Published in final edited form as:

Proceedings (IEEE Int Conf Bioinformatics Biomed). 2016 ; 2016: 707–711. doi:10.1109/BIBM.
2016.7822605.

Sparse Canonical Correlation Analysis via Truncated ℓ_1 -norm with Application to Brain Imaging Genetics

Lei Du^{*}, Tuo Zhang^{*}, Kefei Liu[†], Xiaohui Yao[†], Jingwen Yan[†], Shannon L. Risacher[†], Lei
Guo^{*}, Andrew J. Saykin[†], Li Shen^{†,§}, and for the ADNI

^{*}School of Automation, Northwestern Polytechnical University, Xi'an, China 710072

[†]Indiana University School of Medicine, Indianapolis, USA 46202

Abstract

Discovering bi-multivariate associations between genetic markers and neuroimaging quantitative traits is a major task in brain imaging genetics. Sparse Canonical Correlation Analysis (SCCA) is a popular technique in this area for its powerful capability in identifying bi-multivariate relationships coupled with feature selection. The existing SCCA methods impose either the ℓ_1 -norm or its variants. The ℓ_1 -norm is more desirable, which however remains unexplored since the ℓ_1 -norm minimization is NP-hard. In this paper, we impose the truncated ℓ_1 -norm to improve the performance of the ℓ_1 -norm based SCCA methods. Besides, we propose two efficient optimization algorithms and prove their convergence. The experimental results, compared with two benchmark methods, show that our method identifies better and meaningful canonical loading patterns in both simulated and real imaging genetic analyse.

Keywords

Sparse Canonical Correlation Analysis; Truncated ℓ_1 -norm; Brain Imaging Genetics

I. Introduction

Brain imaging genetics has gained more and more attentions recently [1], [2]. A major task of imaging genetics is to identify bi-multivariate associations between single nucleotide polymorphisms (SNPs) and imaging quantitative traits (QTs). Sparse canonical correlation analysis (SCCA), which is powerful in bi-multivariate relationship discovery coupled with feature selection, has become a popular technique in imaging genetic studies [3], [4], [5], [6], [7].

Witten *et al.* [3] introduced the ℓ_1 -norm (Lasso) to assure sparsity which only selects a small proportion of the features. Since then, many SCCA methods using the ℓ_1 -norm or its variants are proposed [8]. There are two major concerns regarding them. First, the ℓ_1 -norm, which only penalizes those nonzero features, is the most ideal constraint. But it is neither non-

[§]Corresponding to: shenli@iu.edu.

convex nor discontinuous [9]. Second, the ℓ_1 -norm constraint is not a stable feature selector and thus could incur estimation bias [10].

To overcome the problem above, the truncated ℓ_1 -norm penalty (TLP) [10], [11] is proposed.

The TLP is defined as $J_\tau(|x|) = \min(\frac{|x|}{\tau}, 1)$ with τ being a positive tuning parameter. It approximates ℓ_1 -norm and permits desirable sparsity. In addition, TLP can be equivalently transferred to a piecewise linear function, and thus is easy to handle.

In this paper, we propose the TLP based SCCA (TLP-SCCA) which embraces the TLP into the CCA model. The TLP-SCCA has the following advantages [10]. First, the TLP performs as a tradeoff between the ℓ_1 and ℓ_2 functions. This means that it not only has improved feature selection, but also can be solved effectively. Second, it is an adaptive shrinkage method if τ is tuned appropriately. We propose two effective optimization algorithms, both using the alternating direction method of multipliers (ADMM) technique [12], and they are guaranteed to converge. The experimental results, compared with two popular ℓ_1 -norm based SCCA [3], [6], show that both TLP-SCCA exhibit cleaner canonical loading patterns than the ℓ_1 -SCCA.

II. The Truncated ℓ_1 -norm Penalty

In this paper, a boldface lowercase letter denotes a vector, and a boldface uppercase letter denotes a matrix. $\mathbf{X} \in R^{n \times p}$ denotes the SNP data, and $\mathbf{Y} \in R^{n \times q}$ is the QT data.

The truncated ℓ_1 -norm is defined as follows [13]:

$$P^{\text{TLP}}(\mathbf{u}) = \sum_i J_\tau(|u_i|), \text{ where } J_\tau(|u_i|) = \min(\frac{|u_i|}{\tau}, 1). \quad (1)$$

The parameter τ is a threshold. Given an appropriate τ , TLP balances between the ℓ_1 -norm and ℓ_2 -norm according to the magnitude of the coefficients. Fig. 1 presents the norm ball of ℓ_1 -norm, ℓ_2 -norm, and TLP with different τ 's. The TLP ball is more pointy than the ℓ_2 ball, and thus it is more desirable in sparse learning.

Since TLP is non-convex, we employ the Difference of Convex functions (DC) programming technique [14], [11]. Let $f_1(u_i) = |u_i|$ and $f_2(u_i) = \max(|u_i| - \tau, 0)$, we have:

$$J_\tau(|u_i|) = \frac{1}{\tau} (f_1(u_i) - f_2(u_i)) \quad (2)$$

where both $f_1(u_i)$ and $f_2(u_i)$ are convex. The affine minorization of $f_2(\cdot)$ can be denoted as:

$$f_2(u) = f_2(u_i) + \langle u - u_i, d_i \rangle, \quad (3)$$

where $d_i = \text{sign}(u_i)\mathbb{I}(|u_i| > \tau)$ is the sub-gradient of f_2 at u_i with the $\mathbb{I}(\cdot)$ being an indicator function. Then the minimization of P^{TLP} can be transferred to its subproblem,

$$\arg \min P^{\text{TLP}} = \arg \min \frac{1}{\tau} (f_1(\mathbf{u}) - \mathbf{D} \cdot \mathbf{u}), \quad (4)$$

where $\mathbf{D}$ is a diagonal matrix with the i -th element as d_i .

III. The Novel SCCA via Truncated ℓ_1 -norm

A. The TLP-SCCA Model

The TLP-SCCA can be formally defined as follows:

$$\min_{\mathbf{u}, \mathbf{v}} -\mathbf{u}^T \mathbf{X}^T \mathbf{Y} \mathbf{v} \text{ s.t. } \|\mathbf{X} \mathbf{u}\|^2 \leq 1, \|\mathbf{Y} \mathbf{v}\|^2 \leq 1, P_1^{\text{TLP}} \leq c_1, P_2^{\text{TLP}} \leq c_2, \quad (5)$$

where $P_1^{\text{TLP}} = \sum_i J_\tau(|u_i|)$ and $P_2^{\text{TLP}} = \sum_j J_\tau(|v_j|)$ are the TLP penalties. The merits are three-fold. Firstly, it has lower estimation bias than a traditional ℓ_1 -SCCA since it approximates the optimal ℓ_1 -SCCA if given an appropriate τ . Secondly, it has better sparsity-inducing ability than the ℓ_1 -SCCA. Thirdly, it is easier to solve than the ℓ_1 -SCCA.

B. The Optimization Algorithm

The TLP-SCCA is biconvex in $\mathbf{u}$ and $\mathbf{v}$, and we solve it based on the Alternate Convex Search (ACS) strategy [15]. We show how to employ the ADMM algorithm to updating $\mathbf{u}$ with $\mathbf{v}$ fixed.¹ The ADMM combines the advantages of dual decomposition and augmented Lagrangian method [12], which divides a large global problem into a series of small local subproblems [11]. Introducing a slack variable $\mathbf{z}$, TLP-SCCA with respect to $\mathbf{u}$ conveys to:

$$\arg \min_{\mathbf{u}} -\mathbf{u}^T \mathbf{X}^T \mathbf{Y} \mathbf{v} + \frac{\beta}{2} \|\mathbf{X} \mathbf{u}\|^2 + \frac{\lambda}{\tau} \|\mathbf{z}\|_1 - \frac{\lambda}{\tau} \mathbf{D} \cdot \mathbf{u} \text{ s.t. } \mathbf{u} - \mathbf{z} = \mathbf{0} \quad (6)$$

where β and λ are tuning parameters, τ is the parameter of the TLP, and $\mathbf{D}$ is a diagonal matrix with the i -th element as the sub-gradient $\text{sign}(u_i)\mathbb{I}(|u_i| > \tau_1)$ (Recall Section II).

Algorithm 1 The Truncated ℓ_1 -norm Penalty based SCCA (TLP-SCCA)

Require:

$\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n]^T, \mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_n]^T, \Gamma \geq 0, \mathbf{z} = \mathbf{0}, \mu = \mathbf{0}$

Ensure:

¹For ease of description, we only show how to solve $\mathbf{u}$ with $\mathbf{v}$ fixed as a constant. The $\mathbf{v}$ can be solved via swapping $\mathbf{u}$ and $\mathbf{v}$.

Canonical loadings $\mathbf{u}$ and $\mathbf{v}$.

- 1: Initialize $\mathbf{u} \in \mathbb{R}^{p \times 1}$, $\mathbf{v} \in \mathbb{R}^{q \times 1}$;
- 2: Compute Cholesky factorization $\mathbf{F}_u$ and $\mathbf{F}_v$;
- 3: **while** not converged **do**
- 4: Solve $\mathbf{u}$ according to Eq. (10);
- 5: Solve $\mathbf{z}_1$ according to Eq. (13);
- 6: Solve μ_1 according to Eq. (14);
- 7: Solve $\mathbf{v}$ in accordance with Step 4 ~ 6 by swapping $\mathbf{u}$ and $\mathbf{v}$;
- 8: Scale $\mathbf{u}$ so that $\|\mathbf{X}\mathbf{u}\|^2 = 1$, and $\mathbf{v}$ so that $\|\mathbf{Y}\mathbf{v}\|^2 = 1$.
- 9: **end while**

Then we construct the augmented Lagrangian function:

$$\mathcal{L}_{\mathbf{v}, \rho}(\mathbf{u}, \mathbf{z}, \mu) = -\mathbf{u}^T \mathbf{X}^T \mathbf{Y} \mathbf{v} + \frac{\beta}{2} \|\mathbf{X}\mathbf{u}\|^2 + \frac{\lambda}{\tau} \|\mathbf{z}\|_1 - \frac{\lambda}{\tau} \mathbf{D}\mathbf{u} + \mu^T (\mathbf{u} - \mathbf{z}) + \frac{\rho}{2} \|\mathbf{u} - \mathbf{z}\|^2 \quad (7)$$

where μ is the augmented Lagrangian multipliers, and $\rho > 0$ is the dual update step length. According to ADMM, this problem can be solved by sequentially updating $\mathbf{u}$, $\mathbf{z}$ and μ .

1. **u-update:** In each iteration, $\mathbf{u}$ can be obtained from:

$$\tilde{\mathbf{u}} = \arg \min_{\mathbf{u}} -\mathbf{u}^T \mathbf{X}^T \mathbf{Y} \mathbf{v} + \frac{\beta}{2} \|\mathbf{X}\mathbf{u}\|^2 - \frac{\lambda}{\tau} \mathbf{D}\mathbf{u} + \mu^T (\mathbf{u} - \mathbf{z}) + \frac{\rho}{2} \|\mathbf{u} - \mathbf{z}\|^2 \quad (8)$$

where $\mathbf{z}$, μ and $\mathbf{v}$ are all fixed as constant.

Then we can solve $\mathbf{u}$ by $\mathbf{u} = \mathbf{F}^{-1} \mathbf{b}$, where

$$\mathbf{F} = \beta \mathbf{X}^T \mathbf{X} + \rho \mathbf{I}, \mathbf{b} = \mathbf{X}^T \mathbf{Y} \mathbf{v} - \mu + \rho \mathbf{z} + \frac{\lambda}{\tau_1} \text{diag}(\mathbf{D}). \quad (9)$$

Obviously, $\mathbf{F}$ remains unchanged during the iteration, and thus we compute it at the beginning of the algorithm. The Cholesky factorization of the positive definite $\mathbf{F}$ is $\mathbf{F} = \mathbf{R}^T \mathbf{R}$, where $\mathbf{R}$ is an upper triangular matrix. Then $\mathbf{u}$ can be obtained by solving two linear system equations:

$$\tilde{\mathbf{u}} = (\mathbf{R}^T)^{-1} \mathbf{b}, \hat{\mathbf{u}} = \mathbf{R}^{-1} \tilde{\mathbf{u}}. \quad (10)$$

2. **z-update:** Likewise, $\mathbf{z}$ can be obtained from:

$$\tilde{\mathbf{z}} = \arg \min_{\mathbf{z}} \frac{\lambda}{\tau} \|\mathbf{z}\|_1 + \mu(\tilde{\mathbf{u}} - \mathbf{z}) + \frac{\rho}{2} \|\tilde{\mathbf{u}} - \mathbf{z}\|^2, \quad (11)$$

We equivalently reformulate the equation as follows.

$$\tilde{\mathbf{z}} = \arg \min_{\mathbf{z}} \frac{\rho}{2} \|\tilde{\mathbf{u}} + \frac{\mu}{\rho} - \mathbf{z}\|^2 + \frac{\lambda}{\tau} \|\mathbf{z}\|_1 \quad (12)$$

Then the soft-thresholding method can be used:

$$\tilde{\mathbf{z}} = S_{\lambda/(\rho\tau)}(\tilde{\mathbf{u}} + \frac{\mu}{\rho}), \text{ where } S_{\lambda}(a) = \text{sign}(a) \max(|a| - \lambda, 0) \quad (13)$$

3. μ -update: The dual update of μ is shown as follows.

$$\tilde{\mu} = \mu + \rho(\tilde{\mathbf{u}} - \tilde{\mathbf{z}}). \quad (14)$$

According to the ACS strategy, $\mathbf{v}$ can be solved similarly by fixing $\mathbf{u}$. The key steps are exhibited in Algorithm 1.

Algorithm 2 The TLP-SCCA with Restart (rTLP-SCCA)

Require:

$$\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n]^T, \mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_n]^T, \Gamma \geq 0, \mathbf{z} = \mathbf{0}, \mu = \mathbf{0}, \eta \in (0, 1)$$

Ensure:

Canonical loadings $\mathbf{u}$ and $\mathbf{v}$.

1: **while** not converged **do**

2: Compute $\mathbf{u}$, $\mathbf{z}_1$ and μ_1 using Algorithm 1;

3:

if $r_c^k < \eta r_c^{k-1}$ **then**

4:

$$\alpha^{k+1} = \frac{1 + \sqrt{1 + 4\alpha^{k2}}}{2}, \mathbf{z}_1^{k+1} = \mathbf{z}_1^k + \frac{\alpha^k - 1}{\alpha^{k+1}} (\mathbf{z}_1^k - \mathbf{z}_1^{k-1});$$

5:

$$\mu_1^{k+1} = \mu_1^k + \frac{\alpha^k - 1}{\alpha^{k+1}} (\mu_1^k - \mu_1^{k-1});$$

6: **else**

7:

$$\alpha^{k+1} = 1, \mathbf{z}_1^{k+1} = \mathbf{z}_1^{k-1}, \mu_1^{k+1} = \mu_1^{k-1}, r_c^k = \frac{r_c^{k-1}}{\eta};$$

8: **end if**

9: Solve $\mathbf{v}$ in accordance with Step 4 ~ 8 by swapping $\mathbf{u}$ and $\mathbf{v}$;

10: Scale $\mathbf{u}$ so that $\|\mathbf{Xu}\|^2 = 1$, and $\mathbf{v}$ so that $\|\mathbf{Yv}\|^2 = 1$.

11: end while

C. The Optimization Algorithm with Restart

We now propose an optimization algorithm with more stability. The algorithm uses a restart rule depending on the summation of the running primal and dual error [16].

We define the primal residual at the $(k+1)^{th}$ iteration $r_p^{k+1} = \mu^{k+1} - \mu^k = \rho(\mathbf{u}^{k+1} - \mathbf{z}^{k+1})$, and the dual residual as $r_d^{k+1} = \mathbf{z}^{k+1} - \mathbf{z}^k$. Then the combined residual is

$$r_c^{k+1} = \frac{1}{\rho} \|\mu^{k+1} - \mu^k\|^2 + \rho \|\mathbf{z}^{k+1} - \mathbf{z}^k\|^2. \quad (15)$$

The key steps of the rTLP-SCCA (TLP-SCCA with restart) are shown in Algorithm 2. If the residual r_c decreases by a factor of at least η (Step 3), we use the Nesterov's accelerated gradient descent method to accelerate [12] (Steps 4-5). Otherwise, we restart the iteration (Step 7).

D. Computational Analysis and Convergence

In Algorithm 1, the initialization step is simple. Step 2 of Cholesky factorization only needs to be done once. According to Eq. (10), Steps 4 involves solving two linear systems. Thus its time complexity is $\mathcal{O}(pq)$. Step 5 is the soft-thresholding operator whose time consumption is $\mathcal{O}(p)$. Step 6 is quite simple and its time complexity is $\mathcal{O}(p)$. Thus the time complexity of Steps 4-6 is $\mathcal{O}(pq)$. Similarly, the time complexity of solving $\mathbf{v}$ is $\mathcal{O}(pq)$. In summary, the complexity of each iteration is $\mathcal{O}(pq)$.

The difference between TLP-SCCA and rTLP-SCCA is the restart steps which are easy to calculate. Thus the time complexity of each iteration of rTLP-SCCA is also $\mathcal{O}(pq)$.

The TLP-SCCA problem is biconvex in terms of $\mathbf{u}$ and $\mathbf{v}$. According to Theorem 4.5 in [15], the TLP-SCCA converges monotonically depending on two conditions: (1) The set $\mathbf{B} = (\mathbf{u}, \mathbf{v}) \in \mathbb{R}^p \times \mathbb{R}^q$ is bounded. (2) Both $\mathbf{u}$ -update and $\mathbf{v}$ -update are solvable.

We note that $\mathbf{u}$ is bounded by $\|\mathbf{X}\mathbf{u}\|^2 = 1$, and $\mathbf{v}$ is bounded by $\|\mathbf{Y}\mathbf{v}\|^2 = 1$. Therefore, $\mathbf{B}$ is bounded because that $\mathbf{u}$ and $\mathbf{v}$ are bounded. Obviously, the $\mathbf{u}$ -update and $\mathbf{v}$ -update are solvable guaranteed by ADMM method. Therefore, the TLP-SCCA converges during iteration. [15].

IV. Experimental Study

A. Benchmark Methods

In this paper, we mainly focus on the improved performance of SCCA using the TLP regularization. Thus we compare our method with the ℓ_1 -SCCA, other than those structured SCCA methods [5], [17], [7], [18], [19]. We use two different implementations of ℓ_1 -SCCA

as benchmark. The first one is named as L1-SCCA which is contained in the Penalized Multivariate Analysis software toolkit [3], and the second one is named as L1-S2CCA by considering each group consists of only one covariate [6].

B. Parameter Tuning

Take the $\mathbf{u}$ -update as an example, we have four parameters λ , β , τ , and ρ , and two unknown variables $\mathbf{z}$, μ to be decided in advance. $\mathbf{z}$ is a slack variable which is split from $\mathbf{u}$. We simply set $\mathbf{z} = \mathbf{0}$. μ is the dual variable and is set to $\mathbf{0}$ too. For the remaining parameters, we further reduce the computational burden based on some heuristic tricks. First, we set the dual update step size ρ to a conservative value to permit a good accuracy. Second, τ is a key parameter in TLP-SCCA, thus it is reasonable to use an $\mathbf{u}$ -related τ . We set $\tau = E(|\mathbf{u}|) + T^* \sigma(|\mathbf{u}|)$ and update it during the iteration, where $E(|\mathbf{u}|)$ and $\sigma(|\mathbf{u}|)$ are the *mean* and *deviation* of $|\mathbf{u}|$. T is a constant which controls τ , and further controls how many u_i 's are penalized by ℓ_1 -norm. According to normal distribution, $T = 1.96$ indicating that 95% of $|\mathbf{u}|$ are smaller than τ . This will yield a reasonable sparse result. At last, we only have two parameters, i.e. λ_1 and β_1 , left. Likewise, this is the same situation for $\mathbf{v}$ -update. Thus we have four parameters to be tuned and utilize the **nested** 3-fold cross-validation to tune them. In this study, we tune the parameters from the range of 10^θ with $\theta \in [-3, -2, -1, 0, 1, 2, 3]$. All methods utilize the same partition and run on the same platform to ensure a fair comparison. The stopping criteria is $|\text{obj}(\mathbf{u})^k - \text{obj}(\mathbf{u})_{k+1}| / \text{obj}(\mathbf{u})_k \leq \varepsilon$, where $\varepsilon = 10^{-5}$ is a desirable error bound in this paper.

C. Results on Synthetic Data

Two synthetic datasets are used in the experiment. Both datasets are generated by the following steps. 1) We first set up two vectors $\mathbf{u}$ with length p and $\mathbf{v}$ with length q separately. We then let most of elements of both $\mathbf{u}$ and $\mathbf{v}$ be zero. 2) We introduce a latent variable $\mathbf{z}$ with $\mathbf{z} \sim N(\mathbf{0}, \mathbf{I}_{n \times n})$. 3) We generate $\mathbf{X}$ with each element $\mathbf{x}_i \sim N(z_i \mathbf{u}, \sigma)$, where σ stands for the standard deviation of noise. We also create $\mathbf{Y}$ with $\mathbf{y}_i \sim N(z_i \mathbf{v}, \sigma)$.

Fig. 2 shows the estimated canonical loadings. Both TLP-SCCA and rTLP-SCCA obtain cleaner canonical loading patterns than the competing methods. Our methods have smaller estimation bias regarding the canonical loadings. Table I shows results of the estimated correlation coefficients. We observe that three methods, i.e. the L1-SCCA, TLP-SCCA and rTLP-SCCA, perform similarly. The results on synthetic datasets reveal that although TLP-SCCA and rTLP-SCCA do not hold the highest correlation coefficients, they obtain the most unbiased canonical loadings.

D. Results on Real Brain Imaging Genetics Data

We obtained the real imaging genetics data from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database (adni.loni.usc.edu). The primary goal of conducting ADNI was to test whether serial magnetic resonance imaging (MRI), positron emission tomography (PET), other biological markers, and clinical and neuropsychological assessment can be combined to help detect the mild cognitive impairment (MCI) and early Alzheimer's disease (AD). Please see www.adni-info.org for up-to-date information.

The genotyping and baseline amyloid imaging data of 283 non-Hispanic Caucasian participants were downloaded from the ADNI website. The imaging data were preprocessed in accordance with [7], and then the effects of the baseline age, gender, education and handedness were pre-adjusted by regression. After that we extracted 191 ROI level mean amyloid measurements. The genotyping data contain 58 SNP markers from the AD-related genes. Our goal is to verify if TLP-SCCA can find out the meaningful associations between the amyloid data and the SNP data.

We first present the heat map regarding the estimated canonical loadings in Fig. 3. On the genotyping data (left panel), we clearly see that the L1-SCCA, TLP-SCCA and rTLP-SCCA find out 5 out of 58 SNPs. The five SNPs are rs429358 and rs769449 from *APOE*, rs157582 and rs2075650 from *TOMM40*, rs4420638 from *APOC1*. Note that all the five SNPs have been demonstrated to be highly correlated with AD. The L1-S2CCA only discovers the SNP rs429358, which means it loses the useful information carried by other four SNPs. On the imaging data (right panel), both TLP-SCCA and rTLP-SCCA identify only one strong signal, showing very clear and clean canonical loading pattern. The identified frontal region in the brain is known to be related with AD [20]. The results indicate that the TLP-SCCA can identify meaningful associations between genetic data and imaging data.

We also show the estimated correlation coefficients in Table II. There is no significant difference among various methods. This phenomenon reveals that generating a clean canonical pattern does not always correspond to a higher correlation coefficient [10]. Nevertheless, the TLP-SCCA and rTLP-SCCA have the potential to balance between the overpenalization and the overfitting, indicating their promising power in imaging genetic studies.

V. Conclusion

In this paper, we proposed the TLP-SCCA coupled with two effective optimization algorithms, i.e. TLP-SCCA and rTLP-SCCA. They identified cleaner canonical loading patterns which reduced the estimation bias of ℓ_1 -SCCA. We analyzed their convergence and computation complexity.

We evaluated our methods and two ℓ_1 -SCCA methods on both synthetic and real data. Both TLP-SCCA and rTLP-SCCA identified cleaner canonical loadings, indicating they successfully and accurately recognized the signals which were the closest to the ground truth. The results on the real imaging genetic data showed that TLP-SCCA and rTLP-SCCA yielded more meaningful canonical loading patterns than the benchmarks. Possible future directions include (1) investigating the performance trajectory under different τ 's, and (2) using relevant structure information to guide the bivariate association identification.

Acknowledgments

Data used in preparation of this article were obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database (adni.loni.usc.edu). As such, the investigators within the ADNI contributed to the design and implementation of ADNI and/or provided data but did not participate in analysis or writing of this report. A complete listing of ADNI investigators is at: http://adni.loni.usc.edu/wp-content/uploads/how_to_apply/ADNI_Acknowledgement_List.pdf.

Funding: This research was supported by the National Natural Science Foundation of China (No. 61602384, No. 31500798 and No. 31671005), and Scientific Research Start-up Foundation (No. G2016KY0305) at Northwestern Polytechnical University. This work was also supported by NIH R01 EB022574, R01 LM011360, U01 AG024904 (<http://adni.loni.usc.edu>), RC2 AG036535, R01 AG19771, P30 AG10133, UL1 TR001108, R01 AG 042437, and R01 AG046171; DOD W81XWH-14-2-0151, W81XWH-13-1-0259, and W81XWH-12-2-0012; NCAA 14132004; and CTSI SPARC Program.

References

1. Kim S, Swaminathan S, Inlow M, Risacher SL, Nho K, et al. Influence of genetic variation on plasma protein levels in older adults using a multi-analyte panel. *PLoS One*. 2013; 8(7):e70269. [PubMed: 23894628]
2. Saykin AJ, Shen L, Yao X, Kim S, Nho K, et al. Genetic studies of quantitative mci and ad phenotypes in adni: Progress, opportunities, and plans. *Alzheimer's & Dementia*. 2015; 11(7):792–814.
3. Witten DM, Tibshirani R, Hastie T. A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis. *Biostatistics*. 2009; 10(3):515–34. [PubMed: 19377034]
4. Parkhomenko E, Tritchler D, Beyene J. Sparse canonical correlation analysis with application to genomic data integration. *Statistical Applications in Genetics and Molecular Biology*. 2009; 8(1):1–34.
5. Chen X, Liu H. An efficient optimization algorithm for structured sparse cca, with applications to eqtl mapping. *Statistics in Biosciences*. 2012; 4(1):3–26.
6. Du L, Yan J, Kim S, Risacher SL, Huang H, et al. A novel structure-aware sparse learning algorithm for brain imaging genetics. *MICCAI*. 2014:329–336. [PubMed: 25320816]
7. Du L, Huang H, Yan J, Kim S, Risacher SL, et al. Structured sparse canonical correlation analysis for brain imaging genetics: An improved graphnet method. *Bioinformatics*. 2016; 32(10):1544–1551. [PubMed: 26801960]
8. Shen L, Thompson PM, Potkin SG, Bertram L, Farrer LA, et al. Genetic analysis of quantitative phenotypes in ad and mci: imaging, cognition and biomarkers. *Brain imaging and behavior*. 2014; 8(2):183–207. [PubMed: 24092460]
9. Fung G, Mangasarian O. Equivalence of minimal l_0 - and l_1 -norm solutions of linear equalities, inequalities and linear programs for sufficiently small p . *Journal of optimization theory and applications*. 2011; 151(1):1–10.
10. Shen X, Pan W, Zhu Y. Likelihood-based selection and sharp parameter estimation. *Journal of the American Statistical Association*. 2012; 107(497):223–232. [PubMed: 22736876]
11. Yang, S., Yuan, L., Lai, YC., Shen, X., Wonka, P., Ye, J. SIGKDD. ACM; 2012. Feature grouping and selection over an undirected graph; p. 922-930.
12. Boyd S, Parikh N, Chu E, Peleato B, Eckstein J. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*. 2011; 3(1):1–122.
13. Zhang T, et al. Multi-stage convex relaxation for feature selection. *Bernoulli*. 2013; 19(5B):2277–2293.
14. Tao, PD., et al. Trends in Mathematical Optimization. Springer; 1988. Duality in dc (difference of convex functions) optimization. subgradient methods; p. 277-293.
15. Gorski J, Pfeuffer F, Klamroth K. Biconvex sets and optimization with biconvex functions: a survey and extensions. *Mathematical Methods of Operations Research*. 2007; 66(3):373–407.
16. Goldstein T, O'Donoghue B, Setzer S, Baraniuk R. Fast alternating direction optimization methods. *SIAM Journal on Imaging Sciences*. 2014; 7(3):1588–1623.
17. Chen X, Liu H, Carbonell JG. Structured sparse canonical correlation analysis. *AISTATS*. 2012
18. Chen J, Bushman FD, et al. Structure-constrained sparse canonical correlation analysis with an application to microbiome data analysis. *Biostatistics*. 2013; 14(2):244–258. [PubMed: 23074263]
19. Du L, Huang H, Yan J, Kim S, Risacher SL, et al. Structured sparse cca for brain imaging genetics via graph oscar. *BMC Systems Biology*. 10(3)

20. Stuss DT, Gow CA, Hetherington CR. "no longer gage": frontal lobe dysfunction and emotional changes. *Journal of consulting and clinical psychology*. 1992; 60(3):349. [PubMed: 1619089]

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Figure 1.

Visualization of the ℓ_1 -norm ball (left), TLP ball with $\tau = \frac{1}{4}$ and $\tau = \frac{1}{6}$ (middle), and ℓ_2 -norm ball (right).

Figure 2.

Canonical loadings estimated on synthetic datasets. The first two columns are for Dataset 1, and the second two columns are for Dataset2. The first row is the ground truth, and each remaining one corresponds to a method: (1) Ground Truth. (2) L1-SCCA. (3) L1-S2CCA. (4) TLP-SCCA. (5) rTLP-SCCA.

Figure 3. Heat map of estimated canonical loadings of different SCCA methods. The left panel shows the estimated weights of $\mathbf{u}$, and the right one shows the weights of $\mathbf{v}$.

Table I

Results on synthetic data: the estimated correlation coefficients of each fold and their MEAN are shown.

Training Results					
	Dataset 1	MEAN	Dataset 2	MEAN	
L1-SCCA	0.99 0.99 0.98	0.99	0.98 0.99 0.98	0.98	
L1-S2CCA	1.00 1.00 1.00	1.00	1.00 1.00 1.00	1.00	
TLP-SCCA	0.99 0.99 0.99	0.99	0.98 0.97 0.98	0.98	
rTLP-SCCA	0.99 0.99 0.99	0.99	0.98 0.98 0.98	0.98	
Testing Results					
L1-SCCA	0.97 0.99 0.99	0.98	0.98 0.98 0.98	0.98	
L1-S2CCA	0.74 0.85 0.89	0.83	0.92 0.82 0.91	0.89	
TLP-SCCA	0.96 0.99 0.98	0.98	0.91 0.84 0.91	0.89	
rTLP-SCCA	0.96 0.98 0.99	0.98	0.98 0.97 0.98	0.98	

Table II

Results on real imaging genetics data: the estimated correlation coefficients of each individual fold and their MEAN are shown.

Methods	Training Results	MEAN	Testing Results	MEAN
L1-SCCA	0.46 0.49 0.44	0.46	0.45 0.39 0.49	0.44
L1-S2CCA	0.47 0.51 0.46	0.48	0.48 0.43 0.51	0.47
TLP-SCCA	0.45 0.47 0.44	0.45	0.44 0.38 0.50	0.44
rTLP-SCCA	0.45 0.47 0.43	0.45	0.43 0.38 0.49	0.43