MULTI-RANK SPARSE AND FUNCTIONAL PCA
MANIFOLD OPTIMIZATION AND ITERATIVE DEFLATION TECHNIQUES

Michael Weylandt

Department of Statistics
Rice University
Houston, TX 77005
michael.weylandt@rice.edu

ABSTRACT

We consider the problem of estimating multiple principal components using the recently-proposed Sparse and Functional Principal Components Analysis (SFPCA) estimator. We first propose an extension of SFPCA which estimates several principal components simultaneously using manifold optimization techniques to enforce orthogonality constraints. While effective, this approach is computationally burdensome so we also consider iterative deflation approaches which take advantage of existing fast algorithms for rank-one SFPCA. We show that alternative deflation schemes can more efficiently extract signal from the data, in turn improving estimation of subsequent components. Finally, we compare the performance of our manifold optimization and deflation techniques in a scenario where orthogonality does not hold and find that they still lead to significantly improved performance. ***Index Terms*—** regularized PCA, orthogonality, deflation, sparsity, manifold optimization

1. INTRODUCTION

Principal Components Analysis (PCA, [1]) is a widely-used approach to finding low-dimensional patterns in complex data, enabling visualization, dimension reduction (compression), and predictive modeling. While PCA performs well in a wide range of low-dimensional settings, its performance degrades rapidly in high-dimensions, necessitating the use of regularized variants. Recently, Allen and Weylandt [2] proposed Sparse and Functional PCA (SFPCA), a unified regularization scheme that allows for simultaneous smooth (functional) and sparse estimation of both row and column principal components (PCs). The rank-one SFPCA estimator is given by

$$\operatorname*{arg\,max}_{\boldsymbol{u}\in\overline{\mathbb{B}}^n_{\boldsymbol{S}_{\boldsymbol{u}}}, v\in\overline{\mathbb{B}}^p_{\boldsymbol{S}_{\boldsymbol{v}}}} \boldsymbol{u}^T \boldsymbol{X} \boldsymbol{v} - \lambda_{\boldsymbol{u}} P_{\boldsymbol{u}}(\boldsymbol{u}) - \lambda_{\boldsymbol{v}} P_{\boldsymbol{v}}(\boldsymbol{v}) \quad (1)$$

where $P_{\boldsymbol{u}}(\cdot)$ is a regularizer inducing sparsity in the row PCs, with strength controlled by $\lambda_{\boldsymbol{u}}$; $\boldsymbol{\Omega}_{\boldsymbol{u}}$ is a positive semi-definite penalty matrix, typically a second- or fourth-order difference matrix; $\boldsymbol{S}_{\boldsymbol{u}} = \boldsymbol{I} + \alpha_{\boldsymbol{u}} \boldsymbol{\Omega}_{\boldsymbol{u}}$ is a smoothing matrix for the row PCs, with strength controlled by $\alpha_{\boldsymbol{u}}$; and $\overline{\mathbb{B}}^n_{\boldsymbol{S}_{\boldsymbol{u}}}$ is the unit ellipse of the $\boldsymbol{S}_{\boldsymbol{u}}$-norm, *i.e.*, $\overline{\mathbb{B}}^n_{\boldsymbol{S}_{\boldsymbol{u}}} = \{\boldsymbol{u} \in \mathbb{R}^n : \boldsymbol{u}^T \boldsymbol{S}_{\boldsymbol{u}} \boldsymbol{u} \leq 1\}$. (Respectively, $P_{\boldsymbol{v}}(\cdot)$, $\lambda_{\boldsymbol{v}}$, $\boldsymbol{\Omega}_{\boldsymbol{v}}$, $\alpha_{\boldsymbol{v}}$, $\boldsymbol{S}_{\boldsymbol{v}}$, and $\overline{\mathbb{B}}^p_{\boldsymbol{S}_{\boldsymbol{v}}}$ for the column PCs.)

Allen and Weylandt [2] show that SFPCA unifies much of the existing regularized PCA literature [3–8] into a single framework,

---

avoiding many pathologies of other approaches. Finally, they propose an efficient alternating maximization scheme with guaranteed global convergence to solve the bi-concave SFPCA problem (1). The SFPCA estimator only allows for a single pair of PCs to be estimated for a given data matrix $\boldsymbol{X}$. Allen and Weylandt suggest applying SFPCA repeatedly to the deflated data matrix if multiple PCs are desired. While this approach performs acceptably in practice, it loses the interpretable orthogonality properties of standard PCA. In particular, the estimated PCs are no longer guaranteed to be orthogonal to each other, hindering the common interpretation of PCs as statistically independent sources of variance, or to the deflated data matrix, suggesting that additional signal remains uncaptured.

We extend the work of Allen and Weylandt [2] to address these shortcomings: first, in Section 2, we modify the SFPCA estimator to simultaneously estimate several sparse PCs subject to orthogonality constraints. The resulting estimator is constrained to a product of generalized Stiefel mainfolds and we propose three efficient algorithms to solve the resulting manifold optimization problem. Next, in Section 3, we propose improved deflation schemes which provably remove all of the signal from the data matrix, allowing for more accurate iterative estimation of sparse PCs. Finally, we demonstrate the improved performance of our manifold estimators and deflation schemes in Section 4. Supplemental materials for this paper, including proofs, counter-examples, and additional algorithmic details, are available online at https://arxiv.org/abs/1907.12012.

2. MANIFOLD OPTIMIZATION FOR SFPCA

One of the most attractive properties of PCA is the factors it extracts are orthogonal ($\boldsymbol{u}_t^T \boldsymbol{u}_s = \boldsymbol{v}_t^T \boldsymbol{v}_s = 1$ if $t = s$ and 0 otherwise). Because of this, PCs can be interpreted as separate sources of variance and, under an additional Gaussianity assumption, statistically independent. While this follows directly from the properties of eigen-decompositions for standard PCA, it is much more difficult to obtain similar results for sparse PCA. Some authors have suggested that there exists a fundamental tension between orthogonality and sparsity, with Journée *et al.* [9] calling the goal of sparse and orthogonal estimation "questionable." Indeed *ex post* orthogonalization of sparse PCs, *e.g.*, using a Gram-Schmidt step, destroys any sparsity in the estimated PCs. To avoid this, it is necessary to impose orthogonality directly in the estimation step, rather than trying to impose it afterwards.

We modify the SFPCA estimator to simultaneously estimate multiple PCs subject to an orthognality constraint:

$$\operatorname*{arg\,max}_{\boldsymbol{U}\in\mathcal{V}^{\boldsymbol{S}_{\boldsymbol{u}}}_{n\times k}, \boldsymbol{V}\in\mathcal{V}^{\boldsymbol{S}_{\boldsymbol{v}}}_{p\times k}} \operatorname{Tr}(\boldsymbol{U}^T \boldsymbol{X} \boldsymbol{V}) - \lambda_{\boldsymbol{U}} P_{\boldsymbol{U}}(\boldsymbol{U}) - \lambda_{\boldsymbol{V}} P_{\boldsymbol{V}}(\boldsymbol{V}) \quad (2)$$

Method	Two-Way Orthogonality $\boldsymbol{u}_t^T \boldsymbol{X}_t \boldsymbol{v}_t = 0$	One-Way Orthogonality $\boldsymbol{u}_t^T \boldsymbol{X}_t, \boldsymbol{X}_t \boldsymbol{v}_t = \boldsymbol{0}$	Subsequent Orthogonality ($\forall s \geq 0$) $\boldsymbol{u}_t^T \boldsymbol{X}_{t+s}, \boldsymbol{X}_{t+s} \boldsymbol{v}_t = \boldsymbol{0}$	Robust to Scale of $\boldsymbol{u}_t, \boldsymbol{v}_t$
Hotelling's Deflation (HD)	✓	✗	✗	✗
Projection Deflation (PD)	✓	✓	✗	✗
Schur Deflation (SD)	✓	✓	✓	✓

Table 1: Properties of Hotelling's Deflation (HD), Projection Deflation (PD), and Schur Complement Deflation (SD). Only SD captures all of the individual signal of each principal component without re-introducing signal at later iterations. Additionally, only SD allows for the non-unit-norm PCs estimated by SFPCA to be used without rescaling.

where $\mathcal{V}_{n \times k}^{\boldsymbol{S_u}}$ is the generalized Stiefel manifold of order k over $\mathbb{R}^n$, *i.e.*,

$$\boldsymbol{U} \in \mathcal{V}_{n \times k}^{\boldsymbol{S_u}} \iff \boldsymbol{U} \in \mathbb{R}^{n \times k} \text{ and } \boldsymbol{U}^T \boldsymbol{S_u} \boldsymbol{U} = \boldsymbol{I}_k.$$

The generalized Stiefel manifold constraint ensures orthogonality of the estimated PCs, while still allowing us to capture most of the variability in the data.[1] We note that, because we use a generalized Stiefel constraint, the estimated PCs will be *orthogonal with respect to $\boldsymbol{S_u}$, i.e., $\boldsymbol{u}_t^T \boldsymbol{S_u} \boldsymbol{u}_s = 0$ for $t \neq s$*, rather than orthogonal in the standard sense. This is commonly observed for functional PCA variants [3, 5, 6] and can be interpreted as orthogonality under the inner product generating the $\boldsymbol{S_u}$-norm. If no roughness penalty is imposed ($\alpha_{\boldsymbol{u}} = 0$ or $\alpha_{\boldsymbol{v}} = 0$), then our method gives orthogonality in the standard sense.

To solve the Manifold SFPCA problem (2), we employ an alternating maximization scheme, first holding $\boldsymbol{V}$ fixed while we update $\boldsymbol{U}$ and *vice versa*, as described in Algorithm 1. Even with one parameter held fixed, the resulting sub-problems are still difficult manifold optimization problems, which require iterative approaches to obtain a solution [10–17].

Algorithm 1 Manifold SFPCA Algorithm

1. Initialize $\hat{\boldsymbol{U}}, \hat{\boldsymbol{V}}$ to the leading k singular vectors of $\boldsymbol{X}$

2. Repeat until convergence:

   (a) $\boldsymbol{U}$-subproblem. Solve using Algorithm 2 or 3:

   $$\hat{\boldsymbol{U}} = \underset{\boldsymbol{U} \in \mathcal{V}_{n \times k}^{\boldsymbol{S_u}}}{\arg \min} - \text{Tr}(\boldsymbol{U}^T \boldsymbol{X} \hat{\boldsymbol{V}}) + \lambda_{\boldsymbol{U}} P_{\boldsymbol{U}}(\boldsymbol{U})$$

   (b) $\boldsymbol{V}$-subproblem: Solve using Algorithm 2 or 3, with $\boldsymbol{U}$ and $\boldsymbol{V}$ reversed:

   $$\hat{\boldsymbol{V}} = \underset{\boldsymbol{V} \in \mathcal{V}_{n \times p}^{\boldsymbol{S_v}}}{\arg \min} - \text{Tr}(\hat{\boldsymbol{U}}^T \boldsymbol{X} \boldsymbol{V}) + \lambda_{\boldsymbol{V}} P_{\boldsymbol{V}}(\boldsymbol{V})$$

3. Return $\hat{\boldsymbol{U}}$ and $\hat{\boldsymbol{V}}$

Allen and Weylandt [2] developed a custom projected + proximal gradient algorithm to solve the $\boldsymbol{u}$- and $\boldsymbol{v}$ subproblems of the rank-one SFPCA estimator. Assuming $P_{\boldsymbol{U}}$ and $P_{\boldsymbol{V}}$ are positive homogeneous, (*e.g.* $P(\cdot) = \|\boldsymbol{A} \cdot \|_p$ for arbitrary $p \geq 1$ and $\boldsymbol{A}$), they establish convergence to a stationary point. In order to extend this idea to the multi-rank (manifold) case, we use the recently-proposed Manifold Proximal Gradient (ManPG) scheme of Chen *et al.* [14], detailed in Algorithm 2. ManPG proceeds in two-steps: first, we solve for

---

[1]When estimating orthogonal factors, it is common to re-express the problem using a (generalized) Grassmanian manifold constraint to avoid identifiability issues. We cannot use the Grassmanian approach here as sparse estimation (implicitly) fixes a single coordinate system.

a descent direction $\boldsymbol{D}$ of the objective along the tangent space of the generalized Stiefel manifold, subject to the tangency constraint of $\boldsymbol{D}^T \boldsymbol{S_u} \boldsymbol{U}^{(k)}$ being skew-symmetric; secondly, back-tracking line search is used to determine a step-size α, after which the estimate is projected back onto the generalized Stiefel manifold using a retraction anchored at the previous $\boldsymbol{U}^{(k)}$. The retraction, which plays the same role as the projection step in the original SFPCA algorithm [2, Algorithm 1], can be computed using a Cholesky factorization [17, Algorithm 3.1]. Chen *et al.* [15] showed that a single step of ManPG is sufficient to ensure convergence despite the bi-concave objective: we refer to their approach as Alternating ManPG (A-ManPG).

Algorithm 2 Manifold Prox. Gradient (ManPG) for $\hat{\boldsymbol{U}}$-Subproblem

1. Initialize $\boldsymbol{U}^{(k)} = \hat{\boldsymbol{U}}$

2. Repeat until convergence:

   • Solve, subject to $\boldsymbol{D}^T \boldsymbol{S_u} \boldsymbol{U}^{(k)} + (\boldsymbol{U}^{(k)})^T \boldsymbol{S_u} \boldsymbol{D} = \boldsymbol{0}$:

   $$\hat{\boldsymbol{D}} = \underset{\boldsymbol{D} \in \mathbb{R}^{n \times k}}{\arg \min} -\langle \boldsymbol{X} \hat{\boldsymbol{V}}, \boldsymbol{D} \rangle_F + \lambda_{\boldsymbol{U}} P_{\boldsymbol{U}}(\boldsymbol{U}^{(k)} + \boldsymbol{D})$$

   • Select α by Armijo-back-tracking

   • $\boldsymbol{U}^{(k+1)} = \text{Retr}_{\boldsymbol{U}^{(k)}}(\alpha \hat{\boldsymbol{D}})$

3. Return $\hat{\boldsymbol{U}}$

Note that in general manifold proximal gradient schemes [14, 15] impose a maximum step-size to ensure that linearization of the smooth portion of the objective actually leads to descent: because the smooth portion of our objective function in linear in $\boldsymbol{U}$ and $\boldsymbol{V}$, we can omit this term from Algorithm 2.

While efficient when tuned properly, we have found the performance of ManPG on the $\boldsymbol{U}$- and $\boldsymbol{V}$-subproblems quite sensitive to infeasibility in the descent direction. A more robust scheme can be derived by using the Manifold ADMM (MADMM) scheme of Kovnatsky *et al.* [13] to solve the subproblems. Like standard ADMM schemes, MADMM allows us to split a problem into two parts, each of which can be solved more easily than the global problem. When applied to the $\boldsymbol{U}$- and $\boldsymbol{V}$-subproblems, the MADMM allows us to separate the manifold constraint from the sparsity inducing regularizer, thereby side-stepping the orthogonality / sparsity tension at the heart of this paper. After this splitting, the smooth update can be shown to equivalent to the unbalanced Procrustes problem [18, 19] with a closed-form update: $\boldsymbol{U}^{(k+1)} = \boldsymbol{S_u}^{-1/2} \boldsymbol{A} \boldsymbol{B}^T$ where $\boldsymbol{A} \boldsymbol{\Delta} \boldsymbol{B}^T$ is the SVD of $\boldsymbol{S_u}^{-1/2} \boldsymbol{X} \hat{\boldsymbol{V}} + \rho \boldsymbol{S_u}^{1/2}(\boldsymbol{W}^{(k)} - \boldsymbol{Z}^{(k)})$. The sparse update is simply the proximal operator of $P_{\boldsymbol{U}}(\cdot)$, typically a threhsolding step [20, Chapter 6]. To the best of our knowledge, the convergence of MADMM has not yet been established, but we have not observed significant non-convergence problems in our experiments.

Algorithm 3 Manifold ADMM (MADMM) for $\hat{U}$-Subproblem

---

1. Initialize $\boldsymbol{U}^{(k)} = \boldsymbol{W}^{(k)} = \hat{\boldsymbol{U}}$, $\boldsymbol{Z}^{(k)} = \boldsymbol{0}$ and $k = 1$

2. Repeat until convergence:

$$\boldsymbol{U}^{(k+1)} = \underset{\boldsymbol{U} \in \mathcal{V}_{n \times k}^{\boldsymbol{S}_{\boldsymbol{u}}}}{\arg\min} -\operatorname{Tr}(\boldsymbol{U}^T \boldsymbol{X} \boldsymbol{V}) + \frac{\rho}{2}\|\boldsymbol{U} - \boldsymbol{W}^{(k)} + \boldsymbol{Z}^{(k)}\|_F^2$$

$$\boldsymbol{W}^{(k+1)} = \underset{\boldsymbol{W} \in \mathbb{R}^{n \times k}}{\arg\min} \, \lambda_{\boldsymbol{U}} P_{\boldsymbol{U}}(\boldsymbol{W}) + \frac{\rho}{2}\|\boldsymbol{U}^{(k+1)} - \boldsymbol{W} + \boldsymbol{Z}^{(k)}\|_F^2$$

$$= \operatorname{prox}_{\lambda_{\boldsymbol{U}}/\rho P_{\boldsymbol{U}}(\cdot)}\left(\boldsymbol{U}^{(k+1)} + \boldsymbol{Z}^{(k)}\right)$$

$$\boldsymbol{Z}^{(k+1)} = \boldsymbol{Z}^{(k)} + \boldsymbol{U}^{(k+1)} - \boldsymbol{W}^{(k+1)}$$

3. Return $\hat{\boldsymbol{U}}$ and $\hat{\boldsymbol{V}}$

---

3. ITERATIVE DEFLATION FOR SFPCA

We next consider the use of iterative deflation schemes for multi-rank SFPCA. As discussed by Mackey [21], the attractive orthogonality properties of standard (Hotelling's) deflation depend critically on the estimated PCs being exact eigenvectors of the covariance matrix. Because the PCs estimated by sparse PCA schemes are almost surely not eigenvectors, Mackey [21] proposes several alternate deflation schemes which retain some of the attractive properties of Hotelling's deflation even when non-eigenvectors are used. We extend these to the low-rank model and allow for deflation by several PCs, possibly with non-unit norm, simultaneously, *e.g.*, as produced by ManSFPCA. To ease exposition, we first work in the vector setting and consider the general case at the end of this section. The properties of our proposed deflation schemes are summarized in Table 1 above.

The simplest deflation scheme is essentially that used by Hotelling [1], extended to the low-rank model:

$$\boldsymbol{X}_t \coloneqq \boldsymbol{X}_{t-1} - d_t \boldsymbol{u}_t \boldsymbol{v}_t^T \text{ where } d_t = \boldsymbol{u}_t^T \boldsymbol{X}_{t-1} \boldsymbol{v}_t. \qquad \text{(HD)}$$

For two-way sparse PCA variants [7, 8], this deflation gives a deflated matrix which is "two-way" orthogonal to the estimated PCs, *i.e.*, $\boldsymbol{u}_t^T \boldsymbol{X}_t \boldsymbol{v}_t = 0$. We may interpret this as Hotelling's deflation (HD) capturing all the signal *jointly* associated with the pair $(\boldsymbol{u}_t, \boldsymbol{v}_t)$.

We may also ask if HD captures *all* of the signal associated with $\boldsymbol{u}_t$ or only the signal which is also associated with $\boldsymbol{v}_t$. If HD captures all of the signal associated with $\boldsymbol{u}_t$, then we would expect $\boldsymbol{u}_t^T \boldsymbol{X}_t \tilde{\boldsymbol{v}} = 0$ for all $\tilde{\boldsymbol{v}} \in \mathbb{R}^p$, or equivalently, $\boldsymbol{u}_t^T \boldsymbol{X}_t = \boldsymbol{0}_p$. Interestingly, HD does not have this left-orthogonality property, suggesting that it leaves additional $\boldsymbol{u}_t$-signal in the deflated matrix $\boldsymbol{X}_t$.

HD fails to yield left- and right-orthogonality because it is not based on a projection operator. To address this in the covariance model, Mackey [21] proposed a deflation scheme which projects the covariance matrix onto the orthogonal complement of the estimated principal component. We extend this idea to the low-rank model by projecting the column- and row-space of the data matrix into the orthogonal complement of the left- and right-PCs respectively, giving two-way projection deflation (PD):

$$\boldsymbol{X}_t \coloneqq (\boldsymbol{I}_n - \boldsymbol{u}_t \boldsymbol{u}_t^T)\boldsymbol{X}_{t-1}(\boldsymbol{I}_p - \boldsymbol{v}_t \boldsymbol{v}_t^T). \qquad \text{(PD)}$$

Unlike HD, PD captures all of the linear signal associated with $\boldsymbol{u}_t$ and $\boldsymbol{v}_t$ individually ($\boldsymbol{u}_t \boldsymbol{X}_t \tilde{\boldsymbol{v}} = 0, \tilde{\boldsymbol{u}} \boldsymbol{X}_t \boldsymbol{v}_t = 0, \forall \tilde{\boldsymbol{u}} \in \mathbb{R}^n, \tilde{\boldsymbol{v}} \in \mathbb{R}^p$).

If we use PD repeatedly, however, the multiply-deflated matrix will not continue to be orthogonal to the PCs: that is, $\boldsymbol{u}_t^T \boldsymbol{X}_{t+s} \neq \boldsymbol{0}$ for $s \geq 1$. This suggests that repeated application of PD can *reintroduce* signal in the direction of the PCs that we previously removed.

This occurs because PD works by sequentially projecting the data matrix, but in general the compositition of two orthogonal projections is not another orthogonal projection without additional assumptions. To address this, Mackey [21] proposed a Schur complement deflation (SD) technique, which we now extend to the low-rank (two-way) model:

$$\boldsymbol{X}_t \coloneqq \boldsymbol{X}_{t-1} - \frac{\boldsymbol{X}_{t-1}\boldsymbol{v}_t \boldsymbol{u}_t^T \boldsymbol{X}_{t-1}}{\boldsymbol{u}_t^T \boldsymbol{X}_{t-1} \boldsymbol{v}_t}. \qquad \text{(SD)}$$

While Mackey motivates this approach using conditional distributions and a Gaussianity assumption on $\boldsymbol{X}$, it can also be understood as an alternate projection construct which is more robust to scaling and non-orthogonality.

So far, we have only considered the behavior of the proposed deflation schemes for two-way sparse PCA. If we consider SFPCA in generality, however, $\boldsymbol{u}_t$ and $\boldsymbol{v}_t$ are unit vectors under the $\boldsymbol{S}_{\boldsymbol{u}}$- and $\boldsymbol{S}_{\boldsymbol{v}}$-norms, not under the Euclidean norm. Consequently, the projections used by PD may not be actual projections and a PD deflated matrix may fail to be two- or one-way orthogonal. Normalizing the estimated PCs before deflation addresses this problem and is recommended in practice: conversely, because its deflation term is invariant under rescalings of $\boldsymbol{u}_t$ and $\boldsymbol{v}_t$, SD works without renormalization.

The extension of these techniques to the multi-rank case is straightforward. We give the normalized variants here:

$$\boldsymbol{X}_t^{\text{HD}} \coloneqq \boldsymbol{X}_{t-1} - \boldsymbol{U}_t(\boldsymbol{U}_t^T \boldsymbol{U}_t)^{-1}\boldsymbol{U}_t^T \boldsymbol{X}_{t-1} \boldsymbol{V}_t(\boldsymbol{V}_t^T \boldsymbol{V}_t)^{-1}\boldsymbol{V}_t^T$$

$$\boldsymbol{X}_t^{\text{PD}} \coloneqq (\boldsymbol{I}_n - \boldsymbol{U}_t(\boldsymbol{U}_t^T \boldsymbol{U}_t)^{-1}\boldsymbol{U}_t^T)\boldsymbol{X}_{t-1}(\boldsymbol{I}_p - \boldsymbol{V}_t(\boldsymbol{V}_t^T \boldsymbol{V}_t)^{-1}\boldsymbol{V}_t^T)$$

$$\boldsymbol{X}_t^{\text{SD}} \coloneqq \boldsymbol{X}_{t-1} - \boldsymbol{X}_{t-1}\boldsymbol{V}_t(\boldsymbol{U}_t^T \boldsymbol{X}_{t-1}\boldsymbol{V}_t)^{-1}\boldsymbol{U}_t^T \boldsymbol{X}_{t-1}.$$

As in the covariance model, if $\boldsymbol{u}_t$ and $\boldsymbol{v}_t$ are true singular vectors, all three deflation schemes are equivalent.

4. SIMULATION STUDIES

In this section, we compare the performance of Manifold SFPCA and the iterative rank-one deflation schemes proposed above in illustrative simulation studies. Manifold SFPCA using Manifold ADMM (Algorithm 3) to solve the subproblems achieves better solutions than Manifold Proximal Gradient or A-ManPG (Algorithm 2) in less time. Furthermore, despite the additional flexibility of the iterative rank-one variants, Manifold SFPCA achieves both better signal recovery and a higher proportion of variance explained, even when the orthogonality assumptions are violated.

We first consider the relative performance of the three algorithms proposed for solving the Manifold SFPCA problem (2). We generate data in $\boldsymbol{X}^* = \boldsymbol{U}^* \boldsymbol{D}^* \boldsymbol{V}^* \in \mathbb{R}^{250 \times 100}$ with three distinct PCs: the left PCs ($\boldsymbol{U}^*$) are localized sinusoids of varying frequency; the right PCs ($\boldsymbol{V}^*$) are non-overlapping sawtooth waves. (See Figure 1.) We add independent standard Gaussian noise ($\boldsymbol{E}$) to give a signal-to-noise ratio (SNR) of $\|\boldsymbol{X}^*\|/\|\boldsymbol{E}\| \approx 1.2$. We fix $\lambda_{\boldsymbol{u}} = \lambda_{\boldsymbol{v}} = 1$ and $\alpha_{\boldsymbol{u}} = \alpha_{\boldsymbol{v}} = 3$ which is near optimal for all three schemes. This is a favorable setting for Manifold SFPCA as the underlying signals are orthogonal, sparse, smooth, and of comparable magnitude.

Table 2 shows the performance of our three Manifold SFPCA algorithms on several metrics, averaged over 100 replicates. Overall, the Manifold ADMM (MADMM [13]) and Alternating Manifold Proximal Gradient (A-ManPG [15]) variants perform best, handily beating the Manifold Proximal Gradient scheme [14] on all measures. MADMM achieved the best objective value on every replicate.

In terms of signal recovery, MADMM achieves slightly better performance than A-ManPG on the right singular vectors, while A-ManPG is slightly better on the left singular vectors. Computationally,
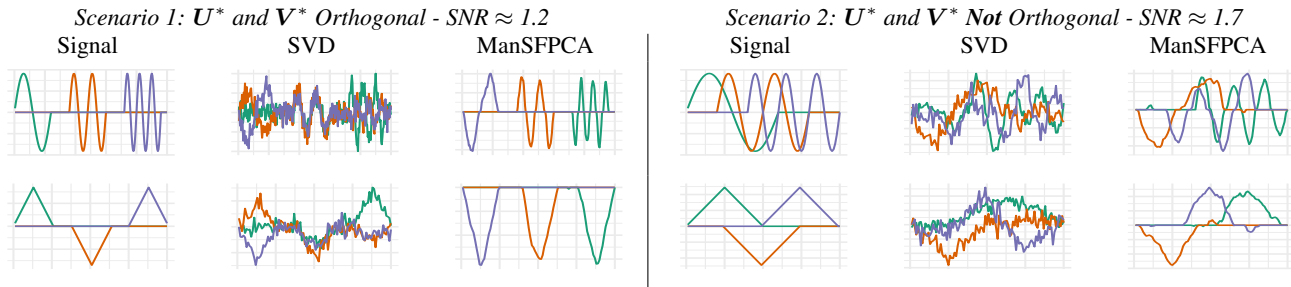
Fig. 1: Simulation Scenarios Used in Section 4. For both the left singular vectors (top row) and the right singular vectors (bottom), ManSFPCA is able to recover the signal far more accurately than unregularized PCA (SVD). ManSFPCA is not identifiable up to change of order or sign.

MADMM dominates both proximal gradient variants even though it requires many more matrix decompositions. The descent direction subproblems of ManPG and A-ManPG are rather expensive to solve repeatedly and their performance is very sensitive to the solver used. Overall, the MADMM variant of Manifold SFPCA achieves the best optimization and statistical performance in far less time than the proximal gradient-based variants.

Next, we compare Manifold SFPCA with the iterative deflation schemes proposed in Section 3 in two different scenarios: the favorable scenario used above and a less-favorable scenario where the true PCs are shifted and no longer orthogonal. ($n = p = 100$ and $\|(\boldsymbol{U}^*)^T(\boldsymbol{U}^*) - \boldsymbol{I}_3\|, \|(\boldsymbol{V}^*)^T\boldsymbol{V}^* - \boldsymbol{I}_3\| \approx 0.37$.) We compare the proportion of variance explained using Manifold SFPCA with iterative rank-one SFPCA using the normalized Hotelling, Projection, and Schur Complement deflation strategies. As can be seen in Table 3, PD and SD consistently dominate HD. Because PD and SD fully remove the signal associated with estimated PCs, the subsequent PCs are able to capture different signals and explain a larger fraction of variance. By ensuring that the signal is never re-introduced, SD does even better than PD as we consider higher ranks. Interestingly, while PD and SD perform about as well when the underlying signals are orthogonal, SD performs much better in the non-orthogonal scenario. By estimating all three PCs simultaneously, Manifold SFPCA is able to find a better set of PCs than any of the greedy deflation methods.

		MADMM	ManPG	A-ManPG
Time (s)		**19.46**	217.40	175.99
Suboptimality		**0**	4.58	12.23
rSS-Error	U	68.66%	74.03%	**64.54%**
	V	**36.85%**	50.69%	43.17%
TPR	U	87.75%	69.72%	**89.99%**
	V	**94.75%**	70.30%	89.87%
FPR	U	12.25%	30.28%	**10.01%**
	V	**5.25%**	29.70%	10.13%

Table 2: Comparison of Manifold-ADMM, Manifold Proximal Gradient, and Alternating Manifold Proximal Gradient approaches for Manifold SFPCA (2). MADMM is consistently more efficient and obtains better solutions, but both MADMM and A-ManPG perform well in terms of signal recovery, as measured by relative subspace recovery error (rSS-Error $= \|\hat{\boldsymbol{U}}\hat{\boldsymbol{U}}^T - \boldsymbol{U}^*(\boldsymbol{U}^*)^T\|/\|\hat{\boldsymbol{U}}_{\text{SVD}}\hat{\boldsymbol{U}}_{\text{SVD}}^T - \boldsymbol{U}^*(\boldsymbol{U}^*)^T\|$), true positive rate (TPR) and false positive rate (FPR).

		HD	PD	SD	ManSFPCA
Scenario 1: $\boldsymbol{U}^*$ and $\boldsymbol{V}^*$ Orthogonal					
CPVE	PC1	15.92%	21.05%	**21.87%**	
	PC2	22.21%	29.42%	**30.59%**	**37.12%**
	PC3	26.80%	35.57%	**37.09%**	
Scenario 2: $\boldsymbol{U}^*$ and $\boldsymbol{V}^*$ Not Orthogonal					
CPVE	PC1	8.85%	19.74%	**29.80%**	
	PC2	13.03%	28.30%	**39.87%**	**50.85%**
	PC3	16.16%	34.22%	**46.48%**	

Table 3: Cumulative Proportion of Variance Explained (CPVE) of Rank-One SFPCA with (normalized) Hotelling, Projection, and Schur Complement Deflation and of (order 3) Manifold SFPCA. SD gives the best CVPE of the iterative approaches and appears to be more robust to violations of orthogonality (Scenario 2). Manifold SFPCA outperforms the iterative methods in both scenarios.

5. DISCUSSION

We have introduced two practical extensions to Sparse and Functional PCA: first, we presented a multi-rank scheme which estimates multiple PCs simultaneously and proposed algorithms to solve the resulting manifold optimization problems. The resulting estimator inherits many of the attractive properties of rank-one SFPCA and is, to the best of our knowledge, the first multi-rank PCA scheme for the low-rank model. ManSFPCA combines both the flexiblity and superior statistical performance of rank-one SFPCA with the superior interpretability of orthogonal (non-regularized) PCA.

Secondly, we re-considered the use of Hotelling's deflation, and developed two additional deflation schemes which have attractive theoretical properties and emprical performance. Our schemes extend the results of Mackey [21] in several ways: they allow for deflation by multiple PCs in a single step, they are applicable to the low-rank model and the covariance model, and they are robust to non-orthogonality and non-unit-scaling common to Functional PCA variants. While developed for SFPCA, these deflation schemes are useful for any regularized PCA model.

We note here that our results can be extended naturally to other multivariate analysis techniques which can be expressed in a regularized SVD framework (*e.g.*, PLS, CCA, *etc.*). Our deflation approaches can be also extended to the higher-order / multi-way context and may be particularly useful in the context of regularized CP decompositions [22, 23]. In the tensor setting, ManSFPCA is a sparse and smooth version of a Tucker decomposition, which suggests several interesting extensions we leave for future work [24, 25].

6. REFERENCES

[1] H. Hotelling, "Analysis of a complex of statistical variables into principal components," *Journal of Educational Psychology*, vol. 24, no. 6, pp. 417–441, 1933. DOI: 10.1037/h0071325.

[2] G. I. Allen and M. Weylandt, "Sparse and functional principal components analysis," in *DSW 2019: Proceedings of the 2nd IEEE Data Science Workshop*, G. Karypis, G. Michailidis, and R. Willett, Eds., Minneapolis, Minnesota: IEEE, 2019, pp. 11–16. DOI: 10.1109/DSW.2019.8755778.

[3] B. W. Silverman, "Smoothed functional principal components analysis by choice of norm," *Annals of Statistics*, vol. 24, no. 1, pp. 1–24, 1996. DOI: 10.1214/aos/1033066196.

[4] H. Shen and J. Z. Huang, "Sparse principal component analysis via regularized low rank matrix approximation," *Journal of Multivariate Analysis*, vol. 99, no. 6, pp. 1015–1034, 2008. DOI: 10.1016/j.jmva.2007.06.007.

[5] J. Z. Huang, H. Shen, and A. Buja, "Functional principal components analysis via penalized rank one approximation," *Electronic Journal of Statistics*, vol. 2, pp. 678–695, 2008. DOI: 10.1214/08-EJS218.

[6] ——, "The analysis of two-way functional data using two-way regularized singular value decompositions," *Journal of the American Statistical Association*, vol. 104, no. 488, pp. 1609–1620, 2009. DOI: 10.1198/jasa.2009.tm08024.

[7] D. M. Witten, R. Tibshirani, and T. Hastie, "A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis," *Biostatistics*, vol. 10, no. 3, pp. 515–534, 2009. DOI: 10.1093/biostatistics/kxp008.

[8] G. I. Allen, L. Grosenick, and J. Taylor, "A generalized least-square matrix decomposition," *Journal of the American Statistical Association*, vol. 109, no. 505, pp. 145–159, 2014. DOI: 10.1080/01621459.2013.852978.

[9] M. Journée, Y. Nesterov, P. Richtárik, and R. Sepulchre, "Generalized power method for sparse principal component analysis," *Journal of Machine Learning Research*, vol. 11, pp. 517–553, 2010. Available at: http://www.jmlr.org/papers/v11/journee10a.html.

[10] P.-A. Absil, R. Mahony, and R. Sepulchre, *Optimization algorithms on matrix manifolds*. Princeton University Press, 2007, ISBN: 978-0-691-13298-3.

[11] Z. Wen and W. Yin, "A feasible method for optimization with orthogonality constraints," *Mathematical Programming*, vol. 142, no. 1-2, pp. 397–434, 2013. DOI: 10.1007/s10107-012-0584-1.

[12] R. Lai and S. Osher, "A splitting method for orthogonality constrained problems," *Journal of Scientific Computing*, vol. 58, no. 2, pp. 431–449, 2014. DOI: 10.1007/s10915-013-9740-x.

[13] A. Kovnatsky, K. Glashoff, and M. M. Bronstein, "MADMM: A generic algorithm for non-smooth optimization on manifolds," in *ECCV 2016: Proceedings of the 14th European Conference on Computer Vision*, B. Leibe, J. Matas, N. Sebe, and M. Welling, Eds., ser. Lecture Notes in Computer Science, vol. 9909, Springer, 2016, pp. 680–696. DOI: 10.1007/978-3-319-46454-1_41.

[14] S. Chen, S. Ma, A. M.-C. So, and T. Zhang, "Proximal gradient method for nonsmooth optimization over the Stiefel manifold," *ArXiv Pre-Print 1811.00980*, Available at: http://arxiv.org/abs/1811.00980.

[15] S. Chen, S. Ma, L. Xue, and H. Zou, "An alternating manifold proximal gradient method for sparse PCA and sparse CCA," *ArXiv Pre-Print 1903.11576*, 2019. Available at: http://arxiv.org/abs/1903.11576.

[16] K. Benidis, Y. Sun, P. Babu, and D. P. Palomar, "Orthogonal sparse PCA and covariance estimation via Procrustes reformulation," *IEEE Transactions on Signal Processing*, vol. 64, no. 23, pp. 6211–6226, 2016. DOI: 10.1109/TSP.2016.2605073.

[17] H. Sato and K. Aihara, "Cholesky QR-based retraction on the generalized Stiefel manifold," *Computational Optimization and Applications*, vol. 72, no. 2, pp. 293–308, 2019. DOI: 10.1007/s10589-018-0046-7.

[18] P. H. Schönemann, "A generalized solution of the orthogonal Procrustes problem," *Psychometrika*, vol. 31, no. 1, pp. 1–10, 1966. DOI: 10.1007/BF02289451.

[19] L. Eldén and H. Park, "A Procrustes problem on the Stiefel manifold," *Numerische Mathematik*, vol. 82, no. 4, pp. 599–619, 1999. DOI: 10.1007/s002110050432.

[20] A. Beck, *First-order methods in optimization*, ser. MOS-SIAM Series on Optimization. 2017. DOI: 10.1137/1.9781611974997.

[21] L. Mackey, "Deflation methods for sparse PCA," in *NIPS 2008: Advances in Neural Information Processing Systems 21*, D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, Eds., 2008, pp. 1017–1024. Available at: https://papers.nips.cc/paper/3575-deflation-methods-for-sparse-pca.

[22] G. I. Allen, "Sparse higher-order principal components analysis," in *AISTATS 2012: Proceedings of the 15th International Conference on Artificial Intelligence and Statistics*, N. D. Lawrence and M. Girolami, Eds., vol. 22, Canary Islands, Spain: PMLR, 2012, pp. 27–36. Available at: http://proceedings.mlr.press/v22/allen12.html.

[23] ——, "Multi-way functional principal components analysis," in *CAMSAP 2013: Proceedings of the 5th IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing*, V. Cevher and S. Gazor, Eds., St. Martin, France: IEEE, 2013, pp. 220–223. DOI: 10.1109/CAMSAP.2013.6714047.

[24] L. R. Tucker, "Some mathematical notes on three-mode factor analysis," *Psychometrika*, vol. 31, no. 3, pp. 279–311, 1966. DOI: 10.1007/BF02289464.

[25] T. G. Kolda and B. W. Bader, "Tensor decompositions and applications," *SIAM Review*, vol. 51, no. 3, pp. 455–500, 2009. DOI: 10.1137/07070111X.

Supplementary Materials

A. PROOFS

A.1. Deflation Schemes

In this section, we give proofs of the claimed properties for the deflation schemes discussed in Section 3 and summarized in Table 1.

We first consider Hotelling's deflation scheme (HD): given estimated left- and right-singular vectors $\boldsymbol{u}_t, \boldsymbol{v}_t$ of a n-by-p matrix $\boldsymbol{X}_{t-1}$, two-way orthogonality is immediate:

$$
\begin{aligned}
\boldsymbol{u}_t^T \boldsymbol{X}_t \boldsymbol{v}_t &= \boldsymbol{u}_t^T (\boldsymbol{X}_{t-1} - d_t \boldsymbol{u}_t \boldsymbol{v}_t^T) \boldsymbol{v}_t \\
&= \boldsymbol{u}_t^T \boldsymbol{X}_{t-1} \boldsymbol{v}_t - d_t \boldsymbol{u}^T \boldsymbol{u}_t \boldsymbol{v}_t^T \boldsymbol{v}_t \\
&= d_t - d_t \|\boldsymbol{u}_t\|^2 \|\boldsymbol{v}_t\|^2
\end{aligned}
$$

If $\|\boldsymbol{u}_t\|^2 = \|\boldsymbol{v}_t\|^2 = 1$, then this gives the desired two-way orthogonality. From here, we clearly see that Hotelling's deflation only gives two-way orthogonality if $\boldsymbol{u}_t$ and $\boldsymbol{v}_t$ are unit-scaled, showing that it is not robust to non-unit scaling. Turning our attention to one-way orthogonality, we give a counter-example using the data matrix

$$
\boldsymbol{X} = \begin{pmatrix} 2 & -4/3 \\ 2 & 2/3 \\ 1 & 4/3 \end{pmatrix} = \underbrace{\begin{pmatrix} 2/3 & -2/3 & 1/3 \\ 2/3 & 1/3 & -2/3 \\ 1/3 & 2/3 & 2/3 \end{pmatrix}}_{=U} \underbrace{\begin{pmatrix} 3 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 1 \end{pmatrix}}_{=D} \underbrace{\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix}^T}_{=V^T}.
$$

If we take take $\boldsymbol{u}_1 = \begin{pmatrix} 1/\sqrt{2} & 1/\sqrt{2} & 0 \end{pmatrix}^T$ as a sparse left singular vector and $\boldsymbol{v}_1 = \begin{pmatrix} 1 & 0 \end{pmatrix}^T$ the true right singular vector, then Hotelling's deflation gives $d_1 = \sqrt{8}$ and

$$
\boldsymbol{X}_1 = \begin{pmatrix} 0 & -4/3 \\ 0 & 2/3 \\ 1 & 4/3 \end{pmatrix}.
$$

The deflated $\boldsymbol{X}_1$ is not left-orthogonal to $\boldsymbol{u}_1$, however, as $\boldsymbol{u}_1^T \boldsymbol{X}_1 = \begin{pmatrix} 0 & -\sqrt{2}/3 \end{pmatrix}$, showing that Hotelling's deflation is not one-way orthogonal.

Next we consider the projection deflation scheme (PD). Left orthogonality can be shown explicitly:

$$
\begin{aligned}
\boldsymbol{u}_t^T \boldsymbol{X}_t &= \boldsymbol{u}_t^T (\boldsymbol{I}_n - \boldsymbol{u}_t \boldsymbol{u}_t^T) \boldsymbol{X}_{t-1} (\boldsymbol{I}_p - \boldsymbol{v}_t \boldsymbol{v}_t^T) \\
&= (\boldsymbol{u}_t^T - \|\boldsymbol{u}_t\|^2 \boldsymbol{u}_t^T) \boldsymbol{X}_{t-1} (\boldsymbol{I}_p - \boldsymbol{v}_t \boldsymbol{v}_t^T)
\end{aligned}
$$

which is clearly zero for arbitrary $\boldsymbol{v}_t$ if and only if $\|\boldsymbol{u}_t\| = 1$. Essentially the same argument shows that projection deflation is right orthogonal if and only if $\|\boldsymbol{v}_t\| = 1$, and two-way orthogonality follows from either left or right orthogonality. Hence projection deflation is both one- and two-way orthogonal but it is also sensitive to the unit scaling of the left and right principal components. Turning to subsequent orthogonality, we now take

$$
\boldsymbol{X} = \begin{pmatrix} -2 & -3/2 & 1 \\ 8/3 & 1/6 & 1/3 \\ 0 & 5/2 & 1 \\ 2/3 & 7/6 & 7/3 \end{pmatrix} = \underbrace{\begin{pmatrix} -1/2 & 1/2 & 1/2 & 1/2 \\ 1/2 & -1/2 & 1/2 & 1/2 \\ 1/2 & 1/2 & -1/2 & 1/2 \\ 1/2 & 1/2 & 1/2 & -1/2 \end{pmatrix}}_{=U} \underbrace{\begin{pmatrix} 4 & 0 & 0 & 0 \\ 0 & 3 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}}_{=D} \underbrace{\begin{pmatrix} 2/3 & 2/3 & 1/3 \\ -2/3 & 1/3 & 2/3 \\ 1/3 & -2/3 & 2/3 \\ 0 & 0 & 0 \end{pmatrix}}_{=V^T}.
$$

We take $\boldsymbol{u}_1 = \begin{pmatrix} 1/2 & 1/2 & 1/2 & 1/2 \end{pmatrix}^T$ and $\boldsymbol{v}_1 = \begin{pmatrix} 1/\sqrt{2} & 1/\sqrt{2} & 0 \end{pmatrix}^T$ as a pair of sparse PCs to yield the deflated matrix:

$$
\boldsymbol{X}_1 = \begin{pmatrix} -1/8 & 1/8 & -1/6 \\ 11/8 & -11/8 & -5/6 \\ -9/8 & 9/8 & -1/6 \\ -1/8 & 1/8 & 7/6 \end{pmatrix}
$$

The leading singular pair of $\boldsymbol{X}_1$ are approximately $\begin{pmatrix} -0.0447 & -0.1673 & 0.8620 & 0.4765 \end{pmatrix}^T$ and $\begin{pmatrix} 0.6731 & 0.2168 & -0.7071 \end{pmatrix}^T$, which we sparsely approximate by $\boldsymbol{u}_2 = \begin{pmatrix} 0 & 0 & 4/5 & 3/5 \end{pmatrix}^T$ and $\boldsymbol{v}_2 = \begin{pmatrix} 1/\sqrt{2} & 0 & 1/\sqrt{2} \end{pmatrix}$ respectively. Another round of PD gives

$$
\boldsymbol{X}_2 \approx \begin{pmatrix} 0.0417 & 0.2917 & 0 \\ 2.2083 & -0.5417 & 0 \\ 0.2750 & 0.9650 & 0 \\ -0.3667 & -1.2867 & 0 \end{pmatrix}
$$

which has $\boldsymbol{u}_1^T \boldsymbol{X}_2 \approx \begin{pmatrix} 1.0792 & -0.2858 & 0 \end{pmatrix}$ which is clearly non-zero.

Finally, we consider Schur complement deflation (SD). As with projection deflation, left- and right-orthogonality (and hence two-way orthogonality), can be shown explicitly:

$$\boldsymbol{u}_t^T \boldsymbol{X}_t = \boldsymbol{u}_t^T \left(\boldsymbol{X}_{t-1} - \frac{\boldsymbol{X}_{t-1} \boldsymbol{v}_t \boldsymbol{u}_t^T \boldsymbol{X}_{t-1}}{\boldsymbol{u}_t^T \boldsymbol{X}_{t-1} \boldsymbol{v}_t} \right)$$

$$= \boldsymbol{u}_t^T \boldsymbol{X}_{t-1} - \frac{(\boldsymbol{u}_t^T \boldsymbol{X}_{t-1} \boldsymbol{v}_t) \boldsymbol{u}_t^T \boldsymbol{X}_{t-1}}{\boldsymbol{u}_t^T \boldsymbol{X}_{t-1} \boldsymbol{v}_t}$$

$$= \boldsymbol{0}$$

which holds without any further restrictions on $\boldsymbol{u}_t$ or $\boldsymbol{v}_t$, showing that Schur complement deflation is robust to unit scaling of estimated principal components. Subsequent left orthogonality can be shown by induction using the above result as the base case: for the inductive step, suppose $\boldsymbol{u}_t^T \boldsymbol{X}_{t+s} = \boldsymbol{0}$ for some s, then

$$\boldsymbol{u}_t^T \boldsymbol{X}_{t+s+1} = \boldsymbol{u}_t^T \left(\boldsymbol{X}_{t+s} - \frac{\boldsymbol{X}_{t+s} \boldsymbol{v}_{t+s+1} \boldsymbol{u}_{t+s+1} \boldsymbol{X}_{t+s}}{\boldsymbol{u}_{t+s+1}^T \boldsymbol{X}_{t+s} \boldsymbol{v}_{t+s+1}} \right)$$

$$= \boldsymbol{u}_t^T \boldsymbol{X}_{t+s} - \frac{(\boldsymbol{u}_t^T \boldsymbol{X}_{t+s}) \boldsymbol{v}_{t+s+1} \boldsymbol{u}_{t+s+1} \boldsymbol{X}_{t+s}}{\boldsymbol{u}_{t+s+1}^T \boldsymbol{X}_{t+s} \boldsymbol{v}_{t+s+1}}$$

$$= \boldsymbol{0} - \frac{\boldsymbol{0} \boldsymbol{v}_{t+s+1} \boldsymbol{u}_{t+s+1} \boldsymbol{X}_{t+s}}{\boldsymbol{u}_{t+s+1}^T \boldsymbol{X}_{t+s} \boldsymbol{v}_{t+s+1}}$$

$$= \boldsymbol{0}$$

as desired.

A.2. Simultaneous Multi-Rank Deflation

In this section, we extend the results of the previous section to allow for simultaneous deflation by several principal components at once. The properties of the multi-rank deflation schemes are essentially the same as those of their counterparts. The normalized decompositions are given by:

$$\boldsymbol{X}_t^{\text{HD}} := \boldsymbol{X}_{t-1} - \boldsymbol{U}_t (\boldsymbol{U}_t^T \boldsymbol{U}_t)^{-1} \boldsymbol{U}_t^T \boldsymbol{X}_{t-1} \boldsymbol{V}_t (\boldsymbol{V}_t^T \boldsymbol{V}_t)^{-1} \boldsymbol{V}_t^T$$

$$\boldsymbol{X}_t^{\text{PD}} := (\boldsymbol{I}_n - \boldsymbol{U}_t (\boldsymbol{U}_t^T \boldsymbol{U}_t)^{-1} \boldsymbol{U}_t^T) \boldsymbol{X}_{t-1} (\boldsymbol{I}_p - \boldsymbol{V}_t (\boldsymbol{V}_t^T \boldsymbol{V}_t)^{-1} \boldsymbol{V}_t^T)$$

$$\boldsymbol{X}_t^{\text{SD}} := \boldsymbol{X}_{t-1} - \boldsymbol{X}_{t-1} \boldsymbol{V}_t (\boldsymbol{U}_t^T \boldsymbol{X}_{t-1} \boldsymbol{V}_t)^{-1} \boldsymbol{U}_t^T \boldsymbol{X}_{t-1}.$$

As before, we see that Hotelling's deflation gives two-way orthogonality:

$$\boldsymbol{U}_t^T \boldsymbol{X}_t^{\text{HD}} \boldsymbol{V}_t = \boldsymbol{U}_t^T \left[\boldsymbol{X}_{t-1} - \boldsymbol{U}_t (\boldsymbol{U}_t^T \boldsymbol{U}_t)^{-1} \boldsymbol{U}_t^T \boldsymbol{X}_{t-1} \boldsymbol{V}_t (\boldsymbol{V}_t^T \boldsymbol{V}_t)^{-1} \boldsymbol{V}_t^T \right] \boldsymbol{V}_t$$

$$= \boldsymbol{U}_t^T \boldsymbol{X}_{t-1} \boldsymbol{V}_t - \boldsymbol{U}_t^T \boldsymbol{U}_t (\boldsymbol{U}_t^T \boldsymbol{U}_t)^{-1} \boldsymbol{U}_t \boldsymbol{X}_{t-1} \boldsymbol{V}_t (\boldsymbol{V}_t^T \boldsymbol{V}_t)^{-1} \boldsymbol{V}_t^T \boldsymbol{V}_t$$

$$= \boldsymbol{U}_t^T \boldsymbol{X}_{t-1} \boldsymbol{V}_t - \boldsymbol{U}_t^T \boldsymbol{X}_{t-1} \boldsymbol{V}_t$$

$$= \boldsymbol{0}.$$

We do not get one-way orthogonality unless $\boldsymbol{V}_t (\boldsymbol{V}_t^T \boldsymbol{V}_t)^{-1} \boldsymbol{V}_t^T = \boldsymbol{I}$:

$$\boldsymbol{U}_t^T \boldsymbol{X}_t^{\text{HD}} = \boldsymbol{U}_t^T \left[\boldsymbol{X}_{t-1} - \boldsymbol{U}_t (\boldsymbol{U}_t^T \boldsymbol{U}_t)^{-1} \boldsymbol{U}_t^T \boldsymbol{X}_{t-1} \boldsymbol{V}_t (\boldsymbol{V}_t^T \boldsymbol{V}_t)^{-1} \boldsymbol{V}_t^T \right]$$

$$= \boldsymbol{U}_t^T \boldsymbol{X}_{t-1} - \boldsymbol{U}_t^T \boldsymbol{U}_t (\boldsymbol{U}_t^T \boldsymbol{U}_t)^{-1} \boldsymbol{U}_t \boldsymbol{X}_{t-1} \boldsymbol{V}_t (\boldsymbol{V}_t^T \boldsymbol{V}_t)^{-1} \boldsymbol{V}_t^T$$

$$= \boldsymbol{U}_t^T \boldsymbol{X}_{t-1} - \boldsymbol{U}_t^T \boldsymbol{X}_{t-1} \boldsymbol{V}_t (\boldsymbol{V}_t^T \boldsymbol{V}_t)^{-1} \boldsymbol{V}_t^T.$$

For projection deflation, we prove (left) one-way orthogonality as before:

$$\boldsymbol{U}_t^T \boldsymbol{X}_t^{\text{PD}} = \boldsymbol{U}_t^T \left[(\boldsymbol{I}_n - \boldsymbol{U}_t (\boldsymbol{U}_t^T \boldsymbol{U}_t)^{-1} \boldsymbol{U}_t^T) \boldsymbol{X}_{t-1} (\boldsymbol{I}_p - \boldsymbol{V}_t (\boldsymbol{V}_t^T \boldsymbol{V}_t)^{-1} \boldsymbol{V}_t^T) \right]$$

$$= (\boldsymbol{U}_t^T - \boldsymbol{U}_t^T \boldsymbol{U}_t (\boldsymbol{U}_t^T \boldsymbol{U}_t)^{-1} \boldsymbol{U}_t^T) \boldsymbol{X}_{t-1} (\boldsymbol{I}_p - \boldsymbol{V}_t (\boldsymbol{V}_t^T \boldsymbol{V}_t)^{-1} \boldsymbol{V}_t^T)$$

$$= (\boldsymbol{U}_t^T - \boldsymbol{U}_t^T) \boldsymbol{X}_{t-1} (\boldsymbol{I}_p - \boldsymbol{V}_t (\boldsymbol{V}_t^T \boldsymbol{V}_t)^{-1} \boldsymbol{V}_t^T)$$

$$= \boldsymbol{0}.$$

An essentially identical argument proves right one-way orthogonality. Either form of one-way orthogonality proves two-way orthogonality *a fortiori*. Considering (left) subsequent orthogonality,

$$\boldsymbol{U}_t^T \boldsymbol{X}_{t+1} = \boldsymbol{U}_t^T \left[(\boldsymbol{I} - \boldsymbol{U}_{t+1} (\boldsymbol{U}_{t+1}^T \boldsymbol{U}_{t+1})^{-1} \boldsymbol{U}_{t+1}^T) \boldsymbol{X}_t (\boldsymbol{I} - \boldsymbol{V}_{t+1} (\boldsymbol{V}_{t+1}^T \boldsymbol{V}_{t+1})^{-1} \boldsymbol{V}_{t+1}^T) \right]$$

we see that it cannot hold unless either $\boldsymbol{U}_t$ lies in the (left) range-space of $\boldsymbol{U}_{t+1}$ or commutes with $(\boldsymbol{I} - \boldsymbol{U}_{t+1}(\boldsymbol{U}_{t+1}^T\boldsymbol{U}_{t+1})^{-1}\boldsymbol{U}_{t+1}^T)$, neither of which hold in general. Interestingly, the matrix form makes the necessity of normalization clear for HD and PD: the unnormalized form is essentially assuming $\boldsymbol{U}_t^T\boldsymbol{U}_t = \boldsymbol{I}$, which will not hold unless we are in the non-smoothed case ($\alpha_{\boldsymbol{u}} = 0$).

Finally, for Schur complement deflation, we get left-orthogonality via the same argument as the vector case:

$$\begin{aligned}
\boldsymbol{U}_t^T\boldsymbol{X}_t^{\text{SD}} &= \boldsymbol{U}_t^T\left[\boldsymbol{X}_{t-1} - \boldsymbol{X}_{t-1}\boldsymbol{V}_t(\boldsymbol{U}_t^T\boldsymbol{X}_{t-1}\boldsymbol{V}_t)^{-1}\boldsymbol{U}_t^T\boldsymbol{X}_{t-1}\right] \\
&= \boldsymbol{U}_t^T\boldsymbol{X}_{t-1} - \boldsymbol{U}_t^T\boldsymbol{X}_{t-1}\boldsymbol{V}_t(\boldsymbol{U}_t^T\boldsymbol{X}_{t-1}\boldsymbol{V}_t)^{-1}\boldsymbol{U}_t^T\boldsymbol{X}_{t-1} \\
&= \boldsymbol{U}_t^T\boldsymbol{X}_{t-1} - \boldsymbol{I}\,\boldsymbol{U}_t^T\boldsymbol{X}_{t-1} \\
&= \boldsymbol{0}.
\end{aligned}$$

Taking this as the base case for subsequent orthogonality, we inductively assume $\boldsymbol{U}_t^T\boldsymbol{X}_{t+s} = \boldsymbol{0}$ and consider $\boldsymbol{U}_t^T\boldsymbol{X}_{t+s+1}$:

$$\begin{aligned}
\boldsymbol{U}_t^T\boldsymbol{X}_{t+s+1} &= \boldsymbol{U}_t^T\left[\boldsymbol{X}_{t+s} - \boldsymbol{X}_{t+s}\boldsymbol{V}_{t+s+1}(\boldsymbol{U}_{t+s+1}\boldsymbol{X}_{t+s}\boldsymbol{V}_{t+s+1})^{-1}\boldsymbol{U}_{t+s+1}^T\boldsymbol{X}_{t+s}\right] \\
&= \boldsymbol{U}_t^T\boldsymbol{X}_{t+s} - \boldsymbol{U}_t^T\boldsymbol{X}_{t+s}\boldsymbol{V}_{t+s+1}(\boldsymbol{U}_{t+s+1}\boldsymbol{X}_{t+s}\boldsymbol{V}_{t+s+1})^{-1}\boldsymbol{U}_{t+s+1}^T\boldsymbol{X}_{t+s} \\
&= \boldsymbol{0} - \boldsymbol{0}\,\boldsymbol{V}_{t+s+1}(\boldsymbol{U}_{t+s+1}\boldsymbol{X}_{t+s}\boldsymbol{V}_{t+s+1})^{-1}\boldsymbol{U}_{t+s+1}^T\boldsymbol{X}_{t+s} \\
&= \boldsymbol{0}
\end{aligned}$$

showing that SD indeed gives subsequent orthogonality. As with the vector case, SD is not sensitive to normalization, as indicated by the fact it has no $(\boldsymbol{U}_t^T\boldsymbol{U}_t)^{-1}$ or $(\boldsymbol{V}_t^T\boldsymbol{V}_t)^{-1}$ terms.

A.3. Solution of the Generalized Unbalanced Procrustes Problem

The $\boldsymbol{U}$-update in the Manifold ADMM scheme for the $\hat{\boldsymbol{U}}$-subproblem (Algorithm 3) requires us to solve the following problem:

$$\hat{\boldsymbol{U}} = \arg\min_{\boldsymbol{U}\in\mathcal{V}_{n\times k}^{\boldsymbol{S_u}}} -\operatorname{Tr}(\boldsymbol{U}^T\boldsymbol{X}\hat{\boldsymbol{V}}) + \frac{\rho}{2}\|\boldsymbol{U} - \boldsymbol{W}^{(k)} + \boldsymbol{Z}^{(k)}\|_{\boldsymbol{S_u}}^2$$

which has the closed-form solution:

$$\hat{\boldsymbol{U}} = \boldsymbol{S_u}^{-1/2}\boldsymbol{A}\boldsymbol{B}^T \text{ where } \boldsymbol{A}, \boldsymbol{\Delta}, \boldsymbol{B}^T = \mathsf{SVD}(\boldsymbol{S_u}^{-1/2}\boldsymbol{X}\hat{\boldsymbol{V}} + \rho\boldsymbol{S_u}^{1/2}(\boldsymbol{W}^{(k)} - \boldsymbol{Z}^{(k)}))$$

as shown by the following theorem.

Theorem 1. *Suppose $\boldsymbol{S_u}$ is a strictly positive definite $n \times n$ matrix and $\boldsymbol{A}, \boldsymbol{B}$ are full (column) rank matrices of size $n \times k$ for $k < n$. Then the solution to*

$$\hat{\boldsymbol{X}} = \arg\min_{\boldsymbol{X}\in\mathcal{V}_{n\times k}^{\boldsymbol{S_u}}} -\operatorname{Tr}(\boldsymbol{X}^T\boldsymbol{A}) + \frac{\rho}{2}\|\boldsymbol{X} - \boldsymbol{B}\|_{\boldsymbol{S_u}}^2$$

is given by

$$\boldsymbol{X} = \boldsymbol{S_u}^{-1/2}\boldsymbol{U}\boldsymbol{V}^T$$

where $\boldsymbol{U}\boldsymbol{D}\boldsymbol{V}^T$ is the (economical) SVD of $\boldsymbol{S_u}^{-1/2}\boldsymbol{A} + \rho\boldsymbol{S_u}^{1/2}\boldsymbol{B}$.

We note that this result is a (slight) generalization of the well-studied Procrustes problem first considered by Schönemann [1] in the orthogonal case and extended to the unbalanced (Stiefel manifold) case by Eldén and Park [2]. We modify their result to the generalized Stiefel manifold, though we do not consider the additional complexities associated with rank-deficient $\boldsymbol{A}, \boldsymbol{B}$, or $\boldsymbol{S_u}$ matrices as they do not apply to our problem.

Proof. Let $\boldsymbol{Y} = \boldsymbol{S_u}^{1/2}\boldsymbol{X}$ so that $\boldsymbol{X}^T\boldsymbol{S_u}\boldsymbol{X} \Leftrightarrow \boldsymbol{Y}^T\boldsymbol{Y} = \boldsymbol{I}$. Then we can rewrite the above problem as:

$$\begin{aligned}
\hat{\boldsymbol{X}} &= \arg\min_{\boldsymbol{X}\in\mathcal{V}_{n\times k}^{\boldsymbol{S_u}}} -\operatorname{Tr}(\boldsymbol{X}^T\boldsymbol{A}) + \frac{\rho}{2}\|\boldsymbol{X} - \boldsymbol{B}\|_{\boldsymbol{S_u}}^2 \\
&= \arg\min_{\boldsymbol{X}\in\mathcal{V}_{n\times k}^{\boldsymbol{S_u}}} -\operatorname{Tr}((\boldsymbol{S_u}^{1/2}\boldsymbol{X})^T(\boldsymbol{S_u}^{-1/2}\boldsymbol{A})) + \frac{\rho}{2}\|\boldsymbol{S_u}^{1/2}\boldsymbol{X} - \boldsymbol{S_u}^{1/2}\boldsymbol{B}\|_F^2 \\
\hat{\boldsymbol{Y}} &= \arg\min_{\boldsymbol{Y}\in\mathcal{V}_{n\times k}} -\operatorname{Tr}(\boldsymbol{Y}^T(\boldsymbol{S_u}^{-1/2}\boldsymbol{A})) + \frac{\rho}{2}\|\boldsymbol{Y} - \boldsymbol{S_u}^{1/2}\boldsymbol{B}\|_F^2
\end{aligned}$$

From here, we apply Lemma 1 to obtain:

$$\hat{\boldsymbol{Y}} = \boldsymbol{U}\boldsymbol{V}^T \text{ where } \boldsymbol{U}, \boldsymbol{D}, \boldsymbol{V}^T = \mathsf{SVD}(\boldsymbol{S_u}^{-1/2}\boldsymbol{A} + \rho\boldsymbol{S_u}^{1/2}\boldsymbol{B})$$

and hence

$$\hat{\boldsymbol{X}} = \boldsymbol{S_u}^{-1/2}\boldsymbol{U}\boldsymbol{V}^T \text{ where } \boldsymbol{U}, \boldsymbol{D}, \boldsymbol{V}^T = \mathsf{SVD}(\boldsymbol{S_u}^{-1/2}\boldsymbol{A} + \rho\boldsymbol{S_u}^{1/2}\boldsymbol{B})$$

$\square$

Lemma 1. *Suppose $\boldsymbol{A}, \boldsymbol{B}$ are full (column) rank matrices of size $n \times k$ for $k < n$. Then the solution to*

$$\hat{\boldsymbol{X}} = \underset{\boldsymbol{X} \in \mathcal{V}_{n \times k}}{\arg \min} - \operatorname{Tr}(\boldsymbol{X}^T \boldsymbol{A}) + \frac{\rho}{2} \|\boldsymbol{X} - \boldsymbol{B}\|_F^2$$

is given by

$$\hat{\boldsymbol{X}} = \boldsymbol{U} \boldsymbol{V}^T$$

where $\boldsymbol{U} \boldsymbol{D} \boldsymbol{V}^T$ is the (economical) SVD of $\boldsymbol{A} + \rho \boldsymbol{B}$.

Proof. Letting $\langle A, B \rangle_F = \operatorname{Tr}(A^T B)$ be the standard (Frobenius) inner product, the above problem becomes:

$$\begin{aligned}
\hat{\boldsymbol{X}} &= \underset{\boldsymbol{X} \in \mathcal{V}_{n \times k}}{\arg \min} - \operatorname{Tr}(\boldsymbol{X}^T \boldsymbol{A}) + \frac{\rho}{2} \|\boldsymbol{X} - \boldsymbol{B}\|_F^2 \\
&= \underset{\boldsymbol{X} \in \mathcal{V}_{n \times k}}{\arg \min} - \langle \boldsymbol{X}, \boldsymbol{A} \rangle_F + \frac{\rho}{2} \langle \boldsymbol{X} - \boldsymbol{B}, \boldsymbol{X} - \boldsymbol{B} \rangle_F \\
&= \underset{\boldsymbol{X} \in \mathcal{V}_{n \times k}}{\arg \min} - \langle \boldsymbol{X}, \boldsymbol{A} \rangle_F + \frac{\rho}{2} \underbrace{\|\boldsymbol{X}\|_F^2}_{=1} - \rho \langle \boldsymbol{X}, \boldsymbol{B} \rangle_F + \frac{\rho}{2} \underbrace{\|\boldsymbol{B}\|_F^2}_{\text{constant}} \\
&= \underset{\boldsymbol{X} \in \mathcal{V}_{n \times k}}{\arg \max} \langle \boldsymbol{X}, \boldsymbol{A} + \rho \boldsymbol{B} \rangle_F
\end{aligned}$$

Let $\boldsymbol{U} \boldsymbol{D} \boldsymbol{V}^T$ be the SVD of $\boldsymbol{A} + \rho \boldsymbol{B}$. Then

$$\langle \boldsymbol{X}, \boldsymbol{A} + \rho \boldsymbol{B} \rangle_F = \langle \boldsymbol{X}, \boldsymbol{U} \boldsymbol{D} \boldsymbol{V}^T \rangle_F = \langle \boldsymbol{U}^T \boldsymbol{X} \boldsymbol{V}, \boldsymbol{D} \rangle_F$$

Since $\boldsymbol{D}$ is a diagonal matrix, this is maximized when the left term is an identity matrix, as can be obtained by taking $\boldsymbol{X} = \boldsymbol{U} \boldsymbol{V}^T$. □

B. ALGORITHMIC DETAILS

In this section, we give additional details of the Algorithms used to solve the Manifold SFPCA (2) problem:

$$\underset{\boldsymbol{U} \in \mathcal{V}_{n \times k}^{\boldsymbol{S_u}}, \boldsymbol{V} \in \mathcal{V}_{p \times k}^{\boldsymbol{S_v}}}{\arg \max} \operatorname{Tr}(\boldsymbol{U}^T \boldsymbol{X} \boldsymbol{V}) - \lambda_{\boldsymbol{U}} P_{\boldsymbol{U}}(\boldsymbol{U}) - \lambda_{\boldsymbol{V}} P_{\boldsymbol{V}}(\boldsymbol{V})$$

B.1. Manifold Proximal Gradient

Manifold proximal gradient [3, 4] proceeds in two steps: first, a descent direction within the tangent space is identified; secondly, a step along along the descent direction is taken, with the step size chosen by Armijo-type [5] back-tracking. (The line search method we use is essentially that of Beck and Teboulle [6]; see also Parikh and Boyd [7, Section 4.2].) Because the step in the descent direction is taken in the ambient space rather than the along the manifold in question (*e.g.*, an unconstrained step in $\mathbb{R}^{n \times k}$ rather than a geodesic move along $\mathcal{V}_{n \times k}$), a retraction step is used to project back onto the manifold and preserve feasibility.

Both steps require further discussion: we first consider identifying the descent direction, which requires solving the following problem

$$\hat{\boldsymbol{D}}_{\boldsymbol{U}} = \underset{\boldsymbol{D}_{\boldsymbol{U}} \in \mathbb{R}^{n \times k}}{\arg \min} - \langle \boldsymbol{X} \hat{\boldsymbol{V}}, \boldsymbol{D}_{\boldsymbol{U}} \rangle_F + \lambda_{\boldsymbol{U}} P_{\boldsymbol{U}}(\boldsymbol{U}^{(k)} + \boldsymbol{D}_{\boldsymbol{U}}) \text{ subject to } \boldsymbol{D}_{\boldsymbol{U}}^T \boldsymbol{S_u} \boldsymbol{U}^{(k)} + (\boldsymbol{U}^{(k)})^T \boldsymbol{S_u} \boldsymbol{D}_{\boldsymbol{U}} = \boldsymbol{0}.$$

The constraint arises from the tangent space of the generalized Stiefel manifold [4, Appendix A.1], while the objective is essentially that of the overall problem. (Note that, because the smooth portion of the objective is already linear, we do not need to linearize it, unlike Chen *et al.* [3, 4].) The tangency constraint makes this problem non-trivial to solve, but it can be reformulated as a linearly constrained quadratic program by splitting $\boldsymbol{D}_{\boldsymbol{U}}$ into positive and negative parts and solved using standard approaches. Chen *et al.* [3, 4] recommend the use of a semi-smooth Newton method to solve this problem [8, 9], though we used the generic SDPT3 solver of Toh *et al.* [10–12] in our timing experiments.

Expanding Algorithm 1 with Algorithm 2 for the subproblems, we obtain Algorithm A1 below. Because Algorithm A1 is initialized at an infeasible pair $(\hat{\boldsymbol{U}}, \hat{\boldsymbol{V}})$ (unless $\alpha_{\boldsymbol{u}} = \alpha_{\boldsymbol{v}} = 0$) the first iteration almost always decreases the objective value: after that, however, each sub-problem typically increases the objective value. (If $\lambda_{\boldsymbol{u}} = \lambda_{\boldsymbol{v}} = 0$, this is guaranteed because the subproblems have a unique global optimum, but we cannot prove monotonicity in general.)

Chen *et al.* [4] suggested an interesting variant of this approach which requires only a single proximal gradient step at each iteration. Applying this approach to our problem yields Algorithm A2. Somewhat surprisingly, they are able to give convergence guarantees for this approach, even though it has both non-convex constraints and a non-convex (bi-convex) objective. As far as we know, this is the only one of our algorithms to have provable convergence.

B.2. Manifold ADMM

If we expand Algorithm 1 using Algorithm 3 (Manifold ADMM [13]) to solve the sub-problems, we obtain Algorithm A3. Note that the result of Section A.3 is used to solve the $U^{(k+1)}$ and $V^{(k+1)}$ updates appearing in Steps 2(a)(ii) and 2(b)(ii) respectively. Because Algorithm A3 is initialized at an infeasible pair $(\hat{U}, \hat{V})$ (unless $\alpha_u = \alpha_v = 0$) the first iteration almost always decreases the objective value: after that, however, each sub-problem typically increases the objective value. (If $\lambda_u = \lambda_v = 0$, this is guaranteed because the subproblems have a unique global optimum, but we cannot prove monotonicity in general.) Unlike Algorithms A1 and A2, Algorithm A3 contains an explicit thresholding step, which we have found improves convergence to exact zeros.

B.3. Additional Note on Identifiability

As written, the Manifold SFPCA problem (2) suffers from two forms of non-identifiability which may impede convergence, at least in the typical case with P_U and P_V elementwise ℓ_1-penalties [14]. In particular, if the columns of U and V are simultaneously permuted (*i.e.*, $U \to U S_k$ and $V \to V S_k$ for some permutation matrix S_k) or if the signs of rows of U and V are simultaneously flipped, the objective is unchanged.

To address the first ambiguity (permutation-invariance), the smooth term may be replaced by $\mathrm{Tr}(U^T X V D)$ where D is a diagonal matrix with elements $(1+\epsilon)^{k-1}, (1+\epsilon)^{k-2}, \ldots, 1$ to ensure that the leading PCs are indeed placed first. Alternatively, both problems may be addressed by post-processing the (U, V) iterates at each step and putting them in a canonical form. A simple canonicalization that we have found works well is to sort the columns of U lexographically by absolute value and set the signs so that the first column of U has as many positive elements as possible.

C. ADDITIONAL EXPERIMENTAL RESULTS

In this section, we give additional details of the simulations performed in Section 4.

The timing comparisons of the first simulation (Table 2) are rather sensitive to the specific sub-problem solvers and matrix decomposition subroutines used. For the problem considered, on average MADMM required 30,659 rank-3 SVDs before convergence; ManPG required 439 descent direction solves and 2,259 retractions; and A-ManPG required 366 descent direction solves and 1,952 retractions. Because the cost of an SVD and a retraction (a QR decomposition) are roughly comparable, it is clear that the cost of solving the descent direction subproblem dominates the proximal gradient schemes and that they will be quite sensitive to the solver used. Recently, Huang and Wei [15] proposed an acceleration scheme for ManPG which could be used here to improve convergence speeds.

Table 4 gives an extended version of Table 3, giving additional details on the true and false positive rates of each deflation scheme, as well as measuring subspace recovery performance. For ManSFPCA, the MADMM approach was used to compute the solution, and tuning parameters were fixed as $\lambda_u = \lambda_v = 1$ and $\alpha_u = \alpha_v = 3$. For the iterative deflation schemes, all four tuning parameters were chosen to maximize the BIC, using the adaptive tuning scheme recommended by Allen and Weylandt [16].

As discussed above, ManSFPCA is able to achieve a greater proportion of variance explained than any of the iterative deflation methods. In terms of subspace recovery, ManSFPCA is far more accurate than the iterative schemes or than an unregularized SVD (indicated by rSS-Error less than 100%). In terms of variable selection, the BIC scheme used to tune the iterative schemes performs as expected, yielding almost no false positives, at the expense of many false negatives (low TPR). ManSFPCA conversely with fixed tuning parameters has a few more false positives but recovers much more of the true signal than any of the deflation schemes, though this is somewhat sensitive to the choice of tuning parameters. In the orthogonal scenario, the different deflation schemes achieve essentially the same results: in the non-orthogonal scenario, the fuller deflation performed by PD and SD gives a better TPR than Hotelling's deflation.

| | | *Scenario 1: U^* and V^* Orthogonal* | | | | *Scenario 2: U^* and V^* Not Orthogonal* | | | |
		HD	PD	SD	ManSFPCA	HD	PD	SD	ManSFPCA
	PC1	15.92%	21.05%	**21.87%**		8.85%	19.74%	**29.80%**	
CPVE	PC2	22.21%	29.42%	**30.59%**	**37.12%**	13.03%	28.30%	**39.87%**	**50.85%**
	PC3	26.80%	35.57%	**37.09%**		16.16%	34.22%	**46.48%**	
rSS-Error	U	129.54%	129.55%	128.35%	**69.32%**	215.73%	206.30%	205.74%	**97.77%**
	V	143.01%	143.72%	141.15%	**36.98%**	211.15%	207.77%	204.38%	**78.26%**
TPR	U	54.57%	54.29%	54.63%	**87.72%**	11.31%	13.11%	15.28%	**88.91%**
	V	59.20%	58.84%	59.41%	**94.72%**	12.89%	15.33%	18.50%	**80.08%**
FPR	U	0.92%	**0.91%**	0.92%	12.28%	1.24%	**1.05%**	1.18%	11.09%
	V	0.61%	**0.60%**	0.60%	5.28%	2.80%	**2.72%**	2.79%	19.92%

Table 4: Extended Version of Table 3, comparing Cumulative Proportion of Variance Explained (CVPE), relative subspace recovery error (rSS-Error $= \|\hat{U}\hat{U}^T - U^*(U^*)^T\| / \|\hat{U}_{\mathrm{SVD}}\hat{U}_{\mathrm{SVD}}^T - U^*(U^*)^T\|$), true positive rate (TPR) and false positive rate (FPR) of Manifold SFPCA with Rank-One SFPCA using (normalized) Hotelling, Projection and Schur Complement Deflation.

D. ADDITIONAL BACKGROUND

The literature on regularized PCA variants is vast and we refer the reader to Appendix C of Allen and Weylandt [16] for a review.[2] As they note, the vast majority of regularized PCA methods use an iterative deflation approach, typically that of Hotelling [17], though the deflations of Mackey [18] can be applied to any method which uses the covariance model and the results of Section 3 can be used for any which uses the low-rank model. While mainly focusing on the rank-one case, Journée *et al.* [19] show how their modified power algorithm can be extended to simultaneously estimate several orthogonal PCs. Benidis *et al.* [20] give a clever MM-type algorithm for finding orthogonal sparse PCs, based on the iteratively reweighted ℓ_1-methods first proposed by Candès *et al.* [21] and extended to orthogonality constraints by Song *et al.* [22].

In the Bayesian context, the sparse probabilistic PCA model of Guan and Dy [23] extends the probabilistic PCA model of Tipping and Bishop [24] and allows for sparsity. This method allows for multiple sparse PCs to be estimated simultaneously, though the authors do not discuss specific difficulties associated with orthogonality. While Gibbs sampling from Stiefel-constrained posteriors poses relatively little additional difficulty, it is less obvious how to use more performant MCMC samplers on the Stiefel manifold, especially the popular NUTS variant of Hamiltonian Monte Carlo (HMC) [25–28]. Girolami and Calderhead [29] and Byrne and Girolami [30] propose geodesic variants of HMC, which extend the strong geometric foundations of (Euclidean) HMC [31] to arbitrary smooth manifolds, but they are not easy to apply in practice and robust software implementations have yet to be developed. Working around this, several authors have proposed methods to reparameterize the Stiefel manifold into an approximately Euclidean coordinate system and apply standard HMC, though the practicality of these methods is limited by the incompatible structures of the Stiefel manifold and Euclidean space [32, 33]. More recently, Jauch *et al.* [34] proposed a promising data augmentation scheme based on the polar decomposition which avoids these topological inconsistencies, at the cost of a slightly larger parameter space.

The use of manifold optimization techniques for estimating multiple regularized PCs simultaneously is a recent development, driven by recent developments in non-smooth manifold optimization, especially the proximal gradient scheme of Chen *et al.* [3, 4] and the Manifold-constrained ADMM of Kovnatsky *et al.* [13], as well as references therein. (Lai and Osher [35] give an interesting splitting method for Stiefel-constrained problems based on Bregman iteration techniques, but to the best of our knowledge, it has not yet been applied to sparse PCA formulations.) Theory for these methods is still under rapid development and several key results remain unproven: in particular, convergence of the Manifold ADMM [13] has not been established, though the analysis of Wang *et al.* [36] comes somewhat close.

On the contrary, smooth manifold optimization techniques are well-established and typically motivated by problems in physics and engineering, where conservation laws are expressed as manifold constraints. These techniques were (re-)popularized by the influential book of Absil *et al.* [37]. Wen and Yin [38] give a particularly nice algorithm which maintains feasibility by using a clever application of the Cayley transform. Ritchie *et al.* [39] used smooth manifold optimization to solve a supervised PCA problem. Interestingly, their approach is Grassmanian- rather than Stiefel-constrained because, in a non-sparse context, the specific coordinate system used is irrelevant. This allows them to use the Grassmanian-based techniques of Edelman *et al.* [40] and of Boumal *et al.* [41].

E. ADDITIONAL REFERENCES

[1] P. H. Schönemann, "A generalized solution of the orthogonal Procrustes problem," *Psychometrika*, vol. 31, no. 1, pp. 1–10, 1966. DOI: 10.1007/BF02289451.

[2] L. Eldén and H. Park, "A Procrustes problem on the Stiefel manifold," *Numerische Mathematik*, vol. 82, no. 4, pp. 599–619, 1999. DOI: 10.1007/s002110050432.

[3] S. Chen, S. Ma, A. M.-C. So, and T. Zhang, "Proximal gradient method for nonsmooth optimization over the Stiefel manifold," *ArXiv Pre-Print 1811.00980*, Available at: http://arxiv.org/abs/1811.00980.

[4] S. Chen, S. Ma, L. Xue, and H. Zou, "An alternating manifold proximal gradient method for sparse PCA and sparse CCA," *ArXiv Pre-Print 1903.11576*, 2019. Available at: http://arxiv.org/abs/1903.11576.

[5] L. Armijo, "Minimization of functions having Lipschitz continuous first partial derivatives," *Pacific Journal of Mathematics*, vol. 16, pp. 1–3, 1966. DOI: 10.2140/pjm.1966.16.1. Available at: https://projecteuclid.org/euclid.pjm/1102995080.

[6] A. Beck and M. Teboulle, "Gradient-based algorithms with applications to signal-recovery problems," in *Convex Optimization in Signal Processing and Communications*, D. P. Palomar and Y. C. Eldar, Eds., Cambridge University Press, 2010, pp. 42–88, ISBN: 978-0-521-76222-9. DOI: 10.1017/CBO9780511804458.003.

[7] N. Parikh and S. Boyd, "Proximal algorithms," *Foundations and Trends in Optimization*, vol. 1, pp. 127–239, 3 2013. DOI: 10.1561/2400000003.

[8] X. Li, D. Sun, and K.-C. Toh, "A highly efficient semismooth newton augmented lagrangian method for solving lasso problems," *SIAM Journal on Optimization*, vol. 28, no. 1, pp. 433–458, 2016. DOI: 10.1137/16M1097572.

[9] A. Ali, E. Wong, and J. Z. Kolter, "A semismooth Newton method for fast, generic convex programming," in *ICML 2017: Proceedings of the 34th International Conference on Machine Learning*, D. Precu and Y. W. Teh, Eds., vol. 70, Sydney, Australia: PMLR, 2017, pp. 70–79. Available at: http://proceedings.mlr.press/v70/ali17a.html.

[10] K.-C. Toh, M. J. Todd, and R. H. Tütüncü, "SDPT3 - a MATLAB software package for semidefinite programming, version 1.3," *Optimization Methods and Software*, vol. 11, no. 1-4, pp. 545–581, 1999. DOI: 10.1080/10556789908805762.

---

[2]Appendix C can be found in the online Supplementary Materials, available at https://arxiv.org/abs/1309.2895.

[11] R. H. Tütüncü, K.-C. Toh, and M. J. Todd, "Solving semidefinite-quadratic-linear programs using SDPT3," *Mathematical Programming*, vol. 95, no. 2, pp. 189–217, 2003. DOI: 10.1007/s10107-002-0347-5.

[12] K.-C. Toh, M. J. Todd, and R. H. Tütüncü, "On the implementation and usage of SDPT3 - a MATLAB software package for semidefinite-quadratic-linear programming, version 4.0," in *Handbook on Semidefinite, Conic and Polynomial Optimization*, ser. International Series in Operations Research & Management Science Frederick S. Hillier, M. F. Anjos and J. B. Lasserre, Eds., 2012, pp. 715–754, ISBN: 978-1-4614-0768-3. DOI: 10.1007/978-1-4614-0769-0_25.

[13] A. Kovnatsky, K. Glashoff, and M. M. Bronstein, "MADMM: A generic algorithm for non-smooth optimization on manifolds," in *ECCV 2016: Proceedings of the 14th European Conference on Computer Vision*, B. Leibe, J. Matas, N. Sebe, and M. Welling, Eds., ser. Lecture Notes in Computer Science, vol. 9909, Springer, 2016, pp. 680–696. DOI: 10.1007/978-3-319-46454-1_41.

[14] R. Tibshirani, "Regression shrinkage and selection via the lasso," *Journal of the Royal Statistical Society, Series B: Methodological*, vol. 58, no. 1, pp. 267–288, 1996. DOI: 10.1111/j.2517-6161.1996.tb02080.x.

[15] W. Huang and K. Wei, "Extending FISTA to Riemannian optimization for sparse PCA," *ArXiv Pre-Print 1909.05485*, 2019. Available at: http://arxiv.org/abs/1909.05485.

[16] G. I. Allen and M. Weylandt, "Sparse and functional principal components analysis," in *DSW 2019: Proceedings of the 2nd IEEE Data Science Workshop*, G. Karypis, G. Michailidis, and R. Willett, Eds., Minneapolis, Minnesota: IEEE, 2019, pp. 11–16. DOI: 10.1109/DSW.2019.8755778.

[17] H. Hotelling, "Analysis of a complex of statistical variables into principal components," *Journal of Educational Psychology*, vol. 24, no. 6, pp. 417–441, 1933. DOI: 10.1037/h0071325.

[18] L. Mackey, "Deflation methods for sparse PCA," in *NIPS 2008: Advances in Neural Information Processing Systems 21*, D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, Eds., 2008, pp. 1017–1024. Available at: https://papers.nips.cc/paper/3575-deflation-methods-for-sparse-pca.

[19] M. Journée, Y. Nesterov, P. Richtárik, and R. Sepulchre, "Generalized power method for sparse principal component analysis," *Journal of Machine Learning Research*, vol. 11, pp. 517–553, 2010. Available at: http://www.jmlr.org/papers/v11/journee10a.html.

[20] K. Benidis, Y. Sun, P. Babu, and D. P. Palomar, "Orthogonal sparse PCA and covariance estimation via Procrustes reformulation," *IEEE Transactions on Signal Processing*, vol. 64, no. 23, pp. 6211–6226, 2016. DOI: 10.1109/TSP.2016.2605073.

[21] E. J. Candès, M. B. Wakin, and S. P. Boyd, "Enhancing sparsity by reweighted ℓ_1 minimization," *Journal of Fourier Analysis and Applications*, vol. 14, no. 5-6, pp. 877–905, 2008. DOI: 10.1007/s00041-008-9045-x.

[22] J. Song, P. Babu, and D. P. Palomar, "Sparse generalized eigenvalue problem via smooth optimization," *IEEE Transactions on Signal Processing*, vol. 63, no. 7, pp. 1627–1642, 2015. DOI: 10.1109/TSP.2015.2394443.

[23] Y. Guan and J. Dy, "Sparse probabilistic principal component analysis," in *AISTATS 2009: Proceedings of 12th International Conference on Artificial Intelligence and Statistics*, D. van Dyk and M. Welling, Eds., vol. 5, Clearwater Beach, Florida, USA: PMLR, 2009, pp. 185–192. Available at: http://proceedings.mlr.press/v5/guan09a.html.

[24] M. E. Tipping and C. M. Bishop, *Journal of the Royal Statistical Society, Series B: Statistical Methodology*, vol. 61, no. 3, pp. 611–622, 1999. DOI: 10.1111/1467-9868.00196.

[25] R. M. Neal, "MCMC using Hamiltonian dynamics," in *Handbook of Markov Chain Monte Carlo*, ser. Handbooks of Modern Statistical Methods, S. Brooks, A. Gelman, G. L. Jones, and X.-L. Meng, Eds., 1st, Chapman & Hall/CRC, 2011, ch. 5, pp. 113–162, ISBN: 978-1-420-07941-8. Available at: http://www.mcmchandbook.net/HandbookChapter5.pdf.

[26] M. D. Hoffman and A. Gelman, "The No-U-Turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo," *Journal of Machine Learning Research*, vol. 15, pp. 1593–1623, 2014. Available at: http://jmlr.org/papers/v15/hoffman14a.html.

[27] B. Carpenter, A. Gelman, M. D. Hoffman, D. Lee, B. Goodrich, M. Betancourt, M. Brubaker, J. Guo, P. Li, and A. Riddell, "Stan: A probabilistic programming language," *Journal of Statistical Software*, vol. 76, no. 1, 2017. DOI: 10.18637/jss.v076.i01.

[28] M. Betancourt, "A conceptual introduction to Hamiltonian Monte Carlo," *ArXiv 1701.02434*, 2017. Available at: https://arxiv.org/abs/1701.02434.

[29] M. Girolami and B. Calderhead, "Riemann manifold Langevin and Hamiltonian Monte Carlo methods," *Journal of the Royal Statistical Society, Series B: Statistical Methodology*, vol. 73, no. 2, pp. 123–214, 2011. DOI: 10.1111/j.1467-9868.2010.00765.x.

[30] S. Byrne and M. Girolami, "Geodesic Monte Carlo on embedded manifolds," *Scandanavian Journal of Statistics*, vol. 40, no. 4, pp. 825–845, 2013. DOI: 10.1111/sjos.12036.

[31] M. Betancourt, S. Byrne, S. Livingstone, and M. Girolami, "The geometric foundations of Hamiltonian Monte Carlo," *Bernoulli*, vol. 23, no. 4A, pp. 2257–2298, 2017. DOI: 10.3150/16-BEJ810.

[32] A. A. Pourzanjani, R. M. Jiang, B. Mitchell, P. J. Atzberger, and L. R. Petzold, "General Bayesian inference over the Stiefel manifold via the Givens representation," *ArXiv Pre-Print 1710.09443*, 2017. Available at: http://arxiv.org/abs/1710.09443.

[33] M. Jauch, P. D. Hoff, and D. B. Dunson, "Random orthogonal matrices and the Cayley transform," *ArXiv Pre-Print 1810.02881*, 2018. Available at: http://arxiv.org/abs/1810.02881.

[34] ——, "Monte Carlo simulation on the Stiefel manifold via polar expansion," *ArXiv Pre-Print 1906.07684*, 2019. Available at: http://arxiv.org/abs/1906.07684.

[35] R. Lai and S. Osher, "A splitting method for orthogonality constrained problems," *Journal of Scientific Computing*, vol. 58, no. 2, pp. 431–449, 2014. DOI: 10.1007/s10915-013-9740-x.

[36] Y. Wang, W. Yin, and J. Zeng, "Global convergence of ADMM in nonconvex nonsmooth optimization," *Journal of Scientific Computing*, vol. 78, no. 1, pp. 29–63, 2019. DOI: 10.1007/s10915-018-0757-z.

[37] P.-A. Absil, R. Mahony, and R. Sepulchre, *Optimization algorithms on matrix manifolds*. Princeton University Press, 2007, ISBN: 978-0-691-13298-3.

[38] Z. Wen and W. Yin, "A feasible method for optimization with orthogonality constraints," *Mathematical Programming*, vol. 142, no. 1-2, pp. 397–434, 2013. DOI: 10.1007/s10107-012-0584-1.

[39] A. Ritchie, C. Scott, L. Balzano, D. Kessler, and C. S. Sripada, "Supervised principal component analysis via manifold optimization," in *DSW 2019: Proceedings of the 2nd IEEE Data Science Workshop*, G. Karypis, G. Michailidis, and R. Willett, Eds., Minneapolis, Minnesota: IEEE, 2019, pp. 6–10. DOI: 10.1109/DSW.2019.8755587.

[40] A. Edelman, T. A. Arias, and S. T. Smith, "The geometry of algorithms with orthogonality constraints," *SIAM Journal on Matrix Analysis and Applications*, vol. 20, no. 2, pp. 303–353, 1998. DOI: 10.1137/S0895479895290954.

[41] N. Boumal, P.-A. Absil, and C. Cartis, "Global rates of convergence for nonconvex optimization on manifolds," *IMA Journal of Numerical Analysis*, vol. 39, no. 1, pp. 1–33, 2019. DOI: 10.1093/imanum/drx080.

Algorithm A1 Alternating Maximization Approach for Manifold SFPCA (2) using Manifold Proximal Gradient for Subproblem Solutions

---

1. Initialize $\hat{U}$ and $\hat{V}$ as the k leading singular vectors of X

2. Repeat Until Convergence:

   (a) Solve U-Subproblem using Manifold Proximal Gradient [3]:

   $$\hat{U} = \arg\min_{U \in \mathcal{V}_{n \times k}^{S_u}} - \text{Tr}(U^T X \hat{V}) + \lambda_U P_U(U)$$

   i. Initialize $U^{(k)} = \hat{U}$

   ii. Repeat Until Convergence:

   A. Determine Descent Direction:

   $$\hat{D}_U = \arg\min_{D_U \in \mathbb{R}^{n \times k}} - \text{Tr}(D_U^T X \hat{V}) + \lambda_U \|U^{(k)} + D_U\|_1 \text{ subject to } D_U^T S_u U^{(k)} + (U^{(k)})^T S_u D_U = 0$$

   B. Perform Backtracking to Determine Step Size:
   - Set $\alpha = 1$
   - While $- \text{Tr}(\text{Retr}_{U^{(k)}}(\alpha D_U)^T X \hat{V}) + \lambda_U \| \text{Retr}_{U^{(k)}}(\alpha D_U)\|_1 > - \text{Tr}((U^{(k)})^T X \hat{V}) + \lambda_U \|U^{(k)}\|_1$:
     - Set $\alpha = 0.8 * \alpha$

   C. Set $U^{(k+1)} = \text{Retr}_{U^{(k)}}(\alpha D_U)$

   iii. Set $\hat{U} = U^{(k)}$

   (b) Solve V-Subproblem using Manifold Proximal Gradient [3]:

   $$\hat{V} = \arg\min_{V \in \mathcal{V}_{p \times k}^{S_v}} - \text{Tr}(\hat{U}^T X V) + \lambda_V P_V(V)$$

   i. Initialize $V^{(k)} = \hat{V}$

   ii. Repeat Until Convergence:

   A. Determine Descent Direction:

   $$\hat{D}_V = \arg\min_{D_V \in \mathbb{R}^{p \times k}} - \text{Tr}(\hat{U}^T X D_V) + \lambda_V \|V^{(k)} + D_V\|_1 \text{ subject to } D_V^T S_v V^{(k)} + (V^{(k)})^T S_v D_V = 0$$

   B. Perform Backtracking to Determine Step Size:
   - Set $\alpha = 1$
   - While $- \text{Tr}(\hat{U}^T X \text{Retr}_{V^{(k)}}(\alpha D_U)) + \lambda_V \| \text{Retr}_{V^{(k)}}(\alpha D_V)\|_1 > - \text{Tr}(\hat{U} X V^{(k)}) + \lambda_V \|V^{(k)}\|_1$:
     - Set $\alpha = 0.8 * \alpha$

   C. Set $V^{(k+1)} = \text{Retr}_{V^{(k)}}(\alpha D_V)$

   (c) Set $\hat{V} = V^{(k)}$

3. Return $\hat{U}$ and $\hat{V}$

Algorithm A2 Alternating Manifold Proximal Gradient Approach [4] for Manifold SFPCA (2)

---

1. Initialize $\hat{U}$ and $\hat{V}$ as the k leading singular vectors of X

2. Repeat Until Convergence:

   (a) U-Update: One Step of Manifold Proximal Gradient [3]:

   $$\hat{U} = \arg\min_{U \in \mathcal{V}_{n \times k}^{S_u}} -\operatorname{Tr}(U^T X \hat{V}) + \lambda_U P_U(U)$$

        i. Determine Descent Direction:

   $$\hat{D}_U = \arg\min_{D_U \in \mathbb{R}^{n \times k}} -\operatorname{Tr}(D_U^T X \hat{V}) + \lambda_U \|\hat{U} + D_U\|_1 \text{ subject to } D_U^T S_u \hat{U} + \hat{U}^T S_u D_U = 0$$

        ii. Perform Backtracking to Determine Step Size:
   - Set $\alpha = 1$
   - While $-\operatorname{Tr}(\operatorname{Retr}_{\hat{U}}(\alpha D_U)^T X \hat{V}) + \lambda_U \|\operatorname{Retr}_{\hat{U}}(\alpha D_U)\|_1 > -\operatorname{Tr}(\hat{U}^T X \hat{V}) + \lambda_U \|\hat{U}\|_1$:
     - Set $\alpha = 0.8 * \alpha$

        iii. Set $\hat{U} = \operatorname{Retr}_{\hat{U}}(\alpha D_U)$

   (b) V-Update: One Step of Manifold Proximal Gradient [3]:

   $$\hat{V} = \arg\min_{V \in \mathcal{V}_{p \times k}^{S_v}} -\operatorname{Tr}(\hat{U}^T X V) + \lambda_V P_V(V)$$

        i. Determine Descent Direction:

   $$\hat{D}_V = \arg\min_{D_V \in \mathbb{R}^{p \times k}} -\operatorname{Tr}(\hat{U}^T X D_V) + \lambda_V \|\hat{V} + D_V\|_1 \text{ subject to } D_V^T S_v \hat{V} + \hat{V}^T S_v D_V = 0$$

        ii. Perform Backtracking to Determine Step Size:
   - Set $\alpha = 1$
   - While $-\operatorname{Tr}(\hat{U}^T X \operatorname{Retr}_{\hat{V}}(\alpha D_V)) + \lambda_V \|\operatorname{Retr}_{\hat{V}}(\alpha D_V)\|_1 > -\operatorname{Tr}(\hat{U}^T X \hat{V}) + \lambda_V \|\hat{V}\|_1$:
     - Set $\alpha = 0.8 * \alpha$

        iii. Set $\hat{V} = \operatorname{Retr}_{\hat{V}}(\alpha D_V)$

3. Return $\hat{U}$ and $\hat{V}$

Algorithm A3 Alternating Maximization Approach for Manifold SFPCA (2) using Manifold ADMM for Subproblem Solutions

1. Initialize $\hat{U}$ and $\hat{V}$ as the k leading singular vectors of X

2. Repeat Until Convergence:

   (a) Solve U-Subproblem using Manifold ADMM [13]:

   $$\hat{U} = \underset{U \in \mathcal{V}_{n \times k}^{S_u}}{\arg\min} - \text{Tr}(U^T X \hat{V}) + \lambda_U P_U(U) = \underset{U, W_U \in \mathcal{V}_{n \times k}^{S_u}}{\arg\min} - \text{Tr}(U^T X \hat{V}) + \lambda_U P_U(W_U) \text{ subject to } U = W_U$$

   i. Initialize $U^{(k)} = \hat{U}$ and restore $W_U^{(k)}, Z_U^{(k)}$ from previous iteration if available else set $W_U^{(k)} = U^{(k)}$ and $Z_U^{(k)} = 0_{n \times k}$

   ii. Repeat Until Convergence:

   $$U^{(k+1)} = \underset{U \in \mathcal{V}_{n \times k}^{S_u}}{\arg\min} - \text{Tr}(U^T X \hat{V}) + \frac{\rho}{2} \|U - W_U^{(k)} + Z_U^{(k)}\|_F^2$$

   $$= S_u^{-1/2} AB^T \text{ where } A, \Delta, B^T = \text{SVD}(S_u^{-1/2} X \hat{V} + \rho S_u^{1/2}(W_U^{(k)} - Z_U^{(k)}))$$

   $$W_U^{(k+1)} = \underset{W_U \in \mathbb{R}^{n \times k}}{\arg\min} \lambda_U P_U(W_U) + \frac{\rho}{2} \|U^{(k+1)} - W_U + Z_U^{(k)}\|_F^2$$

   $$= \text{prox}_{\lambda_U / \rho P_U(\cdot)} \left(U^{(k+1)} + Z_U^{(k)} \right)$$

   $$Z_U^{(k+1)} = Z_U^{(k)} + U^{(k+1)} - W_U^{(k+1)}$$

   iii. Set $\hat{U} = U^{(k)}$

   (b) Solve V-Subproblem using Manifold ADMM [13]:

   $$\hat{V} = \underset{V \in \mathcal{V}_{p \times k}^{S_v}}{\arg\min} - \text{Tr}(\hat{U}^T X V) + \lambda_V P_V(V) = \underset{V, W_V \in \mathcal{V}_{p \times k}^{S_v}}{\arg\min} - \text{Tr}(\hat{U}^T X V) + \lambda_V P_V(W_V) \text{ subject to } V = W_V$$

   i. Initialize $V^{(k)} = \hat{V}$ and restore $W_V^{(k)}, Z_V^{(k)}$ from previous iteration if available else set $W_V^{(k)} = V^{(k)}$ and $Z_V^{(k)} = 0_{p \times k}$

   ii. Repeat Until Convergence:

   $$V^{(k+1)} = \underset{V \in \mathcal{V}_{p \times k}^{S_v}}{\arg\min} - \text{Tr}(\hat{U}^T X V) + \frac{\rho}{2} \|V - W_V^{(k)} + Z_V^{(k)}\|_F^2$$

   $$= S_v^{-1/2} AB^T \text{ where } A, \Delta, B^T = \text{SVD}(S_v^{-1/2} X^T \hat{U} + \rho S_v^{1/2}(W_V^{(k)} - Z_V^{(k)}))$$

   $$W_V^{(k+1)} = \underset{W_V \in \mathbb{R}^{p \times k}}{\arg\min} \lambda_V P_V(V) + \frac{\rho}{2} \|V^{(k+1)} - W_V + Z_V^{(k)}\|_F^2$$

   $$= \text{prox}_{\lambda_V / \rho P_V(\cdot)} \left(V^{(k+1)} + Z_V^{(k)} \right)$$

   $$Z_V^{(k+1)} = Z_V^{(k)} + V^{(k+1)} - W_V^{(k+1)}$$

   (c) Set $\hat{V} = V^{(k)}$

3. Return $\hat{U}$ and $\hat{V}$