

Scalable and Flexible Multiview MAX-VAR Canonical Correlation Analysis

Xiao Fu, *Member, IEEE*, Kejun Huang, *Member, IEEE*, Mingyi Hong, *Member, IEEE*
 Nicholas D. Sidiropoulos, *Fellow, IEEE*, and Anthony Man-Cho So, *Senior Member, IEEE*

Abstract— Generalized canonical correlation analysis (GCCA) aims at finding latent low-dimensional common structure from multiple views (feature vectors in different domains) of the same entities. Unlike principal component analysis (PCA) that handles a single view, (G)CCA is able to integrate information from different feature spaces. Here we focus on MAX-VAR GCCA, a popular formulation which has recently gained renewed interest in multilingual processing and speech modeling. The classic MAX-VAR GCCA problem can be solved optimally via eigen-decomposition of a matrix that compounds the (whitened) correlation matrices of the views; but this solution has serious scalability issues, and is not directly amenable to incorporating pertinent structural constraints such as non-negativity and sparsity on the canonical components. We posit regularized MAX-VAR GCCA as a non-convex optimization problem and propose an alternating optimization (AO)-based algorithm to handle it. Our algorithm alternates between *inexact* solutions of a regularized least squares subproblem and a manifold-constrained non-convex subproblem, thereby achieving substantial memory and computational savings. An important benefit of our design is that it can easily handle structure-promoting regularization. We show that the algorithm globally converges to a critical point at a sublinear rate, and approaches a global optimal solution at a linear rate when no regularization is considered. Judiciously designed simulations and large-scale word embedding tasks are employed to showcase the effectiveness of the proposed algorithm.

Index terms— Canonical correlation analysis (CCA), multiview CCA, MAX-VAR, word embedding, optimization, scalability, feature selection

I. INTRODUCTION

Canonical correlation analysis (CCA) [1] produces low-dimensional representations via finding common structure of two or more views corresponding to the same entities. A view contains high-dimensional representations of the entities in a certain feature space – e.g., the text and audio representations of a given word can be considered as different views of this word. CCA is able to deal with views that have different

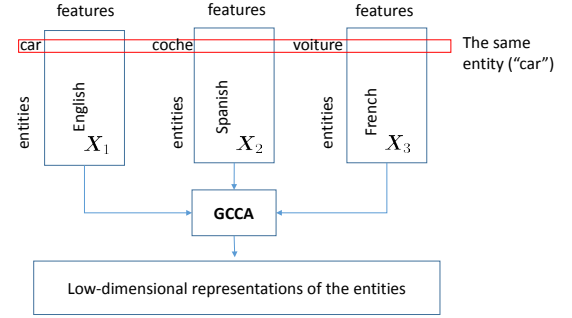


Fig. 1. Word embedding seeks low-dimensional representations of the entities (words) that are well-aligned with human judgment. Different language data (i.e., X_1 - X_3) can be considered as different views / feature spaces of the same entities.

dimensions, and this flexibility is very useful in data fusion, where one is interested in integrating information acquired from different domains. Multiview analysis finds numerous applications in signal processing and machine learning, such as blind source separation [2], [3], direction-of-arrival estimation [4], wireless channel equalization [5], regression [6], clustering [7], speech modeling and recognition [8], [9], and word embedding [10], to name a few. Classical CCA was derived for the two-view case, but *generalized canonical correlation analysis* (GCCA) that aims at handling more than two views has a long history as well [11], [12]. A typical application of GCCA, namely, multilingual word embedding, is shown in Fig. 1. Applying GCCA to integrate multiple languages was shown to yield better embedding results relative to single-view analyses such as principle component analysis (PCA) [10].

Computationally, GCCA poses interesting and challenging optimization problems. Unlike the two-view case that admits an algebraically simple solution (via eigen-decomposition), GCCA is in general not easily solvable. Many prior works considered the GCCA problem with different cost functions [11]–[13] – see a nice summary in [14, Chapter 10]. However, the proposed algorithms often can only extract a single canonical component and then find others through a deflation process, which is known to suffer from error propagation. CCA and GCCA can also pose serious scalability challenges, since they involve auto- and/or cross-correlations of different views and a whitening stage [15]. These procedures can easily lead to memory explosion and require a large number of flops for computation. They also destroy the sparsity of the data, which is usually what one relies upon to deal with large-scale

X. Fu, K. Huang and N.D. Sidiropoulos are supported in part by National Science Foundation under Project NSF ECCS-1608961 and Project NSF IIS-1447788. M. Hong is supported in part by National Science Foundation under Project NSF CCF-1526078 and by the Air Force Office of Scientific Research (AFOSR) under Grant 15RT0767.

X. Fu, K. Huang and N.D. Sidiropoulos are with the Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN 55455, e-mail (xfu,huang663,nikos)@umn.edu.

M. Hong is with Department of Industrial and Manufacturing Systems Engineering, Iowa State University, Ames, Iowa 50011, (515) 294-4111, Email: mingyi@iastate.edu.

A. M.-C. So is with the Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong, Shatin, N.T., Hong Kong, Email: manchoso@se.cuhk.edu.hk

problems. In recent years, effort has been spent on solving these scalability issues, but the focus is mostly on the two-view case [15]–[17].

Among all different formulations of GCCA, there is a particular one that admits a conceptually simple solution, the so-called MAX-VAR GCCA [11], [13], [18]. MAX-VAR GCCA was first proposed in [12], and its solution amounts to finding the ‘directions’ aligned to those exhibiting maximum variance for a matrix aggregated from the (whitened) auto-correlations of the views. It can also be viewed as a problem of enforcing *identical* latent representations of different views as opposed to highly correlated ones, which is the more general goal of (G)CCA. The merit of MAX-VAR GCCA is that it can be solved via eigen-decomposition and finds all the canonical components simultaneously. In practice, MAX-VAR GCCA also demonstrates promising performance in various applications such as word embedding [10] and speech recognition [8]. On the other hand, MAX-VAR GCCA has the same scalability problem as the other GCCA formulations: It involves correlation matrices of different views and their inverses, which is prohibitive to even instantiate when the data dimension is large. The work in [10] provided a pragmatic way to circumvent this difficulty: PCA was first applied to each view to reduce the rank of the views, and then MAX-VAR GCCA was applied to the rank-truncated views. Such a procedure significantly reduces the number of parameters for characterizing the views and is feasible in terms of memory. However, truncating the rank of the views is prone to information loss, and thus leads to performance degradation.

Besides the basic (G)CCA formulations, *structured* (G)CCA [19] that seeks canonical components with pre-specified structure is often considered in applications. Sparse/group-sparse CCA has attracted particular attention, since it has the ability of discarding outlying or irrelevant features when performing CCA [20]–[22]. In multi-lingual word embedding [10], [16], [23], for example, it is known that outlying features (“stop words”), may exist. Gene analysis is another example [20]–[22]. Ideally, CCA seeks a few highly correlated latent components, and so it should naturally be able to identify and down-weight irrelevant features automatically. In practice, however, this ability is often impaired when correlations cannot be reliably estimated, when one only has access to relatively few and/or very noisy samples, or when there is model mismatch due to bad preprocessing. In those cases, performing feature selection jointly with (G)CCA is well-motivated. Some other structure-promoting regularizations may also be of interest: Non-negativity together with sparsity have proven helpful in analyzing audio and video data, since non-negative CCA produces weighted sums of video frames that are interpretable [24]; non-negative CCA has also proven useful in time series analysis [25]. Effective algorithms that tackle large-scale structured GCCA problems are currently missing, to the best of our knowledge.

Contributions In this work, our goal is to provide a scalable and flexible algorithmic framework for handling the MAX-VAR GCCA problem and its variants with structure-promoting regularizers. Instead of truncating the rank of the views as in

[10], we keep the data *intact* and deal with the problem using a two-block alternating optimization (AO) framework. The proposed algorithm alternates between a regularized least squares subproblem and an orthogonality-constrained subproblem. The merit of this framework is that correlation matrices of the views never need to be explicitly instantiated, and the inversion procedure is avoided. Consequently, the algorithm consumes significantly less memory compared to that required by the original solution using eigen-decomposition. The proposed algorithm allows inexact solution to the subproblems, and thus per-iteration computational complexity is also light. In addition, it can easily handle different structure-promoting regularizers (e.g. sparsity, group sparsity and non-negativity) without increasing memory and computational costs, including the feature-selective regularizers that we are mainly interested in.

The AO algorithm alternates between convex and non-convex manifold-constrained subproblems, using possibly inexact updates for the subproblems. Under such circumstances, general convergence analysis tools cannot be directly applied, and thus the associated convergence properties are not obvious. This necessitates custom convergence analysis. We first show that the proposed algorithm *globally* converges to a Karush-Kuhn-Tucker (KKT) point of the formulated problem, even when a variety of regularizers are employed. We also show that the optimality gap shrinks to at most $\mathcal{O}(1/r)$ after r iterations – i.e., at least a sublinear convergence rate can be guaranteed. In addition, we show that when the classic MAX-VAR problem without regularization (or with a minimal energy regularization) is considered, the proposed algorithm approaches a *global optimal solution* and enjoys a *linear* convergence rate.

The proposed algorithm is applied to judiciously designed simulated data, as well as a real large-scale word embedding problem, and promising results are observed.

A conference version of this work appears at ICASSP 2017, New Orleans, USA, Mar. 2017 [26]. This journal version includes detailed convergence analysis and proofs, comprehensive simulations, and a set of experiments using real large-scale multilingual data.

Notation We use \mathbf{X} and \mathbf{x} to denote a matrix and a vector, respectively. $\mathbf{X}(m, :)$ and $\mathbf{X}(:, n)$ denote the m th row and the n th column of \mathbf{X} , respectively; in particular, $\mathbf{X}(:, n_1 : n_2)$ ($\mathbf{X}(n_1 : n_2, :)$) denotes a submatrix of \mathbf{X} consisting of the n_1 – n_2 th columns (rows) of \mathbf{X} (MATLAB notation). $\|\mathbf{X}\|_F$ and $\|\mathbf{X}\|_p$ for $p \geq 1$ denote the Frobenius norm and the matrix-induced p -norm, respectively. $\|\mathbf{X}\|_{p,1} = \sum_{i=1}^m \|\mathbf{X}(i, :)\|_p$ for $p \geq 1$ denotes the ℓ_p/ℓ_1 -mixed norm of $\mathbf{X} \in \mathbb{R}^{m \times n}$. The superscripts “ T ”, “ $+$ ”, and “ -1 ” denote the matrix operators of transpose, pseudo-inverse and inverse, respectively. The operator $\langle \mathbf{X}, \mathbf{Y} \rangle$ denotes the inner product of \mathbf{X} and \mathbf{Y} . $\mathbf{1}_+(\mathbf{X})$ denotes the element-wise indicator function of the nonnegative orthant – i.e., $\mathbf{1}_+(\mathbf{X}) = +\infty$ if any element of \mathbf{X} is negative and $\mathbf{1}_+(\mathbf{X}) = 0$ otherwise.

II. BACKGROUND

Consider a scenario where L entities have different representations in I views. Let $\mathbf{X}_i \in \mathbb{R}^{L \times M_i}$ denote the i th view

with its ℓ th row $\mathbf{X}_i(\ell, :)$ being a feature vector that defines the ℓ th data point (entity) in the i th view (cf. Fig. 1), where M_i is the dimension of the i th feature space. The classic two-view CCA aims at finding common structure of the views via linear transformation. Specifically, the corresponding problem can be expressed in the following form [1]:

$$\min_{\mathbf{Q}_1, \mathbf{Q}_2} \|\mathbf{X}_1 \mathbf{Q}_1 - \mathbf{X}_2 \mathbf{Q}_2\|_F^2 \quad (1a)$$

$$\text{s.t. } \mathbf{Q}_i^T (\mathbf{X}_i^T \mathbf{X}_i) \mathbf{Q}_i = \mathbf{I}, \quad i = 1, 2, \quad (1b)$$

where the columns of $\mathbf{Q}_i \in \mathbb{R}^{M_i \times K}$ correspond to the K canonical components of view \mathbf{X}_i , and K is usually small (i.e., $K \ll \min\{M_i, L\}$). Note that we are essentially maximizing the trace of the estimated cross-correlations between the reduced-dimension views, i.e., $\text{Tr}(\mathbf{Q}_2^T \mathbf{X}_2^T \mathbf{X}_1 \mathbf{Q}_1)$ subject to the normalization in (1b) – which motivates the terminology “correlation analysis”. Problem (1) can be solved via a generalized eigen-decomposition, but this simple solution only applies to the two-view case. To analyze the case with more than two views, one natural thought is to extend the formulation in (1) to a pairwise matching criterion, i.e., $\sum_{i=1}^{I-1} \sum_{j=i+1}^I \|\mathbf{X}_i \mathbf{Q}_i - \mathbf{X}_j \mathbf{Q}_j\|_F^2$ with orthogonality constraints on $\mathbf{X}_i \mathbf{Q}_i$ for all i , where I is the number of views. Such an extension leads to the so-called sum-of-correlations (SUMCOR) generalized CCA [11], which has been shown to be NP-hard [27]. Notice that designing efficient and scalable algorithms for SUMCOR is an interesting topic and it started attracting attention recently [27]–[29]. Another formulation of GCCA is more tractable: Instead of forcing pairwise similarity of the reduced-dimension views, one can seek a common latent representation of different views, i.e., [8], [10], [11], [13], [18]

$$\min_{\{\mathbf{Q}_i\}_{i=1}^I, \mathbf{G}} \sum_{i=1}^I (1/2) \|\mathbf{X}_i \mathbf{Q}_i - \mathbf{G}\|_F^2, \quad (2)$$

$$\text{s.t. } \mathbf{G}^T \mathbf{G} = \mathbf{I},$$

where $\mathbf{G} \in \mathbb{R}^{L \times K}$ is a common latent representation of the different views. Problem (2) also finds highly correlated reduced-dimension views as SUMCOR does. The upshot of Problem (2) is that it “transfers” the multiple difficult constraints $\mathbf{Q}_i^T \mathbf{X}_i^T \mathbf{X}_i \mathbf{Q}_i = \mathbf{I}$ to a single constraint $\mathbf{G}^T \mathbf{G} = \mathbf{I}$, and thus admits a *conceptually* simple algebraic solution, which, as we will show, has the potential to be scaled up to deal with very large problems. In this work, we will focus on Problem (2) and its variants.

Problem (2) is referred to as the MAX-VAR formulation of GCCA since the optimal solution amounts to taking principal eigenvectors of a matrix aggregated from the correlation matrices of the views. To explain, let us first assume that \mathbf{X}_i has full column rank and solve (2) with respect to (w.r.t.) \mathbf{Q}_i , i.e., $\mathbf{Q}_i = \mathbf{X}_i^\dagger \mathbf{G}$, where $\mathbf{X}_i^\dagger = (\mathbf{X}_i^T \mathbf{X}_i)^{-1} \mathbf{X}_i^T$. By substituting it back to (2), we see that an optimal solution \mathbf{G}_{opt} can be obtained via solving the following:

$$\mathbf{G}_{\text{opt}} = \arg \max_{\mathbf{G}^T \mathbf{G} = \mathbf{I}} \text{Tr} \left(\mathbf{G}^T \left(\sum_{i=1}^I \mathbf{X}_i \mathbf{X}_i^\dagger \right) \mathbf{G} \right). \quad (3)$$

Let $\mathbf{M} = \sum_{i=1}^I \mathbf{X}_i \mathbf{X}_i^\dagger$. Then, an optimal solution is $\mathbf{G}_{\text{opt}} = \mathbf{U}_M(:, 1 : K)$, i.e., the first K principal eigenvectors of \mathbf{M}

[30]. Although Problem (2) admits a seemingly easy solution, implementing it in practice has two major challenges:

1) **Scalability Issues:** Implementing the eigen-decomposition based solution for large-scale data is prohibitive. As mentioned, instantiating $\mathbf{M} = \sum_{i=1}^I \mathbf{X}_i (\mathbf{X}_i^T \mathbf{X}_i)^{-1} \mathbf{X}_i^T$ is not doable when L and M_i 's are large. The matrix \mathbf{M} is an $L \times L$ matrix. In applications like word embedding, L and M_i are the vocabulary size of a language and the number of features defining the terms, respectively, which can both easily exceed 100,000. This means that the memory for simply instantiating \mathbf{M} or $(\mathbf{X}_i^T \mathbf{X}_i)^{-1}$ can reach 75GB. In addition, even if the views \mathbf{X}_i are sparse, computing $(\mathbf{X}_i^T \mathbf{X}_i)^{-1}$ will create large dense matrices and make it difficult to exploit sparsity in the subsequent processing. To circumvent these difficulties, Rastogi *et al.* [10] proposed to first apply the singular value decomposition (SVD) to the views, i.e., $\text{svd}(\mathbf{X}_i) = \mathbf{U}_i \mathbf{\Sigma}_i \mathbf{V}_i^T$, and then let $\hat{\mathbf{X}}_i = \mathbf{U}_i(:, 1 : P) \mathbf{\Sigma}_i(1 : P, 1 : P) (\mathbf{V}_i(:, 1 : P))^T \approx \mathbf{X}_i$, where P is much smaller than M_i and L . This procedure enables one to represent the views with significantly fewer parameters, i.e., $(L + M_i + 1)P$ compared to LM_i , and allows the original eigen-decomposition based solution to MAX-VAR GCCA to be applied; see more details in [10]. The drawback, however, is also evident: The procedure truncates the rank of the views significantly (since in practice the views almost always have full column-rank, i.e., $\text{rank}(\mathbf{X}_i) = M_i$), and rank-truncation is prone to information losses. Therefore, it is much more appealing to deal with the intact views.

2) **Structure-Promoting:** Another aspect that is under-addressed by existing approaches is how to incorporate regularizations on \mathbf{Q}_i to multiview large-scale CCA. Note that finding structured \mathbf{Q}_i is well-motivated in practice. Taking multilingual word embedding as an example, $\mathbf{X}_i(:, n)$ represents the n th feature in language i , which is usually defined by the co-occurrence frequency of the words and feature n (also a word in language i). However, many features of \mathbf{X}_i may not be informative (e.g., “the” and “a” in English) or not correlated to data in \mathbf{X}_j . These *irrelevant* or *outlying features* could result in unsatisfactory performance of GCCA if not taken into account. Under such scenarios, a more appealing formulation may include a row-sparsity promoting regularization on \mathbf{Q}_i so that some columns corresponding to the irrelevant features in \mathbf{X}_i can be discounted/downweighted when seeking \mathbf{Q}_i . Sparse (G)CCA is desired in a variety of applications such as gene analytics and fMRI prediction [20]–[22], [31], [32]. Other structure such as nonnegativity of \mathbf{Q}_i was also shown useful in data analytics for maintaining interpretability and enhancing performance; see [24], [25].

III. PROPOSED ALGORITHM

In this work, we consider a scalable and flexible algorithmic framework for handling MAX-VAR GCCA and its variants with structure-promoting regularizers on \mathbf{Q}_i . We aim at offering simple solutions that are memory-efficient, admit light per-iteration complexity, and feature good convergence properties under certain mild conditions. Specifically, we consider the

following formulation:

$$\begin{aligned} \min_{\{Q_i\}, G} \quad & \sum_{i=1}^I (1/2) \|X_i Q_i - G\|_F^2 + \sum_{i=1}^I h_i(Q_i), \\ \text{s.t.} \quad & G^T G = I, \end{aligned} \quad (4)$$

where $h_i(\cdot)$ is a regularizer that imposes a certain structure on Q_i . Popular regularizers include

$$h_i(Q_i) = \mu_i/2 \cdot \|Q_i\|_F^2, \quad (5a)$$

$$h_i(Q_i) = \mu_i \cdot \|Q_i\|_{2,1}, \quad (5b)$$

$$h_i(Q_i) = \mu_i \cdot \|Q_i\|_{1,1}, \quad (5c)$$

$$h_i(Q_i) = \mu_i/2 \cdot \|Q_i\|_F^2 + \beta_i \cdot \|Q_i\|_{2,1}, \quad (5d)$$

$$h_i(Q_i) = \mu_i/2 \cdot \|Q_i\|_F^2 + \beta_i \cdot \|Q_i\|_{1,1}, \quad (5e)$$

$$h_i(Q_i) = \mathbf{1}_+(Q_i), \quad (5f)$$

where $\mu_i, \beta_i \geq 0$ are regularization parameters for balancing the least squares fitting term and the regularization terms. The first regularizer is commonly used for controlling the energy of the dimension-reducing matrix Q_i , which also has an effect of improving the conditioning of the subproblem w.r.t. Q_i . $h_i(Q_i) = \mu_i \|Q_i\|_{2,1}$ that we are mainly interested in has the ability of promoting rows of Q_i to be zeros (or approximately zeros), and thus can suppress the impact of the corresponding columns (features) in X_i – which is effectively feature selection. The function $h_i(Q_i) = \mu_i \|Q_i\|_{1,1}$ also does feature selection, but different columns of the dimension-reduced data, i.e., $X_i Q_i$, may use different features. The regularizers in (5d)–(5e) are sometimes referred to as the *elastic net regularizers* in statistics, which improve conditioning of the Q_i -subproblem and perform feature selection at the same time. $h_i(Q_i) = \mathbf{1}_+(Q_i)$ is for restraining the canonical components to be non-negative so that $X_i Q_i$ maintains interpretability in some applications like video analysis – where the columns of $X_i Q_i$ are weighted combinations of time frames [25]; joint nonnegativity and sparsity regularizers can also be considered [24]. In this section, we propose an algorithm that can deal with the regularized and the original versions of MAX-VAR GCCA under a unified framework.

A. Alternating Optimization

To deal with Problem (4), our approach is founded on alternating optimization (AO); i.e., we solve two subproblems w.r.t. $\{Q_i\}$ and G , respectively. As will be seen, such a simple strategy will lead to highly scalable algorithms in terms of both memory and computational cost.

To begin with, let us assume that after r iterations the current iterate is $(Q^{(r)}, G^{(r)})$ where $Q = [Q_1^T, \dots, Q_I^T]^T$ and consider the subproblem

$$\min_{Q_i} (1/2) \|X_i Q_i - G^{(r)}\|_F^2 + h_i(Q_i), \quad \forall i. \quad (6)$$

The above problem is a regularized least squares problem. When X_i is large and sparse, many efficient algorithms can be considered to solve it. For example, the alternating direction method of multipliers (ADMM) [33] is frequently employed to handle Problem (6) in a scalable manner. However, ADMM

is a primal-dual method that does not guarantee monotonic decrease of the objective value, which will prove useful in later convergence analysis. Hence, we propose to employ the proximal gradient (PG) method for handling Problem (6). To explain, let us denote $Q_i^{(r,t)}$ as the t th update of Q_i when $G^{(r)}$ is fixed. Under this notation, we have $Q_i^{(r,0)} = Q_i^{(r)}$ and $Q_i^{(r,T)} = Q_i^{(r+1)}$. Let us rewrite (6) as

$$\min_{Q_i} f_i(Q_i, G^{(r)}) + g_i(Q_i), \quad (7)$$

where we define $f_i(Q_i, G^{(r)})$ and $g_i(Q_i)$ as the continuously differentiable part and the non-smooth part of the objective function in (6), respectively. We also define $\nabla_{Q_i} f_i(Q_i, G^{(r)})$ as the partial derivative of the differentiable part w.r.t. Q_i . When “single-component” regularizers such as $h_i(Q_i) = \|Q_i\|_{2,1}$ are employed, we have $f_i(Q_i, G^{(r)}) = (1/2) \|X_i Q_i - G^{(r)}\|_F^2$ and $g_i(Q_i) = h_i(Q_i)$; when $h_i(Q_i)$ has multiple components such as $h_i(Q_i) = \mu_i/2 \cdot \|Q_i\|_F^2 + \beta_i \cdot \|Q_i\|_{1,1}$, we have $f_i(Q_i, G^{(r)}) = (1/2) \|X_i Q_i - G^{(r)}\|_F^2 + \mu_i/2 \|Q_i\|_F^2$ and $g_i(Q_i) = \beta_i \cdot \|Q_i\|_{1,1}$.

Per PG, we update Q_i by the following rule:

$$Q_i^{(r,t+1)} \leftarrow \text{prox}_{\alpha_i g_i} \left(Q_i^{(r,t)} - \alpha_i \nabla_{Q_i} f_i(Q_i^{(r,t)}, G_i^{(r)}) \right) \quad (8)$$

$$= \arg \min_{Q_i} \frac{1}{2} \|Q_i - H_i^{(r,t)}\|_F^2 + g_i(Q_i)$$

where $H_i^{(r,t)} = Q_i^{(r,t)} - \alpha_i \nabla_{Q_i} f_i(Q_i^{(r,t)}, G_i^{(r)})$. For many $g_i(\cdot)$'s, the proximity operator in (8) has closed-form or lightweight solutions [34].

For example, if one adopts $g_i(Q_i) = \mu_i \|Q_i\|_{2,1}$, the update rule becomes

$$Q_i^{(r,t)}(m, :) \leftarrow \begin{cases} 0, & \|H_i^{(r,t)}(m, :)\|_2 < \mu_i, \\ \left(1 - \frac{\mu_i}{\|H_i^{(r,t)}(m, :)\|_2}\right) \|H_i^{(r,t)}(m, :)\|_2, & \text{o.w.} \end{cases}$$

For $g_i(Q_i) = \mu_i \|Q_i\|_{1,1}$, the update rule is similar to the above, which is known as the *soft-thresholding operator*. For $g_i(Q_i) = \mathbf{1}_+(Q_i)$, the solution is simply $Q_i^{(r,t)} = \max\{H_i^{(r,t)}, 0\}$. An even simpler case is $h_i(Q_i) = (\mu_i/2) \|Q_i\|_F^2$; for this case, the update of Q_i is simply gradient descent, i.e.,

$$Q_i^{(r,t+1)} \leftarrow Q_i^{(r,t)} - \alpha_i \left((X_i^T X_i + \mu_i I) Q_i^{(r,t)} - X_i^T G^{(r)} \right),$$

since the Q_i -subproblem in (6) does not have a non-smooth part.

By updating Q_i using the rule in (8) for T times where $T \geq 1$, we obtain $Q_i^{(r+1)}$. Next, we consider solving the subproblem w.r.t. G when fixing $\{Q_i\}_{i=1}^I$. The G -subproblem amounts to solving the following:

$$\min_{G^T G = I} \sum_{i=1}^I (1/2) \|X_i Q_i^{(r+1)} - G\|_F^2. \quad (9)$$

Expanding the above and dropping the constants, we come up with the following equivalent problem:

$$\max_{G^T G = I} \text{Tr} \left(G^T \sum_{i=1}^I X_i Q_i^{(r+1)} / I \right).$$

An optimal solution of \mathbf{G} is the so-called Procrustes projection [35], which is implemented as follows: Let $\mathbf{R} = \sum_{i=1}^I \mathbf{X}_i \mathbf{Q}_i^{(r+1)}$. Then, we have

$$\mathbf{G}^{(r+1)} \leftarrow \mathbf{U}_R \mathbf{V}_R^T,$$

where $\mathbf{U}_R \Sigma_R \mathbf{V}_R^T = \text{svd}(\mathbf{R}, \text{'econ'})$, and $\text{svd}(\cdot, \text{'econ'})$ denotes the economy-size SVD that produces $\mathbf{U}_R \in \mathbb{R}^{L \times K}$, $\Sigma_R \in \mathbb{R}^{K \times K}$ and $\mathbf{V}_R^T \in \mathbb{R}^{K \times K}$. The above update is optimal in terms of solving the subproblem. However, since this subproblem has multiple optimal solutions, picking an arbitrary one from the solution set results in difficulties in analyzing some aspects of the algorithm (specifically, the rate of convergence). To overcome this issue, we propose to solve the following

$$\min_{\mathbf{G}^T \mathbf{G} = \mathbf{I}} \sum_{i=1}^I \frac{1}{2} \left\| \mathbf{X}_i \mathbf{Q}_i^{(r+1)} - \mathbf{G} \right\|_F^2 + \omega \cdot \left\| \mathbf{G} - \mathbf{G}^{(r)} \right\|_F^2, \quad (10)$$

where $\omega = (1-\gamma)I/2\gamma$ and $\gamma \in (0, 1]$. Note that $\omega \geq 0$ and the proximal term is added to ensure that $\mathbf{G}^{(r+1)}$ will not wander very far from $\mathbf{G}^{(r)}$. An optimal solution to the above is still simple: The only change to the original \mathbf{G} -solution is to use the following modified \mathbf{R}

$$\mathbf{R} = \gamma \sum_{i=1}^I \mathbf{X}_i \mathbf{Q}_i^{(r+1)} / I + (1-\gamma) \mathbf{G}^{(r)}, \quad (11)$$

and the other operations (e.g., the economy-size SVD) remain the same. While the addition of the proximal term may seem to “degrade” an optimal solution of the \mathbf{G} -subproblem (i.e., Problem (9)) to an inexact one, this simple change helps establish nice convergence rate properties of the overall algorithm, as we will see.

The algorithm is summarized in Algorithm 1, which we call the alternating optimization-based MAX-VAR GCCA (AltMaxVar). As one can see, the algorithm does not instantiate any large dense matrix during the procedure and thus is highly efficient in terms of memory. Also, the procedure does not destroy sparsity of the data, and thus the computational burden is light when the data is sparse – which is often the case in large-scale learning applications. Detailed complexity analysis will be presented in the next subsection.

B. Computational and Memory Complexities

The update rule in (8) inherits the good features from the PG method. First, there is no “heavy computation” if the views \mathbf{X}_i for $i = 1, \dots, I$ are sparse. Specifically, the major computation in the update rule of (8) is computing the partial gradient of the smooth part of the cost function, i.e., $\nabla_{\mathbf{Q}_i} f_i(\mathbf{Q}_i, \mathbf{G}_i)$. To this end, $\mathbf{X}_i \mathbf{Q}_i$ should be calculated first, since if \mathbf{X}_i is sparse, this matrix multiplication step has a complexity order of $\mathcal{O}(\text{nnz}(\mathbf{X}_i) \cdot K)$ flops, where $\text{nnz}(\cdot)$ counts the number of non-zeros. The next multiplication, i.e., $\mathbf{X}_i^T (\mathbf{X}_i \mathbf{Q}_i)$, has the same complexity order. Similarly, the operation of $\mathbf{X}_i^T \mathbf{G}$ has the same complexity. For solving the \mathbf{G} -subproblem, the major operation is the SVD of \mathbf{R} . This step is also not computationally heavy – what we ask for is an economy-size

Algorithm 1: AltMaxVar

```

input :  $\{\mathbf{X}_i, \mu_i, \alpha_i\}_{i=1}^I; \gamma \in (0, 1]; K; T; (\{\mathbf{Q}_i^{(0)}\}_{i=1}^I, \mathbf{G}^{(0)})$ .
1  $r \leftarrow 0$ ;
2 repeat
3    $t \leftarrow 0$ ;
4    $\mathbf{E}_i^{(t)} \leftarrow \mathbf{Q}_i^{(r)}$  for  $i = 1, \dots, I$ ;
5   while  $t \leq T$  and convergence not reached do
6     for all  $i$ , update
7        $\mathbf{H}_i^{(r,t)} \leftarrow \mathbf{Q}_i^{(r,t)} - \alpha_i \nabla_{\mathbf{Q}_i} f_i(\mathbf{Q}_i^{(r,t)}; \mathbf{G}_i^{(r)})$ ;
8        $\mathbf{Q}_i^{(r,t+1)} \leftarrow \text{prox}_{\alpha_i g_i}(\mathbf{H}_i^{(r,t)})$ ;
9      $t \leftarrow t + 1$ ;
10  end
11   $\mathbf{Q}_i^{(r+1)} \leftarrow \mathbf{Q}_i^{(r,T)}$ ;
12   $\mathbf{R} \leftarrow \gamma \sum_{i=1}^I \mathbf{X}_i \mathbf{Q}_i^{(r+1)} / I + (1-\gamma) \mathbf{G}^{(r)}$ ;
13   $\mathbf{U}_R \Sigma_R \mathbf{V}_R^T \leftarrow \text{svd}(\mathbf{R}, \text{'econ'})$ ;
14   $\mathbf{G}^{(r+1)} \leftarrow \mathbf{U}_R \mathbf{V}_R^T$ ;
15   $r \leftarrow r + 1$ ;
16 until Some stopping criterion is reached;
output:  $\{\mathbf{Q}_i^{(r)}\}_{i=1}^I, \mathbf{G}^{(r)}$ 

```

SVD of a very thin matrix (of size $L \times K$, $L \gg K$). This has a complexity order of $\mathcal{O}(LK^2)$ flops [30], which is light.

In terms of memory, all the terms involved (i.e., \mathbf{Q}_i , \mathbf{G}_i , $\mathbf{X}_i \mathbf{Q}_i$, $\mathbf{X}_i^T \mathbf{X}_i \mathbf{Q}_i$ and $\mathbf{X}_i^T \mathbf{G}_i$) only require $\mathcal{O}(LK)$ memory or less, but the eigen-decomposition-based solution needs $\mathcal{O}(M_i^2)$ and $\mathcal{O}(L^2)$ memory to store $(\mathbf{X}_i^T \mathbf{X}_i)^{-1}$ and \mathbf{M} , respectively. Note that K is usually very small (and up to our control) compared to L and M_i , which are approximately of the same large size in applications like word embedding.

IV. CONVERGENCE PROPERTIES

In this section, we study convergence properties of AltMaxVar. Note that the algorithm alternates between a (possibly) non-smooth subproblem and a manifold-constrained subproblem, and the subproblems may or may not be solved to optimality. Existing convergence analysis for exact and inexact block coordinate descent such as those in [36]–[39] can not be directly applied to analyze AltMaxVar, and thus its convergence properties are not obvious. For clarity of exposition, we first define a critical point, or, a KKT point, of Problem (4). A KKT point $(\mathbf{G}^*, \mathbf{Q}^*)$ satisfies the following first-order optimality conditions:

$$\begin{cases} \mathbf{0} \in \nabla_{\mathbf{Q}_i} f_i(\mathbf{Q}_i^*, \mathbf{G}^*) + \partial_{\mathbf{Q}_i} g_i(\mathbf{Q}_i^*), \quad \forall i \\ \mathbf{0} = \nabla_{\mathbf{G}} \sum_{i=1}^I f_i(\mathbf{Q}_i^*, \mathbf{G}^*) + \mathbf{G}^* \mathbf{\Lambda}^*, \quad (\mathbf{G}^*)^T \mathbf{G}^* = \mathbf{I}, \end{cases}$$

where $\mathbf{\Lambda}$ is a Lagrangian multiplier associated with the constraint $\mathbf{G}^T \mathbf{G} = \mathbf{I}$, and $\partial_{\mathbf{Q}_i} g_i(\mathbf{Q}_i)$ denotes a subgradient of the (possibly) non-smooth function $g_i(\mathbf{Q}_i)$. We first show that

Proposition 1 Assume that $\alpha_i \leq 1/L_i$ for all i , where $L_i = \lambda_{\max}(\mathbf{X}_i^T \mathbf{X}_i)$ is the largest eigenvalue of $\mathbf{X}_i^T \mathbf{X}_i$. Also assume that $g_i(\cdot)$ is a closed convex function, $T \geq 1$, and $\gamma \in (0, 1]$. Then, the following holds:

- (a) The objective value of Problem (2) is non-increasing. In addition, every limit point of the solution sequence $\{\mathbf{G}^{(r)}, \{\mathbf{Q}_i^{(r)}\}_{i=1}^I\}_r$ is a KKT point of Problem (2).

- (b) If \mathbf{X}_i and $\mathbf{Q}_i^{(0)}$ for $i = 1, \dots, I$ are bounded and $\text{rank}(\mathbf{X}_i) = M_i$, then, the whole solution sequence converges to the set \mathcal{K} that consists of all the KKT points.

Proposition 1 (a) characterizes the limit points of the solution sequence: Even if only one proximal gradient step is performed in each iteration r , every convergent subsequence of the solution sequence attains a KKT point of Problem (4). As we demonstrate in the proof (relegated to the Appendix), AltMaxVar can be viewed as an algorithm that successively deals with local upper bounds of the two subproblems, which has a similar flavor as *block successive upper bound minimization* (BSUM) [37]. However, the generic BSUM framework does not cover nonconvex constraints such as $\mathbf{G}^T \mathbf{G} = \mathbf{I}$. Hence, the convergence properties of BSUM cannot be applied to show Proposition 1. To fill this gap, careful custom convergence analysis is provided in the appendix. The (b) part of Proposition 1 establishes the convergence of the whole solution sequence – which is a much stronger result. The assumption $\text{rank}(\mathbf{X}_i) = M_i$, on the other hand, is also relatively more restrictive.

It is also meaningful to estimate the number of iterations that is needed for the algorithm to reach a neighborhood of a KKT point. To this end, let us define the following potential function:

$$Z^{(r+1)} = \sum_{t=0}^{T-1} \sum_{i=1}^I \left\| \tilde{\nabla}_{\mathbf{Q}_i} F_i(\mathbf{Q}_i^{(r,t)}, \mathbf{G}^{(r)}) \right\|_F^2 + \left\| \mathbf{G}^{(r)} - \sum_{i=1}^I \mathbf{X}_i \mathbf{Q}_i^{(r+1)} / I + \mathbf{G}^{(r+1)} \mathbf{\Lambda}^{(r+1)} \right\|_F^2,$$

where $F_i(\mathbf{Q}_i, \mathbf{G}) = f_i(\mathbf{Q}_i, \mathbf{G}) + g_i(\mathbf{Q}_i, \mathbf{G})$, $\mathbf{\Lambda}^{(r+1)}$ is the Lagrangian multiplier associated with the solution $\mathbf{G}^{(r+1)}$, and

$$\tilde{\nabla}_{\mathbf{Q}_i} F_i(\mathbf{Q}_i^{(r,t)}, \mathbf{G}^{(r)}) = \frac{1}{\alpha_i} \left(\mathbf{Q}_i^{(r,t)} - \text{prox}_{\alpha_i g_i} \left(\mathbf{H}_i^{(r,t)} \right) \right).$$

Note that the update w.r.t. \mathbf{Q}_i can be written as $\mathbf{Q}_i^{(r,t+1)} = \mathbf{Q}_i^{(r,t)} - \alpha_i \tilde{\nabla}_{\mathbf{Q}_i} F_i(\mathbf{Q}_i^{(r,t)}, \mathbf{G}_i^{(r)})$ [34] – and therefore $\tilde{\nabla}_{\mathbf{Q}_i} F_i(\mathbf{Q}_i^{(r,t)}, \mathbf{G}^{(r)})$ is also called the *proximal gradient* of the \mathbf{Q}_i -subproblem w.r.t. \mathbf{Q}_i at $(\mathbf{Q}_i^{(r,t)}, \mathbf{G}^{(r)})$, as a counterpart of the classic gradient that is defined on smooth functions. One can see that $Z^{(r+1)}$ is a value that is determined by two consecutive outer iterates indexed by r and $r+1$ of the algorithm. $Z^{(r+1)}$ has the following property:

Lemma 1 $Z^{(r+1)} \rightarrow 0$ implies that $(\{\mathbf{Q}_i^{(r)}\}_i, \mathbf{G}^{(r)})$ approaches a KKT point.

The proof of Lemma 1 is in Appendix B. As a result, we can use the value of $Z^{(r+1)}$ to measure how close is the current iterate to a KKT point, thereby estimating the iteration complexity. Following this rationale, we show that

Theorem 1 Assume that $\alpha_i < 1/L_i$, $0 < \gamma < 1$ and $T \geq 1$. Let $\delta > 0$ and J be the number of iterations when $Z^{(r+1)} \leq \delta$ holds for the first time. Then, there exists a constant v such that $\delta \leq v/J-1$; that is, the algorithm converges to a KKT point at least sublinearly.

The proof of Theorem 1 is relegated to Appendix C. By Theorem 1, AltMaxVar reduces the optimality gap (measured by the Z-function) between the current iterate and a KKT point to $\mathcal{O}(1/r)$ after r iterations. One subtle point that is worth mentioning is that the analysis in Theorem 1 holds when $\gamma < 1$ – which corresponds to the case where the \mathbf{G} -subproblem in (9) is *not* optimally solved (to be specific, what we solve is a local surrogate in (10)). This reflects some interesting facts in AO – when the subproblems are handled in a more conservative way using a controlled step size, convergence rate may be guaranteed. On the other hand, more conservative step sizes may result in slower convergence. Hence, choosing an optimization strategy usually poses a trade-off between practical considerations such as speed and theoretical guarantees.

Proposition 1 and Theorem 1 characterize convergence properties of AltMaxVar with a general regularization term $h_i(\cdot)$. It is also interesting to consider the special case where $h_i(\cdot) = (\mu_i/2) \|\cdot\|_F^2$ – which correspond to the original MAX-VAR formulation (when $\mu_i = 0$) and its “diagonally loaded” version ($\mu_i > 0$). The corresponding problem is *optimally solvable* via taking the K leading eigenvectors of $\mathbf{M} = \sum_{i=1}^I \mathbf{X}_i (\mathbf{X}_i^T \mathbf{X}_i + \mu_i \mathbf{I})^{-1} \mathbf{X}_i^T$ [10]. It is natural to wonder if AltMaxVar has sacrificed optimality in dealing with this special case for the sake of gaining scalability? The answer is – thankfully – not really. This is not entirely surprising; to explain, let us denote $\mathbf{U}_1 = \mathbf{U}_M(:, 1 : K)$ and $\mathbf{U}_2 = \mathbf{U}_M(:, K+1 : L)$ as the K principal eigenvectors of \mathbf{M} and the eigenvectors spanning its orthogonal complement, respectively. Recall that our ultimate goal is to find \mathbf{G} that is a basis of the range space of \mathbf{U}_1 , denoted by $\mathcal{R}(\mathbf{U}_1)$. Hence, the speed of convergence can be measured through the distance between $\mathcal{R}(\mathbf{G})$ and $\mathcal{R}(\mathbf{U}_1)$. To this end, we adopt the definition of subspace distance in [30], i.e., $\text{dist}(\mathcal{R}(\mathbf{G}^{(r)}), \mathcal{R}(\mathbf{U}_1)) = \|\mathbf{U}_2^T \mathbf{G}^{(r)}\|_2$ and show that

Theorem 2 Denote the eigenvalues of $\mathbf{M} \in \mathbb{R}^{L \times L}$ by $\lambda_1, \dots, \lambda_L$ in descending order. Consider $h_i(\cdot) = \frac{\mu_i}{2} \|\cdot\|_F^2$ for $\mu_i \geq 0$ and let $\gamma = 1$. Assume that $\text{rank}(\mathbf{X}_i) = M_i$, $\lambda_K > \lambda_{K+1}$, and $\mathcal{R}(\mathbf{G}^{(0)})$ is not orthogonal to any component in $\mathcal{R}(\mathbf{U}_1)$, i.e.,

$$\cos(\theta) = \min_{\mathbf{u} \in \mathcal{R}(\mathbf{U}_1), \mathbf{v} \in \mathcal{R}(\mathbf{G}^{(0)})} \frac{|\mathbf{u}^T \mathbf{v}|}{(\|\mathbf{u}\|_2 \|\mathbf{v}\|_2)} = \sigma_{\min}(\mathbf{U}_1^T \mathbf{G}^{(0)}) > 0. \quad (12)$$

In addition, assume that each subproblem in (6) is solved to accuracy $\epsilon^{(r)}$ at iteration r , i.e., $\|\mathbf{Q}_i^{(r)} - \tilde{\mathbf{Q}}_i^{(r)}\|_2 \leq \epsilon^{(r)}$, where $\tilde{\mathbf{Q}}_i^{(r)} = (\mathbf{X}_i^T \mathbf{X}_i + \mu_i \mathbf{I})^{-1} \mathbf{X}_i^T \mathbf{G}^{(r-1)}$. Assume that $\epsilon^{(r)}$ is sufficiently small, i.e.,

$$\epsilon^{(r)} \leq \frac{\lambda_K - \lambda_{K+1}}{3 \sum_{i=1}^I \lambda_{\max}(\mathbf{X}_i)} \times \min \left\{ \sigma_{\min}(\mathbf{U}_2^T \mathbf{G}^{(r)}), \sigma_{\max}(\mathbf{U}_1^T \mathbf{G}^{(r)}) \right\}. \quad (13)$$

Then, $\text{dist}(\mathcal{R}(\mathbf{G}^{(r)}), \mathcal{R}(\mathbf{U}_1))$ approaches zero at a linear rate; i.e.,

$$\text{dist}(\mathcal{R}(\mathbf{G}^{(r)}), \mathcal{R}(\mathbf{U}_1)) \leq \left(\frac{2\lambda_{K+1} + \lambda_K}{2\lambda_K + \lambda_{K+1}} \right)^r \tan(\theta).$$

Theorem 2 ensures that if a T suffices for the \mathbf{Q} -subproblem to obtain a good enough approximation of the solution of Problem (6), the algorithm converges *linearly* to a *global optimal* solution – this means that we have gained scalability using AltMaxVar without losing optimality. Note that (13) means that the \mathbf{Q}_i -subproblem may require a higher solution accuracy when $\mathcal{R}(\mathbf{G}^{(r)})$ approaches $\mathcal{R}(\mathbf{U}_1)$, since $\sigma_{\min}(\mathbf{U}_2^T \mathbf{G}^{(r)})$ is close to zero under such circumstances. Nevertheless, since the result is based on worst-case analysis, the solution of the \mathbf{Q}_i -subproblem can be far rougher in practice – and one can still observe good convergence behavior of AltMaxVar. In fact, in our simulations, we observe that using $T = 1$ already gives very satisfactory results (as will be shown in the next section), which leads to computationally very cheap updates.

Remark 1 According to the proof of Theorem 2, when dealing with Problem (4) with $h_i(\cdot) = \mu_i/2 \|\cdot\|_F^2$, the procedure of AltMaxVar can be interpreted as a variant of the orthogonal iteration [30]. Based on this insight, many other approaches can be taken; e.g., the \mathbf{Q}_i -subproblem can be handled by conjugate gradient and the SVD step can be replaced by the QR decomposition – which may lead to computationally even cheaper updates. Nevertheless, our interest lies in solving (4) with a variety of regularizations under a unified framework, and the aforementioned alternatives cannot easily handle other regularizations.

V. NUMERICAL RESULTS

In this section, we use synthetic data and real experiments to showcase the effectiveness of the proposed algorithm. Throughout this section, the step size of the \mathbf{Q}_i -subproblem of AltMaxVar is set to be $\alpha_i = 0.99 \times 1/\lambda_{\max}(\mathbf{X}_i^T \mathbf{X}_i)$. When $h_i(\mathbf{Q}_i) = \mu_i/2 \|\mathbf{Q}_i\|_F^2$ is employed, we let $\gamma = 1$ following Theorem 2; otherwise, we let $\gamma = 0.9999$ – so that the convergence rate guarantee in Theorem 1 holds. All the experiments are coded in Matlab and conducted on a Linux server equipped with 32 1.2GHz cores and 128GB RAM.

A. Sanity Check: Small-Size Problems

We first use small-size problem instances to verify the convergence properties that were discussed in the last section.

1) *Classic MAX-VAR GCCA*: We generate the synthetic data in the following way: First, we let $\mathbf{Z} \in \mathbb{R}^{L \times N}$ be a common latent factor of different views, where the entries of \mathbf{Z} are drawn from the zero-mean i.i.d. Gaussian distribution and $L \geq N$. Then, a ‘mixing matrix’ $\mathbf{A}_i \in \mathbb{R}^{N \times M_i}$ is multiplied to \mathbf{Z} , resulting in $\mathbf{Y}_i = \mathbf{Z} \mathbf{A}_i$. We let $M_1 = \dots = M_I = M$ in this section. Finally, we add noise so that $\mathbf{X}_i = \mathbf{Y}_i + \sigma \mathbf{N}_i$. Here, \mathbf{A}_i and \mathbf{N}_i are generated in the same way as \mathbf{Z} . We first apply the algorithm with the regularization term $h_i(\cdot) = \mu_i/2 \|\cdot\|_F^2$ and let $\mu_i = 0.1$. Since L and M are small in this subsection, we employ the optimal solution that is based on eigen-decomposition as a baseline. The multiview latent semantic analysis (MVLSA) algorithm that was proposed in [10] is also employed as a baseline. In this section, we stop AltMaxVar when the absolute change of the objective value is smaller than 10^{-4} .




Fig. 2. Convergence curves of the algorithms.

In Fig. 2, we let $(L, M, N, I) = (500, 25, 20, 3)$. We set $\sigma = 0.1$ in this case, let $P = 8$ and $\gamma = 1$ for MVLSA and AltMaxVar, respectively, and ask for $K = 5$ canonical components. The results are averaged over 50 random trials, where \mathbf{Z} , $\{\mathbf{A}_i\}$, $\{\mathbf{N}_i\}$ are randomly generated in each trial. We test the proposed algorithm under different settings: We let $T = 1$, $T = 10$, and the gradient descent run until the inner loop converges (denoted as ‘solved’ in the figures). We also initialize the algorithm with random initializations (denoted as ‘randn’) and warm starts (denoted as ‘warm’) – i.e., using the solutions of MVLSA as starting points. Some observations from Fig. 2 are in order. First, the proposed algorithm using various T ’s including $T = 1$ and random initialization can reach the global optimum, which supports the analysis in Theorem 2. Second, by increasing T , the overall cost value decreases faster in terms of number of outer iterations – using $T = 10$ already gives very good speed of decreasing the cost value. Third, MVLSA cannot attain the global optimum, as expected. However, it provides good initialization: Using the warm start, the cost value comes close to the optimal value within 100 iterations in this case, even when $T = 1$ is employed. In fact, the combination of MVLSA-based initialization and using $T = 1$ offers the most computationally efficient way of implementing the proposed algorithm – especially for the large-scale case. In the remaining part of this section, we will employ MVLSA as the initialization of AltMaxVar and employ $T = 1$ for the \mathbf{Q} -subproblem.

2) *Feature-Selective MAX-VAR GCCA*: To test the proposed algorithm with non-smooth regularizers, we generate cases where outlying features are present in all views. Specifically, we let $\mathbf{X}_i = [\mathbf{Z} \mathbf{A}_i, \mathbf{O}_i] + \sigma \mathbf{N}_i$, where $\mathbf{O}_i \in \mathbb{R}^{L \times N_o}$ denotes the irrelevant outlying features and the elements of \mathbf{O}_i follow the i.i.d. zero-mean unit-variance Gaussian distribution. We wish to perform MAX-VAR GCCA of the views while discounting \mathbf{O}_i at the same time. To deal with outlying features, we employ the regularizer $g_i(\cdot) = \mu_i \|\cdot\|_{2,1}$ and implement the algorithm with $\mu_i = 0.5$ and $\mu_i = 1$, respectively. Under this setting, the optimal solution to Problem (4) is unknown. Therefore, we evaluate the performance by observing $\text{metric}_1 = 1/I \sum_{i=1}^I \|\mathbf{X}_i(:, \mathcal{S}_i^c) \hat{\mathbf{Q}}_i(\mathcal{S}_i^c, :) - \hat{\mathbf{G}}\|_F^2$ and $\text{metric}_2 = 1/I \sum_{i=1}^I \|\mathbf{X}_i(:, \mathcal{S}_i) \hat{\mathbf{Q}}_i(\mathcal{S}_i, :)\|_F^2$, where \mathcal{S}_i^c and

\mathcal{S}_i denote the index sets of “clean” and outlying features of view i , respectively – i.e., $\mathbf{X}_i(:, \mathcal{S}_i^c) = \mathbf{Z}_i \mathbf{A}_i$ and $\mathbf{X}_i(:, \mathcal{S}_i) = \mathbf{O}_i$ if noise is absent. metric_1 measures the performance of matching $\hat{\mathbf{G}}$ with the relevant part of the views, while metric_2 measures the performance of suppressing the irrelevant part. We wish that our algorithm yields low values of metric_1 and metric_2 simultaneously.

Table I presents the results of a small-size case which are averaged from 50 random trials, where $(L, M, N, I) = (150, 60, 60, 3)$ and $|\mathcal{S}| = \{61, \dots, 120\}$; i.e., 60 out of 120 features of $\mathbf{X}_i \in \mathbb{R}^{150 \times 120}$ are outlying features. The average power of the outlying features is set to be the same as that of the clean features, i.e., $\|\mathbf{Q}_i\|_F^2 / L |\mathcal{S}_i| = \|\mathbf{Z}_i \mathbf{A}_i\|_F^2 / LM$ so that the outlying features are not negligible. We ask for $K = 10$ canonical components. For MVLSA, we let the rank-truncation parameter to be $P = 50$. One can see that the eigen-decomposition based algorithm gives similar high values of both the evaluation metrics since it treats $\mathbf{X}_i(:, \mathcal{S}_i^c)$ and $\mathbf{X}_i(:, \mathcal{S}_i)$ equally. It is interesting to see that MVLSA suppresses the irrelevant features to some extent – although it does not explicitly consider outlying features, our understanding is that the PCA pre-processing on the views can somewhat suppress the outliers. Nevertheless, MVLSA does not fit the relevant part of the views well. The proposed algorithm gives the lowest values of both metrics. In particular, when $\mu_i = 1$ for all i , the irrelevant part is almost suppressed completely. Another observation is that using $\mu_i = 0.5$, the obtained score of metric_1 is slightly lower than that under $\mu_i = 1$, which makes sense since the algorithm pays more attention to feature selection using a larger μ . An illustrative example using a random trial can be seen in Fig. 3. From there, one can see that the proposed algorithm gives \mathbf{Q}_i ’s with almost zero rows over \mathcal{S} , thereby performing feature selection.

TABLE I
PERFORMANCE OF THE ALGORITHMS WHEN IRRELEVANT FEATURES ARE PRESENT. $(L, M, N) = (150, 60, 60)$; $|\mathcal{S}| = 60$; $\mathbf{X}_i \in \mathbb{R}^{150 \times 120}$; $\sigma = 1$.

Algorithm	metric_1	metric_2
eigen-decomp	9.547	9.547
MVLSA	15.506	1.456
proposed ($\mu = .5$)	0.486	9.689×10^{-3}
proposed ($\mu = 1$)	1.074	8.395×10^{-4}

B. Scalability Test: Large-Size Problems

1) *Original MAX-VAR GCCA*: We first test the case where no outlying features are involved and the regularizer $h_i(\cdot) = \mu_i/2 \|\cdot\|_F^2$ is employed. The views $\mathbf{X}_i = \mathbf{Z}_i \mathbf{A}_i + \sigma \mathbf{N}_i$ are generated following a similar way as in the last subsection, but \mathbf{Z} , \mathbf{A}_i and \mathbf{N}_i are sparse so that \mathbf{X}_i are sparse with a density level ρ_i that is defined as $\rho_i = \frac{\text{nnz}(\mathbf{X}_i)}{LM}$. In the simulations, we let $\rho = \rho_1 = \dots = \rho_I$. In the large-scale cases in this subsection, and the results are obtained via averaging 10 random trials.

In Fig. 4, we show the runtime performance of the algorithms for various sizes of the views, where density of the views is controlled so that $\rho \approx 10^{-3}$. The regularization

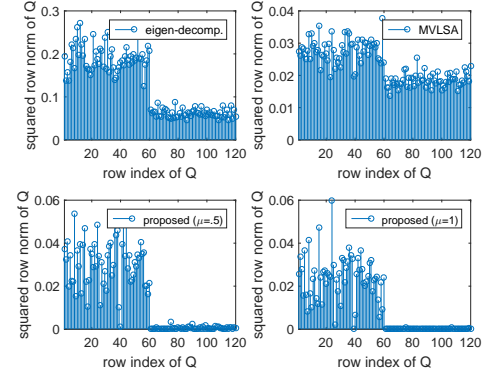


Fig. 3. Average row-norms of \mathbf{Q}_i (i.e., $(1/I) \sum_{i=1}^I \|\mathbf{Q}_i(m, :)\|_2^2$) for all m given by the algorithms.

parameter $\mu_i = 0.1$ is employed by all algorithms. We let $M = L \times 0.8$, $M = N$ and change M from 5,000 to 50,000. To run MVLSA, we truncate the ranks of views to $P = 100$, $P = 500$ and $P = 1,000$, respectively. We use MVLSA with $P = 100$ to initialize AltMaxVar and let $T = 1$ and $\gamma = 1$. We stop the proposed algorithm when the absolute change of the objective value is smaller than 10^{-4} . Ten random trials are used to obtain the results. One can see that the eigen-decomposition based algorithm does not scale well since the matrix $(\mathbf{X}_i^T \mathbf{X}_i + \mu_i \mathbf{I})^{-1}$ is dense. In particular, the algorithm exhausts the memory quota (32GB RAM) when $M = 30,000$. MVLSA with $P = 100$ and the proposed algorithm both scale very well from $M = 5,000$ to $M = 50,000$: When $M = 20,000$, brute-force eigen-decomposition takes almost 80 minutes, whereas MVLSA ($P = 100$) and AltMaxVar both use less than 2 minutes. Note that the runtime of the proposed algorithm already includes the runtime of the initialization time by MVLSA with $P = 100$, and thus the runtime curve of AltMaxVar is slightly higher than that of MVLSA ($P = 100$) in Fig. 4. Another observation is that, although MVLSA exhibits good runtime performance when using $P = 100$, its runtime under $P = 500$ and $P = 1,000$ is not very appealing. The corresponding cost values can be seen in Table II. The eigen-decomposition based method gives the lowest cost values when applicable, as it is an optimal solution. The proposed algorithm gives favorable cost values that are close to the optimal ones, even when only one iteration of the \mathbf{Q} -subproblem is implemented for every fixed $\mathbf{G}^{(r)}$ – this result supports our analysis in Theorem 2. Increasing P helps improve MVLSA. However, even when $P = 1,000$, the cost value given by MVLSA is still higher than that of AltMaxVar, and MVLSA using $P = 1,000$ is much slower than AltMaxVar.

2) *Feature-Selective MAX-VAR GCCA*: Table III presents the simulation results of a large-scale case in the presence of outlying features. Here, we fix $L = 100,000$ and $M = 80,000$ and change the density level ρ . We add $|\mathcal{S}| = 30,000$ outlying features to each view and every outlying feature is a random sparse vector whose non-zero elements follow the zero-mean i.i.d. unit-variance Gaussian distribution. We also scale the outlying features as before so that the energy of the clean and outlying features are comparable. The other settings follow




Fig. 4. Runtime of the algorithms for various problem sizes. $L = M/0.8$, $\rho \leq 10^{-3}$, $\sigma = 0.1$.

TABLE II

COST VALUES OF THE ALGORITHMS FOR DIFFERENT PROBLEM SIZES. $L = M/0.8$, $\rho = 10^{-3}$, $\sigma = 0.1$. † MEANS ‘OUT-OF-MEMORY’.

Algorithm	M					
	5,000	10,000	20,000	30,000	40,000	50,000
Global Opt	0.053	0.033	0.021	†	†	†
MVLSA ($P = 100$)	2.164	3.527	5.065	5.893	6.475	7.058
MVLSA ($P = 500$)	0.280	0.717	1.766	2.582	3.407	3.996
MVLSA ($P = 1,000$)	0.125	0.287	0.854	1.406	2.012	2.513
Proposed	0.092	0.061	0.049	0.043	0.038	0.039

those in the last simulation. One can see from Table III that the proposed algorithm with $\mu_i = 0.05$ gives the most balanced result – both evaluation metrics in with fairly low levels. Using $\mu_i = 0.5$ suppresses the $Q_i(\mathcal{S}, :)$ quite well, but using a larger μ_i also brings some sacrifice to the fitting metric. In terms of runtime, one can see that the proposed algorithm operates within the same order of magnitude of time as MVLSA needs. Note that the proposed algorithm works with the intact views of size $L \times M$, while MVLSA works with heavily truncated data. Therefore, such runtime performance of AltMaxVar is very satisfactory.

Similar results can be seen in Table IV, where we let $\rho = 10^{-4}$ and change I from 3 to 8. One can see that increasing the number of views does not increase the runtime of the proposed algorithm. The reason is that the updates of different Q_i ’s can be easily parallelized since the subproblems w.r.t. Q_i ’s are separable. One can implement the parallel computations using the `parfor` function of Matlab.

C. Real Large-Scale Word Embedding Tasks

We test the algorithms on a large-scale multilingual dataset. The views are extracted from a large word co-occurrence matrix, which is available at <https://sites.google.com/a/umn.edu/huang663/research>. The original data contains words of three languages, namely, English, Spanish, and French, and all the words are defined by the co-occurrences pointwise mutual information (PMI) with other words. We use the English words to form our first view, X_1 , which contains $L = 183,034$

TABLE III

EVALUATION OF THE ALGORITHM FOR DIFFERENT DATA DENSITIES IN THE PRESENCE OF OUTLYING FEATURES. $L = 100,000$, $M = 80,000$, $|S| = 30,000$, $\sigma = 1$, $I = 3$.

Algorithm	measure	ρ (density of views)			
		10^{-5}	5×10^{-4}	10^{-4}	10^{-3}
MVLSA ($P = 100$)	metric1	16.843	13.877	17.159	16.912
	metric2	0.003	0.010	0.009	0.003
	time (min)	0.913	1.019	1.252	3.983
Proposed ($\mu = .05$)	metric1	0.478	0.610	0.565	0.775
	metric2	0.018	0.134	0.034	0.003
	time (min)	3.798	5.425	5.765	24.182
Proposed ($\mu = .1$)	metric1	0.942	1.054	0.941	1.265
	metric2	0.006	0.054	0.004	0.000
	time (min)	2.182	3.791	4.510	16.378
Proposed ($\mu = .5$)	metric1	1.592	1.497	1.306	1.538
	metric2	0.003	0.021	0.000	0.000
	time (min)	1.735	2.714	3.723	13.447

TABLE IV

EVALUATION OF THE ALGORITHM VERSUS THE NUMBER OF VIEWS IN THE PRESENCE OF OUTLYING FEATURES. $L = 100,000$, $M = 80,000$, $|S| = 30,000$, $\sigma = 1$, $\rho = 5 \times 10^{-5}$.

Algorithm	measure	I (no. of views)					
		3	4	5	6	7	8
MVLSA ($P = 100$)	metric1	15.813	15.715	14.667	16.904	17.838	17.691
	metric2	0.008	0.009	0.009	0.009	0.007	0.009
	time (min)	1.087	0.975	0.960	0.958	0.989	1.026
proposed ($\mu = .05$)	metric1	0.731	0.590	0.670	0.611	0.517	0.628
	metric2	0.172	0.078	0.100	0.101	0.065	0.098
	time (min)	5.870	6.064	5.762	5.070	5.895	5.776
proposed ($\mu = .1$)	metric1	1.070	1.057	1.110	1.026	1.042	1.112
	metric2	0.055	0.019	0.018	0.024	0.023	0.023
	time (min)	3.240	2.974	3.313	3.210	3.083	3.529
proposed ($\mu = .5$)	metric1	1.461	1.482	1.578	1.443	1.472	1.561
	metric2	0.018	0.002	0.003	0.001	0.006	0.007
	time (min)	2.700	2.441	2.528	2.569	2.431	2.567

words and each word is defined by $M_i = 100,000$ features (co-occurrences). Note that X_1 is sparse – only 1.21% of its entries are non-zeros. Using a dictionary, we pick out the translations of the English words contained in X_1 in Spanish and French to form X_2 and X_3 , respectively. Note that many English words do not have a corresponding word in Spanish (or French). In such cases, we simply let $X_i(\ell, :) = \mathbf{0}$ for $i = 2$ (or $i = 3$), resulting in sparser X_2 and X_3 . Our objective is to use “side information” provided by Spanish and French to find a G whose rows are low-dimensional embeddings of the English words (cf. the motivating example in Fig. 1).

To evaluate the output, we use the evaluation tool provided at wordvectors.org [40], which runs several word embedding tasks to evaluate a set of given embeddings. Simply speaking, the tasks compare the algorithm-learned embeddings with the judgment of humans and yield high scores if the embeddings are consistent with the humans. The scores are between zero and one, and a score equal to one means a perfect alignment between the learned result and human judgment. We use the result of MVLSA with $P = 640$ as benchmark. The result of applying SVD to X_1 without considering different languages is also presented. We apply the proposed algorithm warm started by MVLSA and set $T = 1$. We run three versions of our algorithm. The first one uses $h_i(\cdot) = \frac{\mu_i}{2} \|\cdot\|_F^2$ with $\mu_i = 1$ for

$i = 1, 2, 3$. The second one uses $h_i(\cdot) = \mu_i \|\cdot\|_{2,1}$ for $i = 2, 3$ where $\mu_i = 0.05$, and we have no regularization on the first view. The third one is $h_i(\cdot) = \mu_i \|\cdot\|_{1,1}$ for $i = 2, 3$ where $\mu_i = 0.05$. The reason for adding ℓ_2/ℓ_1 mixed-norm (and ℓ_1 norm) regularization to the French and Spanish views is twofold: First, the ℓ_2/ℓ_1 norm (ℓ_1 norm) promotes row sparsity (sparsity) of \mathbf{Q}_i and thus performs feature selection on \mathbf{X}_2 and \mathbf{X}_3 – this physically means that we aim at selecting the most useful features from the other languages to help enhance English word embeddings. Second, \mathbf{X}_2 and \mathbf{X}_3 are effectively “fat matrices” and thus a column-selective regularizer can help improve the conditioning. Interestingly, we find that not adding feature-selective regularizations to the English view produces better results for the dataset considered. Our understanding is that \mathbf{X}_1 is a complete view without missing elements, and thus giving \mathbf{Q}_1 more “degrees of freedom” helps improve performance.

Tables V and VI show the word embedding results using $K = 50$ and $K = 100$, respectively. One can see that using the information from multiple views does help in improving the word embeddings: For $K = 50$ and $K = 100$, the multiview approaches perform better relative to SVD in 11 and 12 tasks out of 12 tasks. In addition, the proposed algorithm with the regularizer $h_i(\cdot) = \mu_i/2 \|\cdot\|_F^2$ (denoted by ℓ_2) gives similar or slightly better performance on average in both experiments compared to MVLSA. The proposed algorithm with the feature-selective regularizers (denoted by ℓ_2/ℓ_1 and ℓ_1 , resp.) gives the best evaluation results on both experiments – this suggests that for large-scale multilingual word embedding, feature selection is very meaningful. In particular, we observe that using $h_i(\mathbf{Q}_i) = \|\mathbf{Q}_i\|_{1,1}$ gives the best performance on many tasks. This further suggests that, in this case, different components of the reduced-dimension representations (i.e., columns of $\mathbf{X}_i \mathbf{Q}_i$) may be better learned by using different features of the views.

VI. CONCLUSION AND FUTURE WORK

In this work, we revisited the MAX-VAR GCCA problem with an eye towards scenarios involving large-scale and sparse data. The proposed approach is memory-efficient and has light per-iteration computational complexity if the views are sparse, and is thus suitable for dealing with big data. The algorithm is also flexible for incorporating different structure-promoting regularizers on the canonical components such as feature-selective regularizations. A thorough convergence analysis was presented, showing that the proposed algorithmic framework guarantees a KKT point to be obtained at a sublinear convergence rate in general cases under a variety of structure-promoting regularizers. We also showed that the algorithm approaches a global optimal solution at a linear convergence rate if the original MAX-VAR problem without regularization is considered. Simulations and real experiments with large-scale multi-lingual data showed that the performance of the proposed algorithm is promising in dealing with real-world large and sparse multiview data.

In the future, it is interesting to consider nonlinear operator-based multiview analysis, e.g., kernel (G)CCA [1] or deep neural network-based (G)CCA [41], under large-scale settings.

Nonlinear dimensionality reduction is very well-motivated in practice since it is able to handle more complex models and usually performs well with real-life data. On the other hand, the associated optimization problems are much harder, especially when the data dimension is large – which also promises a fertile research ground ahead. Another interesting direction is to consider constraints on \mathbf{G} (or $\mathbf{X}_i \mathbf{Q}_i$) – in some applications, structured (e.g., sparse and nonnegative) low-dimensional representations of data are desired.

APPENDIX A PROOF OF PROPOSITION 1

To simplify the notation, let us define $\mathbf{Q} = [\mathbf{Q}_1^T, \dots, \mathbf{Q}_I^T]^T$ as a collection of \mathbf{Q}_i ’s. We also rewrite the objective function in (4) as

$$\begin{aligned} F(\mathbf{Q}, \mathbf{G}) &= f(\mathbf{Q}, \mathbf{G}) + g(\mathbf{Q}) \\ &= \sum_{i=1}^I f_i(\mathbf{Q}_i, \mathbf{G}) + \sum_{i=1}^I g_i(\mathbf{Q}_i), \end{aligned}$$

where $f_i(\mathbf{Q}_i, \mathbf{G})$ and $g_i(\mathbf{Q}_i)$ are the smooth and non-smooth parts in (6) as before, and $f(\mathbf{Q}, \mathbf{G}) = \sum_{i=1}^I f_i(\mathbf{Q}_i, \mathbf{G})$ and $g(\mathbf{Q}) = \sum_{i=1}^I g_i(\mathbf{Q}_i)$, respectively. Additionally, let

$$\begin{aligned} \nabla_{\mathbf{Q}} f(\mathbf{Q}, \mathbf{G}) &= [(\nabla_{\mathbf{Q}_1} f(\mathbf{Q}, \mathbf{G}))^T, \dots, (\nabla_{\mathbf{Q}_I} f(\mathbf{Q}, \mathbf{G}))^T]^T, \\ \partial_{\mathbf{Q}} g(\mathbf{Q}) &= [(\partial_{\mathbf{Q}_1} g_1(\mathbf{Q}_1))^T, \dots, (\partial_{\mathbf{Q}_I} g_I(\mathbf{Q}_I))^T]^T. \end{aligned}$$

As the algorithm is essentially a two-block alternating optimization (since \mathbf{Q}_i for all i are updated simultaneously), the above notation suffices to describe the updates. Let

$$\begin{aligned} u_{\mathbf{Q}}(\mathbf{Q}; \hat{\mathbf{G}}, \hat{\mathbf{Q}}) &= f(\hat{\mathbf{G}}, \hat{\mathbf{Q}}) + \langle \nabla_{\mathbf{Q}} f(\hat{\mathbf{Q}}, \hat{\mathbf{G}}), \mathbf{Q} - \hat{\mathbf{Q}} \rangle \\ &\quad + \sum_{i=1}^I \frac{1}{2\alpha_i} \|\mathbf{Q}_i - \hat{\mathbf{Q}}_i\|_F^2 + \sum_{i=1}^I g_i(\mathbf{Q}_i); \end{aligned}$$

i.e., $u_{\mathbf{Q}}(\mathbf{Q}; \hat{\mathbf{G}}, \hat{\mathbf{Q}})$ is an approximation of $F(\mathbf{G}, \mathbf{Q})$ locally at the point $(\hat{\mathbf{G}}, \hat{\mathbf{Q}})$. We further define $\tilde{u}_{\mathbf{Q}}(\mathbf{Q}; \hat{\mathbf{G}}, \hat{\mathbf{Q}}) = u_{\mathbf{Q}}(\mathbf{Q}; \hat{\mathbf{G}}, \hat{\mathbf{Q}}) - \sum_{i=1}^I g_i(\mathbf{Q}_i)$; i.e., $\tilde{u}_{\mathbf{Q}}(\mathbf{Q}; \hat{\mathbf{G}}, \hat{\mathbf{Q}})$ is an approximation of the continuously differentiable part $f(\mathbf{G}, \mathbf{Q})$ locally at the point $(\hat{\mathbf{G}}, \hat{\mathbf{Q}})$. One can see that,

$$\nabla_{\mathbf{Q}} f(\hat{\mathbf{Q}}, \hat{\mathbf{G}}) = \nabla_{\mathbf{Q}} \tilde{u}(\hat{\mathbf{Q}}; \hat{\mathbf{G}}, \hat{\mathbf{Q}}), \quad (14)$$

Since $\nabla_{\mathbf{Q}_i} f_i(\mathbf{Q}_i, \mathbf{G})$ is L_i -Lipschitz continuous w.r.t. \mathbf{Q}_i and $\alpha_i \leq 1/L_i$ for all i , we have the following holds:

$$u_{\mathbf{Q}}(\mathbf{Q}; \hat{\mathbf{G}}, \hat{\mathbf{Q}}) \geq F(\mathbf{Q}, \hat{\mathbf{G}}), \quad \forall \mathbf{Q}, \quad (15)$$

where the equality holds if and only if $\mathbf{Q}_i = \hat{\mathbf{Q}}_i$ for all i , i.e.,

$$u_{\mathbf{Q}}(\hat{\mathbf{Q}}; \hat{\mathbf{G}}, \hat{\mathbf{Q}}) = F(\hat{\mathbf{Q}}, \hat{\mathbf{G}}). \quad (16)$$

Similarly, we define

$$\begin{aligned} u_{\mathbf{G}}(\mathbf{G}; \hat{\mathbf{G}}, \hat{\mathbf{Q}}) &= \sum_{i=1}^I \frac{1}{2} \|\mathbf{X}_i \mathbf{Q}_i^{(r+1)} - \mathbf{G}\|_F^2 \\ &\quad + \omega \cdot \|\mathbf{G} - \mathbf{G}^{(r)}\|_F^2 + \sum_{i=1}^I g_i(\mathbf{Q}_i), \end{aligned}$$

TABLE V
EVALUATION ON 12 WORD EMBEDDING TASKS; $K = 50$.

Task	Algorithm ($K = 50$)				
	SVD	MVLSA	AltMaxVar (ℓ_2)	AltMaxVar (ℓ_2/ℓ_1)	AltMaxVar (ℓ_1)
EN-WS-353-SIM	0.63	0.69	0.67	0.68	0.69
EN-MC-30	0.56	0.63	0.63	0.64	0.66
EN-MTurk-771	0.54	0.58	0.59	0.60	0.59
EN-MEN-TR-3k	0.67	0.66	0.67	0.68	0.69
EN-RG-65	0.51	0.53	0.55	0.58	0.58
EN-MTurk-287	0.65	0.64	0.65	0.64	0.63
EN-WS-353-REL	0.50	0.51	0.53	0.55	0.57
EN-VERB-143	0.21	0.22	0.21	0.21	0.20
EN-YP-130	0.36	0.39	0.38	0.41	0.41
EN-SIMLEX-999	0.31	0.42	0.41	0.39	0.36
EN-RW-STANFORD	0.39	0.43	0.43	0.43	0.43
EN-WS-353-ALL	0.56	0.59	0.59	0.60	0.62
Average	0.49	0.52	0.53	0.54	0.54
Median	0.53	0.56	0.57	0.59	0.59

TABLE VI
EVALUATION ON 12 WORD EMBEDDING TASKS; $K = 100$.

Task	Algorithm ($K = 100$)				
	SVD	MVLSA	AltMaxVar (ℓ_2)	AltMaxVar (ℓ_2/ℓ_1)	AltMaxVar (ℓ_1)
EN-WS-353-SIM	0.68	0.72	0.71	0.72	0.72
EN-MC-30	0.73	0.68	0.72	0.74	0.82
EN-MTurk-771	0.59	0.60	0.60	0.61	0.62
EN-MEN-TR-3k	0.72	0.70	0.70	0.71	0.73
EN-RG-65	0.68	0.63	0.64	0.68	0.70
EN-MTurk-287	0.61	0.66	0.65	0.64	0.64
EN-WS-353-REL	0.57	0.54	0.55	0.56	0.59
EN-VERB-143	0.19	0.28	0.27	0.29	0.28
EN-YP-130	0.42	0.41	0.41	0.45	0.49
EN-SIMLEX-999	0.34	0.42	0.41	0.41	0.39
EN-RW-STANFORD	0.44	0.46	0.45	0.46	0.48
EN-WS-353-ALL	0.62	0.62	0.62	0.62	0.65
Average	0.55	0.56	0.56	0.58	0.59
Median	0.60	0.61	0.61	0.62	0.63

where we recall that $\omega = (1-\gamma)I/2\gamma$ and the last term is a constant if \mathbf{Q} is fixed.

The update rule of \mathbf{G} in Algorithm 1 can be re-expressed as $\mathbf{G} \in \arg \min_{\mathbf{G}^T \mathbf{G} = \mathbf{I}} u_G(\mathbf{G}; \hat{\mathbf{G}}, \hat{\mathbf{Q}})$. It is easily seen that

$$u_G(\mathbf{G}; \hat{\mathbf{G}}, \hat{\mathbf{Q}}) \geq F(\hat{\mathbf{Q}}, \mathbf{G}), \quad (17a)$$

$$u_G(\hat{\mathbf{G}}; \hat{\mathbf{G}}, \hat{\mathbf{Q}}) = F(\hat{\mathbf{Q}}, \hat{\mathbf{G}}). \quad (17b)$$

Hence, Algorithm 1 boils down to

$$\mathbf{Q}_i^{(r,t+1)} = \arg \min_{\mathbf{Q}_i} u_Q(\mathbf{Q}; \mathbf{G}^{(r)}, \mathbf{Q}^{(r,t)}), \quad \forall t \quad (18a)$$

$$\mathbf{G}^{(r+1)} \in \arg \min_{\mathbf{G}^T \mathbf{G} = \mathbf{I}} u_G(\mathbf{G}; \mathbf{G}^{(r)}, \mathbf{Q}^{(r+1)}). \quad (18b)$$

When $\gamma = 1$, (18b) amounts to SVD of $\sum_{i=1}^I \mathbf{X}_i \mathbf{Q}_i / I$ and the \mathbf{G} -subproblem $\min_{\mathbf{G}^T \mathbf{G} = \mathbf{I}} F(\mathbf{Q}^{(r+1)}, \mathbf{G})$ is optimally solved; otherwise, both (18a) and (18b) are local upper bound minimizations.

Note that the following holds:

$$F(\mathbf{Q}^{(r)}, \mathbf{G}^{(r)}) = u_Q(\mathbf{Q}^{(r)}; \mathbf{G}^{(r)}, \mathbf{Q}^{(r)}) \quad (19a)$$

$$\geq u_Q(\mathbf{Q}^{(r+1)}; \mathbf{G}^{(r)}, \mathbf{Q}^{(r,T-1)}) \quad (19b)$$

$$\geq F(\mathbf{Q}^{(r+1)}, \mathbf{G}^{(r)}) \quad (19c)$$

$$= u_G(\mathbf{G}^{(r)}; \mathbf{G}^{(r)}, \mathbf{Q}^{(r+1)}) \quad (19d)$$

$$\geq u_G(\mathbf{G}^{(r+1)}; \mathbf{G}^{(r)}, \mathbf{Q}^{(r+1)}) \quad (19e)$$

$$\geq F(\mathbf{Q}^{(r+1)}, \mathbf{G}^{(r+1)}), \quad (19f)$$

where (19a) holds because of (16), (19b) holds since PG is a descending method when $\alpha_i \leq 1/L_i$ [42], (19c) holds by the property in (16), (19d) holds due to (17), (19e) is due to the fact that (18b) is optimally solved, and (19f) holds also because of the first equation in (17).

Next, we show that every limit point is a KKT point. Assume that there exists a convergent subsequence of $\{\mathbf{G}^{(r)}, \mathbf{Q}^{(r)}\}_{r=0,1,\dots}$, whose limit point is $(\mathbf{G}^*, \mathbf{Q}^*)$ and the subsequence is indexed by $\{r_j\}_{j=1,\dots,\infty}$. We have the follow-

ing chain of inequalities:

$$u_Q(Q; G^{(r_j)}, Q^{(r_j)}) \geq u_Q(Q^{(r_j,1)}; G^{(r_j)}, Q^{(r_j)}) \quad (20a)$$

$$\geq u_Q(Q^{(r_j,T)}; G^{(r_j)}, Q^{(r_j,T-1)}) \quad (20b)$$

$$\geq F(G^{(r_j)}, Q^{(r_j+1)}) \quad (20c)$$

$$\geq F(Q^{(r_j+1)}, G^{(r_j+1)}) \quad (20d)$$

$$\geq F(Q^{(r_j+1)}, G^{(r_j+1)}) \quad (20e)$$

$$= u_Q(Q^{(r_j+1)}; G^{(r_j+1)}, Q^{(r_j+1)}), \quad (20f)$$

where (20a) holds because of the update rule in (18a), (20b) holds, again, by the descending property of PG, (20d) follows (19f), and (20f) is again because of the way that we construct $u_Q(Q; G^{(r_j+1)}, Q^{(r_j+1)})$. Taking $j \rightarrow \infty$, and by continuity of $u_Q(\cdot)$, we have

$$u_Q(Q; G^*, Q^*) \geq u_Q(Q^*; G^*, Q^*), \quad (21)$$

i.e., Q^* is a minimum of $u_Q(Q; G^*, Q^*)$. Consequently, Q^* satisfies the conditional KKT conditions, i.e., $0 \in \nabla_Q \tilde{u}_Q(Q^*; G^*, Q^*) + \partial_Q g(Q^*)$, which, by (14), also means

$$0 \in \nabla_{Q_i} f_i(Q_i^*, G^*) + \partial_{Q_i} g_i(Q_i^*), \quad \forall i. \quad (22)$$

We now show that $Q^{(r_j,t)}$ for $t = 1, \dots, T$ also converges to Q^* . Indeed, we have

$$\begin{aligned} u_Q(Q^{(r_{j+1})}; G^{(r_{j+1})}, Q^{(r_{j+1})}) &\leq u_Q(Q^{(r_j,1)}; G^{(r_j)}, Q^{(r_j)}) \\ &\leq u_Q(Q^{(r_j)}; G^{(r_j)}, Q^{(r_j)}), \end{aligned}$$

where the first inequality was derived from (20). Taking $j \rightarrow \infty$, we see that $u_Q(Q^*; G^*, Q^*) \leq u_Q(Q^{(r_j,1)}; G^*, Q^*) \leq u_Q(Q^*; G^*, Q^*)$, which implies that $u_Q(Q^{(r_j,1)}; G^*, Q^*) = u_Q(Q^*; G^*, Q^*) \leq u_Q(Q; G^*, Q^*)$. On the other hand, the problem in (18a) has a unique minimizer when $g_i(\cdot)$ is a convex closed function [34], which means that $Q^{(r_j,1)} \rightarrow Q^*$. By the same argument, we can show that $Q^{(r_j,t)}$ for $t = 1, \dots, T$ also converges to Q^* . Consequently, we have $Q^{(r_j,T)} = Q^{(r_j+1)} \rightarrow Q^*$. We repeat the proof in (20) to G :

$$\begin{aligned} u_G(G; G^{(r_j)}, Q^{(r_j+1)}) &\geq u_G(G^{(r_j+1)}; G^{(r_j)}, Q^{(r_j+1)}) \\ &\geq F(Q^{(r_j+1)}, G^{(r_j+1)}) \\ &\geq F(Q^{(r_j+1)}, G^{(r_j+1)}) \\ &= u_G(G^{(r_j+1)}; G^{(r_j+1)}, Q^{(r_j+1)}), \end{aligned}$$

Taking $j \rightarrow \infty$ and by $Q^{(r_j+1)} \rightarrow Q^*$, we have

$$u_G(G; G^*, Q^*) \geq u_G(G^*; G^*, Q^*), \quad \forall G^T G = I.$$

The above means that G^* satisfies the partial conditional KKT conditions w.r.t. G . Combining with (22), we see that (G^*, Q^*) is a KKT point of the original problem.

Now, we show the b) part. First, we show that Q_i remains in a bounded set (the variable G is always bounded since we keep it feasible in each iteration). Since the objective value is non-increasing (cf. Proposition 1), if we denote the initial

objective value as V , then $F(G^{(r)}, Q^{(r)}) \leq V$ holds in all subsequent iterations. Note that when $X_i^{(0)}$ and $Q_i^{(0)}$ are bounded, V is also finite. In particular, we have $\|X_i Q_i - G\|_F^2 + 2 \sum_{i=1}^I g_i(Q_i) \leq 2V$ holds, which implies $\|X_i Q_i\|_F \leq \|G\|_F + \sqrt{2V}$ by the triangle inequality. The right-hand side is finite since both terms are bounded. Denote $(\|G\|_F + \sqrt{2V})$ by V' . Then, we have $\|Q_i\|_F = \|(X_i^T X_i)^{-1} X_i^T X_i Q_i\|_F \leq \|(X_i^T X_i)^{-1} X_i^T\|_F \cdot \|X_i Q_i\|_F \leq V' \cdot \|(X_i^T X_i)^{-1} X_i^T\|_F$. Now, by the assumption that $\text{rank}(X_i) = M_i$, the term $\|(X_i^T X_i)^{-1} X_i^T\|_F$ is bounded. This shows that $\|Q_i\|_F$ is bounded. Hence, starting from a bounded $Q_i^{(0)}$, the solution sequence $\{Q^{(r)}, G^{(r)}\}$ remains in a bounded set. Since the constraints of Q_i , i.e., $\mathbb{R}^{M_i \times K}$ and G are also closed sets, $\{Q^{(r)}, G^{(r)}\}$ remains in a compact set.

Now, let us denote \mathcal{K} as the set containing all the KKT points. Suppose the whole sequence does not converge to \mathcal{K} . Then, there exists a convergent subsequence indexed by $\{r_j\}$ such that $\lim_{j \rightarrow \infty} d^{(r)}(\mathcal{K}) \geq \gamma$ for some positive γ , where $d^{(r)}(\mathcal{K}) = \min_{Y \in \mathcal{K}} \|(G^{(r)}, Q^{(r)}) - Y\|$. Since the subsequence indexed by $\{r_j\}$ lies in a closed and bounded set as we have shown, this subsequence has a limit point. However, as we have shown in Theorem 1, every limit point of the solution sequence is a KKT point. This is a contradiction. Therefore, the whole sequence converges to a KKT point.

APPENDIX B PROOF OF LEMMA 1

First, we have the update rule $Q_i^{(r,t+1)} = Q_i^{(r,t)} - \alpha_i \tilde{\nabla}_{Q_i} F(Q_i^{(r,t)}, G^{(r)})$, which leads to the following:

$$\frac{1}{\alpha_i} (Q_i^{(r,t+1)} - Q_i^{(r,t)}) = -\tilde{\nabla}_{Q_i} F(Q_i^{(r,t)}, G^{(r)}). \quad (24)$$

Meanwhile, the updating rule can also be expressed as

$$\begin{aligned} Q_i^{(r,t+1)} &= \arg \min_{Q_i} \left\langle \nabla_{Q_i} f(Q_i^{(r,t)}, G^{(r)}), Q_i - Q_i^{(r,t)} \right\rangle \\ &\quad + g_i(Q_i) + \frac{1}{2\alpha_i} \|Q_i - Q_i^{(r,t)}\|_F^2. \end{aligned} \quad (25)$$

Therefore, there exists a $\partial_{Q_i} g_i(Q_i^{(r,t+1)})$ and a $Q^{(r,t+1)}$ satisfy the following optimality conditions:

$$0 = \nabla_{Q_i} f_i(Q_i^{(r,t)}, G^{(r)}) + \partial_{Q_i} g_i(Q_i^{(r,t+1)}) + \frac{1}{\alpha_i} (Q_i^{(r,t+1)} - Q_i^{(r,t)}).$$

Consequently, we see that

$$\begin{aligned} \sum_{i=1}^I \sum_{t=0}^T \left\| \tilde{\nabla}_{Q_i} F(Q_i^{(r,t)}, G^{(r)}) \right\|_F^2 &\rightarrow 0 \\ \Rightarrow Q_i^{(r,t)} - Q_i^{(r,t+1)} &\rightarrow 0, \quad \forall t = 0, \dots, T-1 \\ \Rightarrow Q_i^{(r)} - Q_i^{(r+1)} &\rightarrow 0, \quad \forall i \\ \Rightarrow \nabla_Q f(Q^{(r)}, G^{(r)}) + \partial_Q g(Q^{(r)}) &\rightarrow 0 \end{aligned}$$

which holds since T is finite. The above means that $0 \in \nabla_Q f(Q^{(r)}, G^{(r)}) + \partial_Q g(Q^{(r)})$ is satisfied when $Z^{(r+1)} \rightarrow 0$.

Recall that $\mathbf{G}^{(r+1)}$ satisfies the optimality condition of Problem (10). Therefore, there exists a $\mathbf{\Lambda}^{(r+1)}$ such that the following optimality condition holds

$$\mathbf{G}^{(r)} - \sum_{i=1}^I \mathbf{X}_i \mathbf{Q}_i^{(r+1)} / I + \frac{1}{\gamma} (\mathbf{G}^{(r+1)} - \mathbf{G}^{(r)}) + \mathbf{G}^{(r+1)} \mathbf{\Lambda}^{(r+1)} = \mathbf{0} \quad (26)$$

Combining (26) and (24), we have

$$Z^{(r+1)} = \frac{1}{\gamma^2} \left\| \mathbf{G}^{(r+1)} - \mathbf{G}^{(r)} \right\|_F^2 + \sum_{i=1}^I \frac{1}{\alpha_i^2} \left\| \mathbf{Q}_i^{(r+1)} - \mathbf{Q}_i^{(r)} \right\|_F^2.$$

We see that $Z^{(r+1)} \rightarrow 0$ implies that a KKT point is reached and this completes the proof of Lemma 1.

APPENDIX C PROOF OF THEOREM 1

We show that every iterate of \mathbf{Q} and \mathbf{G} gives sufficient decreases of the overall objective function. Since $\nabla_{\mathbf{Q}_i} f_i(\mathbf{Q}_i, \mathbf{G})$ is L_i -Lipschitz continuous for all i , we have the following:

$$\begin{aligned} F(\mathbf{Q}^{(r,t+1)}, \mathbf{G}^{(r)}) &\leq u_{\mathbf{Q}}(\mathbf{Q}^{(r,t+1)}; \mathbf{G}^{(r)}, \mathbf{Q}^{(r,t)}) \\ &= f(\mathbf{Q}^{(r,t)}, \mathbf{G}^{(r)}) + \left\langle \nabla_{\mathbf{Q}} f(\mathbf{Q}^{(r,t)}, \mathbf{G}^{(r)}), \mathbf{Q}^{(r,t+1)} - \mathbf{Q}^{(r,t)} \right\rangle \\ &\quad + \sum_{i=1}^I g_i(\mathbf{Q}_i^{(r,t+1)}) + \sum_{i=1}^I \frac{L_i}{2} \left\| \mathbf{Q}_i^{(r,t+1)} - \mathbf{Q}_i^{(r,t)} \right\|_F^2. \end{aligned} \quad (27)$$

Since $\mathbf{Q}^{(r,t+1)}$ is a minimizer of Problem (25), we also have

$$\begin{aligned} &\left\langle \nabla_{\mathbf{Q}} f(\mathbf{Q}^{(r,t)}, \mathbf{G}^{(r)}), \mathbf{Q}^{(r,t+1)} - \mathbf{Q}^{(r,t)} \right\rangle + \sum_{i=1}^I g_i(\mathbf{Q}_i^{(r,t+1)}) \\ &\quad + \sum_{i=1}^I \frac{1}{2\alpha_i} \left\| \mathbf{Q}_i^{(r,t+1)} - \mathbf{Q}_i^{(r,t)} \right\|_F^2 \leq \sum_{i=1}^I g_i(\mathbf{Q}_i^{(r,t)}), \end{aligned} \quad (28)$$

which is obtained by letting $\mathbf{Q}_i = \mathbf{Q}_i^{(r,t)}$. Combining (27) and (28), we have

$$\begin{aligned} &F(\mathbf{Q}^{(r,t+1)}, \mathbf{G}^{(r)}) - F(\mathbf{Q}^{(r,t)}, \mathbf{G}^{(r)}) \\ &\leq - \sum_{i=1}^I \left(\frac{1}{2\alpha_i} - \frac{L_i}{2} \right) \left\| \mathbf{Q}_i^{(r,t+1)} - \mathbf{Q}_i^{(r,t)} \right\|_F^2. \end{aligned} \quad (29)$$

Summing up the above over $t = 0, \dots, T-1$, we have

$$\begin{aligned} &F(\mathbf{Q}^{(r)}, \mathbf{G}^{(r)}) - F(\mathbf{Q}^{(r+1)}, \mathbf{G}^{(r)}) \\ &\geq \sum_{t=0}^{T-1} \sum_{i=1}^I \left(\frac{1}{2\alpha_i} - \frac{L_i}{2} \right) \left\| \mathbf{Q}_i^{(r,t+1)} - \mathbf{Q}_i^{(r,t)} \right\|_F^2. \end{aligned} \quad (30)$$

For the \mathbf{G} -subproblem, we have

$$\begin{aligned} u_{\mathbf{G}}(\mathbf{G}^{(r+1)}; \mathbf{G}^{(r)}, \mathbf{Q}^{(r+1)}) &\leq u_{\mathbf{G}}(\mathbf{G}^{(r)}; \mathbf{G}^{(r)}, \mathbf{Q}^{(r+1)}) \\ &= F(\mathbf{G}^{(r)}, \mathbf{Q}^{(r+1)}) \end{aligned}$$

and thus

$$F(\mathbf{G}^{(r+1)}, \mathbf{Q}^{(r+1)}) + \omega \left\| \mathbf{G}^{(r+1)} - \mathbf{G}^{(r)} \right\|_F^2 \leq F(\mathbf{Q}^{(r+1)}, \mathbf{G}^{(r)}),$$

or, equivalently

$$\begin{aligned} &F(\mathbf{Q}^{(r+1)}, \mathbf{G}^{(r+1)}) - F(\mathbf{Q}^{(r+1)}, \mathbf{G}^{(r)}) \\ &\leq -\omega \left\| \mathbf{G}^{(r+1)} - \mathbf{G}^{(r)} \right\|_F^2, \quad \forall \mathbf{G}^T \mathbf{G} = \mathbf{I}, \end{aligned} \quad (31)$$

where $\omega = \left(\frac{1}{2\gamma} - \frac{I}{2} \right) > 0$ if $\gamma < 1$. Combining (30) and (31), we have

$$\begin{aligned} &F(\mathbf{Q}^{(r)}, \mathbf{G}^{(r)}) - F(\mathbf{Q}^{(r+1)}, \mathbf{G}^{(r+1)}) \\ &\geq \left(\frac{1}{2\gamma} - \frac{I}{2} \right) \left\| \mathbf{G}^{(r+1)} - \mathbf{G}^{(r)} \right\|_F^2 \\ &\quad + \sum_{t=0}^{T-1} \sum_{i=1}^I \left(\frac{1}{2\alpha_i} - \frac{L_i}{2} \right) \left\| \mathbf{Q}_i^{(r,t+1)} - \mathbf{Q}_i^{(r,t)} \right\|_F^2. \end{aligned} \quad (32)$$

Summing up $F(\mathbf{Q}^{(r)}, \mathbf{G}^{(r)})$ over $r = 0, 1, \dots, J-1$, we have the following:

$$\begin{aligned} &F(\mathbf{Q}^{(r)}, \mathbf{G}^{(r)}) - F(\mathbf{Q}^{(r+1)}, \mathbf{G}^{(r+1)}) \\ &\geq \sum_{r=0}^{J-1} \omega \left\| \mathbf{G}^{(r+1)} - \mathbf{G}^{(r)} \right\|_F^2 \\ &\quad + \sum_{r=0}^{J-1} \sum_{t=0}^{T-1} \sum_{i=1}^I \left(\frac{1}{2\alpha_i} - \frac{L_i}{2} \right) \left\| \mathbf{Q}_i^{(r,t+1)} - \mathbf{Q}_i^{(r,t)} \right\|_F^2 \\ &= \sum_{r=0}^{J-1} \omega \gamma^2 \left\| \mathbf{G}^{(r)} - \frac{\sum_{i=1}^I \mathbf{X}_i \mathbf{Q}_i^{(r+1)}}{I} + \mathbf{G}^{(r+1)} \mathbf{\Lambda}^{(r+1)} \right\|_F^2 \\ &\quad + \sum_{r=0}^{J-1} \sum_{i=1}^I \sum_{t=0}^{T-1} \left(\frac{1}{2\alpha_i} - \frac{L_i}{2} \right) \alpha_i^2 \left\| \tilde{\nabla}_{\mathbf{Q}_i} F(\mathbf{Q}^{(r,t)}, \mathbf{G}^{(r)}) \right\|_F^2 \\ &\geq \sum_{r=0}^{J-1} c Z^{(r+1)}, \end{aligned} \quad (33)$$

where $c = \min\{\omega \gamma^2, \{(\frac{1}{2\alpha_i} - \frac{L_i}{2}) \alpha_i^2\}_{i=1, \dots, I}\}$. By the definition of J , we have

$$\begin{aligned} &\frac{F(\mathbf{Q}^{(0)}, \mathbf{G}^{(0)}) - F(\mathbf{Q}^{(J)}, \mathbf{G}^{(J)})}{J-1} \geq \frac{\sum_{r=0}^{J-1} c Z^{(r+1)}}{J-1} \geq c \cdot \delta \\ &\Rightarrow \delta \leq \frac{1}{c} \frac{F(\mathbf{Q}^{(0)}, \mathbf{G}^{(0)}) - \bar{F}}{J-1} \Rightarrow \delta \leq \frac{v}{J-1}, \end{aligned}$$

where \bar{F} is the lower bound of the cost function and $v = (F(\mathbf{Q}^{(0)}, \mathbf{G}^{(0)}) - \bar{F})/c$. This completes the proof.

APPENDIX D PROOF OF THEOREM 2

First consider an easier case where $\epsilon^{(r)} = 0$ for all r . Then, we have $\mathbf{Q}_i^{(r+1)} = (\mathbf{X}_i^T \mathbf{X}_i + \mu_i \mathbf{I})^{-1} \mathbf{X}_i^T \mathbf{G}^{(r)}$. Therefore, the update w.r.t. \mathbf{G} is simply to apply SVD on $\sum_{i=1}^I \mathbf{X}_i \mathbf{Q}_i / I = \mathbf{M} \mathbf{G}^{(r)} / I$. In other words, there exists an invertible $\mathbf{\Theta}^{(r+1)}$ such that

$$\mathbf{G}^{(r+1)} \mathbf{\Theta}^{(r+1)} = \mathbf{M} \mathbf{G}^{(r)}, \quad (34)$$

where $\mathbf{M} = \sum_{i=1}^I \mathbf{X}_i \mathbf{X}_i^\dagger$ as before, since the SVD procedure is nothing but a change of bases. The update rule in (34), is essentially the *orthogonal iteration* algorithm in [30]. Invoking [30, Theorem 8.2.2], one can show that $\|U_2^T \mathbf{G}^{(r)}\|_2$ approaches zero linearly.

The proof of the case where $\epsilon^{(r)} > 0$ can be considered as an extension of round-off error analysis of orthogonal iterations, and can be shown following the insights of [17] and [43] with proper modifications to accommodate the MAX-VAR GCCA case. At the r th iteration, ideally, we have $\tilde{\mathbf{Q}}_i^{(r+1)} = (\mathbf{X}_i^T \mathbf{X}_i + \mu_i \mathbf{I})^{-1} \mathbf{X}_i^T \mathbf{G}^{(r)}$ if the \mathbf{Q} -subproblem is solved to optimality. In practice, what we have is an inexact solution, i.e.,

$$\mathbf{Q}_i^{(r+1)} = (\mathbf{X}_i^T \mathbf{X}_i + \mu_i \mathbf{I})^{-1} \mathbf{X}_i^T \mathbf{G}^{(r)} + \mathbf{W}_i^{(r)},$$

where we have assumed that the largest singular value of $\mathbf{W}_i^{(r)}$ is bounded by ϵ , i.e., $\|\mathbf{W}_i^{(r)}\|_2 \leq \epsilon$. Hence, one can see that

$$\sum_{i=1}^I \mathbf{X}_i \mathbf{Q}_i^{(r+1)} = \mathbf{M} \mathbf{G}^{(r)} + \sum_{i=1}^I \mathbf{X}_i \mathbf{W}_i^{(r)}.$$

Therefore, following the same reason of obtaining (34), we have

$$\mathbf{G}^{(r+1)} \boldsymbol{\Theta}^{(r+1)} = \left(\mathbf{M} \mathbf{G}^{(r)} + \sum_{i=1}^I \mathbf{X}_i \mathbf{W}_i^{(r)} \right),$$

where $\boldsymbol{\Theta}^{(r+1)} \in \mathbb{R}^{K \times K}$ is a full-rank matrix since the solution via SVD is a change of bases. Consequently, we have

$$\begin{bmatrix} \mathbf{U}_1^T \mathbf{G}^{(r+1)} \\ \mathbf{U}_2^T \mathbf{G}^{(r+1)} \end{bmatrix} \boldsymbol{\Theta}^{(r+1)} = \begin{bmatrix} \boldsymbol{\Lambda}_1 \mathbf{U}_1^T \mathbf{G}^{(r)} + \mathbf{U}_1^T \sum_{i=1}^I \mathbf{X}_i \mathbf{W}_i^{(r)} \\ \boldsymbol{\Lambda}_2 \mathbf{U}_2^T \mathbf{G}^{(r)} + \mathbf{U}_2^T \sum_{i=1}^I \mathbf{X}_i \mathbf{W}_i^{(r)} \end{bmatrix}.$$

Now, we denote

$$\Delta_1^{(r)} = \mathbf{U}_1^T \sum_{i=1}^I \mathbf{X}_i \mathbf{W}_i^{(r)}, \quad \Delta_2^{(r)} = \mathbf{U}_2^T \sum_{i=1}^I \mathbf{X}_i \mathbf{W}_i^{(r)},$$

as two error terms at the r th iteration. Next, let us consider the following chain of inequalities:

$$\begin{aligned} & \left\| \mathbf{U}_2^T \mathbf{G}^{(r+1)} \left(\mathbf{U}_1^T \mathbf{G}^{(r+1)} \right)^{-1} \right\|_2 \\ &= \left\| \left(\boldsymbol{\Lambda}_2 \mathbf{U}_2^T \mathbf{G}^{(r)} + \Delta_2^{(r)} \right) \left(\boldsymbol{\Lambda}_1 \mathbf{U}_1^T \mathbf{G}^{(r)} + \Delta_1^{(r)} \right)^{-1} \right\|_2 \\ &\leq \frac{\left\| \left(\boldsymbol{\Lambda}_2 \mathbf{U}_2^T \mathbf{G}^{(r)} + \Delta_2^{(r)} \right) \left(\mathbf{U}_1^T \mathbf{G}^{(r)} \right)^{-1} \right\|_2}{\sigma_K \left(\boldsymbol{\Lambda}_1 + \Delta_1^{(r)} \left(\mathbf{U}_1^T \mathbf{G}^{(r)} \right)^{-1} \right)} \\ &\leq \frac{\lambda_{K+1} \left\| \mathbf{U}_2^T \mathbf{G}^{(r)} \left(\mathbf{U}_1^T \mathbf{G}^{(r)} \right)^{-1} \right\|_2 + \left\| \Delta_2^{(r)} \left(\mathbf{U}_1^T \mathbf{G}^{(r)} \right)^{-1} \right\|_2}{\lambda_K - \left\| \Delta_1^{(r)} \left(\mathbf{U}_1^T \mathbf{G}^{(r)} \right)^{-1} \right\|_2} \\ &\leq \left\| \mathbf{U}_2^T \mathbf{G}^{(r)} \left(\mathbf{U}_1^T \mathbf{G}^{(r)} \right)^{-1} \right\|_2 \left(\frac{\lambda_{K+1} + \frac{\|\Delta_2^{(r)}\|_2}{\sigma_{\min}(\mathbf{U}_2^T \mathbf{G}^{(r)})}}{\lambda_K - \frac{\|\Delta_1^{(r)}\|_2}{\sigma_{\max}(\mathbf{U}_1^T \mathbf{G}^{(r)})}} \right). \end{aligned} \quad (35)$$

Assume that the following holds:

$$\begin{aligned} & \max\{\|\Delta_1^{(r)}\|_2, \|\Delta_2^{(r)}\|_2\} \\ &\leq \frac{\lambda_K - \lambda_{K+1}}{3} \min\left\{ \sigma_{\min}\left(\mathbf{U}_2^T \mathbf{G}^{(r)}\right), \sigma_{\max}\left(\mathbf{U}_1^T \mathbf{G}^{(r)}\right) \right\}. \end{aligned} \quad (37)$$

Then, one can easily show that

$$\begin{aligned} & \left\| \mathbf{U}_2^T \mathbf{G}^{(r+1)} \left(\mathbf{U}_1^T \mathbf{G}^{(r+1)} \right)^{-1} \right\|_2 \\ &\leq \varrho \left\| \mathbf{U}_2^T \mathbf{G}^{(r)} \left(\mathbf{U}_1^T \mathbf{G}^{(r)} \right)^{-1} \right\|_2 \end{aligned} \quad (38)$$

where

$$\varrho = \left(\frac{2\lambda_{K+1} + \lambda_K}{2\lambda_K + \lambda_{K+1}} \right) < 1.$$

One can see that

$$\begin{aligned} \left\| \mathbf{U}_2^T \mathbf{G}^{(r+1)} \right\|_2 &\leq \left\| \mathbf{U}_2^T \mathbf{G}^{(r+1)} \left(\mathbf{U}_1^T \mathbf{G}^{(r+1)} \right)^{-1} \right\|_2 \\ &\leq \varrho^r \left\| \left(\mathbf{U}_2^T \mathbf{G}^{(0)} \right) \left(\mathbf{U}_1^T \mathbf{G}^{(0)} \right)^{-1} \right\|_2 \\ &\leq \varrho^r \tan(\theta), \end{aligned} \quad (39)$$

where the first inequality holds because of $\|\mathbf{U}_1^T \mathbf{G}^{(r+1)}\|_2 \leq 1$. By noticing that $\|\mathbf{U}_2^T \mathbf{G}^{(0)}\|_2 = \sin(\theta)$ and $\|(\mathbf{U}_1^T \mathbf{G}^{(0)})^{-1}\|_2 = 1/\cos(\theta)$ [30, Theorem 8.2.2], we obtain the last inequality.

In addition, we notice that

$$\max\{\|\Delta_1^{(r)}\|_2, \|\Delta_2^{(r)}\|_2\} \leq \sum_{i=1}^I \lambda_{\max}(\mathbf{X}_i) \epsilon^{(r)}.$$

This means that to ensure linear convergence to a global minimal solution, it suffices to have

$$\begin{aligned} \epsilon^{(r)} &\leq \frac{\lambda_K - \lambda_{K+1}}{3 \sum_{i=1}^I \lambda_{\max}(\mathbf{X}_i)} \\ &\times \min\left\{ \sigma_{\min}\left(\mathbf{U}_2^T \mathbf{G}^{(r)}\right), \sigma_{\max}\left(\mathbf{U}_1^T \mathbf{G}^{(r)}\right) \right\} \end{aligned} \quad (40)$$

in the worst case.

The last piece of the proof is to show that $(\mathbf{U}_1^T \mathbf{G}^{(r)})^{-1}$ in (36) always exists. Note that if (39) holds, then $\mathbf{U}_1^T \mathbf{G}^{(r)}$ is always invertible under the condition stated in (12). The reason is that we always have [30]

$$\sigma_{\max}^2(\mathbf{U}_2^T \mathbf{G}^{(r)}) + \sigma_{\min}^2(\mathbf{U}_1^T \mathbf{G}^{(r)}) = 1.$$

Therefore, $\sigma_{\min}^2(\mathbf{U}_1^T \mathbf{G}^{(r)})$ monotonically increases since $\sigma_{\max}^2(\mathbf{U}_2^T \mathbf{G}^{(r)})$ decreases when (40) (and thus (39)) holds. Hence, if $\sigma_{\min}(\mathbf{U}_1^T \mathbf{G}^{(0)}) > 0$, we have $\sigma_{\min}(\mathbf{U}_1^T \mathbf{G}^{(r)}) > 0$ for all $r > 1$.

REFERENCES


- [1] D. R. Hardoon, S. Szedmak, and J. Shawe-Taylor, "Canonical correlation analysis: An overview with application to learning methods," *Neural computation*, vol. 16, no. 12, pp. 2639–2664, 2004.
- [2] Y.-O. Li, T. Adali, W. Wang, and V. D. Calhoun, "Joint blind source separation by multiset canonical correlation analysis," *IEEE Trans. Signal Process.*, vol. 57, no. 10, pp. 3918–3929, 2009.
- [3] A. Bertrand and M. Moonen, "Distributed canonical correlation analysis in wireless sensor networks with application to distributed blind source separation," *IEEE Trans. Signal Process.*, vol. 63, no. 18, pp. 4800–4813, 2015.
- [4] Q. Wu and K. M. Wong, "Un-music and un-cle: An application of generalized correlation analysis to the estimation of the direction of arrival of signals in unknown correlated noise," *IEEE Trans. Signal Process.*, vol. 42, no. 9, pp. 2331–2343, 1994.
- [5] A. Dogandzic and A. Nehorai, "Finite-length mimo equalization using canonical correlation analysis," *IEEE Trans. Signal Process.*, vol. 50, no. 4, pp. 984–989, 2002.

- [6] S. M. Kakade and D. P. Foster, "Multi-view regression via canonical correlation analysis," in *Learning Theory*. Springer, 2007, pp. 82–96.
- [7] K. Chaudhuri, S. M. Kakade, K. Livescu, and K. Sridharan, "Multi-view clustering via canonical correlation analysis," in *Proc. ICML*. ACM, 2009, pp. 129–136.
- [8] R. Arora and K. Livescu, "Multi-view learning with supervision for transformed bottleneck features," in *Proc. ICASSP*. IEEE, 2014, pp. 2499–2503.
- [9] H. Wang, T. Lee, C.-C. Leung, B. Ma, and H. Li, "Acoustic segment modeling with spectral clustering methods," *IEEE/ACM Trans. Audio, Speech, Language Process.*, vol. 23, no. 2, pp. 264–277, 2015.
- [10] P. Rastogi, B. Van Durme, and R. Arora, "Multiview LSA: Representation learning via generalized CCA," in *Proc. NAACL*, 2015.
- [11] J. D. Carroll, "Generalization of canonical correlation analysis to three or more sets of variables," in *Proc. annual convention of the American Psychological Association*, vol. 3, 1968, pp. 227–228.
- [12] P. Horst, "Generalized canonical correlations and their applications to experimental data," *Journal of Clinical Psychology*, vol. 17, no. 4, pp. 331–347, 1961.
- [13] J. R. Kettenring, "Canonical analysis of several sets of variables," *Biometrika*, vol. 58, no. 3, pp. 433–451, 1971.
- [14] N. A. Asendorf, "Informative data fusion: Beyond canonical correlation analysis," Ph.D. dissertation, The University of Michigan, 2015.
- [15] Z. Ma, Y. Lu, and D. Foster, "Finding linear structure in large datasets with scalable canonical correlation analysis," *arXiv preprint arXiv:1506.08170*, 2015.
- [16] L. Sun, S. Ji, and J. Ye, "Canonical correlation analysis for multilabel classification: A least-squares formulation, extensions, and analysis," *IEEE Trans. Pattern Anal. Machine Intel.*, vol. 33, no. 1, pp. 194–200, 2011.
- [17] Y. Lu and D. P. Foster, "Large scale canonical correlation analysis with iterative least squares," in *Proc. NIPS*, 2014, pp. 91–99.
- [18] M. Van De Velden and T. H. A. Bijmolt, "Generalized canonical correlation analysis of matrices with missing rows: a simulation study," *Psychometrika*, vol. 71, no. 2, pp. 323–331, 2006.
- [19] D. R. Hardoon and J. Shawe-Taylor, "Sparse canonical correlation analysis," *Machine Learning*, vol. 83, no. 3, pp. 331–353, 2011.
- [20] D. M. Witten, R. Tibshirani, and T. Hastie, "A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis," *Biostatistics*, p. kxp008, 2009.
- [21] X. Chen, H. Liu, and J. G. Carbonell, "Structured sparse canonical correlation analysis," in *International Conference on Artificial Intelligence and Statistics*, 2012, pp. 199–207.
- [22] D. M. Witten and R. J. Tibshirani, "Extensions of sparse canonical correlation analysis with applications to genomic data," *Statistical applications in genetics and molecular biology*, vol. 8, no. 1, pp. 1–27, 2009.
- [23] M. Faruqui and C. Dyer, "Improving vector space word representations using multilingual correlation." Association for Computational Linguistics, 2014.
- [24] C. Sigg, B. Fischer, B. Ommer, V. Roth, and J. Buhmann, "Nonnegative cca for audiovisual source separation," in *2007 IEEE Workshop on Machine Learning for Signal Processing*. IEEE, 2007, pp. 253–258.
- [25] B. Fischer, V. Roth, and J. M. Buhmann, "Time-series alignment by non-negative multiple generalized canonical correlation analysis," *BMC bioinformatics*, vol. 8, no. Suppl 10, p. S4, 2007.
- [26] X. Fu, K. Huang, M. Hong, N. D. Sidiropoulos, and A. M.-C. So, "Scalable and flexible max-var generalized canonical correlation analysis via alternating optimization," in *Proc. ICASSP 2017*, 2016.
- [27] J. Rupnik, P. Skraba, J. Shawe-Taylor, and S. Guettes, "A comparison of relaxations of multiset canonical correlation analysis and applications," *arXiv preprint arXiv:1302.0974*, 2013.
- [28] L.-H. Zhang, L.-Z. Liao, and L.-M. Sun, "Towards the global solution of the maximal correlation problem," *Journal of Global Optimization*, vol. 49, no. 1, pp. 91–107, 2011.
- [29] X. Fu, K. Huang, E. E. Papalexakis, H. A. Song, P. P. Talukdar, N. D. Sidiropoulos, C. Faloutsos, and T. Mitchell, "Efficient and distributed algorithms for large-scale generalized canonical correlations analysis," in *Proc. IEEE ICDM 2016*, Dec 2016, pp. 871–876.
- [30] G. H. Golub and C. F. V. Loan., *Matrix Computations*. The Johns Hopkins University Press, 1996.
- [31] I. Rustandi, M. A. Just, and T. Mitchell, "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.
- [32] T. M. Mitchell, S. V. Shinkareva, A. Carlson, K.-M. Chang, V. L. Malave, R. A. Mason, and M. A. Just, "Predicting human brain activity associated with the meanings of nouns," *science*, vol. 320, no. 5880, pp. 1191–1195, 2008.
- [33] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends in Machine Learning*, pp. 1–122, 2011.
- [34] N. Parikh and S. Boyd, "Proximal algorithms," *Foundations and Trends in optimization*, vol. 1, no. 3, pp. 123–231, 2013.
- [35] P. Schönemann, "A generalized solution of the orthogonal Procrustes problem," *Psychometrika*, vol. 31, no. 1, pp. 1–10, 1966.
- [36] D. P. Bertsekas, *Nonlinear programming*. Athena Scientific, 1999.
- [37] M. Razaviyayn, M. Hong, and Z.-Q. Luo, "A unified convergence analysis of block successive minimization methods for nonsmooth optimization," *SIAM Journal on Optimization*, vol. 23, no. 2, pp. 1126–1153, 2013.
- [38] Y. Xu and W. Yin, "A block coordinate descent method for regularized multiconvex optimization with applications to nonnegative tensor factorization and completion," *SIAM Journal on imaging sciences*, vol. 6, no. 3, pp. 1758–1789, 2013.
- [39] —, "A globally convergent algorithm for nonconvex optimization based on block coordinate update," *arXiv preprint arXiv:1410.1386*, 2014.
- [40] M. Faruqui and C. Dyer, "Community evaluation and exchange of word vectors at wordvectors.org," in *Proc. 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*. Baltimore, USA: Association for Computational Linguistics, June 2014.
- [41] G. Andrew, R. Arora, J. A. Bilmes, and K. Livescu, "Deep canonical correlation analysis," in *ICML (3)*, 2013, pp. 1247–1255.
- [42] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM journal on imaging sciences*, vol. 2, no. 1, pp. 183–202, 2009.
- [43] R. Ge, C. Jin, S. M. Kakade, P. Netrapalli, and A. Sidford, "Efficient algorithms for large-scale generalized eigenvector computation and canonical correlation analysis," in *ICML 2016*, 2016.




Xiao Fu (S'12-M'15) received his B.Eng and M.Eng degrees in Communications and Information Engineering from the University of Electronic Science and Technology of China, Chengdu, China, in 2005 and 2010, respectively. In 2014, he received his Ph.D. degree in Electronic Engineering from The Chinese University of Hong Kong (CUHK), Shatin, N.T., Hong Kong. From 2005 to 2006, he was an Assistant Engineer at China Telecom Co. Ltd., Shenzhen, China. He is currently a Postdoctoral Associate at the Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN 55455, United States. His research interests include signal processing and machine learning, with a recent emphasis on factor analysis and its applications.

Dr. Fu received a Best Student Paper Award at IEEE ICASSP 2014, and was a finalist of the Best Student Paper Competition at IEEE SAM 2014. He also co-authored a paper that received a Best Student Paper Award at IEEE CAMSAP 2015.



Kejun Huang (S'13-M'17) received the bachelor's degree from Nanjing University of Information Science and Technology, Nanjing, China in 2010. He received the Ph.D. degree in Electrical Engineering from University of Minnesota, Minneapolis, MN, USA in 2016. He is currently a Postdoctoral Associate at the Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, United States. His research interests include signal processing, machine learning, and optimization.




Mingyi Hong received his B.E. degree in Communications Engineering from Zhejiang University, China, in 2005, his M.S. degree in Electrical Engineering from Stony Brook University in 2007, and Ph.D. degree in Systems Engineering from University of Virginia in 2011. From 2011 to 2014 he has been a Post Doctoral Fellow and a Research Assistant Professor with the Department of Electrical and Computer Engineering, University of Minnesota, Twin Cities. He is currently an Assistant Professor with the Department of Industrial and

Manufacturing Systems Engineering (IMSE), Iowa State University, Ames. His research interests include signal processing, wireless communications, large-scale optimization and its applications in compressive sensing, complex networks and high-dimensional data analysis.



Nicholas D. Sidiropoulos (F'09) received the Diploma in Electrical Engineering from the Aristotelian University of Thessaloniki, Greece, and M.S. and Ph.D. degrees in Electrical Engineering from the University of Maryland - College Park, in 1988, 1990 and 1992, respectively. He served as Assistant Professor at the University of Virginia (1997-1999); Associate Professor at the University of Minnesota - Minneapolis (2000-2002); Professor at the Technical University of Crete, Greece (2002-2011); and Professor at the University of Minnesota

- Minneapolis (2011-). His current research focuses primarily on signal and tensor analytics, with applications in cognitive radio, big data, and preference measurement. He received the NSF/CAREER award (1998), the IEEE Signal Processing Society (SPS) Best Paper Award (2001, 2007, 2011), and the IEEE SPS Meritorious Service Award (2010). He has served as IEEE SPS Distinguished Lecturer (2008-2009), and Chair of the IEEE Signal Processing for Communications and Networking Technical Committee (2007-2008). He received the Distinguished Alumni Award of the Department of Electrical and Computer Engineering, University of Maryland, College Park in 2013, and was elected EURASIP Fellow in 2014.



Anthony Man-Cho So (M'12–SM'17) received his BSE degree in Computer Science from Princeton University with minors in Applied and Computational Mathematics, Engineering and Management Systems, and German Language and Culture. He then received his MSc degree in Computer Science and his PhD degree in Computer Science with a PhD minor in Mathematics from Stanford University. Dr. So joined The Chinese University of Hong Kong (CUHK) in 2007. He is currently an Assistant Dean of the Faculty of Engineering and an Associate

Professor in the Department of Systems Engineering and Engineering Management. He also holds a courtesy appointment as an Associate Professor in the CUHK-BGI Innovation Institute of Trans-omics. His research focuses on the interplay between optimization theory and various areas of algorithm design, such as computational geometry, machine learning, signal processing, bioinformatics, and algorithmic game theory.

Dr. So currently serves on the editorial boards of Journal of Global Optimization, Optimization Methods and Software, and SIAM Journal on Optimization. He has also served on the editorial boards of IEEE TRANSACTIONS ON SIGNAL PROCESSING and Mathematics of Operations Research. He received the 2015 IEEE Signal Processing Society Signal Processing Magazine Best Paper Award, the 2014 IEEE Communications Society Asia-Pacific Outstanding Paper Award, the 2010 Institute for Operations Research and the Management Sciences (INFORMS) Optimization Society Optimization Prize for Young Researchers, and the 2010 CUHK Young Researcher Award. He also received the 2008 Exemplary Teaching Award and the 2011, 2013, 2015 Dean's Exemplary Teaching Award from the Faculty of Engineering at CUHK, and the 2013 Vice-Chancellor's Exemplary Teaching Award from CUHK.