

A greedy approach to sparse canonical correlation analysis

Ami Wiesel, Mark Kliger and Alfred O. Hero III
 Dept. of Electrical Engineering and Computer Science
 University of Michigan, Ann Arbor, MI 48109-2122, USA

Abstract— We consider the problem of sparse canonical correlation analysis (CCA), i.e., the search for two linear combinations, one for each multivariate, that yield maximum correlation using a specified number of variables. We propose an efficient numerical approximation based on a direct greedy approach which bounds the correlation at each stage. The method is specifically designed to cope with large data sets and its computational complexity depends only on the sparsity levels. We analyze the algorithm’s performance through the tradeoff between correlation and parsimony. The results of numerical simulation suggest that a significant portion of the correlation may be captured using a relatively small number of variables. In addition, we examine the use of sparse CCA as a regularization method when the number of available samples is small compared to the dimensions of the multivariates.

I. INTRODUCTION

Canonical correlation analysis (CCA), introduced by Harold Hotelling [1], is a standard technique in multivariate data analysis for extracting common features from a pair of data sources [2], [3]. Each of these data sources generates a random vector that we call a multivariate. Unlike classical dimensionality reduction methods which address one multivariate, CCA takes into account the statistical relations between samples from two spaces of possibly different dimensions and structure. In particular, it searches for two linear combinations, one for each multivariate, in order to maximize their correlation. It is used in different disciplines as a stand-alone tool or as a preprocessing step for other statistical methods. Furthermore, CCA is a generalized framework which includes numerous classical methods in statistics, e.g., Principal Component Analysis (PCA), Partial Least Squares (PLS) and Multiple Linear Regression (MLR) [4]. CCA has recently regained attention with the advent of kernel CCA and its application to independent component analysis [5], [6].

The last decade has witnessed a growing interest in the search for sparse representations of signals and sparse numerical methods. Thus, we consider the problem of sparse CCA, i.e., the search for linear combinations with maximal correlation using a small number of variables. The quest for sparsity can be motivated through various reasonings. First is the ability to interpret and visualize the results. A small number of variables allows us to get the “big picture”, while sacrificing some of the small details. Moreover, sparse representations enable the use of computationally efficient

numerical methods, compression techniques, as well as noise reduction algorithms. The second motivation for sparsity is regularization and stability. One of the main vulnerabilities of CCA is its sensitivity to a small number of observations. Thus, regularized methods such as ridge CCA [7] must be used. In this context, sparse CCA is a subset selection scheme which allows us to reduce the dimensions of the vectors and obtain a stable solution.

To the best of our knowledge the first reference to sparse CCA appeared in [2] where backward and stepwise subset selection were proposed. This discussion was of qualitative nature and no specific numerical algorithm was proposed. Recently, increasing demands for multidimensional data processing and decreasing computational cost has caused the topic to rise to prominence once again [8]–[13]. The main disadvantages with these current solutions is that there is no direct control over the sparsity and it is difficult (and non-intuitive) to select their optimal hyperparameters. In addition, the computational complexity of most of these methods is too high for practical applications with high dimensional data sets. Sparse CCA has also been implicitly addressed in [9], [14] and is intimately related to the recent results on sparse PCA [9], [15]–[17]. Indeed, our proposed solution is an extension of the results in [17] to CCA.

The main contribution of this work is twofold. First, we derive CCA algorithms with direct control over the sparsity in each of the multivariates and examine their performance. Our computationally efficient methods are specifically aimed at understanding the relations between two data sets of large dimensions. We adopt a forward (or backward) greedy approach which is based on sequentially picking (or dropping) variables. At each stage, we bound the optimal CCA solution and bypass the need to resolve the full problem. Moreover, the computational complexity of the forward greedy method does not depend on the dimensions of the data but only on the sparsity parameters. Numerical simulation results show that a significant portion of the correlation can be efficiently captured using a relatively low number of non-zero coefficients. Our second contribution is investigation of sparse CCA as a regularization method. Using empirical simulations we examine the use of the different algorithms when the dimensions of the multivariates are larger than (or of the same order of) the number of samples and demonstrate the advantage of sparse CCA. In this context, one of the advantages of the greedy approach is that it generates the full sparsity path in a single run and allows for efficient parameter tuning using

The first two authors contributed equally to this manuscript. This work was supported in part by an AFOSR MURI under Grant FA9550-06-1-0324.

cross validation.

The paper is organized as follows. We begin by describing the standard CCA formulation and solution in Section II. Sparse CCA is addressed in Section III where we review the existing approaches and derive the proposed greedy method. In Section IV, we provide performance analysis using numerical simulations and assess the tradeoff between correlation and parsimony, as well as its use in regularization. Finally, a discussion is provided in Section V.

The following notation is used. Boldface upper case letters denote matrices, boldface lower case letters denote column vectors, and standard lower case letters denote scalars. The superscripts $(\cdot)^T$ and $(\cdot)^{-1}$ denote the transpose and inverse operators, respectively. By \mathbf{I} we denote the identity matrix. The operator $\|\cdot\|_2$ denotes the L2 norm, and $\|\cdot\|_0$ denotes the cardinality operator. For two sets of indices I and J , the matrix $\mathbf{X}^{I,J}$ denotes the submatrix of \mathbf{X} with the rows indexed by I and columns indexed by J . Finally, $\mathbf{X} \succ \mathbf{0}$ or $\mathbf{X} \succeq \mathbf{0}$ means that the matrix \mathbf{X} is positive definite or positive semidefinite, respectively.

II. REVIEW ON CCA

In this section, we provide a review of classical CCA. Let \mathbf{x} and \mathbf{y} be two zero mean random vectors of lengths n and m , respectively, with joint covariance matrix:

$$\Sigma = \begin{bmatrix} \Sigma_x & \Sigma_{xy} \\ \Sigma_{xy}^T & \Sigma_y \end{bmatrix} \succeq \mathbf{0}, \quad (1)$$

where Σ_x , Σ_y and Σ_{xy} are the covariance of \mathbf{x} , the covariance of \mathbf{y} , and their cross covariance, respectively. CCA considers the problem of finding two linear combinations $X = \mathbf{a}^T \mathbf{x}$ and $Y = \mathbf{b}^T \mathbf{y}$ with maximal correlation defined as

$$\frac{\text{cov}\{X, Y\}}{\sqrt{\text{var}\{X\}}\sqrt{\text{var}\{Y\}}}, \quad (2)$$

where $\text{var}\{\cdot\}$ and $\text{cov}\{\cdot\}$ are the variance and covariance operators, respectively, and we define $0/0 = 1$. In terms of \mathbf{a} and \mathbf{b} the correlation can be easily expressed as

$$\frac{\mathbf{a}^T \Sigma_{xy} \mathbf{b}}{\sqrt{\mathbf{a}^T \Sigma_x \mathbf{a}} \sqrt{\mathbf{b}^T \Sigma_y \mathbf{b}}}. \quad (3)$$

Thus, CCA considers the following optimization problem

$$\rho(\Sigma_x, \Sigma_y, \Sigma_{xy}) = \max_{\mathbf{a} \neq \mathbf{0}, \mathbf{b} \neq \mathbf{0}} \frac{\mathbf{a}^T \Sigma_{xy} \mathbf{b}}{\sqrt{\mathbf{a}^T \Sigma_x \mathbf{a}} \sqrt{\mathbf{b}^T \Sigma_y \mathbf{b}}}. \quad (4)$$

Problem (4) is a multidimensional non-concave maximization and therefore appears difficult on first sight. However, it has a simple closed form solution via the generalized eigenvalue decomposition (GEVD). Indeed, if $\Sigma \succ \mathbf{0}$, it is easy to show that the optimal \mathbf{a} and \mathbf{b} must satisfy:

$$\begin{bmatrix} \mathbf{0} & \Sigma_{xy} \\ \Sigma_{xy}^T & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{a} \\ \mathbf{b} \end{bmatrix} = \lambda \begin{bmatrix} \Sigma_x & \mathbf{0} \\ \mathbf{0} & \Sigma_y \end{bmatrix} \begin{bmatrix} \mathbf{a} \\ \mathbf{b} \end{bmatrix}, \quad (5)$$

for some $0 \leq \lambda \leq 1$. Thus, the optimal value of (4) is just the principal generalized eigenvalue λ_{\max} of the pencil (5) and the optimal solution \mathbf{a} and \mathbf{b} can be obtained by appropriately partitioning the associated eigenvector. These solutions are

invariant to scaling of \mathbf{a} and \mathbf{b} and it is customary to normalize them such that $\mathbf{a}^T \Sigma_x \mathbf{a} = 1$ and $\mathbf{b}^T \Sigma_y \mathbf{b} = 1$. On the other hand, if Σ is rank deficient, then choosing \mathbf{a} and \mathbf{b} as the upper and lower partitions of any vector in its null space will lead to full correlation, i.e., $\rho = 1$.

In practice, the covariance matrices Σ_x and Σ_y and cross covariance matrix Σ_{xy} are usually unavailable. Instead, multiple independent observations \mathbf{x}_i and \mathbf{y}_i for $i = 1 \dots N$ are measured and used to construct sample estimates of the (cross) covariance matrices:

$$\hat{\Sigma}_x = \frac{1}{N} \mathbf{X}^T \mathbf{X}, \quad \hat{\Sigma}_y = \frac{1}{N} \mathbf{Y}^T \mathbf{Y}, \quad \hat{\Sigma}_{xy} = \frac{1}{N} \mathbf{X}^T \mathbf{Y}, \quad (6)$$

where $\mathbf{X} = [\mathbf{x}_1 \dots \mathbf{x}_N]$ and $\mathbf{Y} = [\mathbf{y}_1 \dots \mathbf{y}_N]$. Then, these empirical matrices are used in the CCA formulation:

$$\rho(\hat{\Sigma}_x, \hat{\Sigma}_y, \hat{\Sigma}_{xy}) = \max_{\mathbf{a} \neq \mathbf{0}, \mathbf{b} \neq \mathbf{0}} \frac{\mathbf{a}^T \hat{\Sigma}_{xy} \mathbf{b}}{\sqrt{\mathbf{a}^T \hat{\Sigma}_x \mathbf{a}} \sqrt{\mathbf{b}^T \hat{\Sigma}_y \mathbf{b}}}. \quad (7)$$

Clearly, if N is sufficiently large then this sample approach performs well. However, in many applications, the number of samples N is not sufficient. In fact, in the extreme case in which $N < n + m$ the sample covariance is rank deficient and $\rho = 1$ independently of the data. The standard approach in such cases is to regularize the covariance matrices and solve $\rho(\hat{\Sigma}_x + \epsilon_x \mathbf{I}, \hat{\Sigma}_y + \epsilon_y \mathbf{I}, \hat{\Sigma}_{xy})$ where $\epsilon_x > 0$ and $\epsilon_y > 0$ are small tuning ridge parameters [7].

CCA can be viewed as a unified framework for dimensionality reduction in multivariate data analysis and generalizes other existing methods. It is a generalization of PCA which seeks the directions that maximize the variance of \mathbf{x} , and addresses the directions corresponding to the correlation between \mathbf{x} and \mathbf{y} . A special case of CCA is PLS which maximizes the covariance of \mathbf{x} and \mathbf{y} (which is equivalent to choosing $\Sigma_x = \Sigma_y = \mathbf{I}$). Similarly, MLR normalizes only one of the multivariates. In fact, the regularized CCA mentioned above can be interpreted as a combination of PLS and CCA [4].

III. SPARSE CCA

We consider the problem of sparse CCA, i.e., finding a pair of linear combinations of \mathbf{x} and \mathbf{y} with prescribed cardinality which maximize the correlation. Mathematically, we define sparse CCA as the solution to

$$\begin{cases} \max_{\mathbf{a} \neq \mathbf{0}, \mathbf{b} \neq \mathbf{0}} & \frac{\mathbf{a}^T \Sigma_{xy} \mathbf{b}}{\sqrt{\mathbf{a}^T \Sigma_x \mathbf{a}} \sqrt{\mathbf{b}^T \Sigma_y \mathbf{b}}} \\ \text{s.t.} & \|\mathbf{a}\|_0 \leq k_a \\ & \|\mathbf{b}\|_0 \leq k_b. \end{cases} \quad (8)$$

Similarly, sparse PLS and sparse MLR are defined as special cases of (8) by choosing $\Sigma_x = \mathbf{I}$ and/or $\Sigma_y = \mathbf{I}$. In general, all of these problems are difficult combinatorial problems. In small dimensions, they can be solved using a brute force search over all possible sparsity patterns and solving the associated subproblem via GEVD. Unfortunately, this approach is impractical for even moderate sizes of data sets due its exponentially increasing computational complexity. In fact, it is a generalization of sparse PCA which has been proven NP-hard [17]. Thus, suboptimal but efficient approaches are in order and will be discussed in the rest of this section.

A. Existing solutions

We now briefly review the different approaches to sparse CCA that appeared in the last few years. Most of the methods are based on the well known LASSO trick in which the difficult combinatorial cardinality constraints are approximated through the convex L1 norm. This approach has shown promising performance in the context of sparse linear regression [18]. Unfortunately, it is not sufficient in the CCA formulation since the objective itself is not concave. Thus additional approximations are required to transform the problems into tractable form.

Sparse dimensionality reduction of rectangular matrices was considered in [9] by combining the LASSO trick with semidefinite relaxation. In our context, this is exactly sparse PLS which is a special case of sparse CCA. Alternatively, CCA can be formulated as two constrained simultaneous regressions (\mathbf{x} on \mathbf{y} , and \mathbf{y} on \mathbf{x}). Thus, an appealing approach to sparse CCA is to use LASSO penalized regressions. Based on this idea, [8] proposed to approximate the non convex constraints using the infinity norm. Similarly, [10], [11] proposed to use two nested iterative LASSO type regressions.

There are two main disadvantages to the LASSO based techniques. First, there is no mathematical justification for their approximations of the correlation objective. Second, there is no direct control over sparsity. Their parameter tuning is difficult as the relation between the L1 norm and the sparsity parameters is highly nonlinear. The algorithms need to be run for each possible value of the parameters and it is tedious to obtain the full sparsity path.

An alternative approach for sparse CCA is sparse Bayes learning [12], [13]. These methods are based on the probabilistic interpretation of CCA, i.e., its formulation as an estimation problem. It was shown that using different prior probabilistic models, sparse solutions can be obtained. The main disadvantage of this approach is again the lack of direct control on sparsity, and the difficulty in obtaining its complete sparsity path.

Altogether, these works demonstrate the growing interest in deriving efficient sparse CCA algorithms aimed at large data sets with simple and intuitive parameter tuning.

B. Greedy approach

A standard approach to combinatorial problems is the forward (backward) greedy solution which sequentially picks (or drops) the variables at each stage one by one. The backward greedy approach to CCA was proposed in [2] but no specific algorithm was derived or analyzed. In modern applications, the number of dimensions may be much larger than the number of samples and therefore we provide the details of the more natural forward strategy. Nonetheless, the backward approach can be derived in a straightforward manner. In addition, we derive an efficient approximation to the subproblems at each stage which significantly reduces the computational complexity. A similar approach in the context of PCA can be found in [17].

Our goal is to find the two sparsity patterns, i.e., two sets of indices I and J corresponding to the indices of the chosen

variables in \mathbf{x} and \mathbf{y} , respectively. The greedy algorithm chooses the first elements in both sets as the solution to

$$\max_{i,j} \frac{\Sigma_{xy}^{i,j}}{\sqrt{\Sigma_x^{ii} \Sigma_y^{jj}}}. \quad (9)$$

Thus, $I = \{i\}$ and $J = \{j\}$. Next, the algorithm sequentially examines all the remaining indices and computes

$$\max_{i \notin I} \rho(\Sigma_x^{I \cup i, I \cup i}, \Sigma_y^{J, J}, \Sigma_{xy}^{I \cup i, J}) \quad (10)$$

and

$$\max_{j \notin J} \rho(\Sigma_x^{I, I}, \Sigma_y^{J \cup j, J \cup j}, \Sigma_{xy}^{I, J \cup j}). \quad (11)$$

Depending on whether (10) is greater or less than (11), we add index i or j to I or J , respectively. We emphasize that at each stage, only one index is added either to I or J . Once one of the set reaches its required size k_a or k_b , the algorithm continues to add indices only to the other set and terminates when this set reaches its required size as well. It outputs the full sparsity path, and returns $k_a - k_b - 1$ pairs of vectors associated with the sparsity patterns in each of the stages.

The computational complexity of the algorithm is polynomial in the dimensions of the problem. At each of its $i = 1, \dots, k_a + k_b - 1$ stages, the algorithm computes $n + m - i$ CCA solutions as expressed in (10) and (11) in order to select the patterns for the next stage. It is therefore reasonable for small problems, but is still impractical for many applications. Instead, we now propose an alternative approach that computes only one CCA per stage and reduces the complexity significantly.

An approximate greedy solution can be easily obtained by approximating (10) and (11) instead of solving them exactly. Consider for example (10) where index i is added to the set I . Let $\mathbf{a}^{I, J}$ and $\mathbf{b}^{I, J}$ denote the optimal solution to $\rho(\Sigma_x^{I, I}, \Sigma_y^{J, J}, \Sigma_{xy}^{I, J})$. In order to evaluate (10) we need to recalculate both $\mathbf{a}^{I \cup i, J}$ and $\mathbf{b}^{I \cup i, J}$ for each $i \notin I$. However, the previous $\mathbf{b}^{I, J}$ is of the same dimension and still feasible. Thus, we can optimize only with respect to $\mathbf{a}^{I \cup i, J}$ (whose dimension has increased). This approach provides the following bounds:

Lemma 1: Let $\mathbf{a}^{I, J}$ and $\mathbf{b}^{I, J}$ be the optimal solution to $\rho(\Sigma_x^{I, I}, \Sigma_y^{J, J}, \Sigma_{xy}^{I, J})$ in (4). Then,

$$\begin{aligned} \rho^2(\Sigma_x^{I \cup i, I \cup i}, \Sigma_y^{J, J}, \Sigma_{xy}^{I \cup i, J}) &\geq \rho^2(\Sigma_x^{I, I}, \Sigma_y^{J, J}, \Sigma_{xy}^{I, J}) + \delta_i^{I, J} \\ \rho^2(\Sigma_x^{I, I}, \Sigma_y^{J \cup j, J \cup j}, \Sigma_{xy}^{I, J \cup j}) &\geq \rho^2(\Sigma_x^{I, I}, \Sigma_y^{J, J}, \Sigma_{xy}^{I, J}) + \gamma_j^{I, J}, \end{aligned} \quad (12)$$

where

$$\begin{aligned} \delta_i^{I, J} &= \frac{\left([\mathbf{b}^{I, J}]^T [\Sigma_{xy}^{I, J}]^T [\Sigma_x^{I, I}]^{-1} \Sigma_x^{I, i} - [\mathbf{b}^{I, J}]^T [\Sigma_{xy}^{I, J}]^T \right)^2}{\Sigma_x^{i, i} - [\Sigma_x^{I, i}]^T [\Sigma_x^{I, I}]^{-1} \Sigma_x^{I, i}} \\ \gamma_j^{I, J} &= \frac{\left([\mathbf{a}^{I, J}]^T \Sigma_{xy}^{I, J} [\Sigma_y^{J, J}]^{-1} \Sigma_y^{J, j} - [\mathbf{a}^{I, J}]^T \Sigma_{xy}^{I, j} \right)^2}{\Sigma_y^{j, j} - [\Sigma_y^{J, j}]^T [\Sigma_y^{J, J}]^{-1} \Sigma_y^{J, j}}, \end{aligned} \quad (13)$$

and we assume that all the involved matrix inversions are nonsingular.

Before proving the lemma, we note that it provides lower bounds on the increase in cross correlation due to including an additional element in \mathbf{x} or \mathbf{y} without the need of solving a full GEVD. Thus, we propose the following approximate greedy approach. For each sparsity pattern $\{I, J\}$, one CCA is computed via GEVD in order to obtain $\mathbf{a}^{I,J}$ and $\mathbf{b}^{I,J}$. Then, the next sparsity pattern is obtained by adding the element that maximizes $\delta_i^{I,J}$ or $\gamma_j^{I,J}$ among $i \notin I$ and $j \notin J$.

Proof: We begin by rewriting (4) as a quadratic maximization with constraints. Thus, we define $\mathbf{a}^{I,J}$ and $\mathbf{b}^{I,J}$ as the solution to

$$\rho(\Sigma_x^{I,I}, \Sigma_y^{J,J}, \Sigma_{xy}^{I,J}) = \begin{cases} \max_{\mathbf{a}, \mathbf{b}} & \mathbf{a}^T \Sigma_{xy}^{I,J} \mathbf{b} \\ \text{s.t.} & \mathbf{a}^T \Sigma_x^{I,I} \mathbf{a} = 1 \\ & \mathbf{b}^T \Sigma_y^{J,J} \mathbf{b} = 1. \end{cases} \quad (14)$$

Now, consider the problem when we add variable j to the set J :

$$\rho(\Sigma_x^{I,I}, \Sigma_y^{J \cup j, J \cup j}, \Sigma_{xy}^{I, J \cup j}) = \begin{cases} \max_{\mathbf{a}, \mathbf{b}} & \mathbf{a}^T \Sigma_{xy}^{I, J \cup j} \mathbf{b} \\ \text{s.t.} & \mathbf{a}^T \Sigma_x^{I,I} \mathbf{a} = 1 \\ & \mathbf{b}^T \Sigma_y^{J \cup j, J \cup j} \mathbf{b} = 1. \end{cases} \quad (15)$$

Clearly, the vector $\mathbf{a} = \mathbf{a}^{I,J}$ is still feasible (though not necessarily optimal) and yields a lower bound:

$$\rho(\Sigma_x^{I,I}, \Sigma_y^{J \cup j, J \cup j}, \Sigma_{xy}^{I, J \cup j}) \geq \begin{cases} \max_{\mathbf{b}} & [\mathbf{a}^{I,J}]^T \Sigma_{xy}^{I, J \cup j} \mathbf{b} \\ \text{s.t.} & \mathbf{b}^T \Sigma_y^{J \cup j, J \cup j} \mathbf{b} = 1. \end{cases} \quad (16)$$

Changing variables $\bar{\mathbf{b}} = [\Sigma_y^{J \cup j, J \cup j}]^{\frac{1}{2}} \mathbf{b}$ results in:

$$\rho(\Sigma_x^{I,I}, \Sigma_y^{J \cup j, J \cup j}, \Sigma_{xy}^{I, J \cup j}) \geq \begin{cases} \max_{\bar{\mathbf{b}}} & \mathbf{c}^T \bar{\mathbf{b}} \\ \text{s.t.} & \|\bar{\mathbf{b}}\|_2 = 1, \end{cases} \quad (17)$$

where

$$\mathbf{c} = [\Sigma_y^{J \cup j, J \cup j}]^{-\frac{1}{2}} [\Sigma_{xy}^{I, J \cup j}]^T \mathbf{a}^{I,J}. \quad (18)$$

Using the Cauchy Schwartz inequality

$$\mathbf{c}^T \bar{\mathbf{b}} \leq \|\mathbf{c}\|_2 \|\bar{\mathbf{b}}\|_2, \quad (19)$$

we obtain

$$\rho(\Sigma_x^{I,I}, \Sigma_y^{J \cup j, J \cup j}, \Sigma_{xy}^{I, J \cup j}) \geq \left\| [\Sigma_y^{J \cup j, J \cup j}]^{-\frac{1}{2}} [\Sigma_{xy}^{I, J \cup j}]^T \mathbf{a}^{I,J} \right\|_2. \quad (20)$$

Therefore,

$$\rho^2(\Sigma_x^{I,I}, \Sigma_y^{J \cup j, J \cup j}, \Sigma_{xy}^{I, J \cup j}) \geq [\mathbf{a}^{I,J}]^T \begin{bmatrix} \Sigma_{xy}^{I,J} & \Sigma_{xy}^{I,j} \\ \Sigma_{xy}^{J,I} & \Sigma_{xy}^{J,j} \end{bmatrix} \begin{bmatrix} \Sigma_y^{J,J} & \Sigma_y^{J,j} \\ \Sigma_y^{j,J} & \Sigma_y^{j,j} \end{bmatrix}^{-1} \begin{bmatrix} [\Sigma_{xy}^{I,J}]^T \\ [\Sigma_{xy}^{I,j}]^T \end{bmatrix} \mathbf{a}^{I,J}. \quad (21)$$

Finally, (12) is obtained by using the inversion formula for partitioned matrices and simplifying the terms. ■




Fig. 1. Correlation vs. Sparsity, $n = m = 7$.

IV. NUMERICAL RESULTS

We now provide a few numerical examples illustrating the behavior of the greedy sparse CCA methods. In all of the simulations below, we implement the greedy methods using the bounds in Lemma 1. In the first experiment we evaluate the validity of the approximate greedy approach. In particular, we choose $n = m = 7$, and generate 200 independent random realizations of the joint covariance matrix using the Wishart distribution with $7+7$ degrees of freedom. For each realization, we run the approximate greedy forward and backward algorithms and calculate the full sparsity path. For comparison, we also compute the optimal sparse solutions using an exhaustive search. The results are presented in Fig. 1 where the average correlation is plotted as a function of the number of variables (or non-zero coefficients). The greedy methods capture a significant portion of the possible correlation. As expected, the forward greedy approach outperforms the backward method when high sparsity is critical. On the other hand, the backward method is preferable if large values of correlation are required.

In the second experiment we demonstrate the performance of the approximate forward greedy approach in a large scale problem. We present results for a representative (randomly generated) covariance matrix of sizes $n = m = 1000$. Fig. 2 shows the full sparsity path of the greedy method. It is easy to see that about 90 percent of the CCA correlation value can be captured using only half of the variables. Furthermore, if we choose to capture only 80 percent of the full correlation, then about a quarter of the variables are sufficient.

In the third set of simulations, we examine the use of sparse CCA algorithms as regularization methods when the number of samples is not sufficient to estimate the covariance matrix efficiently. For simplicity, we restrict our attention to CCA and PLS (which can be interpreted as an extreme case of ridge CCA). In addition, we show results for an alternative method in which the sample covariance $\hat{\Sigma}_x$ and $\hat{\Sigma}_y$ are approximated as diagonal matrices with the sample variances (which are easier to estimate). We refer to this method as Diagonal CCA (DCCA). In order to assess the regularization properties of CCA we used the following procedure. We randomly generate a single “true” covariance matrix Σ and use it throughout all the simulations. Then, we generate N random Gaussian samples of \mathbf{x} and \mathbf{y} and estimate $\hat{\Sigma}$. We apply the three

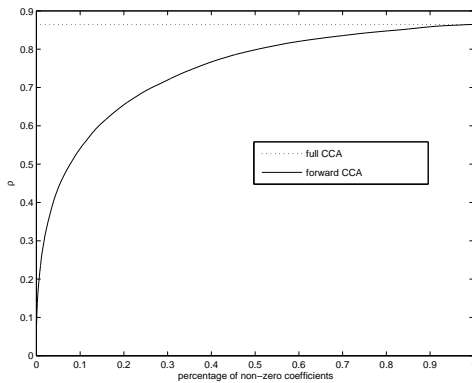


Fig. 2. Correlation vs. Sparsity, $n = m = 1000$.




Fig. 3. Sparsity as a regularized CCA method, $n = 10$, $m = 10$ and $N = 20$.

approximate greedy sparse algorithms, CCA, PLS and DCCA, using the sample covariance and obtain the estimates $\hat{\mathbf{a}}$ and $\hat{\mathbf{b}}$. Finally, our performance measure is the “true” correlation value associated with the estimated weights which is defined as:

$$\frac{\hat{\mathbf{a}}^T \Sigma_{xy} \hat{\mathbf{b}}}{\sqrt{\hat{\mathbf{a}}^T \Sigma_x \hat{\mathbf{a}} \hat{\mathbf{b}}^T \Sigma_y \hat{\mathbf{b}}}}. \quad (22)$$

We then repeat the above procedure (using the same “true” covariance matrix) 500 times and present the average value of (22) over these Monte Carlo trials. Fig. 3 provides these averages as a function of parsimony for two representative realizations of the “true” covariance matrix. Examining the curves reveals that variable selection is indeed a promising regularization strategy. The average correlation increases with the number of variables until it reaches a peak. After this peak, the number of samples are not sufficient to estimate the full covariance and it is better to reduce the number of variables through sparsity. DCCA can also be slightly improved by using fewer variables, and it seems that PLS performs best with no subset selection.

V. DISCUSSION

We considered the problem of sparse CCA and discussed its implementation aspects and statistical properties. In particular, we derived direct greedy methods which are specifically designed to cope with large data sets. Similar to state of the

art sparse regression methods, e.g., Least Angle Regression (LARS) [19], the algorithms allow for direct control over the sparsity and provide the full sparsity path in a single run. We have demonstrated their performance advantage through numerical simulations.

There are a few interesting directions for future research in sparse CCA. First, we have only addressed the first order sparse canonical components. In many applications, analysis of higher order canonical components is preferable. Numerically, this extension can be implemented by subtracting the first components and rerunning the algorithms. However, there remain interesting theoretical questions regarding the relations between the sparsity patterns of the different components. Second, while here we considered the case of a pair of multivariates, it is possible to generalize the setting and address multivariate correlations between more than two data sets.

REFERENCES

- [1] H. Hotelling, “Relations between two sets of variates,” *Biometrika*, vol. 28, pp. 321–377, 1936.
- [2] B. Thompson, *Canonical correlation analysis: uses and interpretation*. SAGE publications, 1984.
- [3] T. Anderson, *An Introduction to Multivariate Statistical Analysis*, 3rd ed. Wiley-Interscience, 2003.
- [4] M. Borge, T. Landelius, and H. Knutsson, “A unified approach to PCA, PLS, MLR and CCA,” *Report LiTH-ISY-R-1992, ISY*, November 1997.
- [5] F. R. Bach and M. I. Jordan, “Kernel independent component analysis,” *J. Mach. Learn. Res.*, vol. 3, pp. 1–48, 2003.
- [6] A. Gretton, R. Herbrich, A. Smola, O. Bousquet, and B. Schölkopf, “Kernel methods for measuring independence,” *J. Mach. Learn. Res.*, vol. 6, pp. 2075–2129, 2005.
- [7] H. D. Vinod, “Canonical ridge and econometrics of joint production,” *Journal of Econometrics*, vol. 4, no. 2, pp. 147–166, May 1976.
- [8] D. R. Hardoon and J. Shawe-Taylor, “Sparse canonical correlation analysis,” University College London, Technical Report, 2007.
- [9] A. d’Aspremont, L. E. Ghaoui, I. Jordan, and G. R. G. Lanckriet, “A direct formulation for sparse PCA using semidefinite programming,” in *Advances in Neural Information Processing Systems 17*. Cambridge, MA: MIT Press, 2005.
- [10] S. Waaijenborg and A. H. Zwinderman, “Penalized canonical correlation analysis to quantify the association between gene expression and DNA markers,” *BMC Proceedings 2007, 1(Suppl 1):S122*, Dec. 2007.
- [11] E. Parkhomenko, D. Trichler, and J. Beyene, “Genome-wide sparse canonical correlation of gene expression with genotypes,” *BMC Proceedings 2007, 1(Suppl 1):S119*, Dec. 2007.
- [12] C. Fyfe and G. Leen, “Two methods for sparsifying probabilistic canonical correlation analysis,” in *ICONIP (1)*, 2006, pp. 361–370.
- [13] L. Tan and C. Fyfe, “Sparse kernel canonical correlation analysis,” in *ESANN*, 2001, pp. 335–340.
- [14] B. K. Sriperumbudur, D. A. Torres, and G. R. G. Lanckriet, “Sparse eigen methods by D.C. programming,” in *ICML ’07: Proceedings of the 24th international conference on Machine learning*. New York, NY, USA: ACM, 2007, pp. 831–838.
- [15] H. Zou, T. Hastie, and R. Tibshirani, “Sparse principal component analysis,” *Journal of Computational; Graphical Statistics*, vol. 15, pp. 265–286(22), June 2006.
- [16] B. Moghaddam, Y. Weiss, and S. Avidan, “Spectral bounds for sparse PCA: Exact and greedy algorithms,” in *Advances in Neural Information Processing Systems 18*, Y. Weiss, B. Schölkopf, and J. Platt, Eds. Cambridge, MA: MIT Press, 2006, pp. 915–922.
- [17] A. d’Aspremont, F. Bach, and L. E. Ghaoui, “Full regularization path for sparse principal component analysis,” in *ICML ’07: Proceedings of the 24th international conference on Machine learning*. New York, NY, USA: ACM, 2007, pp. 177–184.
- [18] R. Tibshirani, “Regression shrinkage and selection via the lasso,” *J. Roy. Statist. Soc. Ser. B*, vol. 58, no. 1, pp. 267–288, 1996.
- [19] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani, “Least angle regression,” *Annals of Statistics (with discussion)*, vol. 32, no. 1, pp. 407–499, 2004.