
A Simple and Provable Algorithm for Sparse Diagonal CCA

Megasthenis Asteris
Anastasios Kyrillidis

Department of Electrical and Computer Engineering, The University of Texas at Austin

MEGAS@UTEXAS.EDU
ANASTASIOS@UTEXAS.EDU

Oluwasanmi Koyejo

Stanford University & University of Illinois at Urbana-Champaign

SANMI@ILLINOIS.EDU

Russell Poldrack

Department of Psychology, Stanford University

POLDRACK@STANFORD.EDU

Abstract

Given two sets of variables, derived from a common set of samples, sparse Canonical Correlation Analysis (CCA) seeks linear combinations of a small number of variables in each set, such that the induced *canonical* variables are maximally correlated. Sparse CCA is NP-hard.

We propose a novel combinatorial algorithm for sparse diagonal CCA, *i.e.*, sparse CCA under the additional assumption that variables within each set are standardized and uncorrelated. Our algorithm operates on a low rank approximation of the input data and its computational complexity scales linearly with the number of input variables. It is simple to implement, and parallelizable. In contrast to most existing approaches, our algorithm administers precise control on the sparsity of the extracted canonical vectors, and comes with theoretical data-dependent global approximation guarantees, that hinge on the spectrum of the input data. Finally, it can be straightforwardly adapted to other constrained variants of CCA enforcing structure beyond sparsity.

We empirically evaluate the proposed scheme and apply it on a real neuroimaging dataset to investigate associations between brain activity and behavior measurements.

1. Introduction

One of the key objectives in cognitive neuroscience is to localize cognitive processes in the brain, and understand

their role in human behavior, as measured by psychological scores and physiological measurements (Posner et al., 1988). This mapping may be investigated by using functional neuroimaging techniques to measure brain activation during carefully designed experimental tasks (Poldrack, 2006). Following the experimental manipulation, a joint analysis of brain activation and behavioral measurements across subjects can reveal associations that exist between the two (Berman et al., 2006).

Similarly, in genetics and molecular biology, several studies involve the joint analysis of multiple assays performed on a single group of patients (Pollack et al., 2002; Morley et al., 2004; Stranger et al., 2007). If DNA variants and gene expression measurements are simultaneously available for a set of tissue samples, a natural objective is to identify correlations between the expression levels of gene subsets and variation in the related genes.

Canonical Correlation Analysis (CCA) (Hotelling, 1936) is a classic method for discovering such linear relationships across two sets of variables and has been extensively used to investigate associations between multiple views of the same set of observations; *e.g.*, see (Deleus & Van Hulle, 2011; Li et al., 2012; Smith et al., 2015) in neuroscience. Given two datasets \mathbf{X} and \mathbf{Y} of dimensions $k \times m$ and $k \times n$, respectively, on k common samples, CCA seeks linear combinations of the original variables of each type that are maximally correlated. More formally, the objective is to compute a pair of *canonical vectors* (or *weights*) \mathbf{u} and \mathbf{v} such that the *canonical variables* $\mathbf{X}\mathbf{u}$ and $\mathbf{Y}\mathbf{v}$ achieve the maximum possible correlation:¹

$$\max_{\mathbf{u}, \mathbf{v} \neq \mathbf{0}} \frac{\mathbf{u}^\top \Sigma_{\mathbf{X}\mathbf{Y}} \mathbf{v}}{(\mathbf{u}^\top \Sigma_{\mathbf{X}\mathbf{X}} \mathbf{u})^{1/2} (\mathbf{v}^\top \Sigma_{\mathbf{Y}\mathbf{Y}} \mathbf{v})^{1/2}}. \quad (1)$$

¹ We assume that the variables in \mathbf{X} and \mathbf{Y} are standardized, *i.e.*, each column has zero mean and has been scaled to have unit standard deviation.

The optimal canonical pair can be computed via a generalized eigenvalue decomposition involving the empirical estimates of the (cross-) covariance matrices in (1).

Imaging and behavioral measurements in cognitive neuroscience, similar to genomic data in bioinformatics, typically involve hundreds of thousands of variables with only a limited number of samples. In that case, the CCA objective in (1) is ill-posed; it is always possible to design canonical variables for which the factors in the denominator vanish, irrespective of the data. Model regularization via constraints such as sparsity, not only improves the interpretability of the extracted canonical vectors, but is critical for enabling the recovery of meaningful results.

Sparse CCA seeks to maximize the correlation between subsets of variables of each type, while performing variable selection. In this work, we consider the following sparse diagonal CCA problem, similar to (Witten et al., 2009)²:

$$\text{(Sparse CCA)} \quad \max_{\mathbf{u} \in \mathcal{U}, \mathbf{v} \in \mathcal{V}} \mathbf{u}^\top \Sigma_{XY} \mathbf{v}, \quad (2)$$

where

$$\begin{aligned} \mathcal{U} &= \{\mathbf{u} \in \mathbb{R}^m : \|\mathbf{u}\|_2 = 1, \|\mathbf{u}\|_0 \leq s_x\}, \\ \mathcal{V} &= \{\mathbf{v} \in \mathbb{R}^n : \|\mathbf{v}\|_2 = 1, \|\mathbf{v}\|_0 \leq s_y\}, \end{aligned} \quad (3)$$

for given parameters s_x and s_y . The $m \times n$ argument $\Sigma_{XY} = \mathbf{X}^\top \mathbf{Y}$ is the empirical estimation of the cross-covariance matrix between the variables in the two views \mathbf{X} and \mathbf{Y} , and it is the input to the optimization problem. Note that besides the introduced sparsity requirement, (2) is obtained from (1) treating the covariance matrices Σ_{XX} and Σ_{YY} as identity matrices, which is common in high dimensions (Dudoit et al., 2002; Tibshirani et al., 2003). Equivalently, we implicitly assume that the original variables within each view are standardized and uncorrelated.

The maximization in (2) is a sparse Singular Value Decomposition (SVD) problem. Disregarding the ℓ_0 cardinality constraint, although the objective is non-convex, the optimal solution of (2) can be easily computed, as it coincides with the leading singular vectors of the input matrix Σ_{XY} . The constraint on the number of nonzero entries of \mathbf{u} and \mathbf{v} , however, renders the problem NP-hard, as it can be shown by a reduction to the closely related sparse PCA problem; see Appendix 6. Several heuristics have been developed to obtain an approximate solution (see Section 1.2).

Finally, we note that sparsity alone may be insufficient to obtain interpretable results; genes participate in groups in biological pathways, and brain activity tends to be localized forming connected components over an underlying network. If higher order structural information is available on a physical system, it is meaningful incorporate that in

²(Witten et al., 2009) consider a relaxation of (2) where the ℓ_0 cardinality constraint on \mathbf{u} and \mathbf{v} is replaced by a threshold on the sparsity inducing ℓ_1 -norm.

the optimization (2). We can then consider *Structured* variants of diagonal CCA (2), by appropriately modifying the feasible regions \mathcal{U}, \mathcal{V} in (3) to reflect the desired structure.

1.1. Our contributions

We present a novel and efficient combinatorial algorithm for sparse diagonal CCA in (2). The main idea is to reduce the exponentially large search space of candidate supports of the canonical vectors, by exploring a low-dimensional principal subspace of the input data. Our algorithm runs in polynomial time –in fact linear– in the dimension of the input. It administers precise control over the sparsity of the extracted canonical vectors and can extract components for multiple sparsity values on a single run. It is simple and trivially parallelizable; we empirically demonstrate that it achieves an almost linear speedup factor in the number of available processing units.

The algorithm is accompanied with theoretical data-dependent global approximation guarantees with respect to the CCA objective (2); this is the first approach with this kind of global guarantees. The latter depend on the rank r of the low-dimensional space and the spectral decay of the input matrix Σ_{XY} . The main weakness is an exponential dependence of the computational complexity on the accuracy parameter r . In practice, however, disregarding the theoretical approximation guarantees, our algorithm can be executed for any allowable time window.

Finally, we note that our approach is similar to that of (Asteris et al., 2014) for sparse PCA. The latter has a similar formulation with (2) but is restricted to a positive semidefinite argument Σ_{XY} . Our main technical contribution is extending those algorithmic ideas and developing theoretical approximation guarantees for the bilinear maximization (2), where the input matrix can be arbitrary.

1.2. Related Work

Sparse CCA is closely related to sparse PCA; the latter can be formulated as in (2) but the argument Σ_{XY} is replaced by a positive semidefinite matrix. There is a large volume of work on sparse PCA –see (Zou et al., 2006; Amini & Wainwright, 2008) and references therein– but these methods cannot be generalized to the CCA problem. One exception is the work of (d’Aspremont et al., 2007) where the authors discuss extensions to the “non-square case”. Their approach relies on a semidefinite relaxation.

References to sparsity in CCA date back to (Thorndike, 1976) and (Thompson, 1984) who identified the importance of sparsity regularization to obtain meaningful results. However, no specific algorithm was proposed. Several subsequent works considered a penalized version of the CCA problem in (1), typically under a Lagrangian formulation

involving a convex relaxation of the ℓ_0 cardinality constraint (Torres et al., 2007; Hardoon & Shawe-Taylor, 2007; 2011). (Chu et al., 2013) characterize the solutions of the unconstrained problem and formulate convex ℓ_1 minimization problems to seek sparse solutions in that set. (Sriperumbudur et al., 2009) consider a constrained generalized eigenvalue problem, which partially captures sparse CCA, and frame it as a difference-of-convex functions program. (Wiesel et al., 2008) proposed an efficient greedy procedure that gradually expands the supports of the canonical vectors. Unlike other methods, this greedy approach allows precise control of the sparsity of the extracted components.

(Witten et al., 2009; Parkhomenko et al., 2009) formulate sparse CCA as the optimization (2) and in particular considered an ℓ_1 relaxation of the ℓ_0 cardinality constraint. They suggest an alternating minimization approach exploiting the bi-convex nature of the relaxed problem, solving a lasso regression in each step. The same approach is followed in (Waaijenborg et al., 2008) combining ℓ_2 and ℓ_1 regularizers similarly to the elastic net approach for sparse PCA (Zou & Hastie, 2005). Similar approaches have appeared in the literature for *sparse SVD* (Yang et al., 2011; Lee et al., 2010). A common weakness in these approaches is the lack of precise control over sparsity: the mapping between the regularization parameters and the number of nonzero entries in the extracted components is highly non-linear. Further, such methods usually lack provable non-asymptotic approximation guarantees. Beyond sparsity, (Witten et al., 2009; Witten & Tibshirani, 2009) discuss alternative penalizations such as fused lasso to impose additional structure, while (Chen et al., 2012) introduce a group-lasso to promote sparsity with structure.

Finally, although a review of CCA applications is beyond the scope of this manuscript, we simply note that CCA is considered a promising approach for scientific research as evidenced by several recent works in the literature, *e.g.*, in neuroscience (Rustandi et al., 2009; Deleus & Van Hulle, 2011; Li et al., 2012; Lin et al., 2014; Smith et al., 2015).

2. SpanCCA: An Algorithm for Sparse Diagonal CCA

We begin this section with a brief discussion of the problem and the key ideas behind our approach. Next, we provide an overview of SpanCCA and the accompanying approximation guarantees and conclude with a short analysis.

2.1. Intuition

The hardness of the sparse CCA problem (2) lies in the detection of the optimal supports for the canonical vectors. In the *unconstrained* problem, where only a unit ℓ_2 -norm constraint is imposed on \mathbf{u} and \mathbf{v} , the optimal CCA pair

coincides with the top singular vectors of the input argument Σ_{XY} . In the sparse variant, if the optimal supports for \mathbf{u} and \mathbf{v} were known, computing the optimal solution would be straightforward: the nonzero subvectors of \mathbf{u} and \mathbf{v} would coincide with the leading singular vectors of the $s_x \times s_y$ submatrix of Σ_{XY} , indexed by the two support sets. Hence, the bottleneck lies in determining the optimal supports for \mathbf{u} and \mathbf{v} .

Exhaustive search A straightforward, brute-force approach is to exhaustively consider all possible supports for \mathbf{u} and \mathbf{v} ; for each candidate pair solve the unconstrained CCA problem on the restricted input, and determine the supports for which the objective (2) is maximized. Albeit optimal, this procedure is intractable as the number of candidate supports $\binom{m}{s_x} \binom{n}{s_y}$ is overwhelming even for small values of s_x and s_y .

Thresholding On the other hand, a feasible pair of sparse canonical vectors \mathbf{u} , \mathbf{v} can be extracted by hard-thresholding the solution to the unconstrained problem, *i.e.*, computing the leading singular vectors \mathbf{u} , \mathbf{v} of Σ_{XY} , suppressing to zero all but the s_x and s_y largest in magnitude entries, respectively, and rescaling to obtain a unit ℓ_2 -norm solution. Essentially, this heuristic resorts to unconstrained CCA for a *guided* selection of the sparse support.

Proposed method Our sparse CCA algorithm covers the ground between these two approaches. Instead of relying on the solution to the unconstrained problem for the choice of the sparse supports, it explores a principal subspace of the input matrix Σ_{XY} , spanned by its leading $r \geq 1$ singular vector pairs. For $r = 1$, its output coincides with that of the thresholding approach, while for $r = \min\{m, n\}$ it approximates that of exhaustive search.

Effectively, we solve (2) on a rank- r approximation of the input Σ_{XY} . The key observation is that the low inner dimension of the argument matrix can be exploited to substantially reduce the search space: our algorithm identifies an (approximately) optimal pair of supports for the low rank sparse CCA problem, without considering the entire collection of possible supports of cardinalities s_x and s_y .

2.2. Overview and Guarantees

SpanCCA is outlined in Algorithm 1. The first step is to compute a rank- r approximation \mathbf{B} of the input Σ_{XY} , where r is an accuracy input parameter, via the truncated singular value decomposition (SVD) of Σ_{XY} .³ From that point on, the algorithm operates exclusively on \mathbf{B} effec-

³ The low-rank approximation can be computed using faster randomized approaches; see (Halko et al., 2011). Here, for simplicity, we consider the exact case.

tively solving a low-rank sparse diagonal CCA problem:

$$\max_{\mathbf{u} \in \mathcal{U}, \mathbf{v} \in \mathcal{V}} \mathbf{u}^\top \mathbf{B} \mathbf{v}, \quad (4)$$

where $\mathcal{U} \subseteq \mathbb{S}_2^{m-1}$ and $\mathcal{V} \subseteq \mathbb{S}_2^{n-1}$ are defined in (3). As we discuss in the next section, we can consider other constrained variants of CCA on potentially arbitrary, non-convex sets. We do require, however, that there exist procedures to (at least approximately) solve the maximizations

$$P_{\mathcal{U}}(\mathbf{a}) \triangleq \arg \max_{\mathbf{u} \in \mathcal{U}} \mathbf{a}^\top \mathbf{u}, \quad (5)$$

$$P_{\mathcal{V}}(\mathbf{b}) \triangleq \arg \max_{\mathbf{v} \in \mathcal{V}} \mathbf{b}^\top \mathbf{v}, \quad (6)$$

for any given vectors $\mathbf{a} \in \mathbb{R}^{m \times 1}$ and $\mathbf{b} \in \mathbb{R}^{n \times 1}$. Fortunately, this is the case for the sets of sparse unit ℓ_2 -norm vectors. Algorithm 2 outlines an efficient $O(m)$ procedure that given $\mathbf{a} \in \mathbb{R}^{m \times 1}$ computes an exact solution to (5) with at most $s \leq m$ nonzero entries: first it determines the s largest (in magnitude) entries of \mathbf{a} (breaking ties arbitrarily), it zeroes out the remaining entries and re-scales the output to meet the ℓ_2 -norm requirement.

The main body of Algorithm 1 consists of a single iteration. In the i th round, it independently samples a point, or equivalently direction, \mathbf{c}_i from the r -dimensional unit ℓ_2 sphere and uses it to appropriately sample a point \mathbf{a}_i in the range of \mathbf{B} . The latter is then used to compute a feasible solution pair $\mathbf{u}_i, \mathbf{v}_i$ via a two-step procedure: first the algorithm computes \mathbf{u}_i by “projecting” \mathbf{a}_i onto \mathcal{U} invoking Alg. 2 as a subroutine to solve maximization (5), and then computes \mathbf{v}_i by projecting $\mathbf{b}_i = \mathbf{B}^\top \mathbf{u}_i$ onto \mathcal{V} in a similar fashion. The algorithm repeats this procedure for T rounds and outputs the pair that achieves the maximum objective value in (4). We emphasize that consecutive rounds are completely independent and can be executed in parallel.

For a sufficiently large number T of rounds (or samples) the procedure guarantees that the output pair will be approximately optimal in terms of the objective for the low-rank problem (4). That, it turn, translates to approximation guarantees for the full-rank sparse CCA problem (2):

Theorem 1. *For any real $m \times n$ matrix Σ_{XY} , $\epsilon \in (0, 1)$, and $r \leq \max\{m, n\}$, Algorithm 1 with input Σ_{XY} , r , and $T = \tilde{O}(2^{r \cdot \log_2(2/\epsilon)})$ outputs $\mathbf{u}_\# \in \mathcal{U}$ and $\mathbf{v}_\# \in \mathcal{V}$ such that*

$$\mathbf{u}_\#^\top \Sigma_{XY} \mathbf{v}_\# \geq \mathbf{u}_*^\top \Sigma_{XY} \mathbf{v}_* - \epsilon \cdot \sigma_1(\Sigma_{XY}) - 2\sigma_{r+1}(\Sigma_{XY}),$$

in time $T_{\text{SVD}}(r) + O(T \cdot (T_{\mathcal{U}} + T_{\mathcal{V}} + r \cdot \max\{m, n\}))$.

Here, \mathbf{u}_* and \mathbf{v}_* denote the unknown optimal pair of canonical vectors satisfying the desired constraints. $T_{\text{SVD}}(r)$ denotes the time to compute the rank- r truncated SVD of the input Σ_{XY} , while $T_{\mathcal{U}}$ and $T_{\mathcal{V}}$ denote the time required to compute the maximizations (5) and (6), respectively, which in the case of Alg. 2 are linear in the dimensions m and n .

Algorithm 1 SpanCCA

input : Σ_{XY} , a real $m \times n$ matrix.

$r \in \mathbb{N}_+$, the rank of the approximation to be used.

$T \in \mathbb{N}_+$, the number of samples/iterations.

output $\mathbf{u}_\# \in \mathcal{U}, \mathbf{v}_\# \in \mathcal{V}$

```

1:  $\mathbf{U}, \Sigma, \mathbf{V} \leftarrow \text{SVD}(\Sigma_{XY}, r)$            {  $\mathbf{B} \leftarrow \mathbf{U}\Sigma\mathbf{V}^\top$  }
2: for  $i = 1, \dots, T$  do
3:    $\mathbf{c}_i \leftarrow \text{randn}(r)$                        {  $\sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{r \times r})$  }
4:    $\mathbf{c}_i \leftarrow \mathbf{c}_i / \|\mathbf{c}_i\|_2$ 
5:    $\mathbf{a}_i \leftarrow \mathbf{U}\Sigma\mathbf{c}_i$                        {  $\mathbf{a}_i \in \mathbb{R}^m$  }
6:    $\mathbf{u}_i \leftarrow \arg \max_{\mathbf{u} \in \mathcal{U}} \mathbf{a}_i^\top \mathbf{u}$        {  $P_{\mathcal{U}}(\cdot)$  }
7:    $\mathbf{b}_i \leftarrow \mathbf{V}\Sigma\mathbf{U}^\top \mathbf{u}_i$              {  $\mathbf{b}_i \in \mathbb{R}^n$  }
8:    $\mathbf{v}_i \leftarrow \arg \max_{\mathbf{v} \in \mathcal{V}} \mathbf{b}_i^\top \mathbf{v}$      {  $P_{\mathcal{V}}(\cdot)$  }
9:    $\text{obj}_i \leftarrow \mathbf{b}_i^\top \mathbf{v}_i$ 
10: end for
11:  $i_0 \leftarrow \arg \max_{i \in [T]} \text{obj}_i$ 
12:  $(\mathbf{u}_\#, \mathbf{v}_\#) \leftarrow (\mathbf{u}_{i_0}, \mathbf{v}_{i_0})$ 
    
```

Algorithm 2 $P_{\mathcal{U}}(\cdot)$ for $\mathcal{U} \triangleq \{\mathbf{u} \in \mathbb{S}_2^{m-1} : \|\mathbf{u}\|_0 \leq s\}$

input : $\mathbf{a} \in \mathbb{R}^{d \times 1}$.

output $\mathbf{u}_0 = \arg \max_{\mathbf{u} \in \mathcal{U}} \mathbf{a}^\top \mathbf{u}$

```

1:  $\mathbf{u}_0 \leftarrow \mathbf{0}_{d \times 1}$ 
2:  $t \leftarrow \text{index of sth order element of } \text{abs}(\mathbf{a})$ 
3:  $\mathcal{I} \leftarrow \{i : |a_i| \geq |a_t|\}$ 
4:  $\mathbf{u}_0[i] \leftarrow \mathbf{a}[i], \forall i \in \mathcal{I}$ 
5:  $\mathbf{u}_0 \leftarrow \mathbf{u}_0 / \|\mathbf{u}_0\|_2$ 
    
```

The first term in the additive error is due to the sampling approach of Alg. 1. The second term is due to the fact that the algorithm operates on the rank- r surrogate matrix \mathbf{B} . Theorem 1 establishes a trade-off between the computational complexity of Alg. 1 and the quality of the approximation guarantees: the latter improves as r increases, but the former depends exponentially in r .

Finally, in the special case where we impose sparsity constraints on only one of the two variables, say \mathbf{u} , while allowing the second variable, here \mathbf{v} , to be any vector with unit ℓ_2 norm, we obtain stronger guarantees.

Theorem 2. *If $\mathcal{V} = \{\mathbf{v} : \|\mathbf{v}\|_2 = 1\}$, i.e., if no constraint is imposed on variable \mathbf{v} besides unit length, then Algorithm 1 under the same configuration as that in Theorem 1 outputs $\mathbf{u}_\# \in \mathcal{U}$ and $\mathbf{v}_\# \in \mathcal{V}$ such that*

$$\mathbf{u}_\#^\top \Sigma_{XY} \mathbf{v}_\# \geq (1 - \epsilon) \cdot \mathbf{u}_*^\top \Sigma_{XY} \mathbf{v}_* - 2 \cdot \sigma_{r+1}(\Sigma_{XY}).$$

Theorem 2 implies that due to the flexibility in the choice of the canonical vector \mathbf{v} , Alg. 1 solves the low-rank problem (4) within a multiplicative $(1 - \epsilon)$ -factor from the optimal; the extra additive error term is once again due to the fact that the algorithm operates on the rank- r approximation \mathbf{B} instead of the original input Σ_{XY} . In this case, the optimal choice of \mathbf{v} in (6) is just a scaled version of the

argument \mathbf{b} . Finally, we note that if constraints need to be applied on \mathbf{v} instead of \mathbf{u} , then the same guarantees can be obtained by applying Algorithm 1 on Σ_{XY}^\top . A formal proof for Theorem 2 is provided in the Appendix, Sec. 7.

Overall, SpanCCA is simple to implement and is trivially parallelizable: the main iteration can be split across an arbitrary number of processing units achieving a potentially linear speedup. It is the first algorithm for sparse diagonal CCA with data-dependent global approximation guarantees. As discussed in Theorem 1, the input accuracy parameter r establishes a trade-off between the running time and the tightness of the theoretical guarantees. Its complexity scales linearly in the dimensions of the input for any constant r , but admittedly becomes prohibitive even for moderate values of r . In practice, if the spectrum of the data exhibits sharp decay, we may be able to obtain useful approximation guarantees even for small values of r such as 2 or 3. Moreover, disregarding the theoretical guarantees, the algorithm can always be executed for any rank r and an arbitrary number of iterations T . In Section 4, we empirically show that moderate values for r and T can potentially achieve better solutions compared to state of the art. We note, however, that the tuning of those parameters needs to be investigated.

2.3. Analysis

Let $\mathbf{B} = \mathbf{U}\Sigma\mathbf{V}^\top$, and $\mathbf{u}_{(B)}, \mathbf{v}_{(B)}$ be a pair that maximizes –not necessarily uniquely– the objective $\mathbf{u}^\top \mathbf{B} \mathbf{v}$ in (4) over all feasible solutions. We assume that $\mathbf{u}_{(B)}^\top \mathbf{B} \mathbf{v}_{(B)} > 0.4$. Define the $r \times 1$ vector $\mathbf{c}_{(B)} \triangleq \mathbf{V}^\top \mathbf{v}_{(B)}$ and let ρ denote its ℓ_2 norm. Then, $0 < \rho \leq 1$; the upper bound follows from the fact that the r columns of \mathbf{V} are orthonormal and $\|\mathbf{v}_{(B)}\|_2 = 1$, while the lower follows by the aforementioned assumption. Finally, define $\bar{\mathbf{c}}_{(B)} = \mathbf{c}_{(B)}/\rho$, the projection of $\mathbf{c}_{(B)}$ on the unit ℓ_2 -sphere \mathbb{S}_2^{r-1} .

Def. 1. For any $\epsilon \in (0, 1)$, an ϵ -net of \mathbb{S}_2^{r-1} is a finite collection N of points in \mathbb{R}^r such that for any $\mathbf{c} \in \mathbb{S}_2^{r-1}$, N contains a point \mathbf{c}' such that $\|\mathbf{c}' - \mathbf{c}\| \leq \epsilon$.

Lemma 2.1 ((Vershynin, 2010), Lemma 5.2). For any $\epsilon \in (0, 1)$, there exists an ϵ -net of \mathbb{S}_2^{r-1} equipped with the Euclidean metric, with at most $(1 + 2/\epsilon)^r$ points.

Algorithm 1 runs in an iteration with T rounds. In each round, it independently samples a point \mathbf{c}_i from \mathbb{S}_2^{r-1} , by randomly generating a vector according to a spherical Gaussian distribution and appropriately scaling its length. Based on Lemma 2.1 and elementary counting arguments, for sufficiently large T the collection of sampled points forms a ϵ -net of \mathbb{S}_2^{r-1} with high probability:

⁴Observe that this is always true for any nonzero argument \mathbf{B} as long as at least one of the two variables \mathbf{u} and \mathbf{v} can take arbitrary signs. It is hence true under vanilla sparsity constraints.

Lemma 2.2. For any $\epsilon, \delta \in (0, 1)$, a set of $T = O(r(\epsilon/4)^{-r} \cdot \ln 4/\epsilon \cdot \delta)$ randomly and independently drawn points uniformly distributed on \mathbb{S}_2^{r-1} suffices to construct an $\epsilon/2$ -net of \mathbb{S}_2^{r-1} with probability at least $1 - \delta$.

It follows that there exists $i_* \in [T]$ such that

$$\|\mathbf{c}_{i_*} - \bar{\mathbf{c}}_{(B)}\|_2 \leq \epsilon/2. \quad (7)$$

In the i_* th round, the algorithm samples the point \mathbf{c}_{i_*} and computes a feasible pair $(\mathbf{u}_{i_*}, \mathbf{v}_{i_*})$ via the two step maximization procedure, that is,

$$\mathbf{u}_{i_*} \triangleq \arg \max_{\mathbf{u} \in \mathcal{U}} \mathbf{u}^\top \mathbf{U} \Sigma \mathbf{c}_{i_*} \quad \text{and} \quad \mathbf{v}_{i_*} \triangleq \arg \max_{\mathbf{v} \in \mathcal{V}} \mathbf{u}_{i_*}^\top \Sigma_{XY} \mathbf{v}.$$

We have:

$$\begin{aligned} \mathbf{u}_{(B)}^\top \mathbf{B} \mathbf{v}_{(B)} &= \rho \cdot \mathbf{u}_{(B)}^\top \mathbf{U} \Sigma \bar{\mathbf{c}}_{(B)} \\ &= \rho \cdot \mathbf{u}_{(B)}^\top \mathbf{U} \Sigma \mathbf{c}_{i_*} + \rho \cdot \mathbf{u}_{(B)}^\top \mathbf{U} \Sigma (\bar{\mathbf{c}}_{(B)} - \mathbf{c}_{i_*}) \\ &\leq \rho \cdot \mathbf{u}_{i_*}^\top \mathbf{U} \Sigma \mathbf{c}_{i_*} + \rho \cdot \mathbf{u}_{(B)}^\top \mathbf{U} \Sigma (\bar{\mathbf{c}}_{(B)} - \mathbf{c}_{i_*}) \\ &\leq \rho \cdot \mathbf{u}_{i_*}^\top \mathbf{U} \Sigma \mathbf{c}_{i_*} + \frac{\epsilon}{2} \cdot \sigma_1(\Sigma_{XY}). \end{aligned} \quad (8)$$

The first step follows by the definition of $\bar{\mathbf{c}}_{(B)}$ and the second by linearity. The first inequality follows from the fact that \mathbf{u}_{i_*} maximizes the first term over all $\mathbf{u} \in \mathcal{U}$. The last inequality follows straightforwardly from the fact that $\|\mathbf{u}_{(B)}\|_2 = 1$ and $\rho \leq 1$ (see Lemma 8.8). Using similar arguments,

$$\begin{aligned} \rho \cdot \mathbf{u}_{i_*}^\top \mathbf{U} \Sigma \mathbf{c}_{i_*} &= \mathbf{u}_{i_*}^\top \mathbf{U} \Sigma \mathbf{V}^\top \mathbf{v}_{(B)} + \rho \cdot \mathbf{u}_{i_*}^\top \mathbf{U} \Sigma (\mathbf{c}_{i_*} - \bar{\mathbf{c}}_{(B)}) \\ &\leq \mathbf{u}_{i_*}^\top \mathbf{B} \mathbf{v}_{i_*} + \frac{\epsilon}{2} \cdot \sigma_1(\Sigma_{XY}). \end{aligned} \quad (9)$$

The inequality follows by the fact that \mathbf{v}_{i_*} maximizes the inner product with $\mathbf{B}^\top \mathbf{u}_{i_*}$ over all $\mathbf{v} \in \mathcal{V}$, as well as that $\|\mathbf{u}_{(B)}\|_2 = 1$ and $\rho \leq 1$. Combining (8) and (9), we obtain

$$\mathbf{u}_{i_*}^\top \mathbf{B} \mathbf{v}_{i_*} \geq \mathbf{u}_{(B)}^\top \mathbf{B} \mathbf{v}_{(B)} - \epsilon \cdot \sigma_1(\Sigma_{XY}). \quad (10)$$

Algorithm 1 computes multiple candidate solution pairs and outputs the pair $(\mathbf{u}_\# , \mathbf{v}_\#)$ that achieves the maximum objective value. The latter is at least as high as that achieved by $(\mathbf{u}_{i_*}, \mathbf{v}_{i_*})$.

Inequality (10) establishes an approximation guarantee for the low-rank problem (4). Those can be translated to guarantees on the original problem with input argument Σ_{XY} . Let \mathbf{u}_* and \mathbf{v}_* denote the (unknown) optimal solution of the sparse CCA problem (2). By the definition of $\mathbf{u}_{(B)}$ and $\mathbf{v}_{(B)}$, it follows that

$$\begin{aligned} \mathbf{u}_{(B)}^\top \mathbf{B} \mathbf{v}_{(B)} &\geq \mathbf{u}_*^\top \mathbf{B} \mathbf{v}_* = \mathbf{u}_*^\top \Sigma_{XY} \mathbf{v}_* - \mathbf{u}_*^\top (\Sigma_{XY} - \mathbf{B}) \mathbf{v}_* \\ &\geq \mathbf{u}_*^\top \Sigma_{XY} \mathbf{v}_* - \sigma_{r+1}(\Sigma_{XY}). \end{aligned} \quad (11)$$

Similarly,

$$\begin{aligned} \mathbf{u}_\#^\top \Sigma_{XY} \mathbf{v}_\# &= \mathbf{u}_\#^\top \mathbf{B} \mathbf{v}_\# - \mathbf{u}_\#^\top (\mathbf{B} - \Sigma_{XY}) \mathbf{v}_\# \\ &\geq \mathbf{u}_\#^\top \mathbf{B} \mathbf{v}_\# - \sigma_{r+1}(\Sigma_{XY}). \end{aligned} \quad (12)$$

Combining (11) and (12) with (10), we obtain the approximation guarantees of Theorem 1.

The running time of Algorithm 1 follows straightforwardly by inspection. The algorithm first computes the truncated singular value decomposition of inner dimension r in time denoted by $T_{\text{SVD}}(r)$. Subsequently, it performs T iterations. The cost of each iteration is determined by the cost of the matrix-vector multiplications and the running times $T_{\mathcal{U}}$ and $T_{\mathcal{V}}$ of the operators $P_{\mathcal{U}}(\cdot)$ and $P_{\mathcal{V}}(\cdot)$. Note that matrix multiplications can exploit the available matrix decomposition and are performed in time $r \cdot \max\{m, n\}$. Substituting the value of T with that specified in Lemma 2.2 completes the proof of Thm. 1. The proof of Theorem 2 follows a similar path; see Appendix Sec. 7.

3. Beyond Sparsity: Structured CCA

While enforcing sparsity results in succinct models, the latter may fall short in capturing the true interactions in a physical system, especially when the number of samples is limited. Incorporating additional prior structural information can improve interpretability⁵; *e.g.*, (Du et al., 2014) argue that a structure-aware sparse CCA incorporating group-like structure obtains biologically more meaningful results, while (Lin et al., 2014) demonstrated that group prior knowledge improved performance compared to standard sparse CCA in a task of identifying brain regions susceptible to schizophrenia. Several works suggest using structure-inducing regularizers to promote smoothness (Witten et al., 2009; Chen et al., 2013; Kobayashi, 2014) or group sparse structure (Chu et al., 2013) in CCA.

Our sparse CCA algorithm and its theoretical approximation guarantees in Theorems 1 and 2 extend straightforwardly to constraints beyond sparsity. The only assumption on the feasible sets \mathcal{U} and \mathcal{V} is that there exist tractable procedures $P_{\mathcal{U}}$ and $P_{\mathcal{V}}$ that solve the constrained maximizations (5) and (6), respectively. The specific structure of the feasible sets only manifests itself through these sub-routines, *e.g.*, Alg. 2 for the case of sparsity constraints. Therefore, Alg. 1 can be straightforwardly adapted to any structural constraint for which the aforementioned conditions are satisfied.

In fact, observe that under the unit ℓ_2 restriction on the feasible vectors, the maximizations in (5) and (6) are equivalent to computing the Euclidean projection of a given real vector on the (nonconvex) sets \mathcal{U} and \mathcal{V} . Such exact or approximate projection procedures exist for several interesting constraints beyond sparsity such as smooth or group sparsity (Huang et al., 2011; Baldassarre et al., 2013;

⁵This is a shared insight in the broader area of sparse approximations (Baraniuk et al., 2010; Huang et al., 2011; Bach et al., 2012; Kyriallidis & Cevher, 2012).

Kyriallidis et al., 2015), sparsity constraints onto norm balls (Kyriallidis et al., 2012), or even sparsity patterns guided by underlying graphs (Hegde et al., 2015; Asteris et al., 2015).

4. Experiments


We empirically evaluate our algorithm on two real datasets: *i*) a publicly available breast cancer dataset (Chin et al., 2006), also used in the evaluation of (Witten et al., 2009), and *ii*) a neuroimaging dataset obtained from the Human Connectome Project (Van Essen et al., 2013) on which we investigate associations between brain activation and behavior measurements.

4.1. Breast Cancer Dataset

The breast cancer dataset (Chin et al., 2006) consists of gene expression and DNA copy number measurements on a set of 89 tissue samples. Among others, it contains a 89×2149 matrix (DNA) with CGH spots for each sample and a 89×19672 matrix (RNA) of genes, along with information for the chromosomal locations of each CGH spot and each gene. As described in (Witten et al., 2009), this dataset can be used to perform integrative analysis of gene expression and DNA copy number data, and in particular to identify sets of genes that have expression that is correlated with a set of chromosomal gains or losses.

We run our algorithm on the breast cancer dataset and compare the output with the PMD algorithm of (Witten et al., 2009); PMD is regarded as state of the art by practitioners and has been used—in its original form or slightly modified—in several neuroscience and biomedical applications; see also Section 1.2. The input to both algorithms is the $m \times n$ matrix $\Sigma_{\text{XY}} = \mathbf{X}^T \mathbf{Y}$ ($m = 2149$, $n = 19672$), where \mathbf{X} and \mathbf{Y} are obtained from the aforementioned DNA and RNA matrices upon feature standardization. Recall that PMD is an iterative, alternating optimization scheme, where the sparsity of the extracted components \mathbf{x} and \mathbf{y} is implicitly controlled by enforcing upper bounds c_1 and c_2 on their ℓ_1 norm, respectively, with $1 \leq c_1 \leq \sqrt{m}$ and $1 \leq c_2 \leq \sqrt{n}$. Here, for simplicity, we set $c_1 = c\sqrt{m}$ and $c_2 = c\sqrt{n}$ and consider multiple values of the constant c in $(0, 1)$. Note that under this configuration, for any given value of c , we expect that the extracted components will be approximately equally sparse, relatively to their dimension.

For each c , we first run the PMD algorithm 10 times with random initializations, determine the pair of components \mathbf{x} and \mathbf{y} that achieves the highest objective value, and count the number of nonzero entries of both components as a percentage of their corresponding dimension. Subsequently, we run SpanCCA (Alg. 1) with parameters $T = 10^4$, $r = 3$, and target sparsity equal to that of the former PMD output.



	Avg Exec. Time	Configuration
PMD	~ 44 seconds	10 rand. restarts
SpanCCA	~ 24 seconds	$T = 10^4, r = 3$.

Figure 1. Comparison of SpanCCA and the PMD algorithm (Witten et al., 2009). We configure PMD with ℓ_1 -norm thresholds $c_1 = c \cdot \sqrt{m}$ and $c_2 = c \cdot \sqrt{n}$, and consider various values of the constant $c \in (0, 1)$. For each c , we run PMD 10 times, select the canonical vectors \mathbf{x}, \mathbf{y} that achieve the highest objective value and count their nonzero entries (depicted as percentage of the corresponding dimension). Finally, we run our SpanCCA algorithm with $T = 10^4$ and $r = 3$, using the latter as target sparsities, and compare the objective values achieved by the two methods. Execution times remain approximately the same for all target sparsity values (equiv. all c).

Recall that our algorithm administers precise control on the number of nonzero entries of the extracted components.

Figure 1 depicts the objective value achieved by the two algorithms, as well as the corresponding sparsity level of the extracted components. SpanCCA achieves a higher objective value in all cases. Finally, note that under the above configuration, both algorithms run for a few seconds per target sparsity, with SpanCCA running approximately half the time of PMD.

4.2. Brain Imaging Dataset

We analyzed functional statistical maps and behavioral variables from 497 subjects available from the Human Connectome Project (HCP) (Van Essen et al., 2013). The HCP consists of high-quality imaging and behavioral data, collected from a large sample of healthy adult subjects, motivated by the goal of advancing knowledge between human brain function and its association to behavior. We apply our algorithm to investigate the shared co-variation between patterns of brain activity as measured by the ex-

perimental tasks, and behavioral variables. We selected the same subset of behavioral variables examined by (Smith et al., 2015), which include scores from psychological tests, physiological measurements, and self reported behavior questionnaires (\mathbf{Y} dataset with dimensions 497×38).

For each subject, we collected statistical maps corresponding to “ n -back” task. These statistical maps summarize the activation of each voxel in response to the experimental manipulation. In the “ n -back” task, designed to measure working memory, items are presented one at a time and subjects identify each item that repeats relative to the item that occurred n items before. Further details on all tasks and variables are available in the HCP documentation (Van Essen et al., 2013).

We used the pre-computed 2back - 0back statistical contrast maps provided by the HCP. Standard preprocessing included motion correction, image registration to the MNI template (for comparison across subjects), and general linear model analysis, resulting in $91 \times 109 \times 91$ voxels. Voxels are then resampled to $61 \times 73 \times 61$ using the `nilearn` python package⁶ and applying standard brain masks, resulting in 65598 voxels after masking non-grey matter regions. (\mathbf{X} dataset with dimensions 497×65598).

We apply our SpanCCA algorithm on the HCP data with arbitrarily selected parameters $T = 10^6$ and $r = 5$. We set the target sparsity at 15% for each canonical vector. Figure 2 depicts the brain regions and the behavioral factors corresponding to the nonzero weights of the extracted canonical pair. The map identifies a set of fronto-parietal regions known to be involved in executive function and working memory, which are the major functions isolated by the 2 back - 0-back contrast. In addition, it identifies deactivation in the default mode areas (medial prefrontal and parietal), which is also associated with engagement of difficult cognitive functions. The behavioral variables associated with activation of this network are all related to various aspects of intelligence; the Penn Matrix Reasoning Test (PMAT24, a measure of fluid intelligence), picture vocabulary (`PicVocab`, a measure of language comprehension), and reading ability (`ReadEng`).

Parallelization To speed up execution, our prototypical Python implementation of SpanCCA exploits the multiprocessing module: N independent worker processes are spawned, and each one independently performs T/N rounds of the main iteration of Alg. 1 returning a single canonical vector pair. The main process collects and compares the candidate pairs to determine the final output.

To demonstrate the parallelizability of our algorithm, we run SpanCCA for the aforementioned task on the brain imaging

⁶<http://nilearn.github.io/>




Figure 2. Brain regions and behavioral factors selected by the sparse left and right canonical vectors extracted by our SpanCCA algorithm. Target sparsity is set at 15% for each canonical vector and SpanCCA is configured to run for $T = 10^6$ samples operating on a rank $r = 5$ approximation of the input data. The map identifies a set of fronto-parietal regions known to be involved in executive function and working memory and deactivation in the default mode areas (medial prefrontal and parietal), which is also associated with engagement of difficult cognitive functions. The behavioral variables identified to be positively correlated with the activation of this network are all related to various aspects of intelligence.

data for various values of the number N of workers on a single server with 36 physical processing cores⁷ and approximately 250Gb of main memory. In Figure 3 (top panel), we plot the run time with respect to the number of workers used. The bottom panel depicts the achieved speedup factor: using the execution time on 5 worker processes as a reference value, the speedup factor is the ratio of the execution time on 5 processes over that on N . As expected, the algorithm achieved a speedup factor that grows almost linearly in the number of available processors.

5. Discussion

We presented a novel combinatorial algorithm for the sparse diagonal CCA problem and other constrained variants, with provable data-dependent global approximation guarantees and several attractive properties: the algorithm is simple, embarrassingly parallelizable, with complexity that scales linearly in the dimension of the input data, while it administers precise control on the sparsity of the extracted canonical vectors. Further it can accommodate additional structural constraints by plugging in a suitable “projection” subroutine.

Several directions remain open. We addressed the question of computing a single pair of sparse canonical vectors. Numerically, multiple pairs can be computed successively employing an appropriate deflation step. However, determining the kind of deflation most suitable for the application at hand, as well as the sparsity level of each component can be a challenging task, leaving a lot of room for research.

References

Amini, Arash A and Wainwright, Martin J. High-dimensional analysis of semidefinite relaxations for sparse principal components. In *Information Theory, 2008. ISIT 2008. IEEE Inter-*

⁷Intel(R) Xeon(R) CPU E5-2699 v3 @ 2.30GHz




Figure 3. Speedup factors and corresponding total execution time, achieved by the prototypical parallel implementation of SpanCCA (Alg. 1) as a function of the number of worker processes or equivalently the number of processors used. Depicted values are medians over 20 executions, each with $T = 10^5$ and $r = 5$, on the 65598×38 example discussed in section 4.2. A speedup factory approximately linear in the number of workers is achieved.

national Symposium on, pp. 2454–2458. IEEE, 2008.

Asteris, Megasthenis, Papailiopoulos, Dimitris, and Dimakis, Alexandros. Nonnegative sparse PCA with provable guarantees. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pp. 1728–1736, 2014.

- Asteris, Megasthenis, Kyrillidis, Anastasios, Dimakis, Alex, Yi, Han-Gyol, and Chandrasekaran, Bharath. Stay on path: Pca along graph paths. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, volume 37. JMLR Workshop and Conference Proceedings, 2015.
- Bach, Francis, Jenatton, Rodolphe, Mairal, Julien, Obozinski, Guillaume, et al. Structured sparsity through convex optimization. *Statistical Science*, 27(4):450–468, 2012.
- Baldassarre, Luca, Bhan, Nirav, Cevher, Volkan, Kyrillidis, Anastasios, and Satpathi, Siddhartha. Group-sparse model selection: Hardness and relaxations. *arXiv preprint arXiv:1303.3207*, 2013.
- Baraniuk, Richard G, Cevher, Volkan, Duarte, Marco F, and Hegde, Chinmay. Model-based compressive sensing. *Information Theory, IEEE Transactions on*, 56(4):1982–2001, 2010.
- Berman, Marc G, Jonides, John, and Nee, Derek Evan. Studying mind and brain with fmri. *Social cognitive and affective neuroscience*, 1(2):158–161, 2006.
- Chen, Jun, Bushman, Frederic D, Lewis, James D, Wu, Gary D, and Li, Hongzhe. Structure-constrained sparse canonical correlation analysis with an application to microbiome data analysis. *Biostatistics*, 14(2):244–258, 2013.
- Chen, Xi, Liu, Han, and Carbonell, Jaime G. Structured sparse canonical correlation analysis. In *International Conference on Artificial Intelligence and Statistics*, pp. 199–207, 2012.
- Chin, Koei, DeVries, Sandy, Fridlyand, Jane, Spellman, Paul T, Roydasgupta, Ritu, Kuo, Wen-Lin, Lapuk, Anna, Neve, Richard M, Qian, Zuwei, Ryder, Tom, et al. Genomic and transcriptional aberrations linked to breast cancer pathophysiology. *Cancer cell*, 10(6):529–541, 2006.
- Chu, Delin, Liao, Li-Zhi, Ng, Michael K, and Zhang, Xiaowei. Sparse canonical correlation analysis: new formulation and algorithm. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 35(12):3050–3065, 2013.
- d’Aspremont, Alexandre, El Ghaoui, Laurent, Jordan, Michael I, and Lanckriet, Gert RG. A direct formulation for sparse pca using semidefinite programming. *SIAM review*, 49(3):434–448, 2007.
- Deleus, Filip and Van Hulle, Marc M. Functional connectivity analysis of fmri data based on regularized multiset canonical correlation analysis. *Journal of Neuroscience methods*, 197(1):143–157, 2011.
- Du, Lei, Yan, Jingwen, Kim, Sungeun, Risacher, Shannon L, Huang, Heng, Inlow, Mark, Moore, Jason H, Saykin, Andrew J, and Shen, Li. A novel structure-aware sparse learning algorithm for brain imaging genetics. In *Medical Image Computing and Computer-Assisted Intervention*, pp. 329–336. Springer, 2014.
- Dudoit, Sandrine, Fridlyand, Jane, and Speed, Terence P. Comparison of discrimination methods for the classification of tumors using gene expression data. *Journal of the American statistical association*, 97(457):77–87, 2002.
- Halko, Nathan, Martinsson, Per-Gunnar, and Tropp, Joel A. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2):217–288, 2011.
- Hardoon, David R and Shawe-Taylor, John. Technical report, university college london (ucl). 2007.
- Hardoon, David R and Shawe-Taylor, John. Sparse canonical correlation analysis. *Machine Learning*, 83(3):331–353, 2011.
- Hegde, Chinmay, Indyk, Piotr, and Schmidt, Ludwig. A nearly-linear time framework for graph-structured sparsity. In *Proceedings of The 32nd International Conference on Machine Learning*, pp. 928–937, 2015.
- Hotelling, Harold. Relations between two sets of variates. *Biometrika*, pp. 321–377, 1936.
- Huang, Junzhou, Zhang, Tong, and Metaxas, Dimitris. Learning with structured sparsity. *The Journal of Machine Learning Research*, 12:3371–3412, 2011.
- Kobayashi, Takumi. S3cca: Smoothly structured sparse cca for partial pattern matching. In *Pattern Recognition (ICPR), 22nd International Conference on*, pp. 1981–1986. IEEE, 2014.
- Kyrillidis, Anastasios and Cevher, Volkan. Combinatorial selection and least absolute shrinkage via the clash algorithm. In *Information Theory Proceedings (ISIT), 2012 IEEE International Symposium on*, pp. 2216–2220. IEEE, 2012.
- Kyrillidis, Anastasios, Puy, Gilles, and Cevher, Volkan. Hard thresholding with norm constraints. In *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, number EPFL-CONF-183061, pp. 3645–3648. Ieee, 2012.
- Kyrillidis, Anastasios, Baldassarre, Luca, El Halabi, Marwa, Tran-Dinh, Quoc, and Cevher, Volkan. Structured sparsity: Discrete and convex approaches. In *Compressed Sensing and its Applications*, pp. 341–387. Springer, 2015.
- Lee, Mihee, Shen, Haipeng, Huang, Jianhua Z, and Marron, JS. Biclustering via sparse singular value decomposition. *Biometrics*, 66(4):1087–1095, 2010.
- Li, Yi-Ou, Eichele, Tom, Calhoun, Vince D, and Adali, Tulay. Group study of simulated driving fmri data by multiset canonical correlation analysis. *Journal of signal processing systems*, 68(1):31–48, 2012.
- Lin, Dongdong, Calhoun, Vince D, and Wang, Yu-Ping. Correspondence between fmri and snp data by group sparse canonical correlation analysis. *Medical image analysis*, 18(6):891–902, 2014.
- Morley, Michael, Molony, Cliona M, Weber, Teresa M, Devlin, James L, Ewens, Kathryn G, Spielman, Richard S, and Cheung, Vivian G. Genetic analysis of genome-wide variation in human gene expression. *Nature*, 430(7001):743–747, 2004.
- Parkhomenko, Elena, Tritchler, David, and Beyene, Joseph. Sparse canonical correlation analysis with application to genomic data integration. *Statistical Applications in Genetics and Molecular Biology*, 8(1):1–34, 2009.
- Poldrack, Russell A. Can cognitive processes be inferred from neuroimaging data? *Trends in cognitive sciences*, 10(2):59–63, 2006.

- Pollack, Jonathan R, Sørlie, Therese, Perou, Charles M, Rees, Christian A, Jeffrey, Stefanie S, Lonning, Per E, Tibshirani, Robert, Botstein, David, Børresen-Dale, Anne-Lise, and Brown, Patrick O. Microarray analysis reveals a major direct role of dna copy number alteration in the transcriptional program of human breast tumors. *Proceedings of the National Academy of Sciences*, 99(20):12963–12968, 2002.
- Posner, Michael I, Petersen, Steven E, Fox, Peter T, and Raichle, Marcus E. Localization of cognitive operations in the human brain. *Science*, 240(4859):1627–1631, 1988.
- Rustandi, Indrayana, Just, Marcel Adam, and Mitchell, Tom. Integrating multiple-study multiple-subject fmri datasets using canonical correlation analysis. In *Proceedings of the MICCAI 2009 Workshop: Statistical modeling and detection issues in intra-and inter-subject functional MRI data analysis*, 2009.
- Smith, Stephen M, Nichols, Thomas E, Vidaurre, Diego, Winkler, Anderson M, Behrens, Timothy EJ, Glasser, Matthew F, Ugurbil, Kamil, Barch, Deanna M, Van Essen, David C, and Miller, Karla L. A positive-negative mode of population covariation links brain connectivity, demographics and behavior. *Nature neuroscience*, 18(11):1565–1567, 2015.
- Sriperumbudur, Bharath, Torres, David, and Lanckriet, Gert. A DC programming approach to the sparse generalized eigenvalue problem. *arXiv preprint arXiv:0901.1504*, 2009.
- Stranger, Barbara E, Forrest, Matthew S, Dunning, Mark, Ingle, Catherine E, Beazley, Claude, Thorne, Natalie, Redon, Richard, Bird, Christine P, de Grassi, Anna, Lee, Charles, et al. Relative impact of nucleotide and copy number variation on gene expression phenotypes. *Science*, 315(5813):848–853, 2007.
- Thompson, Bruce. *Canonical correlation analysis: Uses and interpretation*. Number 47. Sage, 1984.
- Thorndike, Robert M. *Correlational procedures for research*. Wiley, 1976.
- Tibshirani, Robert, Hastie, Trevor, Narasimhan, Balasubramanian, and Chu, Gilbert. Class prediction by nearest shrunken centroids, with applications to dna microarrays. *Statistical Science*, pp. 104–117, 2003.
- Torres, David A, Turnbull, Douglas, Barrington, Luke, and Lanckriet, Gert RG. Identifying words that are musically meaningful. In *ISMIR*, volume 7, pp. 405–410, 2007.
- Van Essen, David C, Smith, Stephen M, Barch, Deanna M, Behrens, Timothy EJ, Yacoub, Essa, Ugurbil, Kamil, Consortium, WU-Minn HCP, et al. The wu-minn human connectome project: an overview. *Neuroimage*, 80:62–79, 2013.
- Vershynin, Roman. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.
- Waaijenborg, Sandra, Verselewe de Witt Hamer, Philip C, and Zwinderman, Aeilko H. Quantifying the association between gene expressions and dna-markers by penalized canonical correlation analysis. *Statistical Applications in Genetics and Molecular Biology*, 7(1), 2008.
- Wiesel, Ami, Kliger, Mark, and Hero III, Alfred O. A greedy approach to sparse canonical correlation analysis. *arXiv preprint arXiv:0801.2748*, 2008.
- Witten, Daniela M and Tibshirani, Robert J. Extensions of sparse canonical correlation analysis with applications to genomic data. *Statistical applications in genetics and molecular biology*, 8(1):1–27, 2009.
- Witten, Daniela M, Tibshirani, Robert, and Hastie, Trevor. A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis. *Biostatistics*, pp. kxp008, 2009.
- Yang, Dan, Ma, Zongming, and Buja, Andreas. A sparse SVD method for high-dimensional data. *arXiv preprint arXiv:1112.2433*, 2011.
- Zou, Hui and Hastie, Trevor. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320, 2005.
- Zou, Hui, Hastie, Trevor, and Tibshirani, Robert. Sparse principal component analysis. *Journal of computational and graphical statistics*, 15(2):265–286, 2006.