

Fast and Efficient MMD-based Fair PCA via Optimization over Stiefel Manifold

Junghyun Lee,¹ Gwangsu Kim*,² Matt Olfat,^{3,4} Mark Hasegawa-Johnson,⁵ Chang D. Yoo*²

¹ Kim Jaechul Graduate School of AI, KAIST, Seoul, Republic of Korea

² School of Electrical Engineering, KAIST, Daejeon, Republic of Korea

³ UC Berkeley IEOR, Berkely, CA, USA

⁴ Citadel, Chicago, IL, USA

⁵ Department of Electrical and Computer Engineering, University of Illinois Urbana-Champaign, IL, USA
{jh_lee00, s88012, cd_yoo}@kaist.ac.kr, molfat@berkeley.edu, jhasegaw@illinois.edu

Abstract

This paper defines fair principal component analysis (PCA) as minimizing the maximum mean discrepancy (MMD) between dimensionality-reduced conditional distributions of different protected classes. The incorporation of MMD naturally leads to an exact and tractable mathematical formulation of fairness with good statistical properties. We formulate the problem of fair PCA subject to MMD constraints as a non-convex optimization over the Stiefel manifold and solve it using the Riemannian Exact Penalty Method with Smoothing (REPMS; Liu and Boumal, 2019). Importantly, we provide local optimality guarantees and explicitly show the theoretical effect of each hyperparameter in practical settings, extending previous results. Experimental comparisons based on synthetic and UCI datasets show that our approach outperforms prior work in explained variance, fairness, and runtime.

1 Introduction

It has become increasingly evident that many widely-deployed machine learning algorithms are biased, yielding outcomes that can be discriminatory across key groupings such as race, gender and ethnicity (Mehrabi et al. 2019). As the applications of these algorithms proliferate in protected areas like healthcare (Karan et al. 2012), hiring (Chien and Chen 2008) and criminal justice (Kirchner et al. 2016), this creates the potential for further exacerbating social biases. To address this, there has recently been a surge of interest in ensuring fairness in resulting machine learning algorithms.

Working in high-dimensional spaces can be undesirable as the curse of dimensionality manifests in the form of data sparsity and computational intractability. Various dimensionality reduction algorithms are deployed to resolve these issues, and Principal Component Analysis (PCA) (Pearson 1901; Hotelling 1933), is one of the most popular methods (Jolliffe and Cadima 2016). One particular advantage of PCA is that there’s no need to train a complex neural network.

In this work, fair PCA is defined as performing dimensionality reduction while minimizing the difference in the conditional distributions of projections of different protected groups. Here, the projected data can be considered as a dimension-reduced fair representation of the input data

(Zemel et al. 2013). We answer the questions of 1) how fairness should be defined for PCA and 2) how to (algorithmically) incorporate fairness into PCA in a *fast* and *efficient* manner. This work takes a different approach from prior studies on PCA fairness (Samadi et al. 2018; Olfat and Aswani 2019), which is discussed in Section 3 and 7.

Our main contributions are as follows:

- We motivate a new mathematical definition of fairness for PCA using the maximum-mean discrepancy (MMD), which can be evaluated in a computationally efficient manner from the samples while guaranteeing asymptotic consistency. Such properties were not available in the previous definition of fair PCA (Olfat and Aswani 2019). This is discussed in detail in Section 3 and 4.
- We formulate the task of performing MMD-based fair PCA as a constrained optimization over the Stiefel manifold and propose using REPMS (Liu and Boumal 2019). For the first time, we prove two general theoretical guarantees of REPMS regarding the local minimality and feasibility. This is discussed in detail in Section 5 and 6.
- Using synthetic and UCI datasets, we verify the efficacy of our approach in terms of explained variance, fairness, and runtime. Furthermore, we verify that using fair PCA does indeed result in a fair representation, as in (Zemel et al. 2013). This is discussed in detail in Section 8.

2 Preliminaries

2.1 Notations

For $b \geq 1$, let \mathcal{P}_b be the set of all Borel probability measures defined on \mathbb{R}^b . For some measurable function $\Pi : \mathbb{R}^p \rightarrow \mathbb{R}^d$ and a measure $P \in \mathcal{P}_p$, the push-forward measure of P via Π is the probability measure $\Pi_{\#}P \in \mathcal{P}_d$, defined as $(\Pi_{\#}P)(S) = P(\Pi^{-1}(S))$ for any Borel set S . Let $\mathbf{0}$ and $\mathbf{1}$ denote matrices (or vectors) of zeros and ones of appropriate size, respectively. In this work, we focus on binary cases, i.e., we assume that the protected attribute A and outcome Y are binary ($A, Y \in \{0, 1\}$). We abbreviate demographic parity, equalized opportunity, and equalized odds as DP, EOP, and EOD, respectively.

2.2 Maximum Mean Discrepancy (MMD)

Let $k : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ be a positive-definite kernel function, and \mathcal{H}_k be a unit-ball in the RKHS generated by k . We

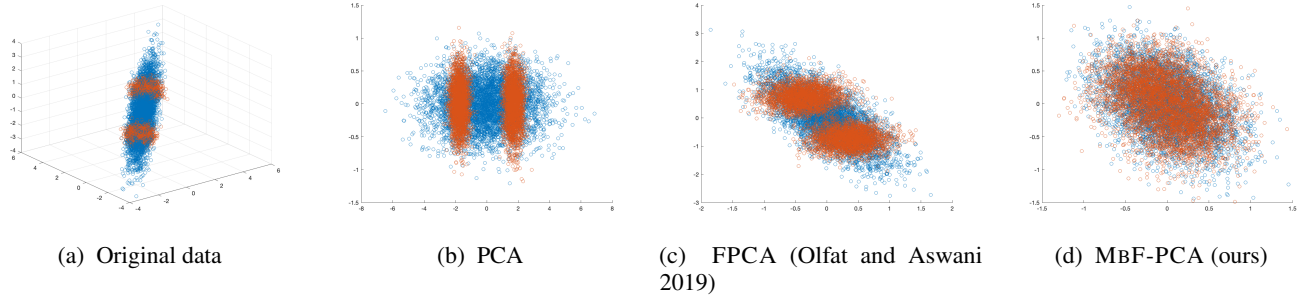


Figure 1: Synthetic data #1: Comparison of PCA, FPCA, and MBF-PCA on data composed of two groups with same mean and covariance, but different distributions. Blue and orange represent different protected groups.

impose some regularity assumptions on k :

Assumption 1. k is measurable, and bounded i.e. $K := \sup_{x,y} k(x,y) < \infty$.

Then one can pseudo-metrize \mathcal{P}_d by the following distance:

Definition 1 (Gretton et al., 2007). Given $\mu, \nu \in \mathcal{P}_d$, their **maximum mean discrepancy (MMD)**, denoted as $\text{MMD}_k(\mu, \nu)$, is a pseudo-metric on \mathcal{P}_d , defined as follows:

$$\text{MMD}_k(\mu, \nu) := \sup_{f \in \mathcal{H}_k} \left| \int_{\mathbb{R}^d} f d(\mu - \nu) \right| \quad (1)$$

As our fairness constraint involves exactly matching the considered distributions using MMD, we require the property of $\text{MMD}_k(\mu, \nu) = 0$ implying $\mu = \nu$. Any kernel k that satisfies such property is said to be **characteristic** (Fukumizu et al. 2008) to \mathcal{P}_d . Furthermore, Sriperumbudur et al. (2008) defined and characterized *stationary* characteristic kernels and identified that well-known kernels such as RBF and Laplace are characteristic. Based on this fact, we set k to be the RBF kernel $k_{r,bf}(x, y) := \exp(-\|x - y\|^2 / 2\sigma^2)$.

For the choice of bandwidth σ , the median of the set of pairwise distances of the samples after *vanilla PCA* is considered following the *median heuristic* of (Schölkopf, Smola, and Müller 1998; Ramdas et al. 2015). For simplicity, we refer to $\text{MMD}_{k_{r,bf}}$ as MMD.

Benefits of MMD There are several reasons for using MMD as the distance on a space of probability measures. First, it can act as a distance between distributions with different, or even disjoint, supports. This is especially crucial as the empirical distributions are often discrete and completely disjoint. Such a property is not generally true, one prominent example being the KL-divergence. Second, since many problems in fairness involve comparing two distributions, MMD has already been used in much of the fairness literature as a metric (Madras et al. 2018; Adel et al. 2019) and as an explicit constraint/penalty (Quadrianto and Sharmanska 2017; Louizos et al. 2016; Prost et al. 2019; Oneto et al. 2020; Jung et al. 2021), among other usages.

2.3 Fairness for supervised learning

The fair PCA discussed above should ultimately lead to fairness in supervised learning tasks based on the dimension-

reduced data with minimal loss in performance. Let us now review three of the most widely-used definitions of fairness in supervised learning, as formulated in (Madras et al. 2018). Let $(Z, Y, A) \in \mathbb{R}^d \times \{0, 1\} \times \{0, 1\}$ be the joint distribution of the dimensionality-reduced data, (downstream task) label, and protected attribute. Furthermore, let $g : \mathbb{R}^d \rightarrow \{0, 1\}$ be a classifier that outputs prediction \hat{Y} for Y from Z . We want to determine the fairness of a well-performing classifier g w.r.t. protected attribute A .

First, let D_s be the probability measure of $Z_s \triangleq Z|A = s$ for $s \in \{0, 1\}$:

Definition 2 (Feldman et al., 2015). g is said to satisfy **demographic parity (DP) up to Δ_{DP}** w.r.t. A with $\Delta_{DP} \triangleq |\mathbb{E}_{x \sim D_0}[g(x)] - \mathbb{E}_{x \sim D_1}[g(x)]|$.

Now, let $D_{s,y}$ be the probability measure of $Z_s \triangleq Z|A = s, Y = y$ for $s, y \in \{0, 1\}$.

Definition 3 (Hardt, Price, and Srebro, 2016). g is said to satisfy **equalized opportunity (EOP) up to Δ_{EOP}** w.r.t. A and Y with $\Delta_{EOP} \triangleq |\mathbb{E}_{x \sim D_{0,1}}[g(x)] - \mathbb{E}_{x \sim D_{1,1}}[g(x)]|$.

Definition 4 (Hardt, Price, and Srebro, 2016). g is said to satisfy **equalized odds (EOD) up to Δ_{EOD}** w.r.t. A and Y with $\Delta_{EOD} \triangleq \max_{y \in \{0,1\}} |\mathbb{E}_{x \sim D_{0,y}}[g(x)] - \mathbb{E}_{x \sim D_{1,y}}[g(x)]|$.

From hereon, we refer to such $\Delta_f(g)$ as the **fairness metric of $f \in \{DP, EOP, EOD\}$ w.r.t. g** , respectively.

3 New definition of fairness for PCA

For $p > d$, let \mathbb{R}^d be the space onto which data will be projected. A dimensionality reduction is a map $\Pi : \mathbb{R}^p \rightarrow \mathbb{R}^d$, and PCA is defined as $\Pi(x) = V^\top x$ for some $V \in \mathbb{R}^{p \times d}$ satisfying¹ $V^\top V = \mathbb{I}_d$ i.e. PCA is a linear, orthogonal dimensionality-reduction. From hereon, we denote a linear PCA as the mapping Π .

Here, we consider a rather canonical notion of fairness; a projection (which includes PCA) is considered to be *fair* when the projected data distributions of two protected classes are the same, or at least similar. This concept was first considered by (Olfat and Aswani 2019). As we discuss later

¹The benefits of pursuing orthogonality in the loading matrix, and thus the resulting PCs, are already well-studied; for example, see (Qi, Luo, and Zhao 2013; Benidis et al. 2016).

(Section 3.1 and 5), their definition does *not* directly translate into a tractable optimization formulation. In fact, their proposed algorithm, FPCA, *necessarily* relies on the Gaussian assumption of the underlying data distributions i.e. FPCA only matches the means and covariances of the projected data distributions, which is problematic. To illustrate this, consider a 3-dimensional synthetic dataset with two protected classes colored as blue and orange, shown in Figure 1a. We’ve designed the data distributions such that their means and covariances are almost the same. We assume that we’re reducing the dimensionality of the dataset to 2. Vanilla PCA results in Figure 1b, where the distribution dissimilarity persists. Due to the Gaussian assumption, FPCA results in Figure 1c where the dissimilarity also persists.

This calls for a new definition of fairness in PCA that can lead to a tractable form of optimization formulation that ensures fairness exactly in the sense that the equality of the two projected data distributions is guaranteed, resulting in Figure 1d. Recall how the necessity for such assumption was derived from the inherent complexity of their definition. Inspired from our previous discussions on MMD, we propose the following **new** definition of fairness for PCA:

Definition 5 (Δ -fairness). *Let P_s be the probability measure of $X_s \triangleq X|A = s$ for $s \in \{0, 1\}$, and let $Q_s := \Pi_{\#} P_s \in \mathcal{P}_d$. Then Π is said to be Δ -fair with $\Delta := \text{MMD}(Q_0, Q_1)$, and we refer to Δ as the **fairness metric**.*

In other words, low Δ means that the discrepancy between the projected conditional distributions of different protected classes, measured in a non-parametric manner using MMD, while retaining as much variance as possible, is low.

Furthermore, Definition 5 ensures that a downstream classification task using Δ -fair dimensionality-reduced data will be fair, as formalized below²:

Proposition 1 (Oneto et al., 2020). *Up to a constant factor, $\text{MMD}(Q_0, Q_1)$ bounds the MMD of the push-forward measures of Q_0, Q_1 via the weight vector of any given downstream task classifier g .*

Remark 1. *The above discussions easily generalize to equal opportunity and equalized odds.*

3.1 Relation with other definitions of fair PCA

The notion of fairness proposed by Olfat and Aswani (2019) is similar to ours in that it measures the predictability of protected group membership in dimensionality-reduced data. However, unlike ours, their definition is explicitly adversarial, which can be a problem.

Definition 6 (Δ_A -fairness; Olfat and Aswani, 2019). *Consider a fixed classifier $h(u, t) : \mathbb{R}^d \times \mathbb{R} \rightarrow \{0, 1\}$ that inputs features $u \in \mathbb{R}^d$ and a threshold t , and predicts the protected class $s \in \{0, 1\}$. Then, Π is $\Delta_A(h)$ -fair if*

$$\sup_{t \in \mathbb{R}} \left| \mathbb{P}[h(\Pi(x), t) = 1 | s = 1] - \mathbb{P}[h(\Pi(x), t) = 1 | s = 0] \right| \leq \Delta_A(h) \quad (2)$$

²See Lemma 3 of (Oneto et al. 2020) for the precise statement.

Moreover, for a family of classifiers \mathcal{F}_c , if Π is $\Delta_A(h)$ -fair for $\forall h \in \mathcal{F}_c$, we say that Π is $\Delta_A(\mathcal{F}_c)$ -fair.

As Δ_A can’t be computed exactly, an estimator of the following form was used:

$$\widehat{\Delta}_A(\mathcal{F}_c) := \sup_{h \in \mathcal{F}_c} \sup_t \left| \frac{1}{|P|} \sum_{i \in P} I_i(\Pi, h_t) - \frac{1}{|N|} \sum_{i \in N} I_i(\Pi, h_t) \right| \quad (3)$$

where $\{x_i\}_{i=1}^n$ are the data points, (P, N) is a partition of the index set $\{1, 2, \dots, n\}$ into two protected groups, $I_i(\Pi, h_t) = \mathbf{1}(h(\Pi(x_i), t) = +1)$, and $\mathbf{1}(\cdot)$ is the indicator function.

Remark 2. *It can be argued that, for some choice of \mathcal{F}_c , Definition 5 and 6 are equivalent: in effect, that these are dual notions. Recognizing this, we proceed with Definition 5, as it has two main advantages in the context of our work:*

- *It ties more directly and intuitively into our optimization formulation; see Section 5.*
- *It can be represented non-variationally which allows for tighter statistical guarantees.*

4 Statistical properties of Δ

4.1 Consistent and efficient estimation of Δ

As defined in Definition 5, let $Q_0, Q_1 \in \mathcal{P}_d$ be the probability measures with respect to the samples of which we want to estimate $\text{MMD}(\cdot, \cdot)$. Let $\{X_i\}_{i=1}^m$ and $\{Y_j\}_{j=1}^n$ be these samples, respectively. Accordingly, we consider the following estimator:

$$\widehat{\Delta} := \text{MMD}(\widehat{Q}_0, \widehat{Q}_1) \quad (4)$$

where \widehat{Q}_s is the usual empirical distribution, defined as the mixture of Dirac measures on the samples.

Unlike other statistical distances (e.g. total variation distance), $\widehat{\Delta}_k$ has several desirable theoretical properties that have important practical implications, some of which we recall here. (See Sriperumbudur et al. (2010) for more comprehensive discussions.)

First, $\widehat{\Delta}_k$ can be computed exactly and efficiently:

Lemma 1 (Gretton et al., 2007). *$\widehat{\Delta}$ is computed as follows:*

$$\widehat{\Delta} = \left[\frac{1}{m^2} \sum_{i,j=1}^m k(X_i, X_j) + \frac{1}{n^2} \sum_{i,j=1}^n k(Y_i, Y_j) - \frac{2}{mn} \sum_{i,j=1}^{m,n} k(X_i, Y_j) \right]^{1/2}. \quad (5)$$

Moreover, it is asymptotically consistent with a convergence rate, depending only on m and n :

Proposition 2 (Gretton et al., 2007). *For any $\delta > 0$, with probability at least $1 - 2 \exp\left(-\frac{\delta^2 mn}{2(m+n)}\right)$ the following holds:*

$$\left| \Delta - \widehat{\Delta} \right| \leq 2 \left(\frac{1}{\sqrt{m}} + \frac{1}{\sqrt{n}} \right) + \delta \quad (6)$$

4.2 Advantages over Δ_A

$\widehat{\Delta}_A$ is known to satisfy the following high probability bound:

Proposition 3 (Olfat and Aswani, 2019). *Consider a fixed family of classifiers \mathcal{F}_c . Then for any $\delta > 0$, with probability at least $1 - \exp\left(-\frac{(n+m)\delta^2}{2}\right)$ the following holds:*

$$\left| \Delta_A(\mathcal{F}_c) - \widehat{\Delta}_A(\mathcal{F}_c) \right| \leq 8\sqrt{\frac{VC(\mathcal{F}_c)}{m+n}} + \delta \quad (7)$$

where $VC(\cdot)$ is the VC dimension.

If \mathcal{F}_c is too expressive in terms of VC-dimension, then the above bound may become void. This is the case, for instance, when \mathcal{F}_c is the set of RBF-kernel SVMs.

In addition, computing $\widehat{\Delta}_A$ requires considering all possible classifiers in the designated family \mathcal{F}_c . This is computationally infeasible, and it forces one to use another approximation (e.g. discretization of \mathcal{F}_c), which incurs additional error that may further inhibit asymptotic consistency.

As exhibited in the previous subsection, our MMD-based approach suffers from none of these issues.

5 Manifold optimization for MBF-PCA

5.1 Improvements over FPCA

Olfat and Aswani (2019) proposed FPCA, an SDP formulation of fair PCA³, in which matching the first and second moments of the protected groups after dimensionality-reduction are approximated as convex constraints. However, this has several shortcomings, which we discuss here and empirically exhibit in a later section.

First, there are cases in which matching the mean and covariance alone is not enough; this was explicitly shown and discussed in Section 3 where we’ve presented a simple “counterexample” in which two protected groups have different distributions with same mean and covariance. (Figure 1) While this previous point may be countered by the application of the kernel trick to FPCA, this raises a second issue: their formulation requires solving⁴ a $p \times p$ -dimensional SDP, motivated by the reparameterization $P = VV^\top$ (Arora, Cotter, and Srebro 2013). Since SDP is known to become inefficient (or even computationally infeasible) in high dimensions, this quickly becomes intractable for high-dimensional data (for linear or polynomial kernels) or for any moderate to large size datasets (for the RBF kernel). Finally, their approach involves a relaxation of a rank constraint ($\text{rank}(P) \leq d$) to a trace constraint ($\text{tr}(P) \leq d$), yielding sub-optimal outputs in presence of (fairness) constraints, even to substantial order in some cases. In Section C of the SP, we discuss in detail why FPCA may lead to such degraded explained variance.

³See Section B of the supplementary material (SP) for its precise description.

⁴In their approach, the final solution V is obtained by taking the first d eigenvectors of P .

5.2 Formulating MBF-PCA

Observing that the shortcomings of FPCA stem from the reparametrization of $P = VV^\top$, we propose a new formulation of fair PCA that solves *directly* for V . This allows for an effective and efficient approach.

We start by noting that the set of all V with orthonormal columns has the intrinsic geometric structure of a *manifold*:

Definition 7. *For $p > d$, the Stiefel manifold, denoted as $St(p, d)$, is an embedded Riemannian sub-manifold of $\mathbb{R}^{p \times d}$ such that each element of $St(p, d)$ has orthonormal columns i.e. $V^\top V = I_d$ for all $V \in St(p, d)$.*

$St(p, d)$ has several desirable properties such as compactness, completeness and smoothness, which we present in Section D of the SP. As $St(p, d)$ is prevalent in various fields of machine learning (most notably PCA), much work has been done that focuses on exploiting this geometric structure for efficient optimization (Hu et al. 2020).

Based on our MMD-based formulation and letting Σ be the sample covariance matrix of the full dataset, we formulate our fair PCA as follows, which we refer to as MBF-PCA:

$$\begin{aligned} & \underset{V \in St(p, d)}{\text{minimize}} && f(V) := -\langle \Sigma, VV^\top \rangle \\ & \text{subject to} && h(V) := \text{MMD}^2(Q_0, Q_1) = 0. \end{aligned} \quad (8)$$

Here, Q_0 and Q_1 are defined as in definition 5. Observe how our definition of fairness *directly* incorporates itself into the optimization problem as a constraint.

Remark 3. *Under Gaussian assumption, our MMD-based formulation indeed amounts to the same constraints as FPCA since $\text{MMD}^2(\cdot, \cdot)$ is a metric and Gaussian distribution is completely characterized by its first and second moments.*

6 REPMS for MBF-PCA, with new theoretical guarantees

6.1 Description of Algorithm 1

One crucial observation is that the constraint function h is always non-negative⁵ and smooth. This motivates the use of the exact penalty method (Han and Mangasarian 1979), recently extended to manifold optimization as the Riemannian Exact Penalty Method via Smoothing (REPMS; Liu and Boumal, 2019). Note that smoothing tricks (Liu and Boumal 2019), which were required to smooth out possible non-differentiable functions emerging from the ℓ_1 -penalty, are *not* necessary. Moreover, by leveraging the kernel trick, there is a *closed* form for $\nabla_V h(V)$, thus alleviating the need for computationally expensive numerical approximations; see Section G of the SP for the derivation. The pseudo-code for the algorithm is shown in Algorithm 1.

For practical concerns that will be addressed in the following subsection, we’ve set the *fairness tolerance level*, τ , to be a fixed and sufficiently small, non-negative value. Formally, we consider the following definition:

Definition 8. *For fixed $\tau \geq 0$, $V \in St(p, d)$ is τ -approximate fair if it satisfies $h(V) \leq \tau$. If $\tau = 0$, we simply say that V is fair.*

⁵This follows from closed form of $\widehat{\Delta}_k$ (Lemma 1). See Section 2 of Gretton et al. (2012) for more details.

Algorithm 1: REPMS for MbF-PCA

Input: $X, K, \epsilon_{min}, \epsilon_0 > 0, \theta_\epsilon \in (0, 1), \rho_0 > 0,$
 $\theta_\rho > 1, \rho_{max} \in (0, \infty), \tau > 0, d_{min} > 0.$

- 1 Initialize V_0 ;
- 2 **for** $k = 0, 1, \dots, K$ **do**
- 3 Compute an approximate solution V_{k+1} for the following sub-problem, with a warm-start at V_k , until $\|\text{grad } \mathcal{Q}\| \leq \epsilon_k$:
$$\min_{V \in St(p,d)} \mathcal{Q}(V, \rho_k) \quad (9)$$

where

$$\mathcal{Q}(V, \rho_k) = f(V) + \rho_k h(V)$$
- 4 **if** $\|V_{k+1} - V_k\|_F \leq d_{min}$ **and** $\epsilon_k \leq \epsilon_{min}$ **then**
- 5 **if** $h(V_{k+1}) \leq \tau$ **then**
- 6 **return** V_{k+1} ;
- 7 **end**
- 8 **end**
- 9 $\epsilon_{k+1} = \max\{\epsilon_{min}, \theta_\epsilon \epsilon_k\}$;
- 10 **if** $h(V_{k+1}) > \tau$ **then**
- 11 $\rho_{k+1} = \min(\theta_\rho \rho_k, \rho_{max})$;
- 12 **else**
- 13 $\rho_{k+1} = \rho_k$;
- 14 **end**
- 15 **end**

6.2 New theoretical guarantees for Algorithm 1

We start by observing that Eq. (9) in Algorithm 1 is smooth, unconstrained manifold optimization problem, which can be solved using conventional algorithms; these include first-order methods like line-search methods (Absil, Mahony, and Sepulchre 2007), or second-order methods like the Riemannian Trust Region (RTR; Absil, Baker, and Gallivan, 2007) method. It is known that, pathological examples excluded, most conventional *unconstrained* manifold optimization solvers produce iterates whose limit points are local minima, and not other stationary points such as saddle point or local maxima: see (Absil, Baker, and Gallivan 2007; Absil, Mahony, and Sepulchre 2007) for more detailed discussions.

Motivated by this, we consider the following assumption:

Assumption 2 (informal; locality assumption). *Each V_{k+1} is sufficiently close to a local minimum of Eq. (9).*

Lastly, we consider the following auxiliary optimization problem:

$$\min_{V \in St(p,d)} h(V) \quad (10)$$

The following theorem, whose proof is deferred to Section F of the SP, provides an *exact* theoretical convergence guarantee of MbF-PCA under the *ideal* hyperparameter setting:

Theorem 1. *Let $K = \infty, \rho_{max} = \infty, \epsilon_{min} = \tau = 0, \{V_k\}$ be the sequence generated by Alg. 1 under Assumption 2, and \bar{V} be any limit point of $\{V_k\}$, whose existence is guaranteed. Then the following holds:*

- (A) \bar{V} is a local minimizer of Eq. (10), which is a necessary condition for \bar{V} to be fair.
- (B) If \bar{V} is fair, then \bar{V} is a local minimizer of Eq. (8)

The assumption of \bar{V} being fair, which is used in (B), is at least partially justified in (A) in the following sense: the ideal hyperparameter setting of $\rho_{max} = \infty, \tau = 0, \epsilon_{min} = 0$ implies the *exact* local minimality of \bar{V} for Eq. (10), which is in turn a *necessary condition* for \bar{V} to be fair.

The next theorem, whose is also deferred to Section F of the SP, asserts that with small τ, ϵ_{min} and large ρ_{max} , the above guarantee can be approximated in rigorous sense:

Theorem 2. *Let $K = \infty, \rho_{max} < \infty, \epsilon_{min}, \tau > 0, \{V_k\}$ be the sequences generated by Alg. 1 under Assumption 2 and \bar{V} be any limit point of $\{V_k\}$, whose existence is guaranteed. Then for any sufficiently small ϵ_{min} and $\tilde{r} = \tilde{r}(\epsilon_{min}) > 0$, the following hold:*

- (A) \bar{V} is an approximate local minimizer of Eq. (10) in the sense that
$$h(\bar{V}) \leq h(V) + \beta \|V - \bar{V}\| + (\beta + L_h)g(\epsilon_{min}) \quad (11)$$

for all $V \in B_{\tilde{r}}(\bar{V}) \cap St(p, d)$, where $\beta = \beta(\rho_{max}, \tau)$ is a function that satisfies the following:

 - $0 < \beta \leq \frac{2\|\Sigma\|}{\rho_0}$
 - $\beta(\rho_{max}, \tau)$ is increasing in ρ_{max} and decreasing in τ .
- (B) If \bar{V} is fair, then it is an approximate local minimizer of Eq. (8) in the sense that it satisfies

$$f(\bar{V}) \leq f(V) + 2\|\Sigma\|g(\epsilon_{min}) \quad (12)$$

for all fair $V \in B_{\tilde{r}}(\bar{V}) \cap St(p, d)$.

In both (A) and (B), g is some continuous, decreasing function that satisfies $g(0) = 0$, and $\tilde{r}(\epsilon_{min}) = r - g(\epsilon_{min})$ for some fixed constant $r > 0$.

Existing optimality guarantee of REPMS (Proposition 4.2; Liu and Boumal, 2019) states that when $\epsilon_{min} = 0, \rho$ is not updated (i.e. line 10-14 is ignored), and the resulting limit point is feasible, then that limit point satisfies the KKT condition (Yang, Zhang, and Song 2014). Comparing Theorem 1 and 2 to the previous result, we see that ours extend the previous result in several ways:

- Our theoretical analyses are much closer to the actual implementation, by incorporating the ρ -update step (line 11) and the *practical* hyperparameter setting.
- Our theoretical analyses are much more stronger in the sense that 1) by *introducing* a reasonable, yet novel locality assumption, we go beyond the existing KKT conditions and prove the *local minimality* of the limit point, and 2) we provide a partial justification of the feasibility assumption in (A) by proving a necessary condition for it.

6.3 Practical implementation

In line 4 in Algorithm 1, we implemented the termination criteria: sufficiently small distance between iterates and sufficiently small tolerance for solving Eq. (9). However, such a heuristic may return some point \bar{V} that is not τ -approximate fair for user-defined level τ in practical hyperparameter setting. To overcome this issue, we've additionally implemented line 5 that forces the algorithm to continue on with the loop until the desired level of fairness is achieved.

7 Related Work

7.1 Fairness in ML

A large body of work regarding fairness in the context of supervised learning (Feldman et al. 2015; Calders, Kamiran, and Pechenizkiy 2009; Dwork et al. 2012; Hardt, Price, and Srebro 2016; Zafar et al. 2017) has been published. This includes key innovations in quantifying algorithmic bias, notably the concepts of *demographic parity* and *equalized odds (opportunity)* that have become ubiquitous in fairness research (Barocas and Selbst 2016; Hardt, Price, and Srebro 2016). More recently, fair machine learning literatures have branched out into a variety of fields, including deep learning (Beutel et al. 2017), regression (Calders et al. 2013), and even hypothesis testing (Olfat et al. 2020).

Among these, one line of research has focused on learning fair representations (Kamiran and Calders 2011; Zemel et al. 2013; Feldman et al. 2015; Calmon et al. 2017), which aims to learn a representation of the given data on which various fairness definitions are ensured for downstream modeling. A growing number of inquiries have been made into highly specialized algorithms for specific unsupervised learning problems like clustering (Chierichetti et al. 2017; Kleindessner, Awasthi, and Morgenstern 2019; Bera et al. 2019), but these lack the general applicability of key dimensionality reduction algorithms such as PCA (Pearson 1901; Hotelling 1933).

To the best of our knowledge, Olfat & Aswani (2019) is the **only** work on incorporating fair representation to PCA, making it the sole comparable approach to ours. Another line of work (Bian and Tao 2011; Samadi et al. 2018; Tantipongpipat et al. 2019; Zalcborg and Wiesel 2021) considers a completely *orthogonal* definition of fairness for PCA: minimizing the discrepancy between reconstruction errors over protected attributes. This doesn't ensure the fairness of downstream tasks, rendering it incomparable to our definition of fairness; see Section A of the SP for more details.

7.2 Manifold Optimization

A constrained problem over Euclidean space can be transformed to an unconstrained problem over a manifold (or at least manifold optimization with less constraints). Many algorithms for solving Euclidean optimization problems have direct counterparts in manifold optimization problems that includes Riemannian gradient descent and Riemannian BFGS. By making use of the geometry of lower dimensional manifold structure, often embedded in potentially very high dimensional ambient space, such Riemannian counterparts are much more computationally *efficient* than algorithms that do not make use of manifold structure. This is shown in numerous literatures (Liu and Boumal 2019; Alsharif et al. 2021; Meng, Chakraborty, and Singh 2021), including this work. We refer interested readers to the standard textbooks Absil, Mahony, and Sepulchre (2007) and Boumal (2020) on this field, along with a survey by Hu et al. (2020).

8 Experiments

MBF-PCA was implemented using ROPTLIB (Huang et al. 2018), a state-of-the-art manifold optimization framework on

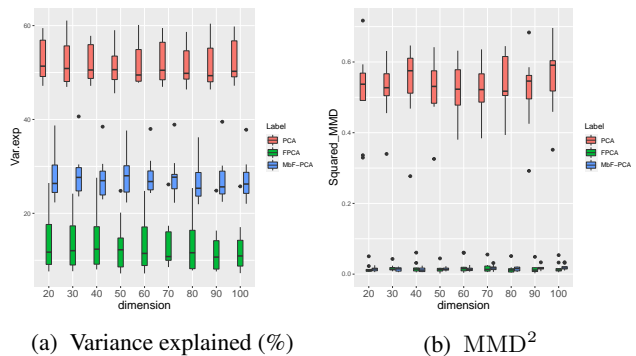


Figure 2: Synthetic data #2: Comparison of PCA, FPCA, and MBF-PCA on the synthetic datasets of increasing dimensions.




Figure 3: Synthetic data #2: Comparison of runtimes of FPCA, and MbF-PCA.

MATLAB. For solving Eq. (9), we use the cautious Riemannian BFGS method (RBFGS; Huang, Absil, and Gallivan, 2018), a memory-efficient quasi-Newton method. As for the hyperparameters, we've set $K = 100$, $\epsilon_{min} = 10^{-6}$, $\epsilon_0 = 10^{-1}$, $\theta_\epsilon = (\epsilon_{min}/\epsilon_0)^{1/5}$, $\rho_{max} = 10^{10}$, $\theta_\rho = 2$, $d_{min} = 10^{-6}$. For FPCA, we use the same Python MOSEK(ApS 2021) implementation as provided by (Olfat and Aswani 2019). (μ, δ) are the hyperparameters of FPCA; see Section B of the SP. Codes are available in our Github repository⁶.

All data is pre-processed to be standardized such that each covariate has zero mean and unit variance. For all experiments, we considered 10 different 70 – 30 train-test splits.

8.1 Synthetic data #1

We consider synthetic data composed of two groups, each of size $n = 150$; one is sampled from $\mathcal{N}_3(\mathbf{0}, 0.1I_3 + \mathbf{1})$ and one is sampled from a (balanced) mixture of $\mathcal{N}_3(\mathbf{1}, 0.1I_3)$ and $\mathcal{N}_3(-\mathbf{1}, 0.1I_3)$. Note how the two groups follow different distributions, yet have the same mean and covariance.

⁶<https://github.com/nick-jhlee/fair-manifold-pca>

Table 1: Comparison of PCA, FPCA, MBF-PCA for UCI datasets. Number in parenthesis for each dataset is its dimension. Also, the parenthesis for each fair algorithm is its hyperparameter setting; (μ, δ) for FPCA and τ for MBF-PCA. Among the fair algorithms considered, results with the best mean values are **bolded**. Results in which our approach terminates improperly in the sense that the maximum iteration is reached before passing the termination criteria are **highlighted**.

d	ALG.	COMPAS (11)				GERMAN CREDIT (57)				ADULT INCOME (97)			
		%VAR	%ACC	MMD ²	Δ_{DP}	%VAR	%ACC	MMD ²	Δ_{DP}	%VAR	%ACC	MMD ²	Δ_{DP}
2	PCA	39.28 _{5.17}	64.53 _{1.45}	0.092 _{0.010}	0.29 _{0.09}	11.42 _{0.47}	76.87 _{1.39}	0.147 _{0.049}	0.12 _{0.06}	7.78 _{0.82}	82.03 _{1.15}	0.349 _{0.027}	0.20 _{0.05}
	FPCA (0.1, 0.01)	35.06 _{5.16}	61.65 _{1.17}	0.012 _{0.007}	0.10 _{0.07}	7.43 _{0.59}	72.17 _{1.09}	0.017 _{0.010}	0.03 _{0.02}	4.05 _{0.98}	77.44 _{2.96}	0.016 _{0.011}	0.04 _{0.04}
	FPCA (0, 0.01)	34.43 _{5.02}	60.86 _{1.09}	0.011 _{0.006}	0.10 _{0.06}	7.33 _{0.57}	71.77 _{1.60}	0.015 _{0.010}	0.03 _{0.03}	3.65 _{0.97}	77.05 _{3.18}	0.005 _{0.004}	0.01 _{0.01}
	MBF-PCA (10 ⁻³)	33.95 _{5.01}	65.37 _{1.11}	0.005 _{0.002}	0.12 _{0.07}	10.17 _{0.57}	74.53 _{1.92}	0.018 _{0.014}	0.05 _{0.04}	6.03 _{0.61}	79.50 _{1.22}	0.005 _{0.004}	0.03 _{0.02}
	MBF-PCA (10 ⁻⁶)	11.83 _{3.59}	57.73 _{1.50}	0.002 _{0.002}	0.06 _{0.08}	9.36 _{0.33}	74.10 _{1.56}	0.016 _{0.010}	0.02 _{0.02}	5.83 _{0.57}	79.12 _{1.14}	0.005 _{0.004}	0.01 _{0.01}
10	PCA	100.00 _{0.00}	73.14 _{1.22}	0.241 _{0.005}	0.21 _{0.07}	38.25 _{0.98}	99.93 _{0.14}	0.130 _{0.019}	0.12 _{0.08}	21.77 _{2.06}	93.64 _{0.92}	0.195 _{0.007}	0.16 _{0.01}
	FPCA (0.1, 0.01)	87.79 _{1.27}	72.25 _{0.93}	0.015 _{0.003}	0.16 _{0.06}	29.85 _{0.87}	99.93 _{0.14}	0.020 _{0.005}	0.12 _{0.08}	15.75 _{1.20}	91.94 _{0.88}	0.006 _{0.003}	0.13 _{0.02}
	FPCA (0, 0.1)	87.44 _{1.35}	72.32 _{0.93}	0.015 _{0.002}	0.16 _{0.07}	29.79 _{0.89}	99.93 _{0.14}	0.020 _{0.006}	0.12 _{0.08}	15.52 _{1.18}	91.66 _{0.97}	0.004 _{0.002}	0.13 _{0.02}
	MBF-PCA (10 ⁻³)	87.75 _{1.36}	72.16 _{0.90}	0.014 _{0.002}	0.16 _{0.07}	34.10 _{1.00}	99.93 _{0.14}	0.020 _{0.008}	0.12 _{0.08}	18.71 _{1.47}	92.81 _{0.84}	0.005 _{0.002}	0.14 _{0.01}
	MBF-PCA (10 ⁻⁶)	87.75 _{1.36}	72.16 _{0.90}	0.014 _{0.002}	0.16 _{0.07}	16.95 _{1.52}	92.70 _{3.00}	0.013 _{0.007}	0.06 _{0.05}	15.49 _{6.44}	86.36 _{3.77}	0.003 _{0.002}	0.07 _{0.03}

Thus, we expect FPCA to project in a similar way as vanilla PCA, while MBF-PCA should find a fairer subspace such that the projected distributions are exactly the same. Hyperparameters are set as follows: $\delta = 0, \mu = 0.01$ for FPCA and $\tau = 10^{-5}$ for MBF-PCA. We’ve set $d = 2$ and Figure 1 displays the results of each algorithm using the top two principal components. Indeed, only MBF-PCA successfully obfuscates the protected group information by merging the two orange clusters with the blue cluster.

8.2 Synthetic data #2

We consider a series of synthetic datasets of dimension p . For each p , the dataset is composed of two groups, each of size $n = 240$ and sampled from two different p -variate normal distributions. We vary $p \in \{20, 30, \dots, 100\}$; see Section H of the SP for a full description of the setting. For the hyperparameters, we’ve set $\delta = 0, \mu = 0.01$ for FPCA and $\tau = 10^{-5}$ for MBF-PCA.

Figure 2 plots the explained variance and fairness metric values. Observe how MBF-PCA achieves better explained variance, while achieving similar level of fairness. In addition, Figure 3 shows a clear gap in runtime between FPCA and MBF-PCA; the runtime of FPCA explodes for even moderate problem sizes, while MBF-PCA scales well. For higher dimensions, conventional computing machine will not be able to handle such computational burden.

8.3 UCI datasets

For target dimensions $d \in \{2, 10\}$, we compare the performance of FPCA and MBF-PCA on 3 datasets from the UCI Machine Learning Repository (Dua and Graff 2017); COMPAS dataset (Kirchner et al. 2016), Adult income dataset, and German credit dataset. See Section I of the SP for complete description of the pre-processing steps. For both algorithms, we consider two different hyperparameter settings, such that one simulates the relaxed fairness while the other simulates a stricter fairness constraints. For computing $\Delta_{DP}(g)$, we trained a RBF SVM g to be the *downstream task classifier* that best classifies the target attribute in the *dimensionality-reduced* data. Table 1 displays the results, from which several observations can be made:

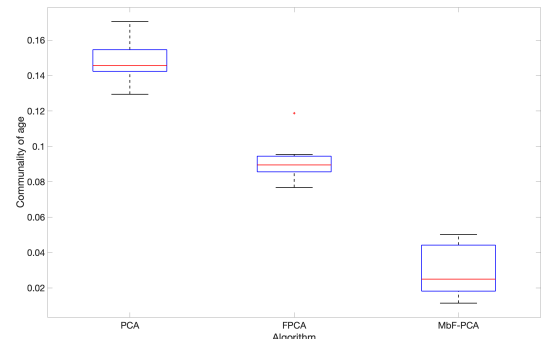


Figure 4: Comparison of communality of “age” of German credit dataset for PCA, FPCA, and MBF-PCA.

- Across all considered datasets, MBF-PCA is shown to outperform FPCA in terms of fairness (both MMD² and Δ_{DP}) with low enough τ .
- For GERMAN CREDIT and ADULT INCOME, MBF-PCA shows a clear trade-off between explained variance and fairness; by relaxing τ , we see that MBF-PCA outperforms FPCA in terms of explained variance and downstream task accuracy.

In addition, to see how correlated are the PCs with the protected attribute, we examine the communalities. For clarity of exposition, we consider the German credit dataset, whose protected attribute is age, and $d = 10$. Here, we again consider PCA, FPCA (0, 0.01), and MBF-PCA (10⁻³). For PCA, communality of a feature is its variance contributed by the PCs (Johnson and Wichern 2008), computed as the sum of squares of the loadings of the considered feature. High communality implies that the correlations between the feature and PCs are strong. Figure 4 displays the boxplot of communality, plotted over the 10 splits. Indeed the amount of variance in age that is accounted for from the loadings of MBF-PCA is much smaller than that of PCA or FPCA. In other words, the PCs resulting from MBF-PCA have the *least* correlations with age, the protected attribute.

9 Conclusion and Future Works

We present a MMD-based definition of fair PCA, and formulate it as a constrained optimization over the Stiefel manifold. Through both theoretical and empirical discussions, we show that our approach outperforms the previous approach (Olfat and Aswani 2019) in terms of explained variance, fairness, and runtime. Many avenues remain for future research. Statistical characterizations of our fair PCA in asymptotic regime, as well as incorporation of sparsity (Johnstone et al. 2009) are important open questions. Incorporating stochastic optimization-type modifications (Shamir 2015; Roh et al. 2021) is also an important direction, as such modifications are expected to result in better scalability and performance.

Acknowledgements

JH, GS, and CD were partly supported by Institute for Information & communications Technology Planning & Evaluation(IITP) grant funded by the Korea government(MSIT) (No. 2019-0-01396, Development of framework for analyzing, detecting, mitigating of bias in AI model and training data), and partly supported by Institute for Information & communications Technology Planning & Evaluation(IITP) grant funded by the Korea government(MSIT) (No. 2021-0-01381, Development of Causal AI through Video Understanding). JH was also supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No. 2019-0-00075, Artificial Intelligence Graduate School Program(KAIST)) JH would like to thank Jaeho Lee and Se-Young Yun for providing helpful comments on the writing and organization of the paper.

References

- Absil, P.-A.; Baker, C. G.; and Gallivan, K. A. 2007. Trust-Region Methods on Riemannian Manifolds. *Foundations of Computational Mathematics*, 7(3): 303–330.
- Absil, P.-A.; Mahony, R.; and Sepulchre, R. 2007. *Optimization Algorithms on Matrix Manifolds*. USA: Princeton University Press.
- Adel, T.; Valera, I.; Ghahramani, Z.; and Weller, A. 2019. One-Network Adversarial Fairness. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 2412–2420.
- Alsharif, M. H.; Douik, A.; Ahmed, M.; Alnaffouri, T.; and Hassibi, B. 2021. Manifold Optimization for High Accuracy Spatial Location Estimation Using Ultrasound Waves. *IEEE Transactions on Signal Processing*, 1–1.
- ApS, M. 2021. *MOSEK Optimizer API for Python. Version 9.2.36*.
- Arora, R.; Cotter, A.; and Srebro, N. 2013. Stochastic Optimization of PCA with Capped MSG. In Burges, C. J. C.; Bottou, L.; Welling, M.; Ghahramani, Z.; and Weinberger, K. Q., eds., *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc.
- Barocas, S.; and Selbst, A. D. 2016. Big Data’s Disparate Impact. *104 California Law Review* 671.
- Bellamy, R. K. E.; Dey, K.; Hind, M.; Hoffman, S. C.; Houde, S.; Kannan, K.; Lohia, P.; Martino, J.; Mehta, S.; Mojsilovic, A.; Nagar, S.; Ramamurthy, K. N.; Richards, J.; Saha, D.; Sattigeri, P.; Singh, M.; Varshney, K. R.; and Zhang, Y. 2018. AI Fairness 360: An Extensible Toolkit for Detecting, Understanding, and Mitigating Unwanted Algorithmic Bias.
- Benidis, K.; Sun, Y.; Babu, P.; and Palomar, D. P. 2016. Orthogonal Sparse PCA and Covariance Estimation via Procrustes Reformulation. *IEEE Transactions on Signal Processing*, 64(23): 6211–6226.
- Bera, S. K.; Chakrabarty, D.; Flores, N.; and Negahbani, M. 2019. Fair Algorithms for Clustering. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019*, 4955–4966. Vancouver, BC, Canada.
- Beutel, A.; Chen, J.; Zhao, Z.; and Chi, E. H. 2017. Data decisions and theoretical implications when adversarially learning fair representations. *arXiv preprint arXiv:1707.00075*.
- Bian, W.; and Tao, D. 2011. Max-Min Distance Analysis by Using Sequential SDP Relaxation for Dimension Reduction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(5): 1037–1050.
- Boumal, N. 2020. An introduction to optimization on smooth manifolds. Available online.
- Boyd, S.; El Ghaoui, L.; Feron, E.; and Balakrishnan, V. 1994. *Linear Matrix Inequalities in System and Control Theory*. Society for Industrial and Applied Mathematics.
- Calders, T.; Kamiran, F.; and Pechenizkiy, M. 2009. Building Classifiers with Independency Constraints. In *ICDM Workshops 2009, IEEE International Conference on Data Mining Workshops*, 13–18. Miami, Florida, USA.
- Calders, T.; Karim, A.; Kamiran, F.; Ali, W.; and Zhang, X. 2013. Controlling attribute effect in linear regression. In *Data Mining (ICDM), 2013 IEEE 13th International Conference on*, 71–80. IEEE.
- Calmon, F. P.; Wei, D.; Vinzamuri, B.; Ramamurthy, K. N.; and Varshney, K. R. 2017. Optimized Pre-Processing for Discrimination Prevention. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017*, 3992–4001. Long Beach, CA, USA.
- Chien, C.-F.; and Chen, L.-F. 2008. Data mining to improve personnel selection and enhance human capital: A case study in high-technology industry. *Expert Systems with applications*, 34(1): 280–290.
- Chierichetti, F.; Kumar, R.; Lattanzi, S.; and Vassilvitskii, S. 2017. Fair Clustering Through Fairlets. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017*, 5029–5037. Long Beach, CA, USA.
- Dua, D.; and Graff, C. 2017. UCI Machine Learning Repository.
- Dwork, C.; Hardt, M.; Pitassi, T.; Reingold, O.; and Zemel, R. S. 2012. Fairness through awareness. In *Innovations in Theoretical Computer Science 2012*, 214–226. Cambridge, MA, USA.

- Eckart, C.; and Young, G. 1936. The approximation of one matrix by another of lower rank. *Psychometrika*, 1: 211–218.
- Feldman, M.; Friedler, S. A.; Moeller, J.; Scheidegger, C.; and Venkatasubramanian, S. 2015. Certifying and Removing Disparate Impact. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 259–268. Sydney, NSW, Australia.
- Fukumizu, K.; Gretton, A.; Sun, X.; and Schölkopf, B. 2008. Kernel Measures of Conditional Dependence. In Platt, J.; Koller, D.; Singer, Y.; and Roweis, S., eds., *Advances in Neural Information Processing Systems*, volume 20. Curran Associates, Inc.
- Gretton, A.; Borgwardt, K.; Rasch, M.; Schölkopf, B.; and Smola, A. 2007. A Kernel Method for the Two-Sample-Problem. In Schölkopf, B.; Platt, J.; and Hoffman, T., eds., *Advances in Neural Information Processing Systems*, volume 19. MIT Press.
- Gretton, A.; Borgwardt, K. M.; Rasch, M. J.; Schölkopf, B.; and Smola, A. 2012. A Kernel Two-Sample Test. *Journal of Machine Learning Research*, 13(25): 723–773.
- Han, S. P.; and Mangasarian, O. L. 1979. Exact penalty functions in nonlinear programming. *Mathematical Programming*, 17(1): 251–269.
- Hardt, M.; Price, E.; and Srebro, N. 2016. Equality of Opportunity in Supervised Learning. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016*, 3315–3323. Barcelona, Spain.
- Hotelling, H. 1933. Analysis of a complex of statistical variables into principal components. *Journal of Educational Psychology*, 24(6): 417–441.
- Hu, J.; Liu, X.; Wen, Z.-W.; and Yuan, Y.-X. 2020. A Brief Introduction to Manifold Optimization. *Journal of Operations Research Society of China*, 8: 199–248.
- Huang, W.; Absil, P.-A.; and Gallivan, K. A. 2018. A Riemannian BFGS Method Without Differentiated Retraction for Nonconvex Optimization Problems. *SIAM Journal on Optimization*, 28(1): 470–495.
- Huang, W.; Absil, P.-A.; Gallivan, K. A.; and Hand, P. 2018. ROPTLIB: An Object-Oriented C++ Library for Optimization on Riemannian Manifolds. *ACM Transactions on Mathematical Softwares*, 44(4).
- Johnson, R. A.; and Wichern, D. W. 2008. *Applied Multivariate Statistical Analysis*. Pearson, 6 edition.
- Johnstone, I. M.; Lu, A. Y.; Nadler, B.; Witten, D. M.; Hastie, T.; Tibshirani, R.; and Ramsay, J. O. 2009. On Consistency and Sparsity for Principal Components Analysis in High Dimensions [with Comments]. *Journal of the American Statistical Association*, 104(486): 682–703.
- Jolliffe, I. T.; and Cadima, J. 2016. Principal component analysis: a review and recent developments. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 374(2065).
- Jung, S.; Lee, D.; Park, T.; and Moon, T. 2021. Fair Feature Distillation for Visual Recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 12115–12124.
- Kamiran, F.; and Calders, T. 2011. Data preprocessing techniques for classification without discrimination. *Knowledge and Information Systems*, 33(1): 1–33.
- Kamishima, T.; Akaho, S.; and Sakuma, J. 2011. Fairness-aware Learning through Regularization Approach. In *2011 IEEE 11th International Conference on Data Mining Workshops*, 643–650.
- Karan, O.; Bayraktar, C.; Gümüşkaya, H.; and Karlık, B. 2012. Diagnosing diabetes using neural networks on small mobile devices. *Expert Systems with Applications*, 39(1): 54–60.
- Kirchner, L.; Larson, J.; Mattu, S.; and Angwin, J. 2016. Machine Bias. ProPublica.
- Kleindessner, M.; Awasthi, P.; and Morgenstern, J. 2019. Fair k-Center Clustering for Data Summarization. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019*, 3448–3457. Long Beach, CA, USA.
- Liu, C.; and Boumal, N. 2019. Simple Algorithms for Optimization on Riemannian Manifolds with Constraints. *Applied Mathematics and Optimization*, 82: 949–981.
- Louizos, C.; Swersky, K.; Li, Y.; Welling, M.; and Zemel, R. S. 2016. The Variational Fair Autoencoder. In Bengio, Y.; and LeCun, Y., eds., *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.
- Madras, D.; Creager, E.; Pitassi, T.; and Zemel, R. 2018. Learning Adversarially Fair and Transferable Representations. In Dy, J.; and Krause, A., eds., *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, 3384–3393. PMLR.
- Mahdi Kamani, M.; Haddadpour, F.; Forsati, R.; and Mahdavi, M. 2019. Efficient Fair Principal Component Analysis. *arXiv e-prints*, arXiv:1911.04931.
- Massart, P. 2007. *Concentration Inequalities and Model Selection*. École d’Été de Probabilités de Saint-Flour. Springer-Verlag Berlin Heidelberg, 1 edition.
- Mehrabi, N.; Morstatter, F.; Saxena, N.; Lerman, K.; and Galstyan, A. 2019. A Survey on Bias and Fairness in Machine Learning. *arXiv e-prints*, arXiv:1908.09635.
- Meng, Z.; Chakraborty, R.; and Singh, V. 2021. An Online Riemannian PCA for Stochastic Canonical Correlation Analysis. *arXiv:2106.07479*.
- Nocedal, J.; and Wright, S. J. 2006. *Numerical optimization*. Springer Series in Operations Research and Financial Engineering. Springer-Verlag New York, 2 edition.
- Olfat, M.; and Aswani, A. 2019. Convex Formulations for Fair Principal Component Analysis. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, 663–670.
- Olfat, M.; Sloan, S.; Hespanhol, P.; Porter, M.; Vasudevan, R.; and Aswani, A. 2020. Covariance-robust dynamic watermarking. In *2020 59th IEEE Conference on Decision and Control (CDC)*, 3793–3799. IEEE.

- Oneto, L.; Donini, M.; Luise, G.; Ciliberto, C.; Maurer, A.; and Pontil, M. 2020. Exploiting MMD and Sinkhorn Divergences for Fair and Transferable Representation Learning. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M. F.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 15360–15370. Curran Associates, Inc.
- Pearson, K. 1901. LIII. On lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 2(11): 559–572.
- Pelegriana, G. D.; Brotto, R. D. B.; Duarte, L. T.; Romano, J. M. T.; and Attux, R. 2020. A multi-objective-based approach for Fair Principal Component Analysis. arXiv:2006.06137.
- Prost, F.; Qian, H.; Chen, Q.; Chi, E. H.; Chen, J.; and Beutel, A. 2019. Toward a better trade-off between performance and fairness with kernel-based distribution matching. In *NeurIPS 2019 Workshop on Machine Learning with Guarantees*. Vancouver, BC, Canada.
- Qi, X.; Luo, R.; and Zhao, H. 2013. Sparse principal component analysis by choice of norm. *Journal of Multivariate Analysis*, 114: 127 – 160.
- Quadrianto, N.; and Sharmanska, V. 2017. Recycling Privileged Learning and Distribution Matching for Fairness. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Ramdas, A.; Reddi, S. J.; Póczos, B.; Singh, A.; and Wasserman, L. 2015. On the Decreasing Power of Kernel and Distance Based Nonparametric Hypothesis Tests in High Dimensions. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, AAAI’15*, 3571–3577. AAAI Press.
- Rockafellar, R. T.; and Wets, R. J.-B. 1998. *Variational Analysis*. Grundlehren der mathematischen Wissenschaften. Springer-Verlag Berlin Heidelberg, 1 edition.
- Roh, Y.; Lee, K.; Whang, S. E.; and Suh, C. 2021. FairBatch: Batch Selection for Model Fairness. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Samadi, S.; Tantipongpipat, U. T.; Morgenstern, J. H.; Singh, M.; and Vempala, S. S. 2018. The Price of Fair PCA: One Extra dimension. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018*, 10999–11010. Montréal, Canada.
- Schölkopf, B.; Smola, A.; and Müller, K. 1998. Nonlinear Component Analysis as a Kernel Eigenvalue Problem. *Neural Computation*, 10(5): 1299–1319.
- Shamir, O. 2015. A Stochastic PCA and SVD Algorithm with an Exponential Convergence Rate. In Bach, F.; and Blei, D., eds., *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, 144–152. Lille, France: PMLR.
- Shumway, R. H.; and Stoffer, D. S. 2017. *Time Series Analysis and Its Applications: With R Examples*. Springer Texts in Statistics. Springer International Publishing, 4 edition.
- Sriperumbudur, B. K.; Gretton, A.; Fukumizu, K.; Lanckriet, G. R.; and Schölkopf, B. 2008. Injective Hilbert Space Embeddings of Probability Measures. In Servedio, R.; and Zhang, T., eds., *Proceedings of the 21st Annual Conference on Learning Theory (COLT)*, 111–222.
- Sriperumbudur, B. K.; Gretton, A.; Fukumizu, K.; Schölkopf, B.; and Lanckriet, G. R. 2010. Hilbert Space Embeddings and Metrics on Probability Measures. *Journal of Machine Learning Research*, 11(50): 1517–1561.
- Tantipongpipat, U.; Samadi, S.; Singh, M.; Morgenstern, J. H.; and Vempala, S. S. 2019. Multi-Criteria Dimensionality Reduction with Applications to Fairness. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019*, 15135–15145. Vancouver, BC, Canada.
- Yang, W. H.; Zhang, L.-H.; and Song, R. 2014. Optimality conditions for the nonlinear programming problems on Riemannian manifolds. *Pacific Journal of Optimization*, 10: 415–434.
- Zafar, M. B.; Valera, I.; Gomez-Rodriguez, M.; and Gummadi, K. P. 2017. Fairness Beyond Disparate Treatment & Disparate Impact: Learning Classification without Disparate Mistreatment. In *Proceedings of the 26th International Conference on World Wide Web, WWW 2017*, 1171–1180. Perth, Australia.
- Zalcborg, G.; and Wiesel, A. 2021. Fair Principal Component Analysis and Filter Design. *IEEE Transactions on Signal Processing*, 69: 4835–4842.
- Zemel, R. S.; Wu, Y.; Swersky, K.; Pitassi, T.; and Dwork, C. 2013. Learning Fair Representations. In *Proceedings of the 30th International Conference on Machine Learning, ICML 2013*, 325–333. Atlanta, GA, USA.




Figure 5: Comparison of fair PCA's as considered in (Samadi et al. 2018), and our work and (Olfat and Aswani 2019).

A Relation to (Samadi et al. 2018)

There is another branch work (Samadi et al. 2018; Tantipongpipat et al. 2019; Mahdi Kamani et al. 2019; Pelegrina et al. 2020) in which the task of solving fair PCA is defined as finding a (linear) subspace in which the reconstruction errors for each protected class are approximately equal. Here, we emphasize again that their fair PCA is *incomparable* with our considered fair PCA.

Consider the task of performing PCA on a two-dimensional data as shown in Figure 5, where $A \in \{0, 1\}$ denotes a protected attribute and the target dimension is set to $d = 1$. We assume that the data distributions for both $A = 0$ and $A = 1$ take an elliptical density, the only difference being their orientation and position.

Fairness in reconstruction error can be achieved by projecting the data onto blue or red subspace. However, fair PCA defined by Samadi et al. (2018) aims to *minimize* the reconstruction error with the fairness constraint, and thus the **blue** subspace is chosen. Under our MMD-based definition of fair PCA, projecting onto the blue subspace is considered unfair because the distributions conditioned on each protected attribute A are not the same; our fair PCA would choose the **red** subspace.

One important observation is that while our definition of fairness ensures that the downstream tasks are fair w.r.t many of existing fairness definitions (see Proposition 1), the reconstruction error-based definition has no such guarantees.

B Brief Description of FPCA

This section follows the discussion in Olfat and Aswani (2019), and by “they”, we refer to its authors. First recall their proposed adversarial definition of fairness for PCA:

Definition 6 (Δ_A -fairness; Olfat and Aswani, 2019). *Consider a fixed classifier $h(u, t) : \mathbb{R}^d \times \mathbb{R} \rightarrow \{0, 1\}$ that inputs features $u \in \mathbb{R}^d$ and a threshold t , and predicts the protected class $s \in \{0, 1\}$. Then, Π is $\Delta_A(h)$ -fair if*

$$\sup_{t \in \mathbb{R}} \left| \mathbb{P}[h(\Pi(x), t) = 1 | s = 1] - \mathbb{P}[h(\Pi(x), t) = 1 | s = 0] \right| \leq \Delta_A(h) \quad (2)$$

Moreover, for a family of classifiers \mathcal{F}_c , if Π is $\Delta_A(h)$ -fair for $\forall h \in \mathcal{F}_c$, we say that Π is $\Delta_A(\mathcal{F}_c)$ -fair.

B.1 Fairness Constraints

By choosing \mathcal{F}_c to be $\mathcal{F}_{lin} = \{h(u, t) = \mathbf{1}(w^\top u - t \leq 0) : w \in \mathbb{R}^d\}$, $\Delta_A(\mathcal{F}_{lin})$ becomes the Kolmogorov distance between $X'_s \triangleq w^\top V^\top X | A = s$, which is upper bounded by their total variation distance. Applying Pinsker's inequality (Massart 2007), we have that

$$\Delta_A(\mathcal{F}_{lin}) \leq \sqrt{\frac{1}{2} \text{KL}(X'_0 || X'_1)}$$

In order to obtain a tractable form of fairness constraints, they make normality assumptions on the data as follows: $X_s = X | A = s \sim \mathcal{N}_p(\mu_s, \Sigma_s)$. Then an upper bound of closed form can be obtained as follows:

$$\Delta_A(\mathcal{F}_{lin}) \leq \frac{1}{2} \sqrt{\frac{|\Sigma_1|}{|\Sigma_0|} + \text{tr}(\Sigma_1^{-1} \Sigma_0) - d + (\mu_1 - \mu_0)^\top \Sigma_1^{-1} (\mu_1 - \mu_0)},$$

and thus, the sufficient condition for the RHS to be 0 is that $\mu_0 = \mu_1$ and $\Sigma_0 = \Sigma_1$.

Denoting $f = \mu_1 - \mu_0$ and $Q = \Sigma_1 - \Sigma_0$, they introduced two constraints:

- Mean constraint: $h_m(V) := \|V^\top f\| = 0$
- Covariance constraint: $h_c(V) = \|V^\top Q V\|_2 = 0$

Based on these, the optimization considered in their work is as follows:

$$\begin{aligned} & \text{minimize} && f(V) = -\langle \Sigma, V V^\top \rangle \\ & \text{subject to} && V^\top V = I_d, \\ & && h_m(V) = 0, \\ & && h_c(V) = 0. \end{aligned} \quad (13)$$

B.2 SDP Formulation

Standard convex relaxation techniques (Boyd et al. 1994) were then used to derive an SDP w.r.t. a *new* variable $P = V V^\top \in \mathbb{R}^{p \times p}$:

$$\max \langle X^\top X, P \rangle - \mu t \quad (14a)$$

$$\text{s.t. } \text{tr}(P) \leq d \quad (14b)$$

$$I_p \succeq P \succeq 0 \quad (14c)$$

$$\langle P, f f^\top \rangle \leq \delta^2 \quad (14d)$$

$$\begin{bmatrix} t I_p & P M_+^\top \\ P M_+ & I_p \end{bmatrix} \succeq 0 \quad (14e)$$

$$\begin{bmatrix} t I_p & M_-^\top P \\ P M_- & I_p \end{bmatrix} \succeq 0 \quad (14f)$$

where $M_s M_s^\top$ is the Cholesky decomposition of $(-1)^s Q + \varphi I_p$, $\varphi \geq \|Q\|_2$, and $\delta, \mu \geq 0$ are the fairness tolerance levels⁷ for mean and covariance constraints, respectively. The final loading matrix V is obtained by extracting the top d eigenvectors from the resulting P^* .

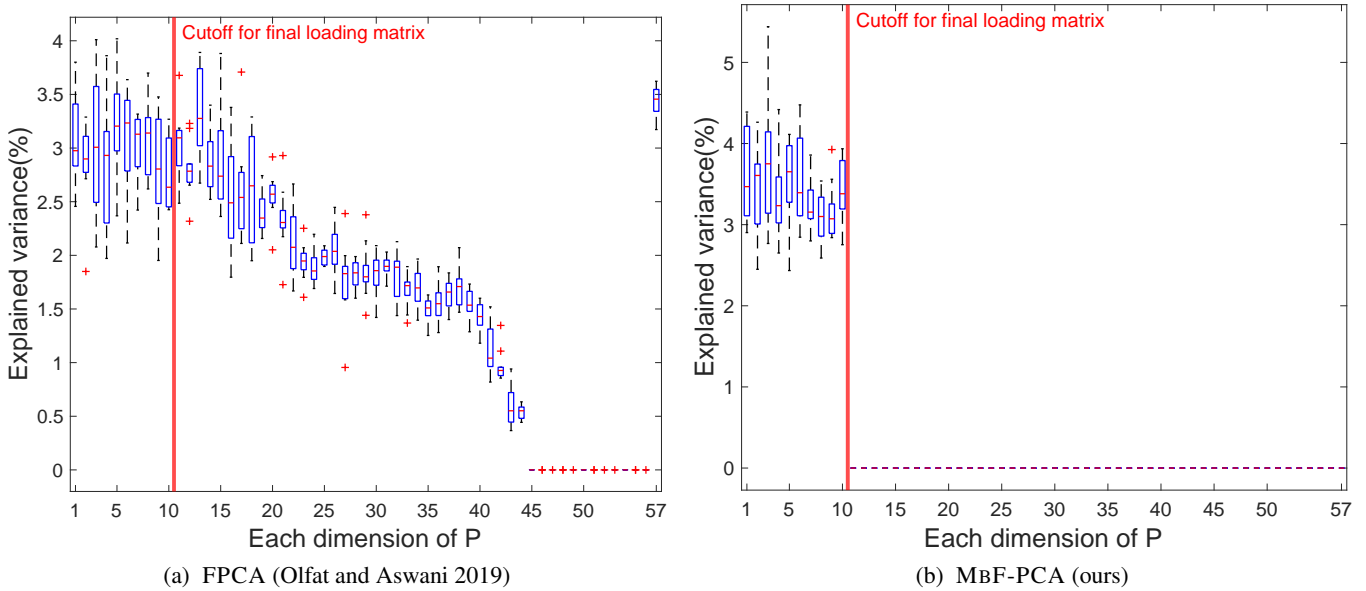


Figure 6: Explained variance of each eigenvector of P^* for GERMAN CREDIT DATASET, over the considered 10 train-test splits. Note how in FPCA’s case, there’s significant “leakage” of explained variance in the latter part (i.e. starting from 11-th eigenvector of P)

C Low Explained Variance of FPCA

C.1 Empirical Exploration

As seen in the Experiments, FPCA generally results in low explained variance than MBF-PCA, when *same(similar)* level of fairness is achieved. For clarity of exposition, we consider GERMAN CREDIT DATASET with $d = 10$ and the following hyperparameter settings: $\tau = 10^{-3}$ for MBF-PCA and $(\delta, \mu) = (0, 0.1)$ for FPCA. The reason for choosing a rather relaxed hyperparameter setting for MBF-PCA is to ensure a fair comparison i.e. to compare the two algorithms under the setting in which they achieve the *same* level of fairness.

To analyze this phenomenon in a detailed manner, we consider the variance explained of each eigenvector of P^* , *in order of the eigenvalues*. Here, P^* is the outputted $p \times p$ matrix from solving the SDP of FPCA. Figure 6 shows the plots of the explained variance of each eigenvector of P^* for FPCA and MBF-PCA, over the considered 10 train-test splits. (Since our algorithm directly outputs the loading matrix V , we simply set P to be VV^\top) Observe how the order of the eigenvalues **do not match** the order of explained variance, and the discrepancy is very high.

C.2 Theoretical Exploration

We now give a sketch of theoretical argument of why such phenomenon of misalignment in ordering of eigenvalues of P^* and explained variance of each eigenvector of P^* may occur. Again, we consider the Gaussian assumption of the data.

Let $\hat{\Sigma}$ be positive definite, and let $\hat{\Sigma} = \sum_{k=1}^p \lambda_k w_k w_k^\top$ be its spectral decomposition with $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_p > 0$ being the eigenvalues and w_k ’s being the corresponding orthonormal eigenvectors.

Also,

$$\text{tr}(V^\top \hat{\Sigma} V) = \sum_{l=1}^d v_l^\top \left\{ \sum_{k=1}^p \lambda_k w_k w_k^\top \right\} v_l \quad (15)$$

$$= \sum_{k=1}^p \sum_{l=1}^d \lambda_k \langle v_l, w_k \rangle^2 \quad (16)$$

$$= \sum_{k=1}^p \sum_{l=1}^d \lambda_k \cos^2 \theta_{v_l, w_k} \quad (17)$$

where $V^\top V = I_d$, v_l is the l -th column of V , and $\cos \theta_{u,v} := \langle u, v \rangle$. Eq. (15) shows that the maximization of $\text{tr}(V^\top \hat{\Sigma} V)$ can be thought of as an optimization problem w.r.t. *orthonormal* v_1, \dots, v_d , given λ_k ’s and w_k ’s. For example, with no constraint, the $\{v_l^*\}_{l=1}^d$ such that $\cos \theta_{v_l^*, w_l} = 1$ achieve the maximization of $\text{tr}(V^\top \hat{\Sigma} V)$ and the variances explained are ordered by the $\{\lambda_l\}_{l=1}^d$, which is precisely Eckart-Young theorem (Eckart and Young 1936).

When the covariance constraint $\|VQV^\top\|_2 = 0$ is imposed on the Eq. (15), we will show the effect of this quadratic constraint to the optimal V^* . The feasible space on the constraint $\|V^\top QV\|_2 = 0$ can be rewritten (Olfat and Aswani 2019): for some fixed $\varphi \geq \|Q\|_2$, the covariance condition $\|V^\top QV\|_2 = 0$, is equivalent to

$$\max \{ \|V^\top QV + \varphi I_d\|_2, \|V^\top QV - \varphi I_d\|_2 \} - \varphi = 0. \quad (18)$$

We can write $Q = \sum_{k=1}^p q_k s_k s_k^\top$, where q_k ’s and s_k ’s are

⁷analogous to τ in our approach.

real eigenvalues and real eigenvectors from the spectral decomposition of Q ; note that all of them are real due to the symmetry of Q , but since Q is not necessarily positive definite, q_k 's can be negative. On the other hand, by construction, we have that $q_k + \varphi, -q_k + \varphi \geq 0$ for all k 's. Here, recall that for symmetric matrices, spectral norm is equivalent to the largest absolute value of eigenvalues.

Rewriting the LHS of Eq. (18), we have that for either choice of $+$ and $-$,

$$\begin{aligned}
\|V^\top(Q \pm \varphi I_d)V\|_2 &= \left\| \sum_{k=1}^p (q_k \pm \varphi) V^\top s_k s_k^\top V \right\|_2 \\
&= \max_{\|x\|=1} \left| \sum_{k=1}^p (q_k \pm \varphi) x^\top V^\top s_k s_k^\top V x \right| \\
&= \max_{\|x\|=1} \left| \sum_{k=1}^p (\varphi \pm q_k) (x^\top V^\top s_k)^2 \right| \\
&= \max_{\|x\|=1} \sum_{k=1}^p (\varphi \pm q_k) (x^\top V^\top s_k)^2 \\
&\leq \sum_{k=1}^p (\varphi \pm q_k) \max_{\|x\|=1} (x^\top V^\top s_k)^2 \\
&= \sum_{k=1}^p (\varphi \pm q_k) \|V^\top s_k\|^2 \\
&= \sum_{k=1}^p (\varphi \pm q_k) \sum_{l=1}^d \cos^2 \theta_{v_l, s_k} \\
&\leq \left(\max_k (\varphi \pm q_k) \right) \sum_{k=1}^p \sum_{l=1}^d \cos^2 \theta_{v_l, s_k}.
\end{aligned}$$

Above and Eq. (18) imply that

$$\begin{aligned}
\varphi &= \max \{ \|V^\top Q V + \varphi I_d\|_2, \|V^\top Q V - \varphi I_d\|_2 \} \\
&\leq \max \left\{ \left(\max_k (\varphi + q_k) \right) \sum_{k=1}^p \sum_{l=1}^d \cos^2 \theta_{v_l, s_k}, \right. \\
&\quad \left. \left(\max_k (\varphi - q_k) \right) \sum_{k=1}^p \sum_{l=1}^d \cos^2 \theta_{v_l, s_k} \right\} \\
&= \max \{ \varphi + q_1, \varphi - q_p \} \sum_{k=1}^p \sum_{l=1}^d \cos^2 \theta_{v_l, s_k} \\
&= (\varphi + \max \{ q_1, -q_p \}) \sum_{k=1}^p \sum_{l=1}^d \cos^2 \theta_{v_l, s_k}
\end{aligned}$$

If $\frac{\varphi}{\varphi + \max\{q_1, -q_p\}}$ is sufficiently large enough, then we would be constraining the sum of square of cosines to be somewhat close to 1 i.e. forcing some of the to-be-found eigenvectors v_l 's into directions such that they are somewhat orthogonal to s_k 's, the eigenvectors of Q . Such restriction becomes more prominent as d increases. All in all, the presence of the *covariance (quadratic)* constraint may cause a misalignment in the ordering of the explained variance of

each loading vector, potentially causing the ordering of the explained variance to be substantially different from the ordering of the eigenvalues.

D Theoretical Minimum

In this section, we provide (minimum) required background for understanding the proofs of propositions/theorems presented in the main paper. All proofs of the results here are deferred to the respective references and therein.

D.1 A Primer on (Matrix) Manifolds, and Stiefel Manifold

Preliminaries (Here, we mainly follow the exposition from the textbook by (Absil, Mahony, and Sepulchre 2007))

Let \mathcal{M} be a smooth, finite-dimensional Riemannian manifold, endowed with a Riemannian metric $\langle \cdot, \cdot \rangle_x$ on each tangent space $T_x \mathcal{M}$ and $x \in \mathcal{M}$. Let $\|\cdot\|_x$ be the associated norm. We omit x when it is clear from context. Let \exp_x be the exponential map at x , and let dist be the Riemannian distance.

Let $f : \mathcal{M} \rightarrow \mathbb{R}$ be a smooth function. We first consider the gradient of f , defined on manifolds::

Definition 9. *The Riemannian gradient at $x \in \mathcal{M}$ of f , denoted as $\text{grad } f(x)$, is the unique element of $T_x \mathcal{M}$ such that*

$$\forall \xi \in T_x \mathcal{M} \quad \langle \text{grad } f(x), \xi \rangle_x = Df(x)[\xi]$$

where $Df(x)[v]$ is the directional derivative of f at x along v .

Finding a closed-form of grad may be difficult in some cases. However, if f can be extended to the ambient Euclidean space $\mathbb{R}^{p \times d}$, then grad can be computed as follows:

$$\text{grad } f(x) = \mathbf{P}_{T_x \mathcal{M}}(\nabla f(x)) \quad (19)$$

where ∇ is the usual Euclidean gradient, and $\mathbf{P}_{T_x \mathcal{M}}$ is the projection operator onto $T_x \mathcal{M}$.

Matrix Manifolds Matrix is a mathematical object of fundamental importance. Often in applications in machine learning and optimization, we are given the task of finding some matrix satisfying certain conditions, such as orthonormality, while optimizing some real-valued function of the matrices (Hu et al. 2020). In the perspective of optimization, it is natural to see whether an intrinsic geometric structure can be found in a set of matrices with certain common properties. It turns out that such structures can be characterized by manifold, hence the term *matrix manifold*.

To start, let us consider $\mathbb{R}^{p \times d}$, a set of all possible real-valued $p \times d$ matrices. Along with the Frobenius inner product, defined as $\langle X, Y \rangle := \text{tr}(X^\top Y)$, $\mathbb{R}^{n \times p}$ can be regarded as a Euclidean space, and thus a Riemannian manifold.

Many of the well-known matrix manifolds are actually *embedded Riemannian submanifold* of $\mathbb{R}^{p \times d}$ i.e. its subspace topology is exactly the given topology on $\mathbb{R}^{p \times d}$. We shall discuss about one specific matrix manifold that is crucial to our MBF-PCA in the next subsection.

Stiefel Manifold We are mainly concerned with the *Stiefel manifold*, defined as follows⁸:

$$St(p, d) = \{V \in St(p, d) \mid V^\top V = I_d\} \quad (20)$$

By considering the mapping $F(V) = V^\top V - I_d$, we can prove the following result:

Theorem 3. *St(p, d), defined as above, is an embedded (Riemannian) submanifold of $\mathbb{R}^{p \times d}$.*

Now let us look at several properties of $St(p, d)$, many of which prove itself to be effective in both theoretical and practical perspective.

Theorem 4. *St(p, d) is (p - d - 1)-connected.*

Corollary 4.1. *When p > d, St(p, d) is always at least path-connected, which implies connectedness.*

Theorem 5. *St(p, d) is compact.*

Lemma 2. *Projection of any X ∈ ℝ^{p×d} onto T_VSt(p, d) is given as*

$$\mathbf{P}_{T_V St(p, d)}(X) = X - V \text{sym}(V^\top X) \quad (21)$$

where $\text{sym}(X) = \frac{1}{2}(X + X^\top)$.

E Feasibility of Fair PCA

Here, we briefly discuss the feasibility problem of fair PCA.

Let us consider Eq. (13), which is a special case of our formulation of fair PCA under Gaussian assumption. For our analysis, we consider Eq. (13) as a constrained manifold optimization over $St(p, d)$; geometric properties of manifolds turn out to be crucial in obtaining to-be-presented results.

When only the mean constraint is in place i.e. assume that $Q = O$, then the feasibility of Eq. (13) is always guaranteed:

Proposition 4. *If Q = O, then Eq. (13) is always feasible.*

Proof. We start by observing that the mean constraint, $V^\top f = 0$, is a linear hyperplane passing through the origin 0, and the manifold constraint⁹, $V^\top V = I_d$, is a manifold that is at least path-connected. Now suppose that there exists some f such that the linear hyperplane does not intersect $St(p, d)$. Let $V \in St(p, d)$ be on the “left” side of the hyperplane. Since $St(p, d)$ is symmetric about the origin ($V^\top V = I_d \Rightarrow (-V)^\top(-V) = I_d$), $-V$ is on the “right” side of the hyperplane. Then there must exist some path on $St(p, d)$ connected V and $-V$, which should intersect the hyperplane, a contradiction. \square

When covariance constraint is also in-place, the problem becomes more complex. However, especially in high dimensions, it is expected that the feasibility assumption is not so far-off. As such discussion is not limited to Gaussian case of Eq. (13), from hereon and forth we consider the following assumption:

Assumption 3. *Eq. (8) is feasible.*

⁸Another characterization of $St(p, d)$ is as a quotient space of set of d linearly independent p -dimensional vectors.

⁹In manifold optimization framework, this shouldn’t be considered as an explicit constraint; we use such words here for clarity.

F Proof of Theoretical Guarantee of Algorithm 1

F.1 Assumptions

Let us start by formally stating Assumption 2: for each outputted solution of each Eq. (9), V_{k+1} , consider the following induced sequence W_{k+1} defined as:

$$W_{k+1} := \underset{\substack{V \in St(p, d) \\ \|\text{grad} \mathcal{Q}(V, \rho_k)\| = 0}}{\text{argmin}} \|V - V_{k+1}\|. \quad (22)$$

Assumption 4 (formal; locality assumption). *For each k, W_{k+1} is a local minimum of Eq. (9) i.e. for some fixed $r > 0$, any $W \in St(p, d) \cap B_r(W_{k+1})$ satisfies $\mathcal{Q}(W_{k+1}, \rho_{k+1}) \leq \mathcal{Q}(W, \rho_{k+1})$.*

For simplicity, we omit the dependency of W_{k+1} on V_{k+1} and we write k instead of $k + 1$. Also, let us define

$$\tilde{r} = \tilde{r}(\epsilon_{min}) := r - g(\epsilon_{min}). \quad (23)$$

F.2 Lemmas

Let us first establish the existence of limit point of $\{V_k\}$:

Lemma 3. *$\{V_k\}$ has at least one limit point.*

Proof. Noting that for all k ’s $V_k \in St(p, d)$, and that $St(p, d)$ is compact i.e. closed and bounded, by Bolzano-Weierstrass theorem, $\{V_k\}$ must have some limit point, which we denote as $\bar{V} \in St(p, d)$. \square

Next lemma establishes an important connection between the two mentioned sequences, namely that there exists a limit point of $\{W_k\}$ that is very close (possibly equal) to \bar{V} :

Lemma 4. *Let \bar{V} be any limit point of $\{V_k\}$. Then there exists some limit point \bar{W} of $\{W_k\}$ and some decreasing, continuous function $g : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ satisfying $g(0) = 0$, such that*

$$\|\bar{V} - \bar{W}\| \leq g(\epsilon_{min}). \quad (24)$$

Proof. Since $W_k \in St(p, d)$ for all k , there exists some subsequence $\{i_k\} \subset \{k\}$ such that $W_{i_k} \rightarrow \bar{W}$ for some $\bar{W} \in St(p, d)$, by Bolzano-Weierstrass theorem. Obviously, $V_{i_k} \rightarrow \bar{V}$. For simplicity, let us denote V_{i_k} and W_{i_k} as V_k and W_k , respectively.

Then we denote

$$\left\{ V \in St(p, d) : \left\| \frac{1}{\rho_k} \text{grad} f(V) + \text{grad} h(V) \right\| \leq \frac{\epsilon}{\rho_k} \right\} \quad (25)$$

by $G_k(\epsilon)$. As $\rho_k < \infty$, above is related to Algorithm 1 as

$$G_k(\epsilon) = \{V \in St(p, d) : \|\text{grad} \mathcal{Q}(V, \rho_k)\| \leq \epsilon\}, \quad (26)$$

where

$$\mathcal{Q}(V, \rho_k) = f(V) + \rho_k h(V). \quad (27)$$

One crucial observation is that $G_k(\epsilon) \subseteq G_k(\epsilon')$ for $\epsilon \leq \epsilon'$.

Let $\epsilon > 0$ be arbitrary. By given, there exists some $K > 0$ such that for any $k > K$, $\|V_k - \bar{V}\| < \epsilon/2$ and $\|W_k - \bar{W}\| < \epsilon/2$. By triangle inequality,

$$\begin{aligned} \|\bar{V} - \bar{W}\| &\leq \|V_k - \bar{V}\| + \|W_k - \bar{W}\| + \|V_k - W_k\| \\ &< \epsilon + \|V_k - W_k\|. \end{aligned}$$

We can further bound the second term as follows:

$$\begin{aligned} \|V_k - W_k\| &= \min_{W \in G_k(0)} \|V_k - W\| \\ &\leq \max_{V \in G_k(\epsilon_k)} \min_{W \in G_k(0)} \|V - W\| \end{aligned}$$

where the first equality follows from the definition of $\{W_k\}$, and the last inequality follows from the fact that our algorithm satisfies $V_k \in G_k(\epsilon_k)$.

Note that since $G_k(0) \subseteq G_k(\epsilon_k)$, above is the variational form of the Hausdorff distance (Rockafellar and Wets 1998) between $G_k(0)$ and $G_k(\epsilon_k)$, which we denote as $d_H(\cdot, \cdot)$. Additionally, we denote

$$\left\{ V \in St(p, d) : \left\| \frac{1}{\rho} \text{grad}f(V) + \text{grad}h(V) \right\| \leq \frac{\epsilon}{\rho} \right\} \quad (28)$$

by $G_\infty(\epsilon)$, where $\rho := \lim_{k \rightarrow \infty} \rho_k \leq \rho_{max} \leq \infty$ and $\frac{1}{\infty} = 0$.

Then it is easy to see that as $k \rightarrow \infty$, $G_k(0) \rightarrow G_\infty(0)$, $G_k(\epsilon_k) \rightarrow G_\infty(\epsilon_{min})$, and $G_\infty(0) \subset G_\infty(\epsilon_{min})$. Now define

$$g(\epsilon) := d_H(G_\infty(0), G_\infty(\epsilon)). \quad (29)$$

Then g is continuous and decreasing with respect to ϵ since $G_\infty(\epsilon) \subseteq G_\infty(\epsilon')$ for $\epsilon \leq \epsilon'$, and it satisfies $g(0) = 0$.

Since $\epsilon > 0$ was arbitrary, the statement follows. \square

Now let us prove the Lipschitzness of the functions considered in our fair PCA:

Lemma 5. f is $2\|\Sigma\|$ -Lipschitz.

Proof. For all $V, W \in St(p, d)$,

$$\begin{aligned} &\|V^\top \Sigma V - W^\top \Sigma W\| \\ &\leq \|V^\top \Sigma (V - W)\| + \|(V - W)^\top \Sigma W\| \\ &\leq \|V\| \|\Sigma\| \|V - W\| + \|V - W\| \|\Sigma\| \|W\| \\ &\leq 2\|\Sigma\| \|V - W\|. \end{aligned}$$

\square

Lemma 6. There exists some $L_h > 0$ such that $h : St(p, d) \rightarrow \mathbb{R}$ is L_h -Lipschitz.

Proof. Suppose not. Then for any $L_h > 0$, there exists some $V, W \in St(p, d)$ such that $|h(V) - h(W)| > L_h \|V - W\|$. Especially for $W \neq V$, we have that $L_h < \frac{|h(V) - h(W)|}{\|V - W\|}$. Letting $W \rightarrow V$, we have that $L_h < |\nabla_V h(V)|$. This implies that $\nabla_x h$ is unbounded on $St(p, d)$, a contradiction since the continuity of $\nabla_V h$ and compactness of $St(p, d)$ implies that $\nabla_V h$ is bounded by the Extreme Value theorem. \square

As done in proof of Lemma 4, let us define $\rho \triangleq \lim_{k \rightarrow \infty} \rho_k$. We now show how ρ_{max}, τ affect ρ :

Lemma 7. $\rho = \rho(\rho_{max}, \tau)$ is characterized as follows: there exists some function $K := K(\tau) \in \mathbb{N}$, satisfying $K < \infty$ for $\tau > h(\bar{V})$ and $K \leq \infty$ otherwise, such that

$$\rho = \min(\theta_\rho^K \rho_0, \rho_{max}). \quad (30)$$

Moreover, K is decreasing in τ and increasing in ρ_{max} .

Proof. Define $K(\tau) := |\{k \in \mathbb{N} : h(V_k) > \tau\}|$, which is the number of ρ -updates in the algorithm. From the algorithm, Eq. (30) is directly implied. Now it remains to characterize the function $K(\tau)$.

Let us first consider the case of $\tau > h(\bar{V})$. Since $h(V_k) \rightarrow h(\bar{V})$ and h is continuous, it must be that $K(\tau) < \infty$ by the algorithm, regardless of the other hyperparameters. Now suppose that $\tau \leq h(\bar{V})$. Then there are two cases to consider:

1. If $h(V_M) \leq \tau$ for some $M < \infty$, then from the algorithm it is clear that $K(\tau) \leq M$.
2. If $h(V_k) > \tau$ for all k , then again, from the algorithm it is clear that $K(\tau) = \infty$.

By construction, $K(\tau)$ is decreasing in τ , and thus, ρ is decreasing in τ and increasing in ρ_{max} . \square

Corollary 5.1. If no k satisfies $h(V_k) \leq \tau$, then $\rho = \infty$ if and only if $\rho_{max} = \infty, \tau = 0$.

Denote $\mathcal{C} \triangleq St(p, d) \cap h^{-1}(0)$ as the set of feasible points. By compactness of $St(p, d)$ and continuity of h , we have that \mathcal{C} is also compact. By Assumption 3, \mathcal{C} is non-empty. Lastly, recall from Eq. (23) $\tilde{r} = r - g(\epsilon_{min})$. The last lemma asserts that after sufficiently many iterations have passed, the sequences $\{V_k\}$ and $\{W_k\}$ are always close to \mathcal{C} .

Lemma 8. If \bar{V} is fair and ϵ_{min} is sufficiently small in the sense that $g(\epsilon_{min}) < r$, then there exists some $K > 0$ such that for all $k > K$, $B_{\tilde{r}}(V_k) \cap \mathcal{C} \neq \emptyset$ and $B_r(W_k) \cap \mathcal{C} \neq \emptyset$.

Proof. First, suppose that for all k , $B_{\tilde{r}}(V_k) \cap \mathcal{C} = \emptyset$. Taking the limit of $k \rightarrow \infty$ on both sides, we have that $\bar{V} \in B_{\tilde{r}}(\bar{V}) \cap \mathcal{C} = \emptyset$ by the given assumption, a contradiction. Thus there exists some K_1 such that for all $k > K_1$, $B_{\tilde{r}}(V_k) \cap \mathcal{C} \neq \emptyset$.

Now suppose that for all k , $B_r(W_k) \cap \mathcal{C} = \emptyset$. Recall from Lemma 4 that

$$\|\bar{W} - \bar{V}\| \leq g(\epsilon_{min}) < r$$

Again taking the limit of $k \rightarrow \infty$, we have that $\bar{V} \in \mathcal{C} \cap B_r(\bar{W}) = \emptyset$, a contradiction. Thus there exists some K_2 such that for all $k > K_2$, $B_r(W_k) \cap \mathcal{C} \neq \emptyset$.

Set $K = \max(K_1, K_2)$, and we are done. \square

F.3 Proof of Theorem 2

(The proof is inspired from the optimality guarantee for the Euclidean quadratic penalty method; see Theorem 17.1 of Nocedal and Wright, 2006) The proof is divided into 3 parts.

Existence of limit point of $\{W_k\}$

By Lemma 4, there exists some limit point \bar{W} of $\{W_k\}$ such that $\|\bar{V} - \bar{W}\| \leq g(\epsilon_{min})$. By considering appropriate subsequence, let $W_k \rightarrow \bar{W}$.

\bar{V} is approximate local minimizer of (10)

By Assumption 4, we have that for all $W \in B_r(W_k) \cap St(p, d)$,

$$\mathcal{Q}(W_k, \rho_k) \leq \mathcal{Q}(W, \rho_k) \quad (31)$$

i.e.

$$f(W_k) + \rho_k h(W_k) \leq f(W) + \rho_k h(W). \quad (32)$$

Rearranging gives

$$h(W_k) \leq h(W) + \frac{1}{\rho_k} (f(W) - f(W_k)) \quad (33)$$

Taking the limit $k \rightarrow \infty$ gives

$$\begin{aligned} h(\bar{W}) &\leq h(W) + \frac{1}{\rho} (f(W) - f(\bar{W})) \\ &\leq h(W) + \frac{2\|\Sigma\|}{\rho} \|W - \bar{W}\|, \end{aligned}$$

for all $W \in B_r(\bar{W}) \cap St(p, d)$.

By Lemma 4, $B_{\bar{r}}(\bar{V}) \cap St(p, d) \subset B_r(\bar{W}) \cap St(p, d)$.

Thus we have that for all $V \in B_{\bar{r}}(\bar{V}) \cap St(p, d)$,

$$\begin{aligned} h(\bar{V}) &\leq h(\bar{W}) + |h(\bar{V}) - h(\bar{W})| \\ &= h(V) + (h(\bar{W}) - h(V)) + L_h \|\bar{V} - \bar{W}\| \\ &\leq h(V) + \frac{2\|\Sigma\|}{\rho} \|V - \bar{W}\| + L_h \|\bar{V} - \bar{W}\| \\ &\leq h(V) + \frac{2\|\Sigma\|}{\rho} (\|V - \bar{V}\| + \|\bar{V} - \bar{W}\|) + L_h \|\bar{V} - \bar{W}\| \\ &= h(V) + \frac{2\|\Sigma\|}{\rho} \|V - \bar{V}\| + \left(L_h + \frac{2\|\Sigma\|}{\rho} \right) g(\epsilon_{min}). \end{aligned}$$

Define $\beta(\rho_{max}, \tau) := \frac{2\|\Sigma\|}{\rho(\rho_{max}, \tau)}$. By Lemma 7, β indeed satisfies $\lim_{\rho_{max} \rightarrow \infty} \rho(\rho_{max}, 0) \geq 0$ and $\lim_{\rho_{max} \rightarrow \infty} \rho(\rho_{max}, \tau) > 0$ for $\tau > 0$. Also, by the same lemma, β is indeed decreasing in τ and increasing in ρ_{max} . (Note how this particular part of the proof does *not* rely on the assumption that \bar{V} is feasible!)

If \bar{V} is fair, then \bar{V} is an approximate local minimizer of Eq. (8)

By Lemma 8, we may only consider k 's such that $B_{\bar{r}}(V_k) \cap \mathcal{C} \neq \emptyset$ and $B_r(W_k) \cap \mathcal{C} \neq \emptyset$. By Assumption 4, we have that for all $W \in B_r(W_k) \cap \mathcal{C}$,

$$\mathcal{Q}(W_k, \rho_k) \leq \mathcal{Q}(W, \rho_k) \quad (34)$$

i.e.

$$f(W_k) \leq f(W_k) + \rho_k h(W_k) \leq f(W) + \rho_k h(W) = f(W), \quad (35)$$

where the last equality follows from $W \in \mathcal{C}$.

Thus we have that

$$\begin{aligned} f(\bar{V}) &\leq f(\bar{W}) + |f(\bar{V}) - f(\bar{W})| \\ &\leq \lim_{k \rightarrow \infty} f(W_k) + 2\|\Sigma\| \|\bar{V} - \bar{W}\| \\ &\leq f(W) + 2\|\Sigma\| g(\epsilon_{min}), \end{aligned}$$

for all $W \in B_r(W_k) \cap \mathcal{C}$. Specifically, since $B_{\bar{r}}(\bar{V}) \cap St(p, d) \subset B_r(\bar{W}) \cap St(p, d)$ by Lemma 4, we have that

$$f(\bar{V}) \leq f(V) + 2\|\Sigma\| g(\epsilon_{min}),$$

for all $V \in B_{\bar{r}}(\bar{V}) \cap \mathcal{C}$.

E.4 Proof of Theorem 1

Set $\epsilon_{min} = 0$, $\tau = 0$, and $\rho_{max} = \infty$ in the above proof. (Note that $g(0) = 0$ implies that $\bar{V} = \bar{W}$) Just as in proof of Lemma 7, if there exists some $M < \infty$ such that $h(V_M) = 0$, then from the algorithm it is clear that \bar{V} is fair and we are done. If not, then we must have that $\rho(\infty, 0) = \infty$ i.e. $\beta(\infty, 0) = 0$ i.e. \bar{V} is a local minimizer of Eq. 10.

G Closed form of $\nabla_V h(V)$

In this section, we derive the closed form of the constraint gradient.

First, we recall the lemma that provides a closed form of $h(V) = \text{MMD}^2(\cdot, \cdot)$:

Lemma 1 (Gretton et al., 2007). $\hat{\Delta}$ is computed as follows:

$$\hat{\Delta} = \left[\frac{1}{m^2} \sum_{i,j=1}^m k(X_i, X_j) + \frac{1}{n^2} \sum_{i,j=1}^n k(Y_i, Y_j) - \frac{2}{mn} \sum_{i,j=1}^{m,n} k(X_i, Y_j) \right]^{1/2}. \quad (5)$$

For simplicity, we use the following abbreviations: $\sum_{i,j=1}^m := \sum_{i=1}^m \sum_{j=1}^m$ and $\sum_{i,j=1}^{m,n} := \sum_{i=1}^m \sum_{j=1}^n$.

After squaring both sides, let us refer to the three terms as h_1, h_2 , and h_3 i.e. $h(V) = \hat{\Delta}^2 = h_1(V) + h_2(V) - 2h_3(V)$ where

$$h_1(V) = \frac{1}{m^2} \sum_{i,j=1}^m k(VX_i, VX_j),$$

$$h_2(V) = \frac{1}{n^2} \sum_{i,j=1}^n k(VY_i, VY_j),$$

$$h_3(V) = \frac{1}{mn} \sum_{i,j=1}^{m,n} k(VX_i, VY_j).$$

G.1 Closed form of $\nabla_V h_1(V)$

Denote

$$g_{ij}(V) = \text{tr}((X_i - X_j)^\top V V^\top (X_i - X_j))$$

$$K_{ij}(V) = \exp(-g_{ij}(V)/(2\sigma^2))$$

$$K_X(V) = \{K_{ij}(V)\} \in \mathbb{R}^{m \times m},$$

and X be the data matrix whose i -th row is X_i^\top . By chain rule, we have that

$$\begin{aligned} \nabla_V h_1(V) &= -\frac{1}{2m^2\sigma^2} \sum_{i,j=1}^m K_{ij}(V) \nabla_V g_{ij}(V) \\ &= -\frac{1}{m^2\sigma^2} \sum_{i,j=1}^m K_{ij}(V) (X_i - X_j)(X_i - X_j)^\top V \end{aligned}$$

Now denote $H_i(V) = \sum_{j=1}^m K_{ij}(V)$ and $H_X(V) = \text{diag}(H_1(V), H_2(V), \dots, H_m(V))$, where ‘diag’ is the block-diagonal operator. We have that

$$\begin{aligned} \sum_{i,j=1}^m K_{ij}(V) X_i X_i^\top &= \sum_{i=1}^m \left(\sum_{j=1}^m K_{ij}(V) \right) X_i X_i^\top \\ &= X^\top H_X(V) X, \end{aligned}$$

and

$$\sum_{i,j=1}^m K_{ij}(V) X_i X_j^\top = X^\top K_X(V) X.$$

Since $K_X(V)$ is symmetric i.e. $K_{ij}(V) = K_{ji}(V)$,

$$\sum_{i,j=1}^m K_{ij}(V) X_j X_j^\top = \sum_{i,j=1}^m K_{ji}(V) X_j X_j^\top = X^\top H_X(V) X,$$

and

$$\sum_{i,j=1}^m K_{ij}(V) X_j X_i^\top = \sum_{i,j=1}^m K_{ji}(V) X_j X_i^\top = X^\top K_X(V) X.$$

Thus we have that

$$\nabla_V h_1(V) = -\frac{2}{m^2 \sigma^2} X^\top (H_X(V) - K_X(V)) X V.$$

G.2 Closed form of $\nabla_V h_2(V)$

The computation is almost exactly the same as above, and thus we only show the final result:

$$\nabla_V h_2(V) = -\frac{2}{n^2 \sigma^2} Y^\top (H_Y(V) - K_Y(V)) Y V,$$

where Y , $H_Y(V)$, and $K_Y(V)$ are defined analogously.

G.3 Closed form of $\nabla_V h_3(V)$

Let us overload the notation a bit and now denote

$$\begin{aligned} g_{ij}(V) &= \text{tr}((X_i - Y_j)^\top V V^\top (X_i - Y_j)), \\ K_{ij}(V) &= \exp(-g_{ij}(V)/(2\sigma^2)), \\ K_{XY}(V) &= \{K_{ij}(V)\} \in \mathbb{R}^{m \times n}. \end{aligned}$$

Again by chain rule, we have that

$$\begin{aligned} \nabla_V h_3(V) &= -\frac{1}{2mn\sigma^2} \sum_{i,j=1}^{m,n} K_{ij}(V) \nabla_V g_{ij}(V) \\ &= -\frac{1}{mm\sigma^2} \sum_{i,j=1}^{m,n} K_{ij}(V) (X_i - Y_j)(X_i - Y_j)^\top V. \end{aligned}$$

Now, again with a slight notation overload, denote

$$\begin{aligned} H_i(V) &= \sum_{j=1}^n K_{ij}(V), \\ H_{XY}(V) &= \text{diag}(H_1(V), H_2(V), \dots, H_m(V)), \\ \tilde{H}_j(V) &= \sum_{i=1}^m K_{ij}(V), \\ \tilde{H}_{XY}(V) &= \text{diag}(\tilde{H}_1(V), \tilde{H}_2(V), \dots, \tilde{H}_n(V)). \end{aligned}$$

We have that

$$\begin{aligned} \sum_{i,j=1}^{m,n} K_{ij}(V) X_i X_i^\top &= \sum_{i=1}^m \left(\sum_{j=1}^n K_{ij}(V) \right) X_i X_i^\top \\ &= X^\top H_{XY}(V) X, \end{aligned}$$

and

$$\begin{aligned} \sum_{i,j=1}^{m,n} K_{ij}(V) Y_j Y_j^\top &= \sum_{j=1}^n \left(\sum_{i=1}^m K_{ij}(V) \right) Y_j Y_j^\top \\ &= Y^\top \tilde{H}_{XY}(V) Y. \end{aligned}$$

Next, we have that

$$\sum_{i,j=1}^{m,n} K_{ij}(V) X_i Y_j^\top = X^\top K_{XY}(V) Y,$$

and

$$\sum_{i,j=1}^{m,n} K_{ij}(V) Y_j X_i^\top = Y^\top K_{XY}(V)^\top X.$$

Thus we have that

$$\begin{aligned} \nabla_V h_3(V) &= -\frac{1}{mn\sigma^2} \left(X^\top H_{XY}(V) X + Y^\top \tilde{H}_{XY}(V) Y \right. \\ &\quad \left. - X^\top K_{XY}(V) Y - Y^\top K_{XY}(V)^\top X \right) V. \end{aligned}$$

H Description of 2nd Part of Experiment 1

With $p_0 = 1000$ and $n = 250$, we first constructed two different p_0 -variate Gaussian distributions $\mathcal{N}_{p_0}(\mu_{p_0}^{(s)}, \Sigma_{p_0}^{(s)})$ and sampled n points for each $s \in \{0, 1\}$. Then for each considered p , we projected the samples by multiplying the constructed data matrices with a Gaussian random matrix of size $p_0 \times p$. This creates 18 pairs of p -dimensional distributions. This allowed for us model the situation in which two protected groups have different *non-Gaussian* distributions.

We now describe how we've set $\mu_{p_0}^{(s)}$ and $\Sigma_{p_0}^{(s)}$ for each s . The mean difference is set as $f_{p_0} = \mu_{p_0}^{(1)} - \mu_{p_0}^{(0)}$ where $2 * \frac{f_{p_0}^{raw}}{\|f_{p_0}^{raw}\|}$ and $f_{p_0}^{raw} = \mathbf{1}$. Then $\mu_{p_0}^{(0)}$ is set as $\mathbf{0}$ and $\mu_{p_0}^{(1)}$ is set as $\mu_{p_0}^{(0)} + f_{p_0}$.

The covariance difference is set as $Q_p^{raw} / \|Q_p^{raw}\|_2$ where

$$Q_p^{raw} = A_p^{(1)} - A_p^{(0)},$$

$$\begin{aligned} \Sigma_p^{(0)} &= \text{diag}(AR_{p/5}(0.99), AR_{p/5}(0.98), AR_{p/5}(0.97), \\ &\quad AR_{p/5}(0.98), AR_{p/5}(0.95)), \end{aligned}$$

and

$$\begin{aligned} \Sigma_p^{(1)} &= \text{diag}(AR_{p/5}(0.99), AR_{p/5}(0.98), AR_{p/5}(0.97), \\ &\quad AR_{p/5}(0.98), AR_{p/5}(0.99)). \end{aligned}$$

Here, ‘diag’ is the block-diagonal operator, and

$$(AR_{p/5}(r))_{ij} = \frac{r^{|i-j|}}{1-r^2}, \quad i, j \in \{1, \dots, p/5\} \quad (36)$$

is the covariance matrix of a $p/5$ -variate Gaussian AR(1) process (Shumway and Stoffer 2017). The $r \in (0, 1)$ is a parameter that controls the strength of correlation. Finally, $\Sigma_p^{(0)}$ is set as $A_p^{(s)} / \|Q_p\|$.

Due to the scaling, $\|f_p\| = 2$ and $\|Q_p\| = 1$ for all p . This is to ensure that $\|f_p\|$ and $\|Q_p\|$ do not explode as p increases, allowing for us to see the effect of high dimensions more clearly under the “same” fairness setting.

I Description of Experiment 2

Here, we describe the full pre-processing steps for the three UCI datasets considered in this paper. All pre-processing steps were done using AI Fairness 360 (AIF360; Bellamy et al., 2018) built-in functionalities. In our Github repository¹⁰, we include the links for downloading the raw datasets, along with the pre-processing codes.

I.1 COMPAS dataset

The raw data was randomly reduced to 40%, resulting in 2468 samples. This was to avoid computational burden in computing the MMD² metric, as it scales quadratically in number of samples. ‘Race’ was set as the protected attribute, and the task is to classify the bank account holders into credit class *Good* or *Bad*. The features ‘sex’ and ‘c.charge_desc’ were dropped, resulting in total of 11 features.

I.2 German credit dataset

The total number of samples was 1000. ‘Age’ was set as the protected attribute as done in (Kamiran and Calders 2011), and the task is to classify whether the individual bank account holder has good or bad credit class. The features ‘sex’ and ‘personal_status’ were dropped, resulting in total of 57 features.

I.3 Adult income dataset

The raw data was randomly reduced to 5%, resulting in 2261 samples; this was due to the same reason as the COMPAS dataset. ‘Sex’ was set as the protected attribute as done in (Kamishima, Akaho, and Sakuma 2011), and the task is to indicate whether the individual’s income is larger than 50K dollars. The features ‘fnlwgt’ and ‘race’ were dropped, resulting in total of 99 features.

¹⁰<https://github.com/nick-jhlee/fair-manifold-pca>

J Full results for EDA of considered UCI datasets

We show two things:

- Boxplots of communality of the protected attribute of each considered UCI dataset, for both $d = 2$ and $d = 10$. (Figures 7, 11, 15)
- PC ($d = 2$) biplot of each considered UCI dataset, for all considered algorithms. (Figures 8, 9, 10, 12, 13, 14, 16, 17, 18)

Particularly for biplots, one train-test split was chosen, and based on that, only the top 10 features with the highest communalities (+ the sensitive attribute, if it is not included) were chosen for plots.




Figure 7: Boxplots of communalities of race for COMPAS dataset




Figure 8: PC biplot of PCA for COMPAS dataset




Figure 9: PC biplot of FPCA for COMPAS dataset






Figure 10: PC biplot of MBF-PCA for COMPAS dataset



(a) $d = 2$



(b) $d = 10$

Figure 11: Boxplots of communalities of age for German credit dataset




Figure 12: PC biplot of PCA for German credit dataset




Figure 13: PC biplot of FPCA for German credit dataset






Figure 14: PC biplot of MBF-PCA for German credit dataset



(a) $d = 2$



(b) $d = 10$

Figure 15: Boxplots of communalities of gender(sex) for Adult income dataset




Figure 16: PC biplot of PCA for Adult income dataset




Figure 17: PC biplot of FPCA for Adult income dataset




Figure 18: PC biplot of MBF-PCA for Adult income dataset