1} \times_2 \mathbf{P}_{\mathbf{G}_2} \times_3 \mathbf{P}_{\mathbf{G}_3}\|_F^2, \quad (2)$$

where $\Delta = \text{E}(\mathbf{X} \circ \mathbf{Y} \circ \mathbf{Z}) \in \mathbb{R}^{p_1 \times p_2 \times p_3}$, $\mathbf{G}_k \in \mathbb{R}^{p_k \times r_k}$ is a semi-orthogonal matrix, $\mathbf{P}_{\mathbf{G}_k} = \mathbf{G}_k (\mathbf{G}_k^\top \mathbf{G}_k)^{-1} \mathbf{G}_k^\top$ is the projection onto the subspace $\text{span}(\mathbf{G}_k)$, and \times_k is the mode- k product between a tensor and a matrix, $k = 1, 2, 3$. We first note that the optimization in (2) is actually the well-known tensor Tucker decomposition (Kolda and Bader, 2009). Let $(\mathbf{G}'_1, \mathbf{G}'_2, \mathbf{G}'_3)$ denote the population minimizer of (2). We next carefully study the connections among $(\mathbf{G}'_1, \mathbf{G}'_2, \mathbf{G}'_3)$, the linear combination coefficient $(\Gamma_1, \Gamma_2, \Gamma_3)$ in model (1), and the GLA measure Φ in Definition 1.

Toward that end, we introduce an intermediate optimization problem,

$$\text{minimize}_{\mathbf{G}_1, \mathbf{G}_2, \mathbf{G}_3} \|\Phi - \Phi \times_1 \mathbf{P}_{\mathbf{G}_1} \times_2 \mathbf{P}_{\mathbf{G}_2} \times_3 \mathbf{P}_{\mathbf{G}_3}\|_F^2. \quad (3)$$

Let $(\mathbf{G}''_1, \mathbf{G}''_2, \mathbf{G}''_3)$ denote the population minimizer of (3). Then $(\mathbf{G}''_1, \mathbf{G}''_2, \mathbf{G}''_3)$ is also the solution to the maximization problem,

$$\text{maximize}_{\mathbf{G}_1, \mathbf{G}_2, \mathbf{G}_3} \|\Phi \times_1 \mathbf{P}_{\mathbf{G}_1} \times_2 \mathbf{P}_{\mathbf{G}_2} \times_3 \mathbf{P}_{\mathbf{G}_3}\|_F^2.$$

In other words, solving (3) helps find the linear combinations of \mathbf{X} and \mathbf{Y} whose generalized liquid association given some linear combination of \mathbf{Z} is maximized. In this sense, it achieves our goal of finding the most contrastive associations of \mathbf{X} and \mathbf{Y} given \mathbf{Z} .

The next theorem characterizes the relations among $(\mathbf{G}''_1, \mathbf{G}''_2, \mathbf{G}''_3)$, $(\mathbf{G}'_1, \mathbf{G}'_2, \mathbf{G}'_3)$, and $(\Gamma_1, \Gamma_2, \Gamma_3)$. Basically, it says minimizing (2) and (3) give the same estimates of Γ_1 and Γ_2 in model (1), in that they span the same subspaces. Furthermore, if \mathbf{Z} is normally distributed, then the estimates of Γ_3 under the two minimizations differ by a rotation.

Theorem 1. *Suppose model (1) holds. Then, (a) $\text{span}(\Gamma_k) = \text{span}(\mathbf{G}''_k) = \text{span}(\mathbf{G}'_k)$, for $k = 1, 2$; and (b) $\text{span}(\Gamma_3) = \text{span}(\mathbf{G}''_3) = \Sigma_Z^{-1}\text{span}(\mathbf{G}'_3)$, if \mathbf{Z} is normally distributed.*

Theorem 1 justifies that we can achieve our goal of finding the linear combinations of \mathbf{X} and \mathbf{Y} that are the most contrastive given \mathbf{Z} through the optimization problem (2), with two crucial implications. First, (2) only involves the three-way tensor Δ , but does not require the inversion of the potentially high-dimensional matrix Σ_Z^{-1} as in Φ in (3). Second, and perhaps more importantly, we do not require the normality of \mathbf{Z} . This is because, regardless of the distribution of \mathbf{Z} , the minimizer $(\mathbf{G}'_1, \mathbf{G}'_2)$ from (2) is the same as the minimizer $(\mathbf{G}''_1, \mathbf{G}''_2)$ from (3), and thus they share the same interpretation. Only if we aim to recover Γ_3 , then we need both Σ_Z^{-1} and the normality of \mathbf{Z} . However, we argue that, in our generalized liquid association analysis, our primary goal is to find the linear combinations of \mathbf{X} and \mathbf{Y} that change the most given \mathbf{Z} . As such, we are more interested in the estimation of Γ_1 and Γ_2 , whereas the estimation of Γ_3 provides additional dimension reduction, but is, relatively speaking, of less interest. Our proposed dimension reduction model (1) essentially serves as a bridge that connects the two optimization problems (2) and (3), which in turn connects the Tucker decomposition formulation in (2) with the generalized liquid association measure Φ in Definition 1.

Finally, we remark that, our proposal is similar in spirit to an SDR method, the principal Hessian directions (Li, 1992). It was also derived based on Stein’s Lemma, but was proven useful for finding low-dimensional representations in graphics (Cook, 1998), and for detecting interaction terms in regressions (Tang et al., 2020), even without the normality.

3 Sparse Tensor Estimation

Tucker decomposition is usually solved by a tensor SVD type algorithm, for example, a higher-order orthogonal iteration (HOOI) algorithm, which was first proposed by De Lathauwer et al. (2000), and later studied in statistical models (e.g., Zhang and Xia, 2018; Luo

Algorithm 1 Generalized liquid association analysis via sparse tensor decomposition.

Input: The centered data $\{\mathbf{X}_i \in \mathbb{R}^{p_1}, \mathbf{Y}_i \in \mathbb{R}^{p_2}, \mathbf{Z}_i \in \mathbb{R}^{p_3}, i = 1, \dots, n\}$, the Tucker ranks $r_k \leq p_k$, and the sparsity parameters $(\eta_k, \tilde{\eta}_k), k = 1, 2, 3$.

Step 1, initialization: Compute the sample estimate $\tilde{\Delta} = n^{-1} \sum_{i=1}^n \mathbf{X}_i \circ \mathbf{Y}_i \circ \mathbf{Z}_i$. Obtain the initial active set $\hat{I}_k^{(0)}$, and the initial basis matrices by,

$$\hat{I}_k^{(0)} = \left\{ j : \|(\tilde{\Delta}_k)_{[j,:]} \|_{\infty} \geq \eta_k \right\}, \hat{\Gamma}_k^{(0)} = \text{SVD}_{r_k} \left\{ \mathbf{D}_{\hat{I}_k^{(0)}} \tilde{\Delta}_k \mathbf{D}_{\hat{I}_k^{(0)}} \right\} \in \mathbb{R}^{p_k \times r_k}, k = 1, 2, 3.$$

repeat

Step 2a: Update the active set: $\hat{I}_k^{(t)} = \left\{ j : \|(\tilde{\Delta}_k)_{[j,:]} \hat{\Gamma}_{-k}^{(t)} \|_2^2 \geq \tilde{\eta}_k \right\}, k = 1, 2, 3$.

Step 2b: Perform SVD: $\hat{\Gamma}_k^{(t)} = \text{SVD}_{r_k} \left\{ \mathbf{D}_{\hat{I}_k^{(t)}} \tilde{\Delta}_k \hat{\Gamma}_{-k}^{(t)} \right\} \in \mathbb{R}^{p_k \times r_k}, k = 1, 2, 3$.

until some stopping criterion is met.

Output: The estimated basis matrices $\hat{\Gamma}_k, k = 1, 2, 3$, and $\hat{\Delta} = \tilde{\Delta} \times_1 \mathbf{P}_{\hat{\Gamma}_1} \times_2 \mathbf{P}_{\hat{\Gamma}_2} \times_3 \mathbf{P}_{\hat{\Gamma}_3}$.

et al., 2020). Next, we develop an iterative algorithm to solve (2). We further introduce sparsity in this decomposition to improve the interpretability of the result.

For n i.i.d. data observations $\{\mathbf{X}_i, \mathbf{Y}_i, \mathbf{Z}_i, i = 1, \dots, n\}$, without loss of generality, we assume the data is centered, so that $\sum_{i=1}^n \mathbf{Z}_i = \mathbf{0}$. The centering of \mathbf{X}_i and \mathbf{Y}_i is not required, but for simplicity, we assume $\sum_{i=1}^n \mathbf{X}_i = \sum_{i=1}^n \mathbf{Y}_i = \mathbf{0}$ as well. Then the sample estimator of Δ is simply $\tilde{\Delta} = n^{-1} \sum_{i=1}^n \mathbf{X}_i \circ \mathbf{Y}_i \circ \mathbf{Z}_i \in \mathbb{R}^{p_1 \times p_2 \times p_3}$. Following the dimension reduction model (1) and the optimization problem (2), $\tilde{\Delta}$ admits a Tucker tensor decomposition structure, which can be solved by some version of the higher-order singular value decomposition algorithm. Specifically, we simplify the STAT-SVD algorithm recently proposed by Zhang and Han (2019) for the tensor denoising problem, and tailor it to our generalized liquid association analysis problem to estimate $\Gamma_k, k = 1, 2, 3$. It consists of two major components, SVD of a matrix, and hard thresholding to identify important variables. We summarize the estimation procedure in Algorithm 1, then discuss each step in detail. We further study the algorithmic convergence of the estimation procedure in Section S2 of the Appendix. It is also noted that, in our formulation, we allow the number of variables $p_k, k = 1, 2, 3$, to be much larger than the sample size n .

We first introduce some notation. For a vector $\mathbf{a} = (a_i) \in \mathbb{R}^p$, define the ℓ_2 -norm as $\|\mathbf{a}\|_2 = (\sum_{i=1}^p a_i^2)^{1/2}$, and the ℓ_∞ -norm as $\|\mathbf{a}\|_\infty = \max_{1 \leq i \leq p} |a_i|$. For a matrix $\mathbf{A} = (a_{ij}) \in \mathbb{R}^{p \times q}$, define the Frobenius norm as $\|\mathbf{A}\|_F = (\sum_{i=1}^p \sum_{j=1}^q a_{ij}^2)^{1/2}$. For an integer p , let $[p]$ denote the set $\{1, \dots, p\}$. For the index sets $I \subseteq [p]$, $J \subseteq [q]$, let $\mathbf{A}_{[I, J]}$ denote the corresponding $|I| \times |J|$ submatrix, while the whole index set $[p]$ is simplified as “:”; e.g., $\mathbf{A}_{[[p], J]} = \mathbf{A}_{[:, J]}$. Let $\text{SVD}_r(\mathbf{A}) \in \mathbb{R}^{p \times r}$ denote the left r singular vectors of \mathbf{A} , with $r \leq q$. Let $\tilde{\Delta}_k$ and Δ_k denote the mode- k matricization of the tensors $\tilde{\Delta}$ and Δ , $k = 1, 2, 3$. Define $\Gamma_{-1} = \Gamma_2 \otimes \Gamma_3$, $\Gamma_{-2} = \Gamma_3 \otimes \Gamma_1$, $\Gamma_{-3} = \Gamma_1 \otimes \Gamma_2$, where \otimes is the Kronecker product. Define the active sets of variables in the generalized liquid association analysis as,

$$I_k = \{j : (\Delta_k)_{[j,:]} \neq \mathbf{0}, 1 \leq j \leq p_k\} \subseteq [p_k], \quad k = 1, 2, 3. \quad (4)$$

As an example, the j th variable X_j in \mathbf{X} corresponds to the j th row of $\Gamma_1 \in \mathbb{R}^{p_1 \times r_1}$, $j = 1, \dots, p_1$. Therefore, variable selection in \mathbf{X} translates to the row-wise sparsity in Γ_1 , and correspondingly, the row-wise sparsity in $\Delta_1 \in \mathbb{R}^{p_1 \times (p_2 p_3)}$. Define the diagonal matrix $\mathbf{D}_{I_k} \in \mathbb{R}^{p_k \times p_k}$ that has one on the i th diagonal element if $i \in I_k$ and zero elsewhere. This matrix represents variable selection along each mode, and is used repeatedly in our estimation algorithm. Define $\mathbf{D}_{I_{-1}} = \mathbf{D}_{I_2} \otimes \mathbf{D}_{I_3}$, whereas I_{-1} denotes the pair of subsets I_2 and I_3 . Define $\mathbf{D}_{I_{-2}}, \mathbf{D}_{I_{-3}}, I_{-2}, I_{-3}$ similarly. Also define $\hat{\Gamma}_{-k}, \hat{I}_{-k}$ in a similar fashion.

We start the algorithm by computing the sample estimate $\tilde{\Delta}$, then perform the initial selection of important variables and initial SVD in Step 1 of Algorithm 1. From (4), we see the selection of important variables can be achieved based on $\|(\tilde{\Delta}_k)_{[j,:]} \|$ for some appropriate norm $\|\cdot\|$. In the initialization step, we employ the ℓ_∞ -norm, and achieve the selection by hard thresholding under the sparsity parameter η_k . The two diagonal matrices $\mathbf{D}_{\hat{I}_k^{(0)}}, \mathbf{D}_{\hat{I}_{-k}^{(0)}}$ operate as the subset selection operator within the SVD operator. Depending on the sparsity parameter η_k , we may keep all the variables in the active set, i.e., $\hat{I}_k = [p_k]$.

Next, we iterate the algorithm by repeatedly selecting important variables and performing SVD in Step 2 of Algorithm 1. We continue to do the selection by hard thresholding,

but we use a different norm, i.e., the ℓ_2 -norm rather than the ℓ_∞ -norm, and a different sparsity parameter $\tilde{\eta}_k$. This change of the norm is practically useful because of the following consideration. In the initialization step, the column dimension of $\tilde{\Delta}_k$ is $\prod_{k' \neq k} p_{k'}$ and is often very large, and thus the ℓ_∞ -norm is more effective in screening out the zero rows in $\tilde{\Delta}_k$. During the iterations, the active variable set I_k is selected based on $\tilde{\Delta}_k \Gamma_{-k}$, which has a much smaller column dimension $\prod_{k' \neq k} r_{k'}$. As such, the ℓ_2 -norm is preferred to being able to pick up potentially weaker signals and to refine the selection from the initialization. Moreover, a different thresholding parameter $\tilde{\eta}_k$ during the iterations gives more flexibility.

We alternate Steps 2a and 2b until some termination criterion is met. That is, we terminate the algorithm if the consecutive estimates are close enough, in that the difference between the squared ℓ_2 -norm of the singular values of the two iterations is smaller than 10^{-6} , or if the algorithm reaches the maximum number of iterations, say, 100. In our numerical experiments, we find that the algorithm converges fast, usually within 10 to 20 iterations. We output the estimated basis matrices $\hat{\Gamma}_k$, $k = 1, 2, 3$, along with the updated estimate $\hat{\Delta} = \tilde{\Delta} \times_1 P_{\hat{\Gamma}_1} \times_2 P_{\hat{\Gamma}_2} \times_3 P_{\hat{\Gamma}_3}$ that follows a Tucker decomposition.

Our algorithm is related to the STAT-SVD algorithm of Zhang and Han (2019), in that we both use hard thresholding SVD iteratively. However, Zhang and Han (2019) targeted a tensor denoising problem involving identically distributed normal errors, and used a double thresholding scheme with a theoretical thresholding value. Their algorithm, after obtaining the variance of the errors, became tuning-free in terms of the thresholding parameter. By contrast, we aim to obtain a low-rank tensor estimator in the context of generalized liquid association analysis. The sample estimator does not have i.i.d. entries, and we use a single thresholding scheme with two data-driven tuning parameters. This leads to a more flexible tuning, and consequently an utterly different approach for the asymptotic analysis.

The thresholding values η_k and $\tilde{\eta}_k$ are treated as tuning parameters, and we propose a prediction-based approach for tuning. That is, we first randomly split the data into a training set and a testing set, and obtain the sample estimates $\tilde{\Delta}^{\text{train}}$ and $\tilde{\Delta}^{\text{test}}$ separately.

We then apply Algorithm 1 to $\tilde{\Delta}^{\text{train}}$ to obtain $\hat{\Gamma}_k^{\text{train}}(\eta)$, $k = 1, 2, 3$, under a given set of tuning parameters $\eta = (\eta_1, \eta_2, \eta_3, \tilde{\eta}_1, \tilde{\eta}_2, \tilde{\eta}_3)$. We choose η to minimize the discrepancy,

$$L(\eta) = \left\| \tilde{\Delta}^{\text{test}} - \tilde{\Delta}^{\text{test}} \times_1 \mathbf{P}_{\hat{\Gamma}_1^{\text{train}}(\eta)} \times_2 \mathbf{P}_{\hat{\Gamma}_2^{\text{train}}(\eta)} \times_3 \mathbf{P}_{\hat{\Gamma}_3^{\text{train}}(\eta)} \right\|_F. \quad (5)$$

Meanwhile, the ranks (r_1, r_2, r_3) take some pre-specified values. In practice, r_k is often taken as 1 or 2 for exploratory analysis and data visualization. This is similar in spirit as canonical correlation analysis. Actually, rank selection is still an open and active topic in CCA and tensor problems, and we leave a full treatment of rank selection as future research.

4 Theoretical Properties

We establish the theoretical guarantees for the estimated Tucker tensor $\hat{\Delta} = \tilde{\Delta} \times_1 \mathbf{P}_{\hat{\Gamma}_1} \times_2 \mathbf{P}_{\hat{\Gamma}_2} \times_3 \mathbf{P}_{\hat{\Gamma}_3}$, the estimated subspace basis matrices $\hat{\Gamma}_k$, and the estimated active sets $\hat{I}_k^{(t)}$, $k = 1, 2, 3$, from Algorithm 1. In our theoretical analysis, we allow both the tensor dimension $p = \prod_{k=1}^3 p_k$ and the sparsity level $s = \prod_{k=1}^3 s_k$ to diverge with the sample size n , while we fix the tensor rank $r = \prod_{k=1}^3 r_k$. We begin with two mild regularity conditions.

- (A1) Suppose $|X_{j_1} Y_{j_2}| \leq M$, for some constant $M > 0$, $j_1 = 1, \dots, p_1, j_2 = 1, \dots, p_2$, and suppose Z_{j_3} is sub-Gaussian with the parameter $\sigma^2 > 0$, $j_3 = 1, \dots, p_3$.
- (A2) Suppose $\lambda \geq \max\{C_1 \sqrt{s \log p/n}, C_2\}$, for some constants $C_1, C_2 > 0$, where $\lambda \equiv \min\{\lambda_1, \lambda_2, \lambda_3\}$, and λ_k is the smallest non-zero singular value of Δ_k , $k = 1, 2, 3$.

Assumption (A1) requires \mathbf{X} and \mathbf{Y} to be bounded, and \mathbf{Z} to be sub-Gaussian, which are necessary to establish the concentration of each element in $\tilde{\Delta}$ to its population counterpart. The sub-Gaussian assumption is weaker than the normality assumption, and is widely used in high-dimensional non-asymptotic analysis (see, e.g., Wainwright, 2019). Besides, it assumes each individual Z_k to be sub-Gaussian, which is weaker than assuming the joint distribution \mathbf{Z} is sub-Gaussian. The constant σ^2 does not require all Z_{j_3} to have the same variance. If Z_{j_3} is sub-Gaussian with the parameter $\sigma_{j_3}^2$, $j_3 = 1, \dots, p_3$, then we can set

$\sigma^2 = \max_{j_3} \sigma_{j_3}^2$. Assumption (A2) ensures that there is a reasonable gap between the zero and nonzero eigenvalues in Δ_k , under which the consistency for the estimator $\widehat{\Gamma}_k$ is ensured. This type of assumption on the eigenvalues is frequently used in high-dimensional singular value decomposition (Yu et al., 2015; Yang et al., 2016; Zhang and Han, 2019).

Next, we derive the non-asymptotic error bound and variable selection property of our estimators. Let $\widehat{\Delta}$, $\widehat{\Gamma}_k$, and \widehat{I}_k , $k = 1, 2, 3$, denote the estimators and the corresponding active sets returned from Algorithm 1 after t_{\max} iterations, under the theoretical thresholding values $\eta_k = \sqrt{\alpha \log p/n}$, and $\tilde{\eta}_k = \alpha s_{-k} \log p/n$, where $\alpha = 513(M + \sigma)^4$, and M, σ are as defined in Assumption (A1). Moreover, since the basis matrix $\widehat{\Gamma}_k$ is identifiable only up to an orthogonal rotation, we characterize its bound in terms of the projection matrix $\mathbf{P}_{\widehat{\Gamma}_k}$.

Theorem 2 (Non-asymptotic properties). *Suppose Assumptions (A1), (A2), and model (1) hold. Then, with probability at least $1 - C \max \left[p^{-\gamma}, p^{-\left\{ \sqrt{n(\gamma+1)/(2 \log p)} - 1 \right\}} \right]$, (a) $\|\widehat{\Delta} - \Delta\|_F \leq (c_1 + c_2 2^{-t_{\max}}) \sqrt{s \log p/n}$; (b) $\|\mathbf{P}_{\widehat{\Gamma}_k} - \mathbf{P}_{\Gamma_k}\|_F \leq (c_3 + c_4 2^{-t_{\max}}) \sqrt{s \log p/n}$, $k = 1, 2$, and $\|\mathbf{P}_{\widehat{\Gamma}_3} - \mathbf{P}_{G'_3}\|_F \leq (c_3 + c_4 2^{-t_{\max}}) \sqrt{s \log p/n}$; and (c) $\widehat{I}_k \subseteq I_k$, $k = 1, 2, 3$, where $\gamma, C, c_1, c_2, c_3, c_4$ are some positive constants.*

We make a few remarks. First, statements (a) and (b) establish the non-asymptotic error bound for the Tucker tensor estimator $\widehat{\Delta}$, as well as the subspace spanned by the basis matrix $\widehat{\Gamma}_k$, $k = 1, 2, 3$. Note that $\widehat{\Gamma}_1, \widehat{\Gamma}_2$ directly target $\mathbf{G}'_1, \mathbf{G}'_2$ from optimization (2), which, by Theorem 1(a), are the same as Γ_1, Γ_2 in our dimension reduction model (1), as well as $\mathbf{G}''_1, \mathbf{G}''_2$ from the generalized liquid association measure Φ , in the sense that they span the same subspaces. Meanwhile, $\widehat{\Gamma}_3$ targets the population minimizer \mathbf{G}'_3 from (2), which, by Theorem 1(b), differs from Γ_3 in model (1) and \mathbf{G}''_3 from Φ by a rotation, when \mathbf{Z} is normally distributed. However, as we have discussed after Theorem 1, our primary interest is to recover Γ_1, Γ_2 , rather than Γ_3 . As such, we do *not* require the normality assumption for Theorem 2. Second, the error bounds in statements (a) and (b) are functions of the maximum number of iterations t_{\max} , and they decrease when t_{\max} increases. As

such, our estimators are to have improved accuracy with more iterations. Third, statement (c) shows that our method avoids selecting inactive variables with a high probability. This result is similar to that in Zhang and Han (2019), and can be viewed as a weaker version of variable selection consistency when compared to Theorem 3 below. Fourth, we treat t_{\max} as a constant in this section, regardless of the tensor dimension or sample size. This is because the algorithm converges fast, often within 10 to 20 iterations. On the other hand, we can easily extend the results by allowing t_{\max} to diverge. For instance, parallel to Zhang and Han (2019), we can let t_{\max} diverge at the rate of $o(p)$. Finally, when the sample size n is sufficiently large, i.e., when $n \gg \log p$, the statements in Theorem 2 hold with probability at least $1 - Cp^{-\gamma}$. Besides, the constants c_1, \dots, c_4 depend on the constants M and σ in Assumption (A1), and their explicit forms are given in Section S5.3 of the Appendix.

We also briefly comment that, in real applications, the modalities may be low-dimensional or have no sparsity. Accordingly, we can modify Algorithm 1 to a non-sparse version, by setting $\eta_k = 0$ in Step 1, and $\tilde{\eta}_k = 0$ in Step 2, for $k = 1, 2, 3$. We give the corresponding non-asymptotic error bounds in Section S3 of the Appendix.

Next, we establish the asymptotic consistency of the tensor parameter estimation, subspace estimation and variable selection as n, p, s diverge to infinity. We allow $s \log p = o(n)$, i.e., each tensor dimension p_k can diverge faster than the sample size n .

Corollary 1 (Asymptotic consistency). *Suppose the conditions in Theorem 2 hold, and $s \log p = o(n)$. Then, as $n, p, s \rightarrow \infty$, with probability tending to one, (a) $\|\widehat{\Delta} - \Delta\|_F \rightarrow 0$; (b) $\|\mathbf{P}_{\widehat{\Gamma}_k} - \mathbf{P}_{\Gamma_k}\|_F \rightarrow 0$, $k = 1, 2$, and $\|\mathbf{P}_{\widehat{\Gamma}_3} - \mathbf{P}_{\mathcal{G}'_3}\|_F \rightarrow 0$; and (c) $\widehat{I}_k \subseteq I_k$, $k = 1, 2, 3$.*

While Theorem 2(c) shows that our method can exclude the inactive variables from the selection, we show in the next theorem that our method can exactly recover the active variables, with a high probability. We need an additional regularity condition.

(A3) Suppose $\delta_{\min} \geq C_3 \sqrt{s \log p/n}$, for some sufficiently large constant C_3 , where $\delta_{\min} \equiv \min_{k \in \{1, 2, 3\}, i \in I_k} \|(\Delta_k)_{[i, :]}\|_2$ denotes the minimal signal strength.

Theorem 3 (Variable selection consistency). *Suppose Assumptions (A1) to (A3), and model (1) hold. Then, with probability at least $1 - C' \max \left[p^{-\gamma}, p^{-\left\{ \sqrt{n(\gamma+1)/(2 \log p)} - 1 \right\}} \right]$, we have, $\widehat{I}_k = I_k$, for $k = 1, 2, 3$, where γ, C' are some positive constants.*

Assumption (A3) ensures the signal of the active variables is of a reasonable strength when n, p, s diverge, which in turn leads to the variable selection consistency in Theorem 3. Note that δ_{\min} is also the minimal Frobenius norm of the non-zero slices in Δ , i.e., the slices of Δ corresponding to those variables $i \in I_k, k = 1, 2, 3$. We feel this assumption is reasonable. Actually, if we allow s, p to diverge with n at the rate of $s \log p = o(n)$, and suppose the nonzero entries of Δ are bounded away from zero, then this assumption is satisfied.

5 Simulation Studies

5.1 Simulation setup

We carry out the simulations to investigate the empirical performance of the proposed generalized liquid association analysis (GLAA) method. We consider three scenarios. In the first scenario, we fix the dimension of \mathbf{Z} at $p_3 = 1$, and increase the dimensions of \mathbf{X} and \mathbf{Y} as $p_1 = p_2 = \{100, 200, 300, 400, 500\}$. In the second scenario, we fix $p_1 = p_2 = 100$, and increase $p_3 = \{20, 40, 60, 80, 100\}$. In both cases, we fix the sample size at $n = 500$. In the third scenario, we fix $p_1 = 100, p_2 = 25, p_3 = 1$, and increase the sample size $n = \{60, 80, 100, 120, 160\}$. We generate the data in the following way. For $i = 1, \dots, n$, we first generate \mathbf{Z}_i from a normal distribution with mean zero and covariance \mathbf{I}_{p_3} . We then generate $(\mathbf{X}_i, \mathbf{Y}_i)$ jointly from a normal distribution with mean zero and covariance,

$$\text{cov}(\mathbf{X}, \mathbf{Y} \mid \mathbf{Z} = \mathbf{Z}_i) = \begin{pmatrix} \Sigma_{\mathbf{X}} & \Gamma_1 \mathbf{f}(\Gamma_3^\top \mathbf{Z}_i) \Gamma_2^\top \\ \Gamma_2 \mathbf{f}^\top(\Gamma_3^\top \mathbf{Z}_i) \Gamma_1^\top & \Sigma_{\mathbf{Y}} \end{pmatrix}.$$

To ensure the positive-definiteness of this covariance matrix, we set $\Gamma_1 = \Sigma_{\mathbf{X}}^{1/2}(\mathbf{O}_1^\top, \mathbf{0})^\top$ and $\Gamma_2 = \Sigma_{\mathbf{Y}}^{1/2}(\mathbf{O}_2^\top, \mathbf{0})^\top$, where $\mathbf{O}_1 = \mathbf{O}_2 \in \mathbb{R}^{5 \times 2}$ with the first column being $(1, 1, 1, 1, 1)/\sqrt{5}$, and the second column being $(0, 0, 0, -1, 1)/\sqrt{2}$. As a result, in this example, for \mathbf{X}

and \mathbf{Y} , the ranks are $r_1 = r_2 = 2$, and the sparsity levels are $s_1 = s_2 = 5$. The marginal covariance matrix $\Sigma_{\mathbf{X}}$ is set as a block diagonal matrix, $\Sigma_{\mathbf{X}} = \text{bdiag}(\Sigma_{\mathbf{X},1}, \Sigma_{\mathbf{X},2})$, where $\Sigma_{\mathbf{X},1} \in \mathbb{R}^{s_1 \times s_1}$ corresponds to the active variables in \mathbf{X} and takes the form of an AR structure such that its (i, j) th entry equals $\sigma_{ij} = 0.3^{|i-j|}$, $i, j = 1, \dots, s_1$, and $\Sigma_{\mathbf{X},2} \in \mathbb{R}^{(p_1-s_1) \times (p_1-s_1)}$ is the identity matrix. The marginal covariance matrix $\Sigma_{\mathbf{Y}}$ is constructed in a similar fashion. The matrix $\mathbf{f}(\Gamma_3^\top \mathbf{Z}_i)$ is set as $\text{diag}\{f_1(\Gamma_3^\top \mathbf{Z}_i), f_2(\Gamma_3^\top \mathbf{Z}_i)\}$, where $f_1(a) = 0.95\text{sign}(a)$ and $f_2(a) = 0.85\text{sign}(a)$. In the Appendix, we consider additional simulations using $f(a; \rho, \xi) = \rho\{2/(1 + e^{-2\xi a}) - 1\}$, with different parameters $0 < \rho \leq 1$ and $\xi > 0$ that control the magnitude and speed of changes in $\text{cov}(\mathbf{X}, \mathbf{Y} | \mathbf{Z})$. For the first and the third scenarios, $p_3 = 1$ and thus $\Gamma_3 = 1$. For the second scenario, where p_3 varies from 20 to 100, we set $\Gamma_3 = (1, 1, 1, 1, 1, 0, \dots, 0)$, with $s_3 = 5$ and $r_3 = 1$. When applying the proposed method, we adopt the theoretical forms for the tuning parameters, i.e., $\eta_k = \sqrt{\zeta \log p/n}$, and $\tilde{\eta}_k = \zeta s_{-k} \log p/n$, and tune ζ following the approach in (5).

There is no existing method designed to directly address our targeting problem. For the purpose of comparison, we consider three relevant solutions. The first solution we consider is a naive and marginal extension of the univariate liquid association (ULA) method from Li (2002). That is, we construct a tensor estimator $\tilde{\Phi}$, each entry of which is the sample univariate LA for the triplet $(X_{j_1}, Y_{j_2}, Z_{j_3})$ as defined in Li (2002), $j_1 = 1, \dots, p_1, j_2 = 1, \dots, p_2, j_3 = 1, \dots, p_3$. We then perform the usual SVD to each mode- k matricization of $\tilde{\Phi}$, denoted by $\tilde{\Phi}_{(k)}$, under the given rank to obtain the estimates of basis matrices; i.e., $\text{SVD}_{r_k}\{\tilde{\Phi}_{(k)}\}$, $k = 1, 2, 3$. The second and third solutions we consider are two different versions of canonical correlation analysis, the penalized matrix decomposition (PMD) method of Witten et al. (2009), and the sparse canonical correlation analysis (SCCA) method of Mai and Zhang (2019). We have chosen these two versions due to their computational simplicity and superior empirical performance. We note that CCA is not designed to incorporate the third set of variables \mathbf{Z} . We thus simply take the first r_1 and r_2 directions of \mathbf{X} and \mathbf{Y} from CCA as the estimated basis matrices corresponding to Γ_1

and Γ_2 . We evaluate the performance of each method in terms of the variable selection accuracy and the subspace estimation accuracy.

For variable selection, we record the true positive rate (TPR) and false positive rate (FPR) for each mode. Recall from (4), the active set of variables is I_k , which is also the index set of nonzero rows in Γ_k . Let \hat{I}_k denote the estimated active set corresponding to $\hat{\Gamma}_k$, then $\text{TPR-}k = |I_k \cap \hat{I}_k|/s_k$ and $\text{FPR-}k = |I_k^c \cap \hat{I}_k|/(p_k - s_k)$, $k = 1, 2, 3$. For GLAA, PMD and SCCA, we estimate the active set as $\hat{I}_k = \{i : \text{there exist non-zero elements in the } i\text{th row of } \hat{\Gamma}_k\}$. For ULA, it does not perform any variable selection. For the purpose of comparison, we simply calculate the ℓ_2 -norm of each row for the k th matricization $\tilde{\Phi}_{(k)}$, arrange the row indices in a descending order by the ℓ_2 -norms, then select the first s_k rows for each mode, $k = 1, 2, 3$. Of course, the information about s_k is generally unknown in practice, and this solution utilizes such knowledge. Even so, as we show later, ULA is still far less effective compared to the proposed GLAA method.

For subspace estimation, we compute the average distance between the true and the estimated subspaces, $D = \sum_{k=1}^{\tilde{k}} D(\Gamma_k, \hat{\Gamma}_k)/\tilde{k}$, where $D(\Gamma_k, \hat{\Gamma}_k) = \|\mathbf{P}_{\Gamma_k} - \mathbf{P}_{\hat{\Gamma}_k}\|_F/\sqrt{2r_k}$. Note that, since \mathbf{Z} is normal with an identity matrix in this example, $\hat{\Gamma}_3$ is estimating $\text{span}(\Gamma_3) = \text{span}(\mathbf{G}'_3) = \text{span}(\mathbf{G}''_3)$. For GLAA and ULA, this distance measure is averaged over all three modes of $\mathbf{X}, \mathbf{Y}, \mathbf{Z}$, so $\tilde{k} = 3$. For PMD and SCCA, it is averaged over the first two modes \mathbf{X}, \mathbf{Y} , so $\tilde{k} = 2$. By definition, this distance measure is between 0 and 1, where 0 indicates a perfect recovery.

5.2 Simulation results

Tables 1–3 summarize the simulation results over 100 replications for the three scenarios.

Table 1 reports the accuracy of variable selection and subspace estimation for Scenario 1 when the dimension $p_1 = p_2$ of \mathbf{X} and \mathbf{Y} increases. It is clearly seen that GLAA dominates all the competing solutions. For ULA, even with the oracle knowledge of the true sparsity level, the naive variable selection of ULA still performs worse, since it only uti-

p_1, p_2	Method	TPR-1	FPR-1	TPR-2	FPR-2	D
100	GLAA	1.000 (0.000)	0.000 (0.000)	1.000 (0.000)	0.000 (0.000)	0.095 (0.002)
	ULA	0.998 (0.002)	0.000 (0.000)	0.996 (0.003)	0.000 (0.000)	0.776 (0.002)
	PMD	0.804 (0.028)	0.735 (0.029)	0.802 (0.029)	0.731 (0.028)	0.971 (0.002)
	SCCA	0.584 (0.030)	0.626 (0.027)	0.634 (0.030)	0.629 (0.027)	0.989 (0.001)
200	GLAA	1.000 (0.000)	0.000 (0.000)	1.000 (0.000)	0.000 (0.000)	0.100 (0.003)
	ULA	0.928 (0.010)	0.002 (0.000)	0.944 (0.009)	0.001 (0.000)	0.873 (0.002)
	PMD	0.766 (0.033)	0.731 (0.029)	0.762 (0.033)	0.727 (0.029)	0.989 (0.001)
	SCCA	0.596 (0.031)	0.611 (0.027)	0.626 (0.035)	0.611 (0.027)	0.993 (0.001)
300	GLAA	1.000 (0.000)	0.000 (0.000)	1.000 (0.000)	0.000 (0.000)	0.100 (0.003)
	ULA	0.808 (0.016)	0.003 (0.000)	0.806 (0.016)	0.003 (0.000)	0.945 (0.003)
	PMD	0.718 (0.032)	0.693 (0.030)	0.730 (0.034)	0.696 (0.029)	0.992 (0.001)
	SCCA	0.670 (0.028)	0.651 (0.019)	0.624 (0.028)	0.651 (0.018)	0.997 (0.000)
400	GLAA	0.932 (0.024)	0.008 (0.004)	0.930 (0.025)	0.008 (0.004)	0.186 (0.025)
	ULA	0.652 (0.018)	0.004 (0.000)	0.678 (0.019)	0.004 (0.000)	0.980 (0.002)
	PMD	0.798 (0.029)	0.766 (0.026)	0.800 (0.029)	0.762 (0.026)	0.994 (0.001)
	SCCA	0.594 (0.022)	0.601 (0.010)	0.590 (0.024)	0.602 (0.010)	0.997 (0.000)
500	GLAA	0.848 (0.032)	0.041 (0.010)	0.848 (0.033)	0.040 (0.009)	0.304 (0.037)
	ULA	0.526 (0.021)	0.005 (0.000)	0.526 (0.020)	0.005 (0.000)	0.989 (0.001)
	PMD	0.788 (0.031)	0.727 (0.027)	0.794 (0.030)	0.729 (0.028)	0.995 (0.001)
	SCCA	0.532 (0.024)	0.518 (0.008)	0.532 (0.024)	0.516 (0.008)	0.997 (0.000)

Table 1: Simulation results for Scenario 1 where $p_1 = p_2$ varies. The reported are the average TPR and FPR for the variable selection accuracy, and D for the subspace estimation accuracy, with the standard errors in the parenthesis. The results are over 100 replications.

lizes the marginal information of each mode. Besides, the estimated subspace is distant away from the true subspace. Moreover, as p_1, p_2 increase, the performance of ULA degrades fast, while GLAA remains competitive. For PMD and SCCA, both suffer large false positive rates in selection, while the estimated subspaces are almost orthogonal to the true subspace with D being almost one. This is because, by design, neither method takes into account the conditioning variable \mathbf{Z} when studying the association between \mathbf{X} and \mathbf{Y} .

Table 2 reports the results for Scenario 2 when the dimension p_3 of \mathbf{Z} increases. In this case, our goal is to estimate Γ_1, Γ_2 accurately, meanwhile select the variables in \mathbf{X} and \mathbf{Y} accurately. It is seen that GLAA outperforms all other methods considerably. Besides, it shows a competitive performance of GLAA even with a relatively large dimension of \mathbf{Z} . This complements our real data example where the dimension of \mathbf{Z} is one.

p_3	Method	TPR-1	FPR-1	TPR-2	FPR-2	D
20	GLAA	1.000 (0.000)	0.000 (0.000)	1.000 (0.000)	0.000 (0.000)	0.132 (0.004)
	ULA	0.642 (0.019)	0.019 (0.001)	0.638 (0.019)	0.019 (0.001)	0.872 (0.003)
	PMD	0.774 (0.031)	0.733 (0.029)	0.776 (0.031)	0.732 (0.029)	0.978 (0.002)
	SCCA	0.518 (0.034)	0.534 (0.029)	0.536 (0.032)	0.532 (0.030)	0.989 (0.001)
40	GLAA	1.000 (0.000)	0.000 (0.000)	1.000 (0.000)	0.000 (0.000)	0.131 (0.004)
	ULA	0.488 (0.020)	0.027 (0.001)	0.498 (0.020)	0.026 (0.001)	0.888 (0.002)
	PMD	0.774 (0.031)	0.695 (0.030)	0.772 (0.032)	0.695 (0.029)	0.973 (0.002)
	SCCA	0.590 (0.034)	0.596 (0.031)	0.624 (0.033)	0.589 (0.031)	0.987 (0.001)
60	GLAA	1.000 (0.000)	0.000 (0.000)	1.000 (0.000)	0.000 (0.000)	0.136 (0.004)
	ULA	0.384 (0.020)	0.032 (0.001)	0.404 (0.020)	0.031 (0.001)	0.898 (0.002)
	PMD	0.748 (0.030)	0.694 (0.030)	0.754 (0.033)	0.701 (0.030)	0.973 (0.002)
	SCCA	0.564 (0.031)	0.537 (0.029)	0.572 (0.031)	0.541 (0.029)	0.985 (0.002)
80	GLAA	0.998 (0.002)	0.000 (0.000)	0.998 (0.002)	0.000 (0.000)	0.145 (0.005)
	ULA	0.378 (0.021)	0.033 (0.001)	0.368 (0.019)	0.033 (0.001)	0.903 (0.001)
	PMD	0.706 (0.033)	0.689 (0.031)	0.760 (0.032)	0.688 (0.031)	0.975 (0.002)
	SCCA	0.624 (0.031)	0.613 (0.027)	0.592 (0.033)	0.609 (0.027)	0.987 (0.002)
100	GLAA	0.996 (0.003)	0.005 (0.005)	0.996 (0.003)	0.005 (0.005)	0.158 (0.010)
	ULA	0.332 (0.020)	0.035 (0.001)	0.360 (0.018)	0.034 (0.001)	0.905 (0.001)
	PMD	0.712 (0.034)	0.633 (0.033)	0.662 (0.038)	0.633 (0.033)	0.974 (0.002)
	SCCA	0.560 (0.033)	0.565 (0.030)	0.586 (0.032)	0.567 (0.030)	0.989 (0.001)

Table 2: Simulation results for Scenario 2 where p_3 varies. The rest is the same as Table 1.

Table 3 reports the results for Scenario 3 when the sample size n increases. Here we examine n that is comparable to the sample size in our multimodal PET example. It is seen that GLAA performs the best, even under a relatively small n . Moreover, the performances of all methods improve as n increases. However, ULA suffers a poor subspace estimation, while both PMD and SCCA continue to suffer both high false positive rates and poor subspace estimation accuracy, even for a relatively large n .

6 Multimodal PET Analysis

6.1 Study and data description

We revisit the multimodal PET study introduced in Section 1.1. It is part of the ongoing Berkeley Aging Cohort Study that targets Alzheimer’s disease (AD) as well as normal aging. AD is an irreversible neurodegenerative disorder and the leading form of dementia.

n	Method	TPR-1	FPR-1	TPR-2	FPR-2	D
60	GLAA	0.606 (0.025)	0.048 (0.008)	0.664 (0.027)	0.154 (0.021)	0.743 (0.015)
	ULA	0.410 (0.022)	0.031 (0.001)	0.362 (0.018)	0.160 (0.004)	0.920 (0.004)
	PMD	0.498 (0.040)	0.486 (0.036)	0.552 (0.039)	0.466 (0.036)	0.957 (0.004)
	SCCA	0.604 (0.023)	0.059 (0.012)	0.688 (0.023)	0.643 (0.017)	0.968 (0.002)
80	GLAA	0.638 (0.027)	0.018 (0.002)	0.738 (0.025)	0.137 (0.027)	0.662 (0.021)
	ULA	0.550 (0.021)	0.024 (0.001)	0.412 (0.018)	0.147 (0.005)	0.895 (0.004)
	PMD	0.494 (0.038)	0.441 (0.034)	0.502 (0.038)	0.429 (0.034)	0.953 (0.004)
	SCCA	0.672 (0.026)	0.665 (0.016)	0.716 (0.025)	0.718 (0.019)	0.966 (0.003)
100	GLAA	0.820 (0.022)	0.013 (0.004)	0.848 (0.020)	0.060 (0.016)	0.483 (0.024)
	ULA	0.686 (0.019)	0.017 (0.001)	0.534 (0.021)	0.117 (0.005)	0.870 (0.004)
	PMD	0.510 (0.035)	0.476 (0.032)	0.560 (0.036)	0.442 (0.033)	0.947 (0.004)
	SCCA	0.702 (0.027)	0.677 (0.021)	0.732 (0.025)	0.714 (0.024)	0.961 (0.003)
120	GLAA	0.896 (0.018)	0.004 (0.001)	0.914 (0.018)	0.010 (0.003)	0.350 (0.021)
	ULA	0.802 (0.017)	0.010 (0.001)	0.622 (0.018)	0.094 (0.004)	0.850 (0.003)
	PMD	0.576 (0.035)	0.564 (0.032)	0.662 (0.035)	0.530 (0.032)	0.941 (0.004)
	SCCA	0.640 (0.030)	0.657 (0.024)	0.708 (0.028)	0.705 (0.022)	0.970 (0.003)
160	GLAA	0.990 (0.004)	0.001 (0.000)	0.984 (0.005)	0.002 (0.001)	0.209 (0.012)
	ULA	0.920 (0.011)	0.004 (0.001)	0.732 (0.016)	0.067 (0.004)	0.812 (0.003)
	PMD	0.610 (0.037)	0.570 (0.034)	0.652 (0.038)	0.544 (0.035)	0.944 (0.004)
	SCCA	0.626 (0.031)	0.651 (0.024)	0.698 (0.029)	0.0691 (0.024)	0.969 (0.003)

Table 3: Simulation results for Scenario 3 where n varies. The rest is the same as Table 1.

It is characterized by progressive impairment of cognitive capabilities, then loss of bodily functions, and ultimately death. AD currently affects more than 10% of adults aged 65 or older, and the prevalence is continuously growing. It has now become an international imperative to understand, diagnose, and treat this disorder (Alzheimer’s Association, 2020).

The data consist of $n = 81$ elderly subjects, with the average age 77.5 years, and the standard deviation 6.2 years. For each subject, three types of neuroimages were acquired, including a Pittsburgh Compound B (PiB) PET scan that measures amyloid-beta protein, an AV-1451 PET scan that measures tau protein, and a 1.5T structural MRI scan for coregistration. MRI and PET images have all been preprocessed, and PET images were both coregistered to each participant’s MRI image. Moreover, a mask representing voxels likely to accumulate cortical amyloid and tau pathology was created (Lockhart et al., 2017). Then a set of MNI-space regions of interest were created, and the amount of amyloid-beta

and tau deposition was summarized for each region. This results in $p_1 = 60$ regions for amyloid-beta PET, and $p_2 = 26$ regions for tau-PET. We note that brain region parcellation is particularly useful to facilitate the interpretation, and has been frequently employed in brain imaging analysis (Fornito et al., 2013; Kang et al., 2016).

6.2 Analyses and results

One of the primary goals of this study is to identify brain regions where the association of amyloid-beta and tau changes the most as age varies, and to further understand this association change. This would offer useful insight about how these two AD pathological proteins interact in the aging brains, which in turn may enable more accurate prediction of individual subjects demonstrating in vivo neuropathology, and allow better design and subject recruitment of clinical trials to potentially slow the spread of AD. For instance, clinical trials that aim at testing anti-amyloid-beta or anti-tau agents would need to know not only that participants have amyloid-beta and tau in the brain, but also how the relative levels of each pathological protein are spatially associated with each other given their ages. We cast this problem in the framework of liquid association analysis. Let $\mathbf{X} \in \mathbb{R}^{60}$, $\mathbf{Y} \in \mathbb{R}^{26}$ denote the amyloid-beta accumulation and tau accumulation in various brain regions, respectively, and $Z \in \mathbb{R}$ denote the subject’s age. We first log-transform each variable in \mathbf{X} and \mathbf{Y} , and standardize \mathbf{X} , \mathbf{Y} and Z marginally. We then apply the proposed generalized liquid association analysis (GLAA) method to this data. We choose the thresholding parameters η_1 and η_2 for the initialization step, so that about half of the variables in \mathbf{X} and in \mathbf{Y} are kept for subsequent iterations. We then tune the thresholding parameters $\tilde{\eta}_1$ and $\tilde{\eta}_2$ used in iterative sparse SVD by cross-validation over a grid of candidate values. We choose the ranks r_1, r_2 , i.e., the numbers of linear combinations for \mathbf{X} and \mathbf{Y} , to be one, which is most common in canonical correlation analysis.

After obtaining the two estimated linear combinations $\hat{\Gamma}_1^\top \mathbf{X}$ and $\hat{\Gamma}_2^\top \mathbf{Y}$, we plot them as the value of Z changes. We divide the interval of Z into six equal-sized intervals with

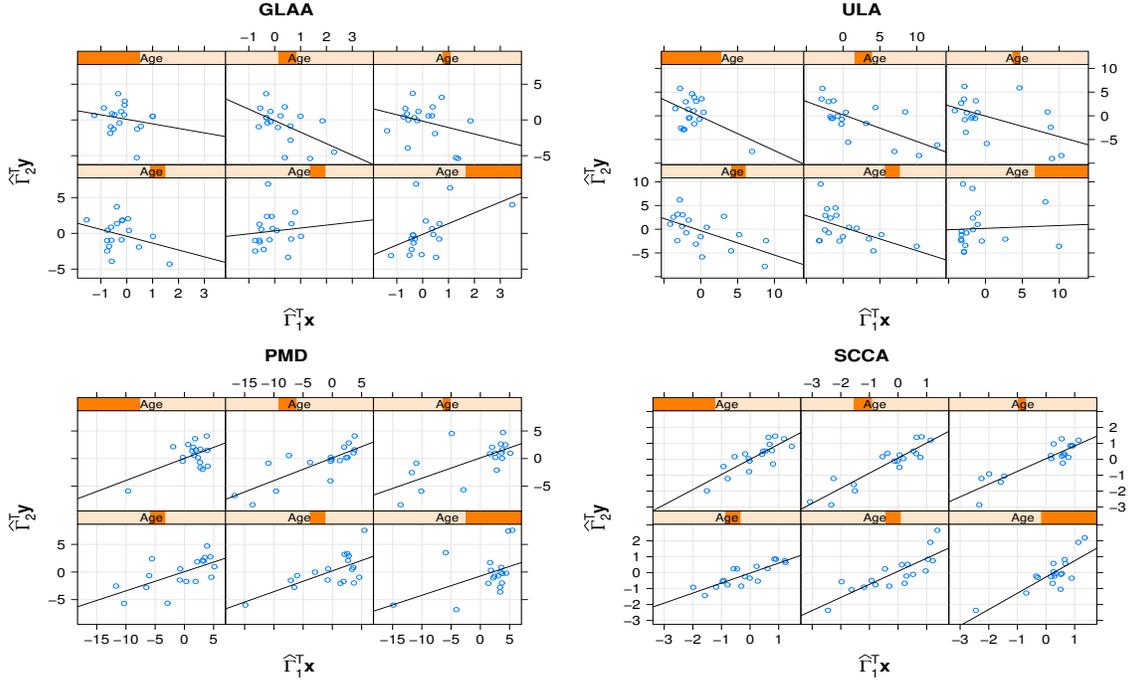


Figure 1: Trellis plots of the estimated linear combinations $\hat{\Gamma}_2^T \mathbf{Y}$ versus $\hat{\Gamma}_1^T \mathbf{X}$ as Z varies. Each panel represents an interval of Z , with a linear line added. The methods include: generalized liquid association analysis (GLAA), univariate liquid association (ULA), penalized matrix decomposition (PMD), and sparse canonical correlation analysis (SCCA).

overlaps, then draw the scatterplot of $\hat{\Gamma}_2^T \mathbf{Y}$ versus $\hat{\Gamma}_1^T \mathbf{X}$ within each interval. We also add a fitted linear regression line in each panel to reflect the correlation between $\hat{\Gamma}_1^T \mathbf{X}$ and $\hat{\Gamma}_2^T \mathbf{Y}$. Figure 1 shows the trellis plots, where the stripe at the top of each panel represents the range of Z it covers. It is interesting to see from the GLAA estimation, the correlation between $\Gamma_1^T \mathbf{X}$ and $\Gamma_2^T \mathbf{Y}$ changes from negative to positive gradually, as the age variable Z increases. This may be due to different deposition patterns of amyloid-beta and tau. In particular, amyloid-beta plaques are detectable in the brain many years before dementia onset, while tau neurofibrillary tangles aggregate specifically in the medial temporal lobes in normal aging. The spread of tau out of medial temporal lobes and into the surrounding isocortex at elder age coincides with cognitive impairment, and the process is hypothesized to be potentiated or accelerated by the presence of amyloid-beta (He et al., 2018; Vogel

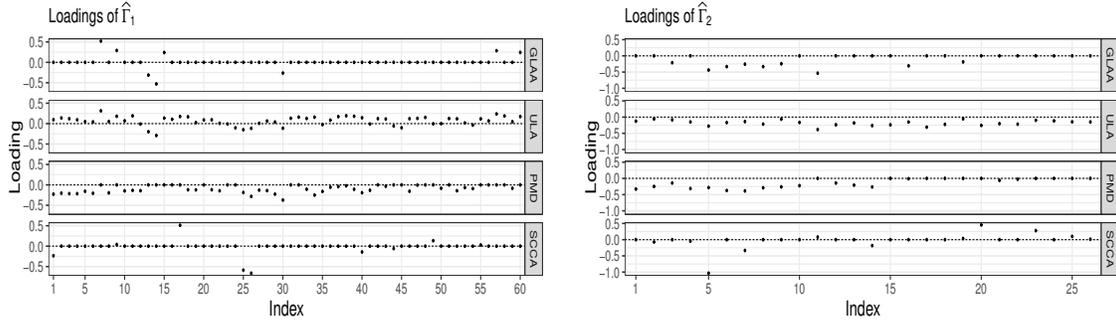


Figure 2: Estimated loadings in $\hat{\Gamma}_1$ and $\hat{\Gamma}_2$. The number of non-zero loading entries estimated by GLAA, ULA, PMD, SCCA are 8, 60, 37, 9 for $\hat{\Gamma}_1$, and 9, 26, 16, 11 for $\hat{\Gamma}_2$.

et al., 2020). The change from a negative association in early years to a positive association in later years between amyloid-beta and tau found by our GLAA method may offer some support to this hypothesis. As a comparison, no clear changing pattern is observed from the other three estimation methods.

Next, we examine more closely the brain regions identified by GLAA that demonstrate dynamic association patterns. Figure 2 plots the loadings of the estimated $\hat{\Gamma}_1$ and $\hat{\Gamma}_2$, where the indices of non-zero loadings correspond to the selected regions. The number of non-zero loading entries estimated by GLAA, ULA, PMD, SCCA are 8, 60, 37, 9 for $\hat{\Gamma}_1$, and 9, 26, 16, 11 for $\hat{\Gamma}_2$, respectively. Note that, the ULA method does not do variable selection, and for the real data, no information on the true sparsity level is known, so its estimated loadings are non-sparse. Moreover, the PMD method yields a large number of non-zero estimates, making the interpretation difficult. The SCCA method selects about the same number of non-zero regions as GLAA, but the selected regions are less meaningful.

Modality	Identified regions				
amyloid-beta	Entorhinal R	Entorhinal L	Hippocampus R	Hippocampus L	Amygdala R
	Orbitofrontal L	Posterior Cingulate L	Middle Frontal R		
tau	Entorhinal R	Entorhinal L	Hippocampus R	Parahippocampal R	Fusiform L
	Middle Temporal R	Middle Temporal L	Insula L	Rostral Anterior Cingulate R	

Table 4: Identified brain regions for amyloid-beta and tau by GLAA. Regions in the left hemisphere are denoted by “L”, and regions in the right hemisphere are denoted by “R”.

Table 4 reports the identified brain regions by GLAA for amyloid-beta and tau, respectively, while Figure 3 visualizes those regions on a template brain using BrainNet Viewer (Xia et al., 2013). Many of these regions are known to be closely related to AD, and the dynamic associations between amyloid-beta and tau of those regions reveal interesting and new insights. Particularly, for both amyloid-beta and tau, the identified regions include hippocampus and entorhinal cortex. Hippocampus is a major component functionally involved in response inhibition, episodic memory, and spatial cognition. It is one of the first brain regions to suffer damage from AD (Jack et al., 2011). Entorhinal cortex is a brain region that functions as a hub in a widespread network for memory and navigation. Entorhinal cortex and hippocampus together play an important role in memories. Atrophy in entorhinal cortex has been consistently reported in AD (Pini et al., 2016). Moreover, animal models have suggested that neurofibrillary tangles of tau first appear in entorhinal cortex, then spread to hippocampus (Cho et al., 2016). For amyloid-beta, other identified regions include amygdala, orbitofrontal cortex, posterior cingulate cortex and areas of middle frontal cortices. Amygdala is responsible for memory processing, decision-making and emotional responses. Amygdala atrophy is found prominent in early AD (Poulin et al., 2011). Orbitofrontal cortex is involved in decision-making, while posterior cingulate cortex is one of the most metabolically active brain regions, and is linked to emotion and memory. Atrophy of both regions and middle frontal cortices have been found associated with AD (Pini et al., 2016). For tau, other identified regions include parahippocampal gyrus, middle temporal gyrus, fusiform, insula, and rostral anterior cingulate cortex. Parahippocampal gyrus is central for memory encoding and retrieval. Atrophy in parahippocampal gyrus has been identified as an early biomarker of AD (Echavarri et al., 2011). Middle temporal gyrus is connected with recognition of known faces and accessing word meaning. Fusiform is linked with various neural pathways related to recognition. Insula is involved in consciousness and emotion. Rostral anterior cingulate cortex is involved in attention allocation, decision-making and emotion. There have been numerous evidences suggesting

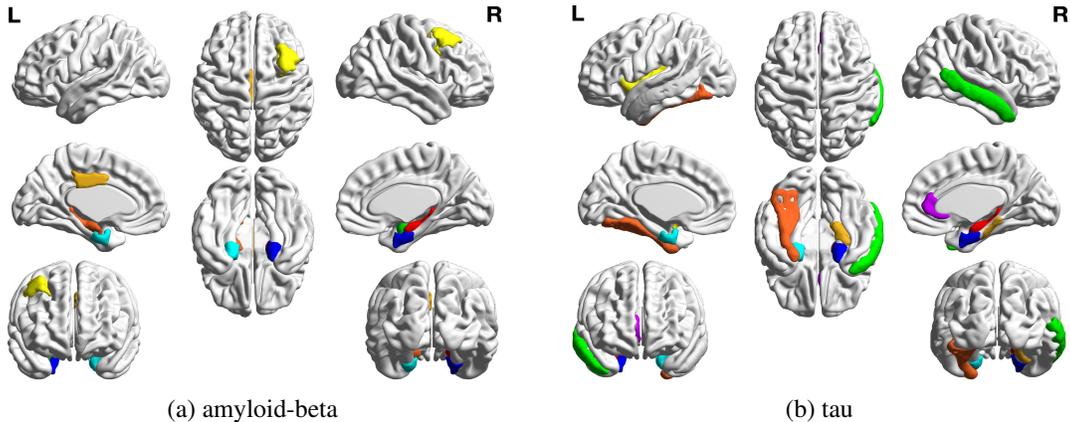


Figure 3: Identified brain regions for amyloid-beta and tau by GLAA.

associations between these regions and AD (Convit et al., 2000; Pini et al., 2016).

In summary, GLAA identifies interesting dynamic association patterns among a number of important brain regions between amyloid-beta and tau as age increases. Moreover, GLAA provides a useful dimension reduction tool to help visualize such patterns.

7 Discussion

In this article, we have proposed generalized liquid association analysis, which offers a new angle to study three-way associations among random variables, and is particularly useful for multimodal integrative data analysis. We have illustrated with a multimodal neuroimaging study of Alzheimer’s disease in detail. Meanwhile, the proposal is potentially applicable to other multimodal data problems, e.g., to understand gene co-expressions given single nucleotide polymorphisms or under varying physiological states (Chen et al., 2011; Yu, 2018), or to understand interactions between gene expressions and microRNA levels or comparative genomic hybridizations given cancer states and demographics (Cai et al., 2016; Mai and Zhang, 2019). Next, we discuss some potential extensions.

First, we begin with the situation when there is a univariate and categorical Z , whereas the analysis so far has primarily concentrated on the case when each variable in Z is con-

tinous. In general, it remains an open question on how to define liquid association for a categorical variable, since the function $g(z)$ is no longer differentiable for a categorical Z . For a binary $Z \in \{0, 1\}$, we propose to replace the derivative of the conditional mean function with the absolute change in the conditional means across the two groups, i.e., $\text{LA}(X, Y|Z) = |\text{E}(XY|Z = 1) - \text{E}(XY|Z = 0)|$, where the absolute value is used because the class labels are interchangeable. This naturally fits the original interpretation of LA. Similarly, for a categorical or ordinal $Z \in \{1, \dots, K\}$, we can use the weighted sum of pairwise absolute mean difference between the pairs of groups. Accordingly, the liquid association of \mathbf{X} and \mathbf{Y} given Z is defined as a $p_1 \times p_2$ matrix.

Next, for a multivariate mixed type \mathbf{Z} , we first organize $\mathbf{Z} = (\mathbf{Z}_1, \mathbf{Z}_2)^\top$ to separate the continuous variables, $\mathbf{Z}_1 = (Z_1, \dots, Z_q)^\top \in \mathbb{R}^q$, from the categorical variables, $\mathbf{Z}_2 = (Z_{q+1}, \dots, Z_{p_3})^\top \in \mathbb{R}^{p_3-q}$. Directly imposing a low-dimensional structure on entire \mathbf{Z} would lead to difficulty in interpretation. Alternatively, we propose a dimension reduction approach, by recognizing the reduction on \mathbf{Z} in model (1) is indeed a sufficient dimension reduction model. Specifically, when \mathbf{Z} is continuous, by model (1), we have $\mathbf{g}(\mathbf{Z}) = \mathbf{g}(\mathbf{P}_\mathcal{S}\mathbf{Z})$, where $\mathcal{S} = \text{span}(\mathbf{\Gamma}_3)$. Therefore, $\mathbf{g}(\mathbf{Z}) \perp\!\!\!\perp \mathbf{Z} | \mathbf{P}_\mathcal{S}\mathbf{Z}$. This leads to a sufficient dimension reduction model of \mathbf{Z} for the conditional mean function $\mathbf{g}(\mathbf{Z}) = \text{E}(\mathbf{X}\mathbf{Y}^\top | \mathbf{Z})$, in the sense that all the mean information of the regression of the matrix response $\mathbf{X}\mathbf{Y}^\top$ given the predictor vector \mathbf{Z} is fully captured by the linear combinations $\mathbf{\Gamma}_3^\top \mathbf{Z}$. As such, model (1) can be viewed as a generalization of the notion of sufficient mean reduction (Cook and Li, 2002). Now for the mixed type $\mathbf{Z} = (\mathbf{Z}_1^\top, \mathbf{Z}_2^\top)^\top$, we adopt the idea of partial dimension reduction (Chiaromonte et al., 2002), or groupwise dimension reduction (Li et al., 2010), and estimate the subspace $\mathcal{S} \subseteq \mathbb{R}^q$ such that $\mathbf{g}(\mathbf{Z}) \perp\!\!\!\perp \mathbf{Z} | (\mathbf{P}_\mathcal{S}\mathbf{Z}_1, \mathbf{Z}_2)$.

We have so far focused on studying the associations between two sets of random variables conditioning on another set. It is possible to generalize to the associations of more than two sets of variables conditioning on another set. It involves a higher-order tensor than an order-3 tensor. Nevertheless, both the estimation algorithm and the theory can be

extended to a general order tensor in a relatively straightforward fashion. Moreover, from the simulation results in Table 2, we observe that, the higher the dimension of \mathbf{Z} , the more challenging the problem becomes. We believe it is possible to employ an alternative modeling strategy such as Chen et al. (2011) when the dimension of \mathbf{Z} is ultrahigh. This is also true when the scientific interest is to select important variables in \mathbf{Z} , while our current interest concentrates on the selection of variables in \mathbf{X} and \mathbf{Y} , but not in \mathbf{Z} . We leave the pursuit of these lines of research as our future work.

Acknowledgement

Li’s research was supported by NSF grant CIF-2102227, and NIH grants R01AG061303, R01AG062542, and R01AG034570. Zeng’s research was supported by NSF grant CCF-1908969. Zhang’s research was supported by NSF grant DMS-2053697, and NIH grant DE-030509. The authors thank the Editor, the Associate Editor, and three referees for their constructive comments and suggestions.

References

- Abid, A., Zhang, M. J., Bagaria, V. K., and Zou, J. (2018). Exploring patterns enriched in a dataset with contrastive principal component analysis. *Nature communications*, 9:1–7.
- Alzheimer’s Association (2020). 2020 Alzheimer’s disease facts and figures. *Alzheimer’s & Dementia*, 16(3):391–460.
- Bi, X., Qu, A., and Shen, X. (2018). Multilayer tensor factorization with applications to recommender systems. *The Annals of Statistics*, 46:3308–3333.
- Bi, X., Tang, X., Yuan, Y., Zhang, Y., and Qu, A. (2021). Tensor in statistics. *Annual Review of Statistics and Its Application*, 8:345–368.

- Braak, H. and Braak, E. (1991). Neuropathological staging of alzheimer-related changes. *Acta Neuropathologica*, 82(4):239–259.
- Cai, T., Cai, T. T., and Zhang, A. (2016). Structured matrix completion with applications to genomic data integration. *Journal of the American Statistical Association*, 111:621–633.
- Chen, J., Xie, J., and Li, H. (2011). A penalized likelihood approach for bivariate conditional normal models for dynamic co-expression analysis. *Biometrics*, 67(1):299–308.
- Chiaromonte, F., Cook, R. D., and Li, B. (2002). Sufficient dimensions reduction in regressions with categorical predictors. *The Annals of Statistics*, 30(2):475–497.
- Cho, H., Choi, J. Y., Hwang, M. S., Kim, Y. J., Lee, H. M., Lee, H. S., Lee, J. H., Ryu, Y. H., Lee, M. S., and Lyoo, C. H. (2016). In vivo cortical spreading pattern of tau and amyloid in the alzheimer disease spectrum. *Annals of Neurology*, 80(2):247–258.
- Convit, A., de Asis, J., et al. (2000). Atrophy of the medial occipitotemporal, inferior, and middle temporal gyri in non-demented elderly predict decline to alzheimer’s disease. *Neurobiology of Aging*, 21(1):19–26.
- Cook, R. D. (1998). Principal hessian directions revisited. *Journal of the American Statistical Association*, 93:84–94.
- Cook, R. D. (2007). Fisher lecture: Dimension reduction in regression. *Statistical Science*, 22(1):1–26.
- Cook, R. D. and Li, B. (2002). Dimension reduction for conditional mean in regression. *The Annals of Statistics*, 30(2):455–474.
- De Lathauwer, L., De Moor, B., and Vandewalle, J. (2000). On the best rank-1 and rank- (r_1, r_2, \dots, r_m) approximation of higher-order tensors. *SIAM journal on Matrix Analysis and Applications*, 21(4):1324–1342.

- Echavarri, C., Aalten, P., Uylings, H., Jacobs, H., Visser, P., Gronenschild, E., Verhey, F., and Burgmans, S. (2011). Atrophy in the parahippocampal gyrus as an early biomarker of alzheimer’s disease. *Brain Structure & Function*, 215:265–271.
- Fornito, A., Zalesky, A., and Breakspear, M. (2013). Graph analysis of the human connectome: Promise, progress, and pitfalls. *NeuroImage*, 80:426–444.
- Gao, C., Ma, Z., Ren, Z., and Zhou, H. H. (2015). Minimax estimation in sparse canonical correlation analysis. *The Annals of Statistics*, 43:2168–2197.
- Hao, B., Zhang, A., and Cheng, G. (2020). Sparse and low-rank tensor estimation via cubic sketchings. *IEEE Transactions on Information Theory*, 66(9):5927–5964.
- He, Z., Guo, J. L., et al. (2018). Amyloid-beta plaques enhance alzheimer’s brain tau-seeded pathologies by facilitating neuritic plaque tau aggregation. *Nature Medicine*, 24(1):29–38.
- Ho, Y.-Y., Parmigiani, G., Louis, T. A., and Cope, L. M. (2011). Modeling liquid association. *Biometrics*, 67(1):133–141.
- Jack, C. R., Barkhof, F., et al. (2011). Steps to standardization and validation of hippocampal volumetry as a biomarker in clinical trials and diagnostic criterion for alzheimer’s disease. *Alzheimer’s & Dementia*, 7(4):474–485.e4.
- Kang, J., Bowman, F. D., Mayberg, H., and Liu, H. (2016). A depression network of functionally connected regions discovered via multi-attribute canonical correlation graphs. *NeuroImage*, 141:431–441.
- Kolda, T. G. and Bader, B. W. (2009). Tensor decompositions and applications. *SIAM review*, 51(3):455–500.
- Li, B. (2018). *Sufficient Dimension Reduction: Methods and Applications with R*. Chapman & Hall/CRC Monographs on Statistics and Applied Probability. CRC Press.

- Li, G. and Gaynanova, I. (2018). A general framework for association analysis of heterogeneous data. *The Annals of Applied Statistics*, 12(3):1700–1726.
- Li, G. and Jung, S. (2017). Incorporating covariates into integrated factor analysis of multi-view data. *Biometrics*, 73(4):1433–1442.
- Li, K.-C. (1992). On principal hessian directions for data visualization and dimension reduction: Another application of stein’s lemma. *Journal of the American Statistical Association*, 87(420):1025–1039.
- Li, K.-C. (2002). Genome-wide coexpression dynamics: theory and application. *Proceedings of the National Academy of Sciences*, 99(26):16875–16880.
- Li, K.-C., Liu, C.-T., Sun, W., Yuan, S., and Yu, T. (2004). A system for enhancing genome-wide coexpression dynamics study. *Proceedings of the National Academy of Sciences*, 101(44):15561–15566.
- Li, L., Li, B., and Zhu, L.-X. (2010). Groupwise dimension reduction. *Journal of the American Statistical Association*, 105(491):1188–1201.
- Liu, J. S. (1994). Siegel’s formula via stein’s identities. *Statistics & Probability Letters*, 21(3):247–251.
- 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(1):523.
- Lockhart, S. N., Schöll, M., Baker, S. L., Ayakta, N., Swinnerton, K. N., Bell, R. K., Mellinger, T. J., Shah, V. D., O’Neil, J. P., Janabi, M., and Jagust, W. J. (2017). Amyloid and tau PET demonstrate region-specific associations in normal older people. *NeuroImage*, 150:191–199.

- Luo, Y., Raskutti, G., Yuan, M., and Zhang, A. R. (2020). A sharp blockwise tensor perturbation bound for orthogonal iteration. *arXiv preprint arXiv:2008.02437*.
- Mai, Q. and Zhang, X. (2019). An iterative penalized least squares approach to sparse canonical correlation analysis. *Biometrics*, 75(3):734–744.
- Nathoo, F. S., Kong, L., and Zhu, H. (2017). Inference on high-dimensional differential correlation matrix. *arXiv preprint arXiv:1707.07332*.
- Pini, L., Pievani, M., Bocchetta, M., Altomare, D., Bosco, P., Cavedo, E., Galluzzi, S., Marizzoni, M., and Frisoni, G. B. (2016). Brain atrophy in alzheimer’s disease and aging. *Ageing Research Reviews*, 30:25–48. Brain Imaging and Aging.
- Poulin, S. P., Dautoff, R., Morris, J. C., Barrett, L. F., and Dickerson, B. C. (2011). Amygdala atrophy is prominent in early alzheimer’s disease and relates to symptom severity. *Psychiatry Research: Neuroimaging*, 194(1):7–13.
- Richardson, S., Tseng, G. C., and Sun, W. (2016). Statistical methods in integrative genomics. *Annual Review of Statistics and Its Application*, 3:181–209.
- Shu, H., Wang, X., and Zhu, H. (2019). D-cca: A decomposition-based canonical correlation analysis for high-dimensional datasets. *Journal of the American Statistical Association*, pages 1–29.
- Stein, C. M. (1981). Estimation of the mean of a multivariate normal distribution. *The annals of Statistics*, pages 1135–1151.
- Sun, W., Lu, J., Liu, H., and Cheng, G. (2017). Provable sparse tensor decomposition. *Journal of the Royal Statistical Society, Series B*, 79:899–916.
- Tang, C. Y., Fang, E. X., and Dong, Y. (2020). High-dimensional interactions detection with sparse principal hessian matrix. *Journal of Machine Learning Research*, 21:1–25.

- Tang, X., Bi, X., and Qu, A. (2019). Individualized multilayer tensor learning with an application in imaging analysis. *Journal of the American Statistical Association*, pages 1–26.
- Uludag, K. and Roebroek, A. (2014). General overview on the merits of multimodal neuroimaging data fusion. *NeuroImage*, 102:3 – 10.
- Vogel, J. W., Iturria-Medina, Y., and et al. (2020). Spread of pathological tau proteins through communicating neurons in human alzheimer’s disease. *Nature Communications*, 11(1):2612.
- Wainwright, M. J. (2019). *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press.
- Witten, D. M., Tibshirani, R., and Hastie, T. (2009). A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis. *Biostatistics*, 10(3):515–534.
- Xia, M., Wang, J., and He, Y. (2013). Brainnet viewer: A network visualization tool for human brain connectomics. *PLOS ONE*, 8(7):1–15.
- Xia, Y., Li, L., Lockhart, S. N., and Jagust, W. J. (2019). Simultaneous covariance inference for multimodal integrative analysis. *Journal of the American Statistical Association*, 115:1279–1291.
- Yang, D., Ma, Z., and Buja, A. (2016). Rate optimal denoising of simultaneously sparse and low rank matrices. *The Journal of Machine Learning Research*, 17(1):3163–3189.
- Yu, T. (2018). A new dynamic correlation algorithm reveals novel functional aspects in single cell and bulk rna-seq data. *PLOS Computational Biology*, 14:1–22.
- Yu, Y., Wang, T., and Samworth, R. J. (2015). A useful variant of the davis–kahan theorem for statisticians. *Biometrika*, 102(2):315–323.

- Zhang, A. and Han, R. (2019). Optimal sparse singular value decomposition for high-dimensional high-order data. *Journal of the American Statistical Association*, pages 1–34.
- Zhang, A. and Xia, D. (2018). Tensor svd: Statistical and computational limits. *IEEE Transactions on Information Theory*, 64(11):7311–7338.
- Zhou, H., Li, L., and Zhu, H. (2013). Tensor regression with applications in neuroimaging data analysis. *Journal of the American Statistical Association*, 108:540–552.