

Coordinate Descent Without Coordinates: Tangent Subspace Descent on Riemannian Manifolds

David Huckleberry Gutman¹ and Nam Ho-Nguyen²

¹Texas Tech University

²University of Sydney

June 13, 2020

Abstract

We extend coordinate descent to manifold domains, and provide convergence analyses for geodesically convex and non-convex smooth objective functions. Our key insight is to draw an analogy between coordinate blocks in Euclidean space and tangent subspaces of a manifold. Hence, our method is called tangent subspace descent (TSD). The core principle behind ensuring convergence of TSD is the appropriate choice of subspace at each iteration. To this end, we propose two novel conditions, the gap ensuring and C -randomized norm conditions on deterministic and randomized modes of subspace selection respectively, that promise convergence for smooth functions and that are satisfied in practical contexts. We propose two subspace selection rules of particular practical interest that satisfy these conditions: a deterministic one for the manifold of square orthogonal matrices, and a randomized one for the Stiefel manifold. Our proof-of-concept numerical experiments on the orthogonal Procrustes problem demonstrate TSD’s efficacy.

1 Introduction

The block coordinate descent (BCD) method [2, 8, 14, 18, 19] is a popular choice of algorithm for solving

$$\min_{x \in \mathbb{R}^n} f(x),$$

where $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is a continuously differentiable function. At each iteration, BCD employs a single gradient step to f in a block of variables, chosen at random or in a cyclic fashion, while holding all other variables constant. Symbolically, we can write this key update step as

$$y^k = y^{k-1} - P_{i_k} \nabla f(y^{k-1}), \quad k = 1, 2, \dots$$

where $\{P_1, \dots, P_m\}$ is a set of matrices each with n rows such that $[P_1 \mid \dots \mid P_m] = I_{n \times n}$, and $i_k \in \{1, \dots, m\}$ is selected in a random or cyclic fashion. One may regard this as an alternating minimization scheme where exact minimization with respect to a block of variables is replaced with a gradient descent step. When n is extremely large, BCD is advantageous compared to standard gradient-based algorithms, since computing $P_{i_k} \nabla f(y^{k-1})$ only requires computing partial derivatives in the P_{i_k} block, and updating only these variables, is relatively easy compared with computing the full gradient $\nabla f(y^{k-1})$ and updating all variables.

Due to the increasing importance of optimization on manifolds to large-scale data science problems, there has been a flurry of activity analyzing first-order methods adapted to this problem class

[3, 7, 11, 16, 17, 20, 24, 25, 26]. In contrast to many other methods depending on first-order information, such as gradient descent and stochastic gradient descent, it is not obvious how to extend BCD to capture problems where the feasible set is an arbitrary Riemannian manifold M . The primary challenge in adapting BCD to an arbitrary manifold is its inherent dependence on the linear structure of \mathbb{R}^n , which allows us to reuse the same block coordinate decomposition $\{P_1, \dots, P_m\}$ at different iterations.

This paper extends BCD methods to the problem

$$\min_{x \in M} f(x)$$

where M is an n -dimensional Riemannian manifold and $f : M \rightarrow \mathbb{R}$ is a continuously differentiable function. Our key insight is that inexact alternating minimization over blocks of coordinates in \mathbb{R}^n generalizes to inexact alternating minimization over subspaces of the manifold’s tangent space at the current iterate. The resultant algorithm, *Tangent Subspace Descent* (TSD), is the focus of our paper. TSD reformulates the key update step of BCD as

$$y^k = \text{Exp}_{y^{k-1}} \left(-\eta \mathcal{P}(y^{k-1}, k) \nabla f(y^{k-1}) \right), \quad k = 1, 2, \dots$$

where Exp_{y^k} is the manifold’s exponential map at y^k , ∇f denotes the Riemannian gradient of f , $\mathcal{P}(y^{k-1}, k)$ is an orthogonal (with respect to the Riemannian metric) projection onto a subspace of the tangent space $T_{y^{k-1}}M$ at y^{k-1} . We formally define these diffeo-geometric objects in Section 2.

TSD is a generalization of BCD, since a block of coordinates can also be interpreted as a tangent subspace. In the Euclidean setting $M = \mathbb{R}^n$, tangent spaces at different points are isomorphic to \mathbb{R}^n itself, and are easily identified with each other. This is the key property that allows us to reuse the same block coordinate decomposition $\{P_1, \dots, P_m\}$ at different iterations. By using the same decomposition, BCD ensures that all possible directions of descent lay in the span of the projections, and that the collection of possible descent directions do not progressively cluster near a lower dimensional subspace. Such clustering can destroy the convergence of an iterative method if the gradient remains roughly orthogonal to said lower dimensional subspace. Examples of this phenomenon are presented in Theorems 3.2 and 3.3.

1.1 Contributions and Related Work

Different tangent spaces on a general manifold M , however, are not so easily related as in the Euclidean case, which necessitates the introduction of \mathcal{P} in TSD, what we call a *subspace selection rule*. Unlike the Euclidean setting, the subspace selection rule \mathcal{P} can choose subspace decompositions that may vary across iterations t , and may be non-orthogonal. A key goal of our paper is to precisely define new concepts which quantify how \mathcal{P} affects the span and clustering of descent directions on a general manifold, and show how these can be used to analyze the convergence of TSD. Our key contributions are as follows.

- Through counterexamples, we demonstrate that care must be taken in specifying \mathcal{P} . Section 3.1 describes subspace selection rules which satisfy simple conditions one might believe ensure convergence, but in fact, do not. This motivates the introduction of the novel and mild conditions, the (γ, r) -gap ensuring (Assumption 4.2) and the C -randomized norm (Assumption 5.1) conditions for deterministic and randomized rules \mathcal{P} respectively. These conditions quantify the ‘information loss’ of \mathcal{P} .
- Our analysis, namely Theorems 4.5 and 5.2, prove that these assumptions ensure convergence of TSD. The convergence analysis for deterministic rules with Assumption 4.2 requires a

non-trivial modification of the existing analysis for BCD (see Remark 4.1). The convergence analysis for randomized rules is sufficiently general to allow selection from an infinite collection of subspaces, unlike in previous work.

- Our new conditions restrict the choice of \mathcal{P} , but are still sufficiently flexible to capture a number of useful settings. We provide examples of different manifolds and subspace selection rules in Sections 4.2 and 5.2. In particular, Section 5.2.1 presents a new randomized selection rule for Stiefel manifolds, which satisfies the C -randomized norm condition and thus ensures convergence of TSD. In addition, Section 6 elaborates on the special case of the manifold of orthogonal $n \times n$ real matrices, for which we provide a class of deterministic subspace selection rules, and show that it satisfies the (γ, r) -gap ensuring condition for convergence of TSD. Furthermore, we show that this class of rules result in computationally cheap and sparse updates for TSD, of order $O(n)$ cost per update as opposed to $O(n^3)$ for standard Riemannian gradient descent updates.
- We conduct proof-of-concept numerical experiments on the well-known orthogonal Procrustes problem $\min_{X \in O_n} \|AX - B\|_F^2$, with $A \in \mathbb{R}^{d \times n}$, $B \in \mathbb{R}^{d \times n}$, and O_n being the manifold of square orthogonal matrices, which demonstrate the efficacy of TSD over standard Riemannian gradient descent schemes.

Frongillo and Reid [8] considered randomized and not necessarily orthogonal subspace descent in the Euclidean setting. We significantly generalize this to cyclic and randomized subspace descent on manifolds. To our knowledge, the only papers which have considered coordinate descent in a non-Euclidean setting are Gao et al. [9], Shalit and Chechik [21]. Both of these examine specifically the Stiefel manifold of orthogonal matrices. The algorithm by Gao et al. [9] updates a single column of the matrix at each step, but then adds a so-called “correction” which involves computing a singular value decomposition, and is not a direct generalization of BCD. The algorithm by Shalit and Chechik [21] is a specialization of our proposed one to O_n , for a particular randomized subspace selection rule.

In contrast to [8, 9, 21], our work examines a more general setting on any manifold, with any geometry, for both deterministic and randomized (valid) subspace selection rules. Our analysis also gives global convergence guarantees for geodesically convex objectives, in addition to convergence to stationary points for non-convex objectives. Of particular note, the randomized selection rule in Theorem 5.6 for general Stiefel manifolds broadly generalizes the randomized rule in [21] for O_n . Furthermore, to our knowledge, the deterministic selection rule in Section 6 is the first for optimization over O_n .

Note also that our gap ensuring condition is closely related to principal angles between subspaces, which are known to control the convergence rates of cyclic projection schemes as far back as [4]. Principal angles between subspaces have also appeared in convergence rates for alternating projection schemes on manifolds [15].

1.2 Outline

In Section 2, we provide the language of, and convergence analysis tools for, manifold-constrained optimization. In Section 3, we present Tangent Subspace Descent and two theorems (Theorems 3.3 and 3.2) explaining the need for the judicious selection of subspace selection rules. In Section 4, we elaborate the (γ, r) -gap ensuring condition, delineate examples of deterministic subspace selection rules satisfying it (including cyclic BCD), and analyze the convergence TSD equipped with such a rule for smooth (both g -convex and non g -convex) functions. In Section 5, we repeat this program for randomized subspace selection rules C -randomized norm. We show that randomized BCD can

be seen as an instance of TSD with a rule satisfying this condition. This section also provides the new randomized subspace selection rule for general Stiefel manifolds (Theorem 5.6). In Section 6, we propose the first deterministic rule for subspace selection over the set of square orthogonal matrices and apply TSD to the orthogonal Procrustes problem.

2 Optimization on Manifolds

In this section, we present some of the basic terminology and convergence analysis tools for optimization on manifolds. We will regularly use the standard notation and known results of Riemannian manifolds, cf. Absil et al. [1], Lee [12, 13], Vishnoi [22]. We consider the following general optimization problem

$$f^* = \min_x \{f(x) : x \in M\}, \quad (1)$$

where M is an n -dimensional Riemannian manifold, and $f : M \rightarrow \mathbb{R}$ is a differentiable function.

For each point $x \in M$, we denote the *tangent space* to M at x as $T_x M$. Recall that $T_x M$ is an n -dimensional real vector space. We denote the tangent bundle of M , the union of all tangent spaces of M , as $TM := \bigcup_{x \in M} T_x M$; it is a smooth manifold in its own right. A *vector field* is a smooth map $V : M \rightarrow TM$ satisfying $V(x) \in T_x M$ for each $x \in M$. The *Riemannian metric* on M is a collection of inner products $\{\langle \cdot, \cdot \rangle_x\}_{x \in M}$ on each $T_x M$ which is smoothly varying in the sense that for any two vector fields $V, W : M \rightarrow TM$, $x \mapsto \langle V(x), W(x) \rangle_x$ is a smooth function on M . The norm on $T_x M$ induced by $\langle \cdot, \cdot \rangle_x$ is denoted $\|\cdot\|_x$.

The *length* of a piecewise smooth curve $\gamma : [0, 1] \rightarrow M$ with $\gamma(0) = x$, $\gamma(1) = y$, and derivative $\gamma' : [0, 1] \rightarrow TM$ is defined as

$$L(\gamma) = \int_0^1 \|\gamma'(t)\|_{\gamma(t)} dt.$$

The *Riemannian distance* $d(x, y)$ between points x and y is the infimum of the lengths of all piecewise smooth curves joining x and y . The manifold M is a metric space when endowed with the Riemannian distance. If M is complete and connected under this metric then the distance can be attained by some smooth curve. We say that a curve $\gamma : [0, 1] \rightarrow M$ with $\gamma(0) = x$ and $\gamma(1) = y$ is a *geodesic between x and y* if it is locally distance minimizing. Given any two points $x, y \in M$, we can find a geodesic γ between them such that $L(\gamma) = d(x, y)$.

The *exponential map* at a point $x \in M$ is denoted $\text{Exp}_x : T_x M \rightarrow M$ gives us a way to generalize the operation $x + v$ in Euclidean space, where $v \in T_x M$. We have $\text{Exp}_x(0) = x$ for any $x \in M$. The exponential map gives rise to geodesics: the curve $t \mapsto \gamma(t) := \text{Exp}_x(tv)$ is always a geodesic between x and $\text{Exp}_x(v)$. We assume that $\text{Exp}_x(v)$ is well defined for all $x \in M$, $v \in T_x M$, that is, M is *geodesically complete*. The Hopf-Rinow theorem states that geodesic completeness and completeness under the Riemannian distance are equivalent. Thus, this is a decidedly mild assumption. We define a set-valued *inverse exponential map*, $\text{Exp}_x^{-1} : M \rightrightarrows T_x M$ as

$$\text{Exp}_x^{-1}(y) = \arg \min_{v \in T_x M} \{\|v\|_x : \text{Exp}_x(v) = y\}.$$

Furthermore, we have that for any $v \in \text{Exp}_x^{-1}(y)$, $d(x, y) = \|v\|_x$, and for general $v \in T_x M$, we have $\|v\|_x \geq d(x, \text{Exp}_x(v))$.

For any smooth curve $\gamma : I \rightarrow M$ defined on an interval $I \subset \mathbb{R}$, there is a well-defined collection of invertible linear maps called *parallel transport operators along γ* , denoted $\Gamma_{\gamma, t^0}^{t^1} : T_{\gamma(t^0)} M \rightarrow T_{\gamma(t^1)} M$, for $t^0, t^1 \in I$. These maps satisfy $(\Gamma_{\gamma, t^0}^{t^1})^{-1} = \Gamma_{\gamma, t^1}^{t^0}$ and $\Gamma_{\gamma, t^0}^{t^0} = \text{id}_{T_{\gamma(t^0)} M}$. Furthermore, a key property is that they preserve the metric, i.e., $\langle v, v' \rangle_{\gamma(t^0)} = \langle \Gamma_{\gamma, t^0}^{t^1} v, \Gamma_{\gamma, t^0}^{t^1} v' \rangle_{\gamma(t^1)}$ for any $t^0, t^1 \in I$, $v, v' \in T_{\gamma(t^0)} M$. In this paper, we will only be interested in parallel transport operators along geodesic curves. When $y = \text{Exp}_x(v)$, we use the shorthand $\Gamma_x^y : T_x M \rightarrow T_y M$ to denote $\Gamma_{\gamma, 0}^1$, where

$\gamma(t) := \text{Exp}_x(tv)$ is a geodesic between x and y . Note that while Γ_x^y depends on v , we will suppress this in the notation, since the dependence will be implied from $y = \text{Exp}_x(v)$, which will be clear from context. We will also use the notation Γ_x^y when there is a *unique* geodesic curve between x and y , i.e., when $\text{Exp}_x^{-1}(y) = \{v\}$ is a singleton. When $y = \text{Exp}_x(v)$ is not clear from context, the fact that $\text{Exp}_x^{-1}(y) = \{v\}$ will be made clear. Note that if $y = \text{Exp}_x(v)$, then $x = \text{Exp}_y(-\Gamma_x^y v)$ and $(\Gamma_x^y)^{-1} = \Gamma_y^x$.

For a differentiable function $f : M \rightarrow \mathbb{R}$, the *Riemannian gradient* of f is defined as the element $\nabla f(x) \in T_x M$ such that $\langle \nabla f(x), v \rangle_x = df_x v$ for all $v \in T_x M$, where $df_x : T_x M \rightarrow \mathbb{R}$ is the differential of f at $x \in M$. The gradient is key to the analysis of optimization algorithms on manifolds as the necessary first-order stationarity condition $\nabla f(x^*) = 0$ for minimizers of (1) still holds. In this case, we say that x^* is a *stationary point* of f . This condition is also sufficient when f is *g-convex*: a condition defined shortly. When f is continuously differentiable, it holds that if $x_n \rightarrow x^*$ and $\|\nabla f(x_n)\|_{x_n} \rightarrow 0$ then $\nabla f(x^*) = 0$.

We are interested in two classes of functions $f : M \rightarrow \mathbb{R}$ on geodesically complete Riemannian manifolds defined by two different conditions on the gradient:

$$\forall (x, v) \in TM, y = \text{Exp}_x(v), \|\Gamma_x^y \nabla f(x) - \nabla f(y)\|_y \leq L_f d(x, y). \quad (2)$$

$$\forall x, y \in M, v \in \text{Exp}_x^{-1}(y), f(y) \geq f(x) + \langle \nabla f(x), v \rangle_x. \quad (3)$$

We say that a function f satisfying (2) is *L_f -smooth*, and a function satisfying (3) is *geodesically convex*, or *g-convex*. Note that $\|\Gamma_x^y \nabla f(x) - \nabla f(y)\|_y = \|\nabla f(x) - \Gamma_y^x \nabla f(y)\|_x$, and that L_f -smoothness of f implies the following upper bound:

$$\forall x \in M, v \in T_x M, f(\text{Exp}_x(v)) \leq f(x) + \langle \nabla f(x), v \rangle_x + \frac{L_f}{2} \|v\|_x^2. \quad (4)$$

In this paper, we use the term smoothness of the function to mean both infinite differentiability and Lipschitz continuity of the gradient. Which definition we mean will be clear from context. Furthermore, Ferreira and Oliveira [6, Theorem 2.3] show that g-convexity is equivalent to the following: for any $x, y \in M$, $v \in \text{Exp}_x^{-1}(y)$, and $\beta \in [0, 1]$, $f(\text{Exp}_x(\beta v)) \leq (1 - \beta)f(x) + \beta f(y)$.

The convergence analyses of (non-accelerated) gradient-based methods for smooth convex optimization problems is often based on establishing the following so-called ‘sufficient decrease’ relation between iterates:

$$f(x^{t-1}) - f(x^t) \geq \eta \|\nabla f(x^{t-1})\|_{x^{t-1}}^2. \quad (\text{Dec})$$

Often, (Dec) is derived only from smoothness of f rather than convexity: if f is L_f -smooth then by selecting $v = -\frac{1}{L_f} \nabla f(x)$ in (4),

$$f(x) - f\left(\text{Exp}_x\left(-\frac{1}{L_f} \nabla f(x)\right)\right) \geq \frac{1}{L_f} \|\nabla f(x)\|_x^2. \quad (5)$$

Thus, when f only satisfies smoothness, we can get a stationarity guarantee when this sufficient decrease condition is met. The analyses in Sections 4 and 5 are based on providing conditions for which (Dec), or a suitable modification thereof, is satisfied.

3 Tangent Subspace Descent on General Riemannian Manifolds

In this section, we introduce the template for tangent subspace descent and elaborate the issues that arise when choosing an appropriate subspace selection rule \mathcal{P} . With the key TSD step (1) in hand, we state the complete algorithmic template in Algorithm 1.

Algorithm 1: Tangent Subspace Descent

Data: Initial point $x^0 \in M$, subspace selection rule \mathcal{P} , $m \in \mathbb{N}$

Result: Sequence $\{x^t\}_{t \geq 1} \subset M$.

for $t = 1, 2, \dots$ **do**

 Set $y^{t,0} := x^{t-1}$;

for $k \in [m]$ **do**

 Pick $P_k^{y^{t,k-1}} := \mathcal{P}(\{y^{t,j}\}_{j=0}^{k-1}, k)$ and $\eta_{t,k} > 0$;

 Update $y^{t,k} := \text{Exp}_{y^{t,k-1}}(-\eta_{t,k} P_k^{y^{t,k-1}} \nabla f(y^{t,k-1}))$;

end

 Update $x^t := y^{t,m}$;

end

At iteration $t \geq 1$ of Algorithm 1, given iterates $\{y^{t,k-1}\}_{k \in [m]}$ and projections $\{P_k^{y^{t,k-1}}\}_{k \in [m]}$, we define

$$P_{t,k} := \Gamma_{y^{t,k-1}}^{y^{t,0}} P_k^{y^{t,k-1}} \Gamma_{y^{t,0}}^{y^{t,k-1}}, \quad k \in [m], \quad \|\cdot\|_{y^{t,0}, \mathcal{P}} := \sqrt{\sum_{k \in [m]} \|P_{t,k}(\cdot)\|_{y^{t,0}}^2}.$$

Note that $\|\cdot\|_{y^{t,0}, \mathcal{P}}$ is, in general, a seminorm on $T_{y^{t,0}}M$, and is a norm if and only if $\text{Span}\left(\bigcup_{k \in [m]} \text{Im}(P_{t,k})\right) = T_{y^{t,0}}M$. Furthermore, if $P_{t,k}$ are projections onto orthogonal subspaces which span $T_{y^{t,0}}M$, then $\|\cdot\|_{y^{t,0}, \mathcal{P}} = \|\cdot\|_{y^{t,0}}$. This is the case for BCD, and leads to sufficient decrease bounds like (Dec), yet is not true in general for TSD. Instead, for TSD, further conditions must be imposed and a new analysis is required to relate the two norms. The choice of the subspace selection rule \mathcal{P} is a critical element of this. Note that \mathcal{P} can cover both deterministic and randomized rules. For randomized rules, we will always take $m = 1$, and denote for convenience $P_t := \mathcal{P}(y^{t,0}, 1) = P_1^{y^{t,0}}$ (this is a randomly chosen projection on $T_{x^{t-1}}M$). We spend the remainder of this section discussing issues regarding \mathcal{P} that can inhibit the convergence of Algorithm 1 and that motivate the conditions we place upon \mathcal{P} in Sections 4 and 5.

3.1 The Problem of Subspace Selection

Since M is not (in general) a vector space, it cannot be decomposed into a collection of subspaces, so we need to pick different subspaces to project onto at each point and iteration. While the tangent space T_xM is a vector space for each $x \in M$, the spaces T_xM and T_yM can be very different for $y \neq x$. Consider the example of the tangent space to the north pole of the sphere, \mathbb{S}^2 . This space is nothing more than an affine plane parallel to the xy -plane that lies flat on the sphere at the pole. However, if we translate this plane to any point on the equator then we see the plane cuts through the sphere so it no longer lies tangent to the sphere. The differing structures of the tangent spaces at different points on M thus disallows us from picking a fixed decomposition of a single tangent space that we use at every subsequent iteration, as in BCD.

Furthermore, the choice of subspaces dictates not only the convergence rate of this algorithm, but whether it converges at all. Even simple examples on \mathbb{R}^n show how badly the algorithm can breakdown when \mathcal{P} is chosen poorly. When V is a Hilbert space and $S \subseteq V$ is a subspace, we will let Π_S denote the orthogonal projection onto S .

Example 3.1. Let $M = \mathbb{R}^n$, $\mathcal{P} \equiv \Pi_{\text{Span}\{\{e_1\}\}}$, $m = 1$, and $f(x) = \frac{1}{2}\|x\|_2^2$. If $x^0 = e_1 + e_2$ and $\eta_{t,k} = 1$ then $x^t = e_2$ for $t \geq 1$, but $\arg \min_{x \in \mathbb{R}^n} f(x) = \{0\}$. ■

The primary issue in the above example is that \mathcal{P} produces a spartan collection of subspaces. Specifically, $\text{Span}(\{e_1\}) \neq \mathbb{R}^2$. Recasting this in terms of tangent spaces, $\text{Span}(\text{Im } \Gamma_y^x \Pi_{\text{Span}(\{e_1\})} \Gamma_x^y) \neq T_x \mathbb{R}^2 \cong \mathbb{R}^2$ for any $x, y \in \mathbb{R}^2$.

One is tempted to believe for deterministic rules it might suffice simply to pick a \mathcal{P} that promises a “spanning condition” such as

$$\text{Span} \left(\bigcup_{k \in [m]} \text{Im}(P_{t,k}) \right) = T_{y^t,0} M \quad (6)$$

for any sequence $y^{t,0}, \dots, y^{t,m}$ generated by one execution of the inner loop in Algorithm 1. Note that this spanning condition is readily seen to be equivalent to the requirement that $\|\cdot\|_{y^t,0,\mathcal{P}}$ is a norm (as opposed to just a seminorm). Thus, for randomized rules, we are similarly tempted to conjecture that if

$$v \mapsto \sqrt{\mathbb{E}[\|P_t v\|_{x^{t-1}}^2 \mid x^{t-1}]} \quad (7)$$

is a norm on $T_{x^{t-1}} M$ then convergence is guaranteed. Unfortunately, these conditions are insufficient. We now exhibit deterministic and randomized selection rules \mathcal{P} satisfying (6) and (7) for which convergence of TSD fails.

For both examples, we take $M = \mathbb{R}^n$ and $f(x) = \frac{1}{2}\|x\|_2^2$, which is a 1-smooth function. Note also that $T_x M = \mathbb{R}^n$ for each $x \in \mathbb{R}^n$, $\text{Exp}_x(v) = x + v$, and the Riemannian metric $\langle \cdot, \cdot \rangle_x$ is just the usual inner product on \mathbb{R}^n . For the deterministic setting, we take $m = n$ in Algorithm 1, and denote for convenience $\|\cdot\|_{t,\mathcal{P}} := \|\cdot\|_{y^t,0,\mathcal{P}}$. As mentioned above, for the randomized setting, we have $m = 1$ and $P_t := \mathcal{P}(y^{t,0}, 1)$ is a random projection on $T_{x^{t-1}} M$. Furthermore, the stepsizes in Algorithm 1 will be taken as 1, since $\nabla f(x) = x$ and for any orthogonal projection $\Pi : \mathbb{R}^n \rightarrow \mathbb{R}^n$, we have

$$\min_{\eta} f(x - \eta \Pi \nabla f(x)) = \frac{1}{2} \|x - \Pi x\|_2^2 = f(x) - f(\Pi x).$$

We start with the randomized setting.

Theorem 3.2. *There exists a randomized projection selection rule \mathcal{P} satisfying (7) and constant $\epsilon > 0$ such that, at any $t \geq 1$ in Algorithm 1, $f(x^t) > \epsilon$. Furthermore,*

$$\limsup_{t \rightarrow \infty} \sup_{v \neq 0} \frac{\|v\|_2}{\sqrt{\mathbb{E}[\|P_t v\|_2^2 \mid x^{t-1}]}} = \infty$$

almost surely.

Proof of Theorem 3.2. We define the following projection selection rule. Given a starting point $x^0 \neq 0$, denote $\epsilon := f(x^0)/2 > 0$. For $t \geq 1$, let $H_t = \left\{ v \in \mathbb{R}^n : \langle v, x^{t-1} \rangle = \sqrt{f(x^{t-1}) - \epsilon} \right\} \cap \{v \in \mathbb{R}^n : \|v\|_2 = 1\}$. We choose $\{v_{t,i}\}_{i \in [n]} \subset H_t$ such that $\text{Span}(\{v_{t,i}\}_{i \in [n]}) = \mathbb{R}^n$ (this is possible due to a simple dimension counting argument). We then choose v_t uniformly at random from $\{v_{t,i}\}_{i \in [n]}$, and set P_t to be the projection onto the subspace spanned by v_t , i.e., $x \mapsto \langle v_t, x \rangle v_t$.

We thus have $x^t = x^{t-1} - P_t x^{t-1}$ and

$$f(x^t) = f(x^{t-1}) - f(P_t x^{t-1}) = f(x^{t-1}) - \frac{1}{2} \langle v_t, x^{t-1} \rangle^2 = \frac{f(x^{t-1}) + \epsilon}{2}.$$

Solving this recurrence relation with the initial condition $f(x^0) = 2\epsilon$, we get $f(x^t) = \epsilon(1 + 2^{-t}) > \epsilon$.

Note that this also implies that for any $t \geq 1$, $\langle v_t, x^{t-1} \rangle = \sqrt{f(x^{t-1}) - \epsilon} = \sqrt{\epsilon 2^{-t}}$. Consequently, we see that

$$x^t = x^0 + \sum_{s \in [t]} \langle v_s, x^{s-1} \rangle v_s = x^0 + \sqrt{\epsilon} \sum_{s \in [t]} 2^{-s/2} v_s$$

must converge to some $\bar{x} \in \mathbb{R}^n$, since $\{2^{-s/2}\}_{s \geq 1}$ is an absolutely convergent sequence and $\{v_s\}_{s \geq 1}$ are from the unit ball. It follows that for any $t \geq 1$ and $i \in [n]$, $\langle v_{t,i}, \bar{x} \rangle \rightarrow 0$. Thus,

$$\mathbb{E} [\|P_t \bar{x}\|_2^2 \mid x^{t-1}] = \frac{1}{n} \sum_{i \in [n]} \langle v_{t,i}, \bar{x} \rangle^2 \rightarrow 0,$$

so

$$\limsup_{t \rightarrow \infty} \sup_{v \neq 0} \frac{\|v\|_2}{\sqrt{\mathbb{E}[\|P_t v\|_2^2 \mid x^{t-1}]}} \geq \lim_{t \rightarrow \infty} \frac{\|\bar{x}\|_2}{\sqrt{\mathbb{E}[\|P_t \bar{x}\|_2^2 \mid x^{t-1}]}} = \infty.$$

□

A non-convergent deterministic rule can be constructed in a similar manner, with small modifications.

Theorem 3.3. *There exists a deterministic projection selection rule \mathcal{P} satisfying (6) with $m = n$ and constant $\epsilon > 0$ such that, at any $t \geq 1$ in Algorithm 1, $f(x^t) > \epsilon$. Furthermore,*

$$\limsup_{t \rightarrow \infty} \sup_{v \neq 0} \frac{\|v\|_2}{\|v\|_{t, \mathcal{P}}} = \infty.$$

Proof of Theorem 3.3. We define the following projection selection rule. Given a starting point $x^0 \neq 0$, denote $\epsilon := f(x^0)/2 > 0$. Fix $t \geq 1$, and set $y^{t,0} = x^{t-1}$, $S_{t,0} = \{0\}$. For $k \in [m]$, we do the following:

- Let $H_{t,k} := \{v : \langle v, y^{t,k-1} \rangle = \sqrt{(f(x^{t-1}) - \epsilon)/m}\}$.
- Choose $v_{t,k} \in H_{t,k} \cap \{v : \|v\|_2 = 1\}$ such that $\dim(S_{t,k-1} \cup \{v_{t,k}\}) = \dim(S_{t,k-1}) + 1$, where $\dim(S)$ is the affine dimension of S (this is possible due to a simple dimension counting argument).
- Set $S_{t,k} := S_{t,k-1} \cup \{v_{t,k}\}$.
- Define $\mathcal{P}(y^{t,k-1}, k)$ as the projection onto the subspace spanned by $v_{t,k}$.

Set $x^t = y^{t,n}$. Also observe that $\dim(\{v_{t,k}\}_{k \in [m]}) = m = n$ so $\text{Span}(\{v_{t,k}\}_{k \in [m]}) = \mathbb{R}^n$. Then for each $k \in [m]$ we have

$$\begin{aligned} f(y^{t,k}) &= f(y^{t,k-1}) - \frac{1}{2} \langle v_{t,k}, y^{t,k-1} \rangle^2 = \dots = f(y^{t,0}) - \frac{1}{2} \sum_{j \in [k]} \langle v_{t,j}, y^{t,j-1} \rangle^2 = f(x^{t-1}) - \frac{k}{2m} (f(x^{t-1}) - \epsilon) \\ &= \frac{(2m - k)f(x^{t-1}) + k\epsilon}{2m}. \end{aligned}$$

In particular, $f(x^t) = (f(x^{t-1}) + \epsilon)/2$, so as in the proof of Theorem 3.2, $f(x^t) = \epsilon(1 + 2^{-t}) > \epsilon$. Similar arguments to the proof of Theorem 3.2 show that $x^t \rightarrow \bar{x}$, $\langle v_{t,k}, \bar{x} \rangle \rightarrow 0$ for any $k \in [m]$, and $\frac{\|\bar{x}\|_2}{\|\bar{v}\|_{t, \mathcal{P}}} \rightarrow \infty$ (recognize that $\|\bar{x}\|_{t, \mathcal{P}} = \sqrt{\sum_{k \in [m]} \langle v_{t,k}, \bar{x} \rangle^2}$). □

The primary issue that prevents convergence in Theorems 3.2 and 3.3 is that, while the $\{v_{t,k}\}_{k \in [m]}$ span \mathbb{R}^n , they eventually cluster near a lower-dimensional subspace $\{v : \langle v, \bar{x} \rangle = 0\}$ as the algorithm progresses. As a result, the projections cause some information loss in the direction \bar{x} .

It seems reasonable to conjecture that if the projections maintain a certain level of ‘orthogonality’ (i.e., do not cluster near some lower-dimensional subspace), then the degeneracy of Theorems 3.2 and 3.3 cannot occur. Constructing viable deterministic and randomized subspace selection rules that respect this observation is the thrust of the main two sections (Sections 4 and 5) of the paper. In the respective cases, the orthogonality of the selected subspaces will be reflected by the quantities $\sup_{v \in T_{y^t,0} M \setminus \{0\}} \frac{\|v\|_{y^t,0}}{\|v\|_{y^t,0,\mathcal{P}}}$ and $\sup_{v \in T_{x^{t-1}} M \setminus \{0\}} \frac{\|v\|_{x^{t-1}}}{\sqrt{\mathbb{E}[\|P_t v\|_{x^{t-1}}^2 | x^{t-1}]}}$. As we will see,

providing finite bounds on these quantities is sufficient to guarantee convergence of TSD, so the main challenge will be to provide conditions to guarantee this.

4 Deterministic Subspace Selection

In this section, we analyze the convergence of TSD when it is equipped with a deterministic subspace selection rule \mathcal{P} . In Section 4.1, we motivate and present a novel and mild condition on \mathcal{P} , the (γ, r) -gap ensuring condition, which guarantees convergence of tangent subspace descent for smooth functions. We then give a convergence analysis for (γ, r) -gap ensuring rules. In Section 4.2, we present a series of important examples of selection rules \mathcal{P} that satisfy this condition including product manifolds, which includes block cyclic coordinate descent in Euclidean space as a special case, as well as an important non-Euclidean application, tensor PCA. Note that we also construct a practical (γ, r) -gap ensuring \mathcal{P} for the manifold of orthogonal square matrices, but we defer its construction to Section 6.

4.1 A Convergent Deterministic \mathcal{P} via the (γ, r) -Gap Ensuring Condition

Given a Hilbert space V , we will say that a set of orthogonal projections $\{P'_k\}_{k \in [m]}$ is an *orthogonal decomposition* of V if $\text{Im}(P'_i) \perp \text{Im}(P'_j)$ for $j \neq i$ and $\text{Span}\left(\bigcup_{k \in [m]} \text{Im}(P'_k)\right) = V$. The foundation of our analysis is Assumption 4.2 below, the (γ, r) -gap ensuring condition. We now describe why the norm equivalence terms from Theorems 3.2 and 3.3 arise in our analysis, as well as why Assumption 4.2 is required for bounding them.

First, recall that we assume that our objective function f is L_f -smooth, in the sense of (4). Similar to (4), smoothness will facilitate stepsize selection in order to achieve sufficient decrease as in (Dec). We refine the definition of smoothness slightly to individual subspaces $P_k^{y^t, k-1} T_{y^t, k-1} M$ in an inner loop.

Assumption 4.1 (Generalized Block Smoothness). There exists constants $0 < L_1, \dots, L_m < \infty$ such that for any $t \geq 1$ and $k \in [m]$, the iterates of Algorithm 1 satisfy

$$f\left(\text{Exp}_{y^t, k-1}(P_k^{y^t, k-1} v)\right) \leq f(y^{t, k-1}) + \left\langle \nabla f(y^{t, k-1}), P_k^{y^t, k-1} v \right\rangle_{y^t, k-1} + \frac{L_k}{2} \left\| P_k^{y^t, k-1} v \right\|_{y^t, k-1}^2 \quad (8)$$

for all $v \in T_{y^t, k-1} M$.

In practice, each L_k may not be known, but can be implemented via backtracking line search (to achieve the condition (12)), which will terminate as long as L_k is finite; this will hold since one can simply take $L_k = L_f$. From Assumption 4.1 and another mild condition, we obtain the following.

Lemma 4.1 (Deterministic Sufficient Decrease Template). *Suppose Assumption 4.1 holds, and that at each $t \geq 1$, $k \in [m]$ in Algorithm 1, the stepsize is chosen as $\eta_{t,k} := 1/L_k$. Then*

$$f(y^{t,0}) - f(y^{t,m}) \geq \sum_{k \in [m]} \frac{1}{2L_k} \left\| P_k^{y^{t,k-1}} \nabla f(y^{t,k-1}) \right\|_{y^{t,k-1}}^2. \quad (9)$$

Furthermore, suppose that there exists some $C > 0$ such that for each $k \in [m]$, \mathcal{P} ensures the iterates of Algorithm 1 satisfy

$$\left\| P_k \nabla f(y^{t,0}) - P_k \Gamma_{y^{t,k-1}}^{y^{t,0}} \nabla f(y^{t,k-1}) \right\|_{y^{t,0}}^2 \leq C \sum_{j \in [k]} \left\| P_j^{y^{t,j-1}} \nabla f(y^{t,j-1}) \right\|_{y^{t,j-1}}^2. \quad (10)$$

Then for any $t \geq 1$,

$$f(y^{t,0}) - f(y^{t,m}) \geq \frac{1}{4L_{\max}(1 + Cm)} \left\| \nabla f(y^{t,0}) \right\|_{y^{t,0}, \mathcal{P}}^2. \quad (11)$$

Proof of Lemma 4.1. For convenience, we fix a $t \geq 1$ and suppress it in the notation for the proof. Our stepsize selection, $\eta_{t,k} = \frac{1}{L_k}$, minimizes the right-hand side of (8), which gives

$$f(y^{k-1}) - f(y^k) \geq \frac{1}{2L_k} \left\| P_k^{y^{k-1}} \nabla f(y^{k-1}) \right\|_{y^{k-1}}^2. \quad (12)$$

We get (9) by summing this over $k \in [m]$.

For every $k \in [m]$,

$$\begin{aligned} \left\| P_k \nabla f(y^0) \right\|_{y^0}^2 &= \left\| P_k \nabla f(y^0) - P_k \Gamma_{y^{k-1}}^{y^0} \nabla f(y^{k-1}) + P_k \Gamma_{y^{k-1}}^{y^0} \nabla f(y^{k-1}) \right\|_{y^0}^2 \\ &\leq \left(\left\| P_k \nabla f(y^0) - P_k \Gamma_{y^{k-1}}^{y^0} \nabla f(y^{k-1}) \right\|_{y^0} + \left\| P_k \Gamma_{y^{k-1}}^{y^0} \nabla f(y^{k-1}) \right\|_{y^0} \right)^2 \\ &\leq 2 \left\| P_k \nabla f(y^0) - P_k \Gamma_{y^{k-1}}^{y^0} \nabla f(y^{k-1}) \right\|_{y^0}^2 + 2 \left\| P_k \Gamma_{y^{k-1}}^{y^0} \nabla f(y^{k-1}) \right\|_{y^0}^2 \\ &= 2 \left\| P_k \nabla f(y^0) - P_k \Gamma_{y^{k-1}}^{y^0} \nabla f(y^{k-1}) \right\|_{y^0}^2 + 2 \left\| \Gamma_{y^{k-1}}^{y^0} P_k^{y^{k-1}} \nabla f(y^{k-1}) \right\|_{y^0}^2 \\ &= 2 \left\| P_k \nabla f(x) - P_k \Gamma_{y^{k-1}}^{y^0} \nabla f(y^{k-1}) \right\|_{y^0}^2 + 2 \left\| P_k^{y^{k-1}} \nabla f(y^{k-1}) \right\|_{y^{k-1}}^2 \\ &\leq 2 \left\| P_k^{y^{k-1}} \nabla f(y^{k-1}) \right\|_{y^{k-1}}^2 + 2C \sum_{j \in [k]} \left\| P_j^{y^{j-1}} \nabla f(y^{j-1}) \right\|_{y^{j-1}}^2. \end{aligned}$$

where we apply the triangle inequality in line 2, the Cauchy-Schwarz inequality in line 3, the definition of P_i in the rightmost term of line 4, and the fact that $\Gamma_{y^{k-1}}^{y^0}$ is an isometry in line 5.

Summing this over $k \in [m]$ we have

$$\begin{aligned} \left\| \nabla f(y^0) \right\|_{y^0, \mathcal{P}}^2 &= \sum_{k \in [m]} \left\| P_k \nabla f(y^0) \right\|_{y^0}^2 \leq 2 \sum_{k \in [m]} (1 + (m - k)C) \left\| P_k^{y^{k-1}} \nabla f(y^{k-1}) \right\|_{y^{k-1}}^2 \\ &\leq 2(1 + Cm) \sum_{k \in [m]} \left\| P_k^{y^{k-1}} \nabla f(y^{k-1}) \right\|_{y^{k-1}}^2. \end{aligned}$$

We then get (11) by simple rearrangement. \square

We will address the condition (10) later. For now, we examine the result (11). Since $x^{t-1} = y^{t,0}$, $x^t = y^{t,m}$, this is almost of the same form as (Dec), except the norm on the right hand side is $\|\cdot\|_{y^{t,0},\mathcal{P}}$ instead of the required $\|\cdot\|_{y^{t,0}} = \|\cdot\|_{x^{t-1}}$. This also clarifies why the projection selection rule of Theorem 3.3 forces non-convergence of tangent subspace descent. In order to achieve the sufficient decrease condition (Dec), we need to replace $\|\nabla f(y^{t,0})\|_{y^{t,0},\mathcal{P}}$ with $\|\nabla f(y^{t,0})\|_{y^{t,0}}$ which can be done if there exists a uniform (over $t \geq 1$) finite bound on the norm-equivalence term $\sup_{v \in T_{y^{t,0}}M \setminus \{0\}} \|v\|_{y^{t,0}} / \|v\|_{y^{t,0},\mathcal{P}}$. For the example in Theorem 3.3, it was shown that this quantity diverges to ∞ as $t \rightarrow \infty$.

As we will see in Example 4.6, when \mathcal{P} is selected in a particular manner, this can be guaranteed. However, unless the manifold M has special structure (e.g., product manifold, which includes Euclidean space) using this particular \mathcal{P} can be costly since it requires repeated computations of parallel transports. In order to circumvent this, we now introduce the (γ, r) -gap ensuring condition.

Assumption 4.2. Let $r > 0$ and $\gamma \in (0, 1]$. Suppose that for any $t \geq 1$ and $\{y^{t,0}, \dots, y^{t,m}\} \subset M$ generated by Algorithm 1, we have that whenever

$$\max_{k \in [m]} d(y^{t,0}, y^{t,k}) \leq r,$$

there exists an orthogonal decomposition $\{P'_{t,k}\}_{k \in [m]}$ of $T_{y^{t,0}}M$ such that for all $k \in [m]$, $P_{t,k}$ and $P'_{t,k}$ have the same rank and $\|P'_{t,k} - P_{t,k}\|_{y^{t,0}} \leq \gamma$ (the induced operator norm on $T_{y^{t,0}}M$). We say such a projection selection rule \mathcal{P} is (γ, r) -gap ensuring.

Remark 4.1. As we will prove in Proposition 4.2, we can bound the norm-equivalence term if $\max_{k \in [m]} d(y^{t,0}, y^{t,k}) \leq r$. One may ask, however, whether the distance parameter r is necessary for our analysis. As it turns out, Section 6.3 shows that on the manifold of orthogonal matrices O_n , there exists a \mathcal{P} such that Assumption 4.2 is satisfied, but the norm-equivalence term is unbounded in general. Thus, the parameter r is critical for bounding the norm equivalence term.

Note that in Euclidean space, r is not required. Furthermore, it is known that small neighbourhoods around any point in M behave like Euclidean space. Thus, an intuitive reason why the norm-equivalence term is well-behaved when $\max_{k \in [m]} d(y^{t,0}, y^{t,k}) \leq r$ is because the neighbourhood around $y^{t,0}$ becomes more like usual Euclidean space.

In terms of the analysis, when the distance condition is violated, in Lemma 4.4 we will show that another sufficient decrease inequality holds, similar to (Dec). Thus, we still make progress for each step. This is the main departure from the traditional analysis of BCD in Euclidean space. ■

Remark 4.2. Golub and Van Loan [10, p. 331] explain that the quantity $\|P'_{t,k} - P_{t,k}\|_{y^{t,0}}$ is equal to $\sqrt{1 - \cos^2(\theta_{\min})}$ where θ_{\min} denotes the smallest principal angle between the subspaces $\text{Im}(P'_{t,k})$ and $\text{Im}(P_{t,k})$. The smallest principal angle between tangent subspaces has also been used by Lewis and Malick [15] to describe the convergence rates of alternating projection methods on manifolds. ■

We now show how Assumption 4.2 leads to bounded norm-equivalence terms.

Proposition 4.2. *If \mathcal{P} is (γ, r) -gap ensuring then at any iteration $t \geq 1$ of Algorithm 1, it holds that*

$$(1 - \sqrt{m}\gamma) \cdot \|\cdot\|_{y^{t,0}} \leq \|\cdot\|_{y^{t,0},\mathcal{P}} \leq (1 + \sqrt{m}\gamma) \cdot \|\cdot\|_{y^{t,0}}.$$

Proof of Proposition 4.2. We suppress the dependence on t in the notation throughout this proof. Observe that

$$\|v\|_{y^0}^2 = \sum_{k \in [m]} \|P'_k v\|_{y^0}^2 \leq \sum_{k \in [m]} (\|P_k v\|_{y^0} + \|(P'_k - P_k)v\|_{y^0})^2$$

$$\begin{aligned}
&\leq \sum_{k \in [m]} (\|P_k v\|_{y^0} + \gamma \|v\|_{y^0})^2 \\
&= \sum_{k \in [m]} \left[\|P_k v\|_{y^0}^2 + 2\gamma \|v\|_{y^0} \|P_k v\|_{y^0} + \gamma^2 \|v\|_{y^0}^2 \right] \\
&= \|v\|_{y^0, \mathcal{P}}^2 + \sum_{k \in [m]} \left[2\gamma \|v\|_{y^0} \|P_k v\|_{y^0} + \gamma^2 \|v\|_{y^0}^2 \right] \\
&\leq \|v\|_{y^0, \mathcal{P}}^2 + 2\sqrt{m}\gamma \|v\|_{y^0} \|v\|_{y^0, \mathcal{P}} + m\gamma^2 \|v\|_{y^0}^2 \\
&= (\|v\|_{y^0, \mathcal{P}} + \sqrt{m}\gamma \|v\|_{y^0})^2 \\
\implies (1 - \sqrt{m}\gamma) \|v\|_{y^0} &\leq \|v\|_{y^0, \mathcal{P}}.
\end{aligned}$$

A similar computation applied to the inequality

$$\|v\|_{y^0, \mathcal{P}}^2 = \sum_{k \in [m]} \|P_k v\|_{y^0}^2 \leq \sum_{k \in [m]} (\|P'_k v\|_{y^0} + \|(P_k - P'_k)v\|_{y^0})^2$$

shows that

$$\|v\|_{y^0, \mathcal{P}} \leq (1 + \sqrt{m}\gamma) \|v\|_{y^0}$$

which completes our chain of inequalities. \square

We are now in a position to complete our analysis. We first verify that (10) holds for any projection selection rule \mathcal{P} with $C = L_f^2(m-1)/L_{\min}^2$.

Lemma 4.3. *Suppose f is L_f -smooth and Assumption 4.1 holds. Then for any $t \geq 1$, the iterates of Algorithm 1*

$$\left\| P_{t,k} \nabla f(y^{t,0}) - P_{t,k} \Gamma_{y^{t,k-1}}^{y^{t,0}} \nabla f(y^{t,k-1}) \right\|_{y^{t,0}}^2 \leq \frac{L_f^2(m-1)}{L_{\min}^2} \sum_{j \in [k-1]} \|P_{t,j}^{y^{t,j-1}} \nabla f(y^{t,j-1})\|_{y^{t,j-1}}^2.$$

Proof of Lemma 4.3. We suppress the dependence of t in the notation throughout the proof. Since P_k is an orthogonal projection,

$$\left\| \nabla f(y^0) - \Gamma_{y^{k-1}}^{y^0} \nabla f(y^{k-1}) \right\|_{y^0}^2 \geq \left\| P_k \nabla f(y^0) - P_k \Gamma_{y^{k-1}}^{y^0} \nabla f(y^{k-1}) \right\|_{y^0}^2$$

Thus, it suffices to examine $\left\| \nabla f(y^0) - \Gamma_{y^{k-1}}^{y^0} \nabla f(y^{k-1}) \right\|_{y^0}^2$ to get the bound. For $k > 1$,

$$\left\| \nabla f(y^0) - \Gamma_{y^{k-1}}^{y^0} \nabla f(y^{k-1}) \right\|_{y^0}^2 \leq L_f d(y^{k-1}, y^0)^2$$

by L_f -smoothness (2). Then

$$d(y^{k-1}, y^0)^2 \leq (k-1) \sum_{j \in [k-1]} d(y^{j-1}, y^j)^2$$

by the bound $\|\cdot\|_2 \geq \sqrt{k-1} \|\cdot\|_1$ where we regard both as norms on \mathbb{R}^{k-1} . Thus

$$\left\| \nabla f(y^0) - \Gamma_{y^{k-1}}^{y^0} \nabla f(y^{k-1}) \right\|_{y^0}^2 \leq L_f d(y^{k-1}, y^0)^2$$

$$\begin{aligned}
&\leq L_f^2(k-1) \sum_{j \in [k-1]} d(y^{j-1}, y^j)^2 \\
&\leq L_f^2(k-1) \sum_{j \in [k-1]} \frac{1}{L_j^2} \|P_j^{y^{j-1}} \nabla f(y^{j-1})\|_{y^{j-1}}^2 \\
&\leq \frac{L_f^2(k-1)}{L_{\min}^2} \sum_{j \in [k-1]} \|P_j^{y^{j-1}} \nabla f(y^{j-1})\|_{y^{j-1}}^2
\end{aligned}$$

In the third inequality, we used the fact that $y^j = \text{Exp}_{y^{j-1}} \left(-\frac{1}{L_j} P_j^{y^{j-1}} \nabla f(y^{j-1}) \right)$. \square

We now present a refined version of (Dec) in light of Assumption 4.2.

Lemma 4.4. *Suppose that Assumption 4.1 holds and that \mathcal{P} is (γ, r) -gap ensuring (Assumption 4.2). Let*

$$\eta := \frac{L_{\min}^2(1 - \sqrt{m}\gamma)^2}{4L_{\max} \left(L_{\min}^2 + L_f^2(m-1)m \right)}, \quad \eta' := \frac{L_{\min} r^2}{2m}.$$

Then for any iteration $t \geq 1$ of Algorithm 1,

1. If $\sum_{k \in [m]} d(y^{t,k-1}, y^{t,k}) \leq r$, then

$$f(y^{t,0}) - f(y^{t,m}) \geq \eta \|\nabla f(y^{t,0})\|_{y^{t,0}}^2. \quad (13)$$

2. If $\sum_{k \in [m]} d(y^{t,k-1}, y^{t,k}) > r$, then

$$f(y^{t,0}) - f(y^{t,m}) \geq \eta'. \quad (14)$$

Proof of Lemma 4.4. We fix some $t \geq 1$ and suppress the dependence on t in the notation throughout the proof.

1. If $\sum_{k \in [m]} d(y^{k-1}, y^k) \leq r$, then $\max_{k \in [m]} d(y^k, y^0) \leq r$ also by the triangle inequality. The result now follows from (11) in Lemma 4.1, Proposition 4.2 and Lemma 4.3.
2. If $\sum_{k \in [m]} d(y^{k-1}, y^k) > r$, then

$$\begin{aligned}
r^2 &< \left(\sum_{k \in [m]} d(y^{k-1}, y^k) \right)^2 \leq m \sum_{k \in [m]} d(y^{k-1}, y^k)^2 \leq m \sum_{k \in [m]} \frac{1}{L_k^2} \left\| P_k^{y^{k-1}} \nabla f(y^{k-1}) \right\|_{y^{k-1}}^2 \\
&\leq \frac{m}{L_{\min}} \sum_{k \in [m]} \frac{1}{L_k} \left\| P_k^{y^{k-1}} \nabla f(y^{k-1}) \right\|_{y^{k-1}}^2 \\
&\leq \frac{2m}{L_{\min}} (f(y^0) - f(y^m)),
\end{aligned}$$

where the second inequality follows since $\|\cdot\|_1 \leq \sqrt{m}\|\cdot\|_2$ for ℓ_1 - and ℓ_2 -norms in \mathbb{R}^m , the third inequality follows from the definition of y^k in Algorithm 1, and the final inequality follows from (9) in Lemma 4.1. \square

Theorem 4.5 (Convergence with Deterministic \mathcal{P}). *Let $f : M \rightarrow \mathbb{R}$ be a L_f -smooth function such that Assumption 4.1 holds. and that \mathcal{P} is (γ, r) -gap ensuring (Assumption 4.2). Suppose that the sequence $\{x^t\}_{t \geq 1}$ is generated from Algorithm 1 with a projection selection rule \mathcal{P} satisfying Assumption 4.2. Let*

$$\eta := \frac{L_{\min}^2(1 - \sqrt{m}\gamma)^2}{4L_{\max} \left(L_{\min}^2 + L_f^2(m-1)m \right)}, \quad \eta' := \frac{L_{\min}r^2}{2m}.$$

Then the following hold:

1. For $t > (f(x^1) - f^*)/\eta'$, we have

$$\lim_{t \rightarrow \infty} \|\nabla f(x^t)\|_{x^t} = 0,$$

$$\min_{s \in [t]} \|\nabla f(x^s)\|_{x^s} \leq \max_{k \leq (f(x^1) - f^*)/\eta'} \sqrt{\frac{f(x^1) - f^* - \eta'k}{\eta(t-k)}} = \mathcal{O}\left(m/\sqrt{t}\right).$$

2. Any limit point $x^* \in M$ of $\{x^t\}_{t \in \mathbb{N}}$ is a stationary point $\nabla f(x^*) = 0$.
3. If in addition f is g -convex and the diameter of $\{x \in M : f(x) \leq f(x^1)\}$ is R , then

$$f(x^{t+1}) - f^* \leq \frac{f(x^1) - f^*}{1 + \min \left\{ \frac{(f(x^1) - f^*)\eta}{R^2}, \frac{\eta'}{f(x^1) - f^*} \right\} t} = \mathcal{O}(m^2/t).$$

Note that the rates in Theorem 4.5 are worse than their Euclidean counterparts [2]. However, we show that in Example 4.7 that when M is a product manifold (which generalizes the Euclidean setting, see Example 4.8), Lemma 4.3 can be improved, which results in improved rates in Theorem 4.5. These match the ones from the Euclidean setting [2].

Proof of Theorem 4.5. Fix some $t \geq 1$. For $s \in [t-1]$, we know from Lemma 4.4 that one of the following must hold:

$$f(x^s) - f(x^{s+1}) \geq \eta \|\nabla f(x^s)\|_{x^s}^2, \quad \text{or} \quad f(x^s) - f(x^{s+1}) \geq \eta'.$$

Let $R(t) := \{s \in [t-1] : f(x^s) - f(x^{s+1}) \geq \eta'\}$. Summing the inequalities over $s \in [t-1]$ gives

$$f(x^1) - f(x^t) \geq \eta \sum_{s \in [t-1] \setminus R(t)} \|\nabla f(x^s)\|_{x^s}^2 + |R(t)|\eta'.$$

Note that we must have $|R(t)| \leq (f(x^1) - f(x^t))/\eta' \leq (f(x^1) - f^*)/\eta'$.

1. Fix $\delta > 0$. Since $|R(t)|$ remains bounded, if there exists a subsequence of $\{\|\nabla f(x^t)\|_{x^t} : t \geq 1\}$ which is bounded below by δ , then $\eta \sum_{s \in [t-1] \setminus R(t)} \|\nabla f(x^s)\|_{x^s}^2 \rightarrow \infty$. Therefore, every subsequence of $\{\|\nabla f(x^t)\|_{x^t} : t \geq 1\}$ must be less than δ eventually, so $\lim_{t \rightarrow \infty} \|\nabla f(x^t)\|_{x^t} = 0$.

The bound above implies

$$f(x^1) - f(x^t) - \eta'|R(t)| \geq \eta \sum_{s \in [t] \setminus R(t)} \|\nabla f(x^s)\|_{x^s}^2 \geq (t - |R(t)|)\eta \min_{s \in [t]} \|\nabla f(x^s)\|_{x^s}^2,$$

which gives the second result.

2. Since $\|\nabla f(x^t)\|_{x^t} \rightarrow 0$, any limit point x^* is a stationary point.
3. For $s \in [t]$, if we have $f(x^s) - f(x^{s+1}) \geq \eta \|\nabla f(x^s)\|_{x^s}^2$ then we also have

$$\begin{aligned} f(x^s) - f(x^*) &\leq \langle \nabla f(x^s), \text{Exp}_{x^s}^{-1}(x^*) \rangle_{x^s} \leq \|\nabla f(x^s)\|_{x^s} \|\text{Exp}_{x^s}^{-1}(x^*)\|_{x^s} \leq \|\nabla f(x^s)\|_{x^s} d(x^s, x^*) \\ &\leq \sqrt{(f(x^s) - f(x^{s+1}))/\eta} d(x^s, x^*) \\ &\leq R \sqrt{(f(x^s) - f(x^{s+1}))/\eta}. \end{aligned}$$

Otherwise, we get $f(x^s) - f(x^{s+1}) \geq \eta'$. Denoting $A_s := f(x^s) - f(x^*)$, these two inequalities can be written as

$$\frac{\eta}{R^2} A_s^2 \leq A_s - A_{s+1} \quad \text{or} \quad A_s - A_{s+1} \geq \eta'.$$

Dividing by $A_s A_{s+1}$ and recognizing that $1 \leq A_s/A_{s+1}$, $A_s A_{s+1} \leq A_1^2$, we have

$$\frac{\eta}{R^2} \leq \frac{1}{A_{s+1}} - \frac{1}{A_s} \quad \text{or} \quad \frac{\eta'}{A_1^2} \leq \frac{1}{A_{s+1}} - \frac{1}{A_s} \implies \frac{1}{A_{s+1}} - \frac{1}{A_s} \geq \min \left\{ \frac{\eta}{R^2}, \frac{\eta'}{A_1^2} \right\}.$$

Summing this over $s \in [t]$, we get

$$\frac{1}{A_{t+1}} - \frac{1}{A_1} \geq t \min \left\{ \frac{\eta}{R^2}, \frac{\eta'}{A_1^2} \right\} \implies A_{t+1} \leq \frac{A_1}{1 + \min \left\{ \frac{A_1 \eta}{R^2}, \frac{\eta'}{A_1} \right\} t},$$

which gives the result. □

4.2 Examples

We now exhibit a series of selection rules satisfying the gap ensuring condition. The first rule applies to any manifold.

Example 4.6 (Parallel Transport Rule). For any $t \geq 1$, given $y^{t,0}$, fix an orthogonal decomposition $\{P'_{t,k}\}_{k \in [m]}$ of $T_{y^{t,0}}M$. We will define \mathcal{P} recursively via the rule

$$\mathcal{P}(y^{t,k-1}, k) = \Gamma_{y^{t,0}}^{y^{t,k-1}} P'_{t,k} \Gamma_{y^{t,k-1}}^{y^{t,0}}.$$

It follows that $P_{t,k} = P'_{t,k}$ so \mathcal{P} is a $(0, \infty)$ -gap ensuring rule. Note that since $r = \infty$, the convergence rate bounds are obtained by taking the bounds in Theorem 4.5 and setting $\eta' = 0$. ■

While \mathcal{P} applies generally to any manifold, the main difficulty in its practical implementation is that on many manifolds it is unclear how, or incredibly expensive, to compute parallel transport maps. For example, on every Stiefel manifold $\text{St}(p, n)$ with the exception to that of the manifold of orthogonal matrices $O_n = \text{St}(n, n)$, there is no known closed form for the parallel transport operator. We provide a non-trivial, and very useful, example of a (γ, r) -gap ensuring rule for the orthogonal manifold O_n . While the construction is conceptually simple, it requires some background knowledge of the manifold structure of O_n , so we defer its presentation to Section 6.

For the rest of this section, we examine the class of product manifolds, for which the rule described in Example 4.6 can be efficiently computed.

Example 4.7 (Product Manifolds). Let $M := M_1 \times \dots \times M_m$ be a Cartesian product of Riemannian manifolds with the endowed product structure. Write $y \in M$ as $y = (y_1, \dots, y_m)$. The tangent space is $T_y M = \bigoplus_{k \in [m]} T_{y_k} M_k$. Denote the zero element of $T_{y_k} M$ as 0_{y_k} . Due to the product structure of M , it can be easily shown that the exponential map and parallel transport also decomposes into components M_k : for any $x = (x_1, \dots, x_m) \in M$, $v = (v_1, \dots, v_m) \in T_x M$,

$$\begin{aligned} y &= \text{Exp}_x(v) = (\text{Exp}_{x_1}(v_1), \dots, \text{Exp}_{x_m}(v_m)) \\ \Gamma_x^y &= (\Gamma_{x_1}^{y_1}, \dots, \Gamma_{x_m}^{y_m}). \end{aligned}$$

For $y = (y_1, \dots, y_m) \in M$, $k \in [m]$, we define

$$\mathcal{P}(y, k) : v = (v_1, \dots, v_m) \mapsto (0_{y_1}, \dots, 0_{y_{k-1}}, v_k, 0_{y_{k+1}}, \dots, 0_{y_m}).$$

Observe that $\{\mathcal{P}(y, k)\}_{k \in [m]}$ is an orthogonal decomposition of $T_y M$ for any $y \in M$. The projection selection rule we use in Algorithm 1 is the same as the one from Example 4.6: at a point $y^{t,0} \in M$, use the orthogonal decomposition $\{\mathcal{P}(y^{t,0}, k)\}_{k \in [m]}$ and define

$$P_k^{y^{t,k-1}} := \Gamma_{y^{t,0}}^{y^{t,k-1}} \mathcal{P}(y^{t,0}, k) \Gamma_{y^{t,k-1}}^{y^{t,0}}.$$

As before, this is a $(0, \infty)$ -gap ensuring rule. Due to the decomposability of the exponential map and parallel transport operators on the product manifold M , it is easily shown that

$$P_k^{y^{t,k-1}} = \mathcal{P}(y^{t,k-1}, k).$$

Thus, due to the product structure, unlike Example 4.6, $P_k^{y^{t,k-1}}$ does not require computing parallel transports.

In fact, the product manifold setting offers two more simplifications. Note that if $y = \text{Exp}_x(\mathcal{P}(x, k)v)$, then $y_j = x_j$ for $j \neq k$, and $y_k = \text{Exp}_{x_k}(v_k)$. Thus, at any inner iteration $k \in [m]$ of Algorithm 1, computing projections of a tangent vector and updating the current iterate is easy; we only need to update the component corresponding to M_k .

Furthermore, due to the decomposability of the parallel transport, the constant in Lemma 4.3 can be improved to L_f^2/L_{\min}^2 (removing a factor of $m - 1$). This means that the η, η' terms in the bounds from Lemma 4.4 and Theorem 4.5 become $\eta = \frac{L_{\min}^2}{4L_{\max}(L_{\min}^2 + L_f^2 m)}$, $\eta' = \infty$. A closer inspection shows that these bounds exactly match those of Beck and Tetrushvili [2] for BCD in the Euclidean setting. \blacksquare

As a special case of this subspace selection rule, we have the selection rules for BCD.

Example 4.8 (Block Coordinate Descent). Let e_i denote the i -th standard coordinate vector in \mathbb{R}^n and $\{S_k\}_{k \in [m]}$ be a partition of the set $\{e_i\}_{i \in [n]}$. It follows that \mathbb{R}^n is equal to the direct sum of $\bigoplus_{k \in [m]} \text{Span}(S_k)$, which is a product manifold. Furthermore, letting P_k be the projection onto $\text{Span}(S_k)$, the rule from the previous example is $\mathcal{P}(x, k) = P_k$, which is exactly block coordinate descent. This shows that our analysis can recover the standard analysis for BCD by Beck and Tetrushvili [2].

Another important application of the product setting is Tensor principle component analysis (PCA).

Example 4.9 (Tensor PCA). Given a tensor $T \in \mathbb{R}^{d_1 \times \dots \times d_m}$, we wish to compute a decomposition of T into tensor products of orthogonal matrices. We do this by solving the following problem:

$$\min_{C, U_1, \dots, U_m} \|T - C \times_1 U_1 \times \dots \times_m U_m\|_F^2$$

$$\begin{aligned} \text{s.t. } C &\in \mathbb{R}^{n_1 \times \dots \times n_m} \\ U_k &\in \text{St}(d_k, n_k), \quad k \in [m]. \end{aligned} \tag{15}$$

In the objective, $\|\cdot\|_F$ is the Frobenius norm on tensors, and \times_k is the mode- k tensor product. The objective is smooth in the variables (C, U_1, \dots, U_m) ; we refer to Xu [23] for detailed definitions and smoothness results.

Of particular interest to us is the domain. Here, $\text{St}(p, n)$ is the *Stiefel manifold* of $n \times p$ matrices (with $p \leq n$) with orthogonal columns:

$$\text{St}(p, n) = \left\{ U \in \mathbb{R}^{n \times p} : U^\top U = I_p \right\}.$$

Thus, the domain is actually a product manifold:

$$(C, U_1, \dots, U_m) \in M := \mathbb{R}^{n_1 \times \dots \times n_m} \times \text{St}(d_1, n_1) \times \dots \times \text{St}(d_m, n_m).$$

To describe the geometry of M , it thus suffices to describe the geometry of $\text{St}(p, n)$. We give the key identities here, and refer to Edelman et al. [5] for full details. The tangent space at a matrix $U \in \text{St}(p, n)$ is

$$T_U \text{St}(p, n) := \left\{ \begin{array}{l} A = -A^\top \in \mathbb{R}^{p \times p}, \\ UA + U_\perp B : B \in \mathbb{R}^{(n-p) \times p}, \\ U_\perp \in \mathbb{R}^{n \times (n-p)} \text{ s.t. } UU^\top + U_\perp U_\perp^\top = I_n \end{array} \right\},$$

which is a subset of $\mathbb{R}^{n \times p}$. Since $\text{St}(p, n)$ is embedded in $\mathbb{R}^{n \times p}$, the gradient of a function $f : \text{St}(p, n) \rightarrow \mathbb{R}$ at U can be computed as the projection of the Euclidean gradient onto $T_U \text{St}(p, n)$ (see Zhang et al. [26, Section 2]). The exponential map at some $UA + U_\perp B \in T_U \text{St}(p, n)$ is

$$\text{Exp}_U(UA + U_\perp B) := [U \quad U_\perp] \exp \left(\begin{bmatrix} A & -B^\top \\ B & 0 \end{bmatrix} \right) \begin{bmatrix} I_p \\ 0 \end{bmatrix}.$$

Computing this exponential map requires a computing the matrix exponential of a $n \times n$ matrix. If we attempt to run (the generalization of) gradient descent to solve (15), we must compute the exponential map for each $\text{St}(d_k, n_k)$ at each iteration. On the other hand, using tangent subspace descent, we need only compute one of these each iteration. ■

5 Randomized Subspace Selection

In this section, we analyze Tangent Subspace Descent when it is equipped with a randomized subspace selection rule \mathcal{P} . Recall that a randomized subspace selection rule \mathcal{P} can be instantiated in Algorithm 1 by setting $m = 1$, and $P_t := \mathcal{P}(y^{t,0}, 1)$ is a projection onto a randomly selected subspace of $T_{x^{t-1}}M$. Thus, at iteration $t \geq 1$ in Algorithm 1, we compute

$$x^t = \text{Exp}_{x^{t-1}}(-\eta_t P_t \nabla f(x^{t-1})), \tag{16}$$

noting that x^t itself is a random variable due to the randomness in P_t .

Most existing rules of practical interest draw P_t from a finite collection of projections, such as the ones for randomized BCD [8, 18] and the orthogonal manifold [21]. In Section 5.1, we propose a new and relatively unrestrictive condition on \mathcal{P} , the C -randomized norm condition, that guarantees convergence for smooth functions. This condition is satisfied for the schemes of [8, 18, 21] as well as other important classes, demonstrated in Section 5.2. However, we note that our condition is sufficiently general that it covers the case when \mathcal{P} draws from an infinite collection of projections. We exploit this in Section 5.2.1, where we present a new randomized \mathcal{P} for optimization over the Stiefel manifold $\text{St}(p, n)$ for $p < n$.

5.1 A Convergent Randomized \mathcal{P} via the C -Randomized Norm Condition

As before, the aim is to achieve a sufficient decrease inequality similar to (Dec), from which the convergence rate follows. In Section 4, we saw that Assumption 4.1 facilitated the selection of stepsizes $\eta_{t,k}$ and achieving sufficient decrease. Since P_t is a random projection drawn from a possibly infinite collection, we will assume a uniform constant $L_f < \infty$ exists such that for any $t \geq 1$ and realization P_t ,

$$f(\text{Exp}_{x^{t-1}}(P_t v)) \leq f(x^{t-1}) + \langle \nabla f(x^{t-1}), \mathcal{P}(x^{t-1}, k)v \rangle_{x^{t-1}} + \frac{L_f}{2} \|P_t v\|_{x^{t-1}}^2. \quad (17)$$

Since we are considering smooth objective functions f , such an upper bound exists from (4). From this, we can choose stepsizes η_t which give a preliminary sufficient decrease guarantee, similar to Lemma 4.1. (In practice, we can implement η_t via backtracking line search, to achieve the condition (18), which can result in larger steps and quicker convergence.)

Lemma 5.1 (Randomized Sufficient Decrease Template). *Suppose f is an L_f -smooth function so that (17) holds, and that at each $t \geq 1$ in Algorithm 1, the stepsize is chosen as $\eta_t := 1/L_f$. Then*

$$f(x^{t-1}) - \mathbb{E}[f(x^t) \mid x^{t-1}] \geq \frac{1}{2L_f} \mathbb{E} \left[\|P_t \nabla f(x^{t-1})\|_{x^{t-1}}^2 \mid x^{t-1} \right].$$

Proof of Lemma 5.1. The stepsize selection $\eta_t = 1/L_f$ minimizes the right hand side of (17) to get

$$f(x^{t-1}) - f(x^t) \geq \frac{1}{2L_f} \|P_t \nabla f(x^{t-1})\|_{x^{t-1}}^2. \quad (18)$$

Taking the expectation over the randomness in P_t , we have

$$f(x^{t-1}) - \mathbb{E}[f(x^t) \mid x^{t-1}] \geq \frac{1}{2L_f} \mathbb{E} \left[\|P_t \nabla f(x^{t-1})\|_{x^{t-1}}^2 \mid x^{t-1} \right].$$

□

In order to get a sufficient decrease guarantee similar to (Dec), we must lower bound the right hand side with $\|\nabla f(x^{t-1})\|_{x^{t-1}}$. This motivates the C -randomized norm condition for \mathcal{P} .

Assumption 5.1 (C -randomized norm). Suppose there exists $C > 0$ such that at any $t \geq 1$ in Algorithm 1 and for any $v \in T_{x^{t-1}}M$, we have

$$\sqrt{\mathbb{E} \left[\|P_t v\|_{x^{t-1}}^2 \mid x^{t-1} \right]} \geq C \|v\|_{x^{t-1}}.$$

Then we say that \mathcal{P} satisfies the C -randomized norm condition.

Assumption 5.1 precludes the progressive degeneration of $v \mapsto \mathbb{E}[\|P_t v\|_{x^{t-1}}^2 \mid x^{t-1}]$ that undergirds the pathology of Theorem 3.2. We are now ready to present the convergence analysis.

Theorem 5.2. *Let $f : M \rightarrow \mathbb{R}$ be a L_f -smooth function satisfying (17). Suppose that the sequence $\{x^t\}_{t \geq 1}$ is generated from Algorithm 1 with \mathcal{P} satisfying Assumption 5.1 and stepsizes chosen according to Lemma 5.1. Then the following hold:*

1. We have

$$\begin{aligned} \lim_{t \rightarrow \infty} \mathbb{E} [\|\nabla f(x^t)\|_{x^t}] &= 0, \\ \min_{s \in [t]} \mathbb{E} [\|\nabla f(x^{s-1})\|_{x^{s-1}}] &\leq \sqrt{\frac{2L_f(f(x^0) - f^*)}{C^2 t}}. \end{aligned}$$

2. Almost surely, $\|\nabla f(x^t)\|_{x^t} \rightarrow 0$ and any limit point $x^* \in M$ of $\{x^t\}_{t \in \mathbb{N}}$ is a stationary point $\nabla f(x^*) = 0$.

3. If in addition f is g -convex and the diameter of $\{x \in M : f(x) \leq f(x^1)\}$ is R , then

$$\mathbb{E}[f(x^t)] - f^* \leq \frac{2L_f R^2 (f(x^0) - f^*)}{2L_f R^2 + C^2 (f(x^0) - f^*) t}.$$

Proof of Theorem 5.2.

$$f(x^{t-1}) - \mathbb{E}[f(x^t) \mid x^{t-1}] \geq \eta \|\nabla f(x^{t-1})\|_{x^{t-1}}^2. \quad (19)$$

1. Take expectations in (19) to get $\eta \mathbb{E}[\|\nabla f(x^{s-1})\|_{x^{s-1}}^2] \leq \mathbb{E}[f(x^{s-1})] - \mathbb{E}[f(x^s)]$. Summing this for $s \in [t]$ gives

$$\eta \sum_{s \in [t]} \mathbb{E}[\|\nabla f(x^{s-1})\|_{x^{s-1}}^2] \leq \mathbb{E}[f(x^0)] - \mathbb{E}[f(x^t)] \leq f(x^0) - f^*.$$

This holds for any t , and since $\eta > 0$, we must have $\mathbb{E}[\|\nabla f(x^t)\|_{x^t}^2] \rightarrow 0$. Furthermore, observe that

$$\eta t \min_{s \in [t]} \mathbb{E}[\|\nabla f(x^{s-1})\|_{x^{s-1}}^2] \leq \eta \sum_{s \in [t]} \mathbb{E}[\|\nabla f(x^{s-1})\|_{x^{s-1}}^2] \leq f(x^0) - f^*,$$

from which the second result follows by Jensen's inequality.

2. Almost sure convergence of $\|\nabla f(x^t)\|_{x^t} \rightarrow 0$ follows by taking \mathcal{F}_t to be the filtration generated by $\{x^1, \dots, x^t\}$, $a_t = f(x_t)$ and $b_t = \eta \|\nabla f(x^t)\|_{x^t}^2$. Then $a_t \geq f^*$ and integrability of a_t holds since $a_t = f(x^t) \leq f(x^1)$, which holds since by construction $f(x^t)$ is monotonically decreasing. Then (19) is rewritten as $b_t \leq a_t - \mathbb{E}[a_t \mid \mathcal{F}_t]$, which also shows integrability of b_t . Then apply Proposition A.1 (see Appendix A) to deduce $\|\nabla f(x^t)\|_{x^t} \rightarrow 0$ almost surely. Then any limit point is a stationary point since f is continuously differentiable.

3. Since f is g -convex, we have for any minimizer x^* , $f(x^{s-1}) - f^* \leq R \|\nabla f(x^t)\|_{x^t}$. This and (19) give $\frac{\eta}{R^2} (f(x^{s-1}) - f^*)^2 \leq f(x^{s-1}) - \mathbb{E}[f(x^s) \mid x^{s-1}]$. By Jensen's inequality, the expectation of the left hand side is greater than or equal to $\frac{\eta}{R^2} (\mathbb{E}[f(x^{s-1})] - f^*)^2$, hence $\frac{\eta}{R^2} (\mathbb{E}[f(x^{s-1})] - f^*)^2 \leq \mathbb{E}[f(x^{s-1})] - \mathbb{E}[f(x^s)]$. Write $A_s = \mathbb{E}[f(x^s)] - f^*$, hence we have $\frac{\eta}{R^2} A_s^2 \leq A_{s-1} - A_s$. Dividing by $A_{s-1} A_s$ and recognizing that $1 \leq A_{s-1}/A_s$, $A_{s-1} A_s \leq A_0^2$, we have

$$\frac{\eta}{R^2} \leq \frac{1}{A_s} - \frac{1}{A_{s-1}}.$$

Summing this over $s \in [t]$, we get

$$\frac{1}{A_t} - \frac{1}{A_0} \geq t \frac{\eta}{R^2} \implies A_t \leq \frac{A_0}{1 + \frac{A_0 \eta}{R^2} t},$$

which gives the result.

□

Note that there is no explicit dependence on m in the bounds of Theorem 5.2. In fact, this dependence is hidden in the constant C . The constant C is affected by how the probabilities p_k are chosen, as well as m . As we will see, for natural choices of probabilities, the dependence on m will be similar to deterministic TSD for product manifolds (Example 4.7).

5.2 Examples

We now elaborate a few important examples of C -randomized norm ensuring subspace selection rules as well as our new subspace selection rule for a general Stiefel manifold. We begin by providing a general randomized selection rule which works for any manifold M .

Example 5.3. At each $x^{t-1} \in M$, take any orthogonal decomposition $\{\mathcal{P}(x^{t-1}, k)\}_{k \in [m]}$ of $T_x M$. Let $P_t = \mathcal{P}(x^{t-1}, k)$ with probability $p_k > 0$, where $\sum_{k \in [m]} p_k = 1$. Then we can take $C = \sqrt{\min_{k \in [m]} p_k}$ in Assumption 5.1: for any $v \in T_{x^{t-1}} M$, we see

$$\mathbb{E} [\|P_t v\|_{x^{t-1}}^2 \mid x^{t-1}] = \sum_{k \in [m]} p_k \|\mathcal{P}(x^{t-1}, k)v\|_{x^{t-1}}^2 \geq \left(\min_{k \in [m]} p_k \right) \sum_{k \in [m]} \|\mathcal{P}(x^{t-1}, k)v\|_{x^{t-1}}^2 = \left(\min_{k \in [m]} p_k \right) \|v\|_{x^{t-1}}^2.$$

To optimize the bounds in Theorem 5.2, we make C as large as possible, which is done by taking $p_k = 1/m$, so $C = \sqrt{1/m}$. The bounds then become $\mathcal{O}(\sqrt{m}/t)$ and $\mathcal{O}(m/t)$ respectively. These bounds match the rates from Example 4.7 on product manifolds (but with slightly better hidden constants). ■

Example 5.4 (Product Manifolds and Randomized Coordinate Descent). Consider again $M = M_1 \times \dots \times M_m$, the product manifold setting from Example 4.7. Since for any $x = (x_1, \dots, x_m)$, $T_x M = \bigoplus_{k \in [m]} T_{x_k} M$, so this yields a natural orthogonal decomposition

$$\mathcal{P}(x^{t-1}, k) : v \mapsto \left(0_{x_1^{t-1}}, \dots, 0_{x_{k-1}^{t-1}}, v_k, 0_{x_{k+1}^{t-1}}, \dots, 0_{x_m^{t-1}} \right), \quad k \in [m].$$

Randomized BCD is instantiated in the following way: set $M = \mathbb{R}^n$, $\{S_k\}_{k \in [m]}$ is a partition of $\{e_i\}_{i \in [n]}$, and $M_k = \text{Span}(S_k)$ for $k \in [m]$. ■

Example 5.5 (Orthogonal Manifold). We consider $O_n = \text{St}(n, n)$, which has tangent space that can be nicely expressed as

$$T_X O_n = \text{Span}\{X H_{ij} : 1 \leq i, j \leq n\}, \quad H_{ij} = e_i e_j^\top - e_j e_i^\top,$$

for any $X \in O_n$. The collection of tangent vectors $\{X H_{ij} : 1 \leq i, j \leq n\}$ is clearly orthogonal under the Riemannian metric $\langle X H_{ij}, X H_{i'j'} \rangle_X = \frac{1}{2} \text{Tr}((X H_{ij})^\top X H_{i'j'})$. We let $P_t = \mathcal{P}(X^{t-1}, ij)$ be the projection onto $\text{Span}(\{X^{t-1} H_{ij}\})$ for each $1 \leq i < j \leq n$, which is chosen with probability $p_{ij} > 0$, where $\sum_{1 \leq i < j \leq n} p_{ij} = 1$. Then \mathcal{P} satisfies Assumption 5.1 with $C = \sqrt{\min_{1 \leq i < j \leq n} p_{ij}}$. This is the subspace selection rule of [21]. ■

5.2.1 Randomized rule for the Stiefel manifold

We can significantly generalize the latter example to $\text{St}(p, n)$ with $p < n$. Given a matrix $U \in \mathbb{R}^{n \times p}$, denote $\text{Ker}(U^\top) \subseteq \mathbb{R}^n$ to be the kernel of U^\top . The tangent space $T_U \text{St}(p, n)$ from Example 4.9 can be written as

$$T_U \text{St}(p, n) = \left\{ UA + [b_1 \ \dots \ b_p] : \begin{array}{l} A = -A^\top \in \mathbb{R}^{p \times p}, \\ b_\ell \in \text{Ker}(U^\top), \ell \in [p] \end{array} \right\}.$$

Furthermore, given $V, V' \in T_U \text{St}(p, n)$, the Riemannian metric is $\langle V, V' \rangle_U = \text{Tr}(V^\top (I_n - UU^\top / 2) V')$.

We consider the non-orthogonal subspace decomposition of $T_U \text{St}(p, n)$:

$$D = \{UH_{ij} : 1 \leq i < j \leq p\} \cup \left(\bigcup_{\ell \in [p]} \{ve_{\ell,p}^\top : v \in \text{Ker}(U^\top), \|v\|_2 = 1\} \right),$$

where $e_{\ell,p}$ is the ℓ -th standard basis vector in \mathbb{R}^p . Note that $\text{Span}(D) = T_U \text{St}(p, n)$. Our main idea is draw a random tangent vector V from D , use it to define a projection onto the subspace spanned by V , then show the random projection satisfies Assumption 5.1.

We define the following distribution over D . Suppose we have a distribution over the indices $\{(i, j) : 1 \leq i < j \leq p\} \cup [p]$ where a pair (i, j) is drawn with probability $p_{ij} > 0$ and ℓ is drawn with probability $p_\ell > 0$, so $\sum_{1 \leq i < j \leq p} p_{ij} + \sum_{\ell \in [p]} p_\ell = 1$. We now draw a random element from $\{(i, j) : 1 \leq i < j \leq p\} \cup [p]$. If it is a pair (i, j) , then the element drawn from D is $V = UH_{ij}$. If it is a single index ℓ , then we randomly draw from the set $\{ve_{\ell,p}^\top : v \in \text{Ker}(U^\top)\}$ by drawing v uniformly at random over $\{v : \|v\|_2 = 1, v \in \text{Ker}(U^\top)\}$ and setting $V = ve_{\ell,p}^\top$. In practice, drawing v can be implemented as $z \sim N(0, I_n)$, $v' = (I_n - UU^\top)z$, $v = v'/\|v'\|_2$. Since U has orthonormal columns, $I_n - UU^\top$ is simply the projection matrix onto $\text{Ker}(U^\top)$.

Theorem 5.6. Fix any $U \in \text{St}(p, n)$, and let $\mathcal{P}(U)$ be the random projection onto the subspace $\{\alpha V : \alpha \in \mathbb{R}\} \subset T_U \text{St}(p, n)$, where V is drawn randomly according to the distribution described above. Let $C := \sqrt{\min\{\min_{1 \leq i < j \leq p} p_{ij}, \min_{\ell \in [p]} p_\ell / (n - p)\}}$. Then for any $W \in T_U \text{St}(p, n)$,

$$\mathbb{E} \left[\|\mathcal{P}(U)W\|_U^2 \right] \geq C^2 \|W\|_U^2.$$

Proof of Theorem 5.6. Denote $W = UA + [b_1 \ \dots \ b_p]$ and observe that

$$\|W\|_U^2 = \sum_{1 \leq i < j \leq p} a_{ij}^2 + \sum_{\ell \in [p]} \|b_\ell\|_2^2.$$

Given $V \in T_U \text{St}(p, n)$, the projection of W onto $\{\alpha V : \alpha \in \mathbb{R}\}$ is given by

$$\frac{\langle V, W \rangle_U}{\|V\|_U^2} V.$$

When $V = UH_{ij}$, $\|V\|_U^2 = 1$, so the projection is

$$\text{Tr} \left(H_{ij}^\top U^\top (I - UU^\top / 2) W \right) UH_{ij} = \frac{1}{2} \text{Tr} \left(H_{ij}^\top U^\top W \right) UH_{ij} = \frac{1}{2} \text{Tr} \left(H_{ij}^\top A \right) UH_{ij} = a_{ij} UH_{ij}.$$

When $V = ve_{\ell,p}^\top$ and $v \in \text{Ker}(U^\top)$, $\|v\|_2 = 1$, $\|V\|_U^2 = 1$, so the projection is

$$\text{Tr} \left(e_{\ell,p} v^\top (I - UU^\top / 2) W \right) ve_{\ell,p}^\top = \text{Tr} \left(e_{\ell,p} v^\top W \right) ve_{\ell,p}^\top = (v^\top b_\ell) ve_{\ell,p}^\top.$$

We then have

$$\mathbb{E} \left[\|\mathcal{P}(U)W\|_U^2 \right] = \sum_{1 \leq i < j \leq p} p_{ij} a_{ij}^2 + \sum_{\ell \in [p]} p_\ell \mathbb{E} \left[(v^\top b_\ell)^2 \right],$$

where the expectations are over the randomness in v . To evaluate the expectation, for convenience write $b := b_\ell$, and letting $\bar{u}_1, \dots, \bar{u}_{n-p}$ be an orthonormal basis for $\text{Ker}(U^\top)$, write $b = \sum_{k \in [n-p]} \beta_k \bar{u}_k$. Clearly $\|b\|_2^2 = \sum_{k \in [n-p]} \beta_k^2$. Then

$$\mathbb{E} \left[\left(v^\top b \right)^2 \right] = \mathbb{E} \left[\left(\sum_{k \in [n-p]} \beta_k v^\top \bar{u}_k \right)^2 \right] = \sum_{k \in [n-p]} \beta_k^2 \mathbb{E} \left[\left(v^\top \bar{u}_k \right)^2 \right] + 2 \sum_{1 \leq k < k' \leq n-p} \beta_k \beta_{k'} \mathbb{E} \left[\left(v^\top \bar{u}_k \right) \left(v^\top \bar{u}_{k'} \right) \right].$$

Since $\bar{u}_1, \dots, \bar{u}_{n-p}$ are an orthonormal basis for $\text{Ker}(U^\top)$ and $v \in \text{Ker}(U^\top)$, we have

$$1 = \mathbb{E}[\|v\|_2^2] = \sum_{k \in [n-p]} \mathbb{E} \left[\left(v^\top \bar{u}_k \right)^2 \right],$$

and since v is spherically symmetric in $\text{Ker}(U^\top)$, $v^\top \bar{u}_k$ are identically distributed for $k \in [n-p]$. Thus $\sum_{k \in [n-p]} \beta_k^2 \mathbb{E} \left[\left(v^\top \bar{u}_k \right)^2 \right] = \frac{1}{n-p} \|b\|_2^2$. Furthermore, due to symmetry, the pairs $(v^\top \bar{u}_k, v^\top \bar{u}_{k'})$ and $(-v^\top \bar{u}_k, v^\top \bar{u}_{k'})$ have the same distribution, thus

$$\mathbb{E} \left[\left(v^\top \bar{u}_k \right) \left(v^\top \bar{u}_{k'} \right) \right] = \mathbb{E} \left[\left(-v^\top \bar{u}_k \right) \left(v^\top \bar{u}_{k'} \right) \right] = 0.$$

Therefore,

$$\mathbb{E} \left[\left(v^\top b \right)^2 \right] = \frac{1}{n-p} \|b\|_2^2 \implies \mathbb{E} \left[\|\mathcal{P}(U)W\|_U^2 \right] = \sum_{1 \leq i < j \leq p} p_{ij} a_{ij}^2 + \frac{1}{n-p} \sum_{\ell \in [p]} p_\ell \|b\|_2^2 \geq C^2 \|W\|_U^2.$$

□

In Algorithm 1, at each $t \geq 1$ we can choose the random projection P_t according to the distribution $\mathcal{P}(U)$ outlined above with $U = U^{t-1}$, which will satisfy Assumption 5.1 by Theorem 5.6

6 Tangent Subspace Descent on the Orthogonal Manifold

In this section, we consider optimization on the manifold of orthogonal $n \times n$ matrices

$$M := O_n = \left\{ Y \in \mathbb{R}^{n \times n} : Y^\top Y = Y Y^\top = I_n \right\}.$$

We first briefly describe the manifold structure (see Edelman et al. [5] for full details). At a given $Y \in O_n$, the tangent space is

$$T_Y O_n := \{ Y A : A \in \text{Skew}_n \}, \quad \text{Skew}_n := \left\{ A \in \mathbb{R}^{n \times n} : A = -A^\top \right\}.$$

Since M is an embedded submanifold of $\mathbb{R}^{n \times n}$, it inherits the usual Euclidean metric

$$g_Y(YA, YB) := \text{Tr}((YA)^\top (YB)) = \text{Tr}(A^\top B).$$

Geodesics originating from $Y \in M$ in the direction $YC \in T_Y M$ are defined by

$$t \mapsto \gamma(t) := Y \text{Exp}_m(tC),$$

where Exp_m is the matrix exponential operator. Consequently, the exponential map is

$$\text{Exp}_Y(YC) := Y \text{Exp}_m(C).$$

Parallel transport of a tangent vector $YA \in T_Y M$ along the geodesic $t \mapsto Y \text{Exp}_m(tC)$ is given by

$$\Gamma_{Y,C}(t; A) := Y \text{Exp}_m(tC) \text{Exp}_m(tC/2)^\top A \text{Exp}_m(tC/2) \in T_{Y \text{Exp}_m(tC)} O_n.$$

Given a smooth function $\bar{f} : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}$, let $f : O_n \rightarrow \mathbb{R}$ be its restriction to O_n , and $X \in O_n \subset \mathbb{R}^{n \times n}$. Absil et al. [1, Eqs. (3.35), (3.37)] show us how to compute $\nabla f(X)$:

$$\nabla f(X) = \frac{1}{2} X \left(X^\top \nabla \bar{f}(X) - \nabla \bar{f}(X)^\top X \right), \quad (20)$$

where $\nabla \bar{f}(X)$ is the usual Euclidean gradient of \bar{f} at X . Note that $\frac{1}{2} (X^\top \nabla \bar{f}(X) - \nabla \bar{f}(X)^\top X) \in \text{Skew}_n$.

6.1 Deterministic Tangent Subspace Descent on O_n

The goal of this section is to provide a deterministic subspace selection rule \mathcal{P} that satisfies Assumption 4.2, and therefore tangent subspace descent with this rule converges according to Theorem 4.5. Our subspace decompositions will be based on partitions of the following basis for Skew_n :

$$\mathcal{B} := \left\{ \frac{1}{\sqrt{2}} (e_i e_j^\top - e_j e_i^\top) : 1 \leq i < j \leq n \right\}.$$

Let $\{\mathcal{B}_k\}_{k \in [m]}$ be some partition of \mathcal{B} . Define

$$\text{Skew}_{n,k} := \text{Span}(\mathcal{B}_k), \quad B_k := \text{Proj}_{\text{Skew}_{n,k}} : \text{Skew}_n \rightarrow \text{Skew}_{n,k},$$

where Proj is the usual Euclidean projection. Then, given $Y \in O_n$ and $k \in [m]$, define projections $P_k^Y : T_Y O_n \rightarrow T_Y O_n$ as the projection onto the subspace,

$$S_k^Y := \{YA : A \in \text{Skew}_{n,k}\} \subset T_Y O_n,$$

which is computed as

$$YA \mapsto P_k^Y(YA) := YB_k(A).$$

We propose the following subspace selection rule \mathcal{P} for the orthogonal manifold O_n : given Y^0, \dots, Y^{k-1} obtained at step k of the inner loop of Algorithm 1, define

$$\mathcal{P} \left(\{Y^i\}_{i=0}^{k-1}, k \right) := P_k^{Y^{k-1}}. \quad (21)$$

Our main result in this section is to show that \mathcal{P} satisfies Assumption 4.2.

Theorem 6.1. *Fix any $\beta \in (0, 1)$. The projection selection rule $\mathcal{P} := \{P_k\}_{k \in [m]}$ defined in (21) is $\left(\max_{k \in [m]} |\mathcal{B}_k| \sqrt{1 - \beta^2}, 2 \log(1 + \sqrt{1 - \beta}) \right)$ -gap ensuring. Thus, when $\beta > \sqrt{1 - \frac{1}{m \max_{k \in [m]} |\mathcal{B}_k|}}$, we have*

$$\|Y^0 A\|_{Y^0} \leq \frac{1}{1 - \sqrt{m(1 - \beta^2)} \max_{k \in [m]} |\mathcal{B}_k|} \|Y^0 A\|_{Y^0, \mathcal{P}},$$

provided $d(Y^0, Y^k) \leq 2 \log(1 + \sqrt{1 - \beta})$ for each $k \in [m]$.

Proof of Theorem 6.1. The second part follows from the first part and Proposition 4.2, thus it suffices to prove the first part.

To do this, we assume that $d(Y^0, Y^i) \leq 2 \log(1 + \sqrt{1 - \beta})$ for all $i \in [k-1]$, and we provide an orthogonal subspace decomposition $\{P'_k\}_{k \in [m]}$ of $T_{Y^0} O_n$ such that $\|P_k - P'_k\| \leq \max_{k \in [m]} |\mathcal{B}_k| \sqrt{1 - \beta^2}$, where $\|P_k - P'_k\|$ is the operator norm. A natural choice is to set

$$P'_k := P_k^{Y^0}, \quad k \in [m].$$

Let $C_k \in \text{Skew}_n$ be matrices such that $\text{Exp}_{Y^{k-1}}(Y^{k-1}C_k) = Y^0$, and $\Gamma_{Y^{k-1}}^{Y^0}$ be the parallel transport along the geodesic $Y^{k-1} \text{Exp}_m(tC_k)$. According to Lemma A.2, P_k is the projection onto the subspace

$$S_k := \left\{ \Gamma_{Y^{k-1}}^{Y^0}(Y^{k-1}A) : A \in \text{Skew}_{n,k} \right\} = \left\{ Y^0 \text{Exp}_m(C_k/2)^\top A \text{Exp}_m(C_k/2) : A \in \text{Skew}_{n,k} \right\}.$$

We compare this with P'_k , which is the projection onto the subspace $S'_k := \{Y^0 A : A \in \text{Skew}_{n,k}\}$. The result now follows from Proposition A.3, provided we can show that for any $A \in \mathcal{B}_k$,

$$\text{Tr} \left(A^\top \text{Exp}_m(C_k/2)^\top A \text{Exp}_m(C_k/2) \right) \geq \beta.$$

For convenience, denote $C = C_k$. We have

$$\begin{aligned} \text{Tr} \left(A^\top \text{Exp}_m(C/2)^\top A \text{Exp}_m(C/2) \right) &= \text{Tr} \left((\text{Exp}_m(C/2)A)^\top A \text{Exp}_m(C/2) \right) \\ &= \text{Tr} \left(\left(\left(I + \sum_{i=1}^{\infty} \frac{C^i}{2^i i!} \right) A \right)^\top A \left(I + \sum_{i=1}^{\infty} \frac{C^i}{2^i i!} \right) \right) \\ &= \text{Tr}(A^\top A) + \sum_{i,j=1}^{\infty} \frac{1}{2^{i+j} i! j!} \text{Tr} \left((C^i A)^\top A C^j \right). \end{aligned}$$

Notice that using Cauchy-Schwarz and the fact that $\|XX'\|_F \leq \|X\|_F \|X'\|_F$, we have

$$-\|A\|_F^2 \|C\|_F^{i+j} \leq -\|C^i A\|_F \|A C^j\|_F \leq \text{Tr} \left((C^i A)^\top A C^j \right) \leq \|C^i A\|_F \|A C^j\|_F \leq \|A\|_F^2 \|C\|_F^{i+j},$$

which implies

$$\begin{aligned} -\|A\|_F^2 (\exp(\|C\|_F/2) - 1)^2 &= -\|A\|_F^2 \sum_{i,j=1}^{\infty} \frac{1}{2^{i+j} i! j!} \|C\|_F^{i+j} \leq \sum_{i,j=1}^{\infty} \frac{1}{2^{i+j} i! j!} \text{Tr} \left((C^i A)^\top A C^j \right) \\ &\leq \|A\|_F^2 \sum_{i,j=1}^{\infty} \frac{1}{2^{i+j} i! j!} \|C\|_F^{i+j} \\ &= \|A\|_F^2 (\exp(\|C\|_F/2) - 1)^2. \end{aligned}$$

Therefore

$$\begin{aligned} \text{Tr} \left(A^\top \text{Exp}_m(C/2)^\top A \text{Exp}_m(C/2) \right) &= \text{Tr}(A^\top A) + \sum_{i,j=1}^{\infty} \frac{1}{2^{i+j} i! j!} \text{Tr} \left((C^i A)^\top A C^j \right) \\ &= \|A\|_F^2 + \sum_{i,j=1}^{\infty} \frac{1}{2^{i+j} i! j!} \text{Tr} \left((C^i A)^\top A C^j \right) \\ &\geq \|A\|_F^2 - \|A\|_F^2 (\exp(\|C\|_F/2) - 1)^2 \\ &= \|A\|_F^2 \left(1 - (\exp(\|C\|_F/2) - 1)^2 \right). \end{aligned}$$

Note that $\|A\|_F = 1$ by definition, and since $d(Y^0, Y^{k-1}) \leq 2 \log(1 + \sqrt{1 - \beta})$, we can choose C such that $\|C\|_F \leq 2 \log(1 + \sqrt{1 - \beta})$, or equivalently, $1 - (\exp(\|C\|_F/2) - 1)^2 \geq \beta$, which gives the result. \square

Convergence of tangent subspace descent with \mathcal{P} from (21) now follows from Theorem 4.5.

6.2 Givens Rotations and Computational Complexity

We now show that, when the partition $\{\mathcal{B}_k\}_{k \in [m]}$ is chosen appropriately, tangent subspace descent updates take a very simple form. Given Y^0 , an inner loop of Algorithm 1 computes the sequence Y^1, \dots, Y^m recursively as

$$Y^k := \text{Exp}_{Y^{k-1}} \left(-\eta_{t,k} P_k^{Y^{k-1}} \nabla f(Y^{k-1}) \right), \quad k \in [m].$$

Let $A^k = -\eta_{t,k} B_k(\nabla f(Y^{k-1}))$. Then

$$Y^k = Y^{k-1} \text{Exp}_m(A^k).$$

We now describe the complexity of computing

$$\text{Exp}_m(A^k) = I_n + \sum_{p=1}^{\infty} \frac{1}{p!} (A^k)^p.$$

Define $H_{ij} := e_i e_j^\top - e_j e_i^\top$. Write each index set (i, j) defining \mathcal{B}_k as I_k , i.e., $\mathcal{B}_k = \{H_{ij}/\sqrt{2} : (i, j) \in I_k\}$. Now write

$$A^k = \sum_{(i,j) \in I_k} a_{ij}^k H_{ij}.$$

In order to efficiently compute $\text{Exp}_m(A^k)$, we make the following assumption on B_k .

Assumption 6.1. For each $k \in [m]$, if pairs $(i, j), (i', j') \in I_k$ are distinct, then i, i', j, j' are also distinct. Consequently, $H_{ij} H_{i'j'}$ is the zero matrix.

Assumption 6.1 ensures that $(A^k)^p = \sum_{(i,j) \in I_k} (a_{ij}^k)^p H_{ij}^p$, and H_{ij}^p takes an easily computable form, which results in the following expression for each (i, j) -submatrix of $\text{Exp}_m(A^k)$ where $(i, j) \in I_k$:

$$\begin{bmatrix} \text{Exp}_m(A^k)_{ii} & \text{Exp}_m(A^k)_{ij} \\ \text{Exp}_m(A^k)_{ji} & \text{Exp}_m(A^k)_{jj} \end{bmatrix} = \begin{bmatrix} \cos(a_{ij}^k) & \sin(a_{ij}^k) \\ -\sin(a_{ij}^k) & \cos(a_{ij}^k) \end{bmatrix}. \quad (22)$$

This submatrix is known as a *Givens rotation matrix*. All other off-diagonal entries of $\text{Exp}_m(A^k)$ are 0, and all other diagonal entries are 1.

Proposition 6.2. Under Assumption 6.1, the complexity of one inner loop of Algorithm 1 is $O(n^2(n-1)/2)$.

Proof of Proposition 6.2. Under Assumption 6.1, when $A^k \in \text{Span}(\mathcal{B}_k)$, computation of $\text{Exp}_m(A^k)$ takes time $O(|\mathcal{B}_k|)$. Updating $Y^k = Y^{k-1} \text{Exp}_m(A^k)$ takes time $O(n|\mathcal{B}_k|)$, since for each pair $(i, j) \in I_k$, we must perform two vector additions with two columns of Y^{k-1} . Overall, the complexity of one inner loop of Algorithm 1 is $O\left(n \sum_{k \in [m]} |\mathcal{B}_k|\right) = O(n^2(n-1)/2)$. \square

Note that while the $O(n^3)$ complexity in Proposition 6.2 is of the same order as computing a matrix exponential for a general matrix $\text{Exp}_m(A)$, which is required for performing Riemannian gradient descent in O_n , the advantage of tangent subspace descent is the fact that progress can be made every $O(n|\mathcal{B}_k|)$ flops, whereas Riemannian gradient descent must wait until $O(n^3)$ flops have elapsed to make progress. This is particularly advantageous when n is large and the \mathcal{B}_k are chosen to be small.

6.3 Necessity of Distance-Dependent Bounds

We now show by a counter-example in O_n why Assumption 4.1, the (γ, r) gap-ensuring condition, and the bound in Proposition 4.2 in general requires the dependence on a distance parameter r , as promised in Remark 4.1.

Suppose we are given iterates Y^0, Y^1, \dots, Y^m generated from an inner loop of Algorithm 1, and that we can write $Y^0 = \text{Exp}_{Y^{k-1}}(Y^{k-1}C_k) = Y^{k-1} \text{Exp}_m(C_k)$ for some $C_k \in \text{Skew}_n$. The proof of Theorem 6.1 (namely, the application of Lemma A.2) showed that the (non-orthogonal) subspace decomposition $\{P_k\}_{k \in [m]}$ of $T_{Y^0}O_n$ is given by

$$P_k = \text{Proj}_{S_k}, \quad S_k := \left\{ Y^0 \text{Exp}_m(C_k/2)^\top A \text{Exp}_m(C_k/2) : A \in \text{Skew}_{n,k} \right\}.$$

The norm under consideration is

$$\|Y^0 A\|_{Y^0, \mathcal{P}} := \sqrt{\sum_{k \in [m]} \|P_k Y^0 A\|_{Y^0}^2}.$$

The proof of Theorem 6.1 showed that when each C_k is sufficiently small, we can show that Assumption 4.2 is satisfied, and hence bound $\sup_{A \in \text{Skew}_n} \|Y^0 A\|_{Y^0} / \|Y^0 A\|_{Y^0, \mathcal{P}}$.

Our goal in this section is to show that when there is no bound on the skew-symmetric matrices C_k , $\sup_{A \in \text{Skew}_n} \|Y^0 A\|_{Y^0} / \|Y^0 A\|_{Y^0, \mathcal{P}}$ can be arbitrarily large. For convenience (although this is not necessary), we set each \mathcal{B}_k to be a singleton member of \mathcal{B} , so $m = n(n-1)/2$. We show that there exists $k > 1$ and some adversarial choice of C_k that can make $S_k = S_1$, hence $\text{Span}\left(\bigcup_{k \in [m]} S_k\right) \subsetneq T_{Y^0}O_n$, hence $\sup_{A \in \text{Skew}_n} \|Y^0 A\|_{Y^0} / \|Y^0 A\|_{Y^0, \mathcal{P}} = \infty$.

Note that for $k = 1$, $Y^{k-1} = Y^0$, so we can take $C_1 = 0$, hence $S_1 = \{\alpha Y^0 H_{ij} : \alpha \in \mathbb{R}\}$ for some pair (i, j) . Consider some $k \neq 1$, where

$$S_k = \left\{ \alpha Y^0 \text{Exp}_m(C_k/2)^\top H_{i'j'} \text{Exp}_m(C_k/2) : \alpha \in \mathbb{R} \right\}$$

for some other pair (i', j') such that i, j, i', j' are all unique. We show that we can choose C_k in such a way that

$$\text{Exp}_m(C_k/2)^\top H_{i'j'} \text{Exp}_m(C_k/2) = H_{ij}.$$

This can be done by setting

$$C_k = \pi(H_{ii'} + H_{jj'}).$$

According to the arguments in Section 6.2, particularly (22), $\text{Exp}_m(C_k/2)$ is the identity matrix except for the following entries:

$$\begin{aligned} \text{Exp}_m(C_k/2)_{ii} &= \text{Exp}_m(C_k/2)_{jj} = \text{Exp}_m(C_k/2)_{i'i'} = \text{Exp}_m(C_k/2)_{j'j'} = 0 \\ \text{Exp}_m(C_k/2)_{ii'} &= \text{Exp}_m(C_k/2)_{jj'} = 1 \\ \text{Exp}_m(C_k/2)_{i'i} &= \text{Exp}_m(C_k/2)_{j'j} = -1. \end{aligned}$$

It is then easily checked that $\text{Exp}_m(C_k/2)^\top H_{i'j'} \text{Exp}_m(C_k/2) = H_{ij}$.

6.4 Numerical Study: Linear Optimization over O_n

In this section, we conduct a numerical study comparing Riemannian gradient descent to tangent subspace descent for linear optimization over O_n :

$$\min_{Y \in O_n} \left\{ f(Y) := \text{Tr}(D^\top Y) \right\}.$$

This is equivalent to the well-known orthogonal Procrustes problem of minimizing a quadratic objective $\|AY - B\|_F^2$ over O_n , which has numerous applications in science and engineering. Since $Y \in O_n$, it is easily seen that this is equivalent to maximizing $\text{Tr}((A^\top B)^\top Y)$.

For Riemannian gradient descent, we use a backtracking line search rule to determine the step size at each iteration. For tangent subspace descent, we show that when \mathcal{B}_k are chosen to consist of a single pair (i, j) , an exact line search over $\eta > 0$ can be computed in closed form to optimize the objective of $Y^{k-1} \text{Expm}(-\eta A^k)$. We use this rule to perform our updates in Algorithm 1.

Proposition 6.3. *Fix some $Y \in O_n$ and $H_{ij} \in \mathcal{B}$. Let $G = D^\top Y$. Then*

$$\min_{\eta \geq 0} \text{Tr} \left(D^\top Y \text{Expm}(-\eta H_{ij}) \right) = \sum_{i' \neq i, j} g_{i'i'} - \sqrt{(g_{ii} + g_{jj})^2 + (g_{ij} - g_{ji})^2}$$

and the (i, j) -submatrix of $\text{Expm}(-\eta^* H_{ij})$ is

$$\begin{bmatrix} \text{Expm}(-\eta^* H_{ij})_{ii} & \text{Expm}(-\eta^* H_{ij})_{ij} \\ \text{Expm}(-\eta^* H_{ij})_{ji} & \text{Expm}(-\eta^* H_{ij})_{jj} \end{bmatrix} = -\frac{1}{\sqrt{(g_{ii} + g_{jj})^2 + (g_{ij} - g_{ji})^2}} \begin{bmatrix} g_{ii} + g_{jj} & g_{ij} - g_{ji} \\ -(g_{ij} - g_{ji}) & g_{ii} + g_{jj} \end{bmatrix}.$$

Proof of Proposition 6.3. Observe that from (22),

$$\text{Tr} \left(D^\top Y \text{Expm}(-\eta H_{ij}) \right) = \sum_{i' \neq i, j} g_{i'i'} + (g_{ii} + g_{jj}) \cos(-\eta) + (g_{ij} - g_{ji}) \sin(-\eta).$$

Since $(\cos(-\eta), \sin(-\eta))$ belong to the unit ball in \mathbb{R}^2 , the result follows from linear optimization over the unit ball. \square

Test instances were generated synthetically, inspired by the orthogonal Procrustes problem. We generated $A \in \mathbb{R}^{n \times n}$ with i.i.d. $N(0, 4)$ entries, a random orthogonal matrix $X \in O_n$, and $B = AX + \epsilon$ where $\epsilon \in \mathbb{R}^{n \times n}$ with i.i.d. $N(0, 1)$ entries, then set $D = -A^\top B$. This is equivalent to minimizing $\|AY - B\|_F^2$. We tested $n = 50, 100, 150, 200$. For each n , we generated 10 instances, and ran both algorithms. Experiments were run on a 2.5GHz personal machine with 16GB memory. We plotted the results in Figure 1. Generally, tangent subspace descent takes much fewer cycles to close the gap than Riemannian gradient descent.

7 Conclusion

In this paper we have presented and analysed a generalization of coordinate descent to manifold domains called tangent subspace descent. We found that the crux of ensuring convergence is a prudent choice of subspace selection rule \mathcal{P} . For deterministic rules, are precisely those that satisfy the (γ, r) -gap ensuring condition while appropriate randomized rules satisfy the \mathcal{C} -randomized norm condition. We proposed a new deterministic rule for optimization over the manifold of square orthogonal matrices, and applied TSD with this rule to the Orthogonal Procrustes Problem. Our numerical experiments indicate the TSD's promise in application. We also provided a new randomized subspace selection rule for optimization over a general Stiefel manifold.

There are several interesting directions that can stem from this work, we give two here. First, it is known that accelerated rates for coordinate descent can be achieved with a simple modification; it is unknown whether something similar can be done by appealing to recent work [24] on extending Nesterov's accelerated gradient algorithm to manifolds. Second, it is well-known that replacing the exponential map with a retraction operator can result in computational savings; it would be interesting to see how retractions can be incorporated into the TSD framework.

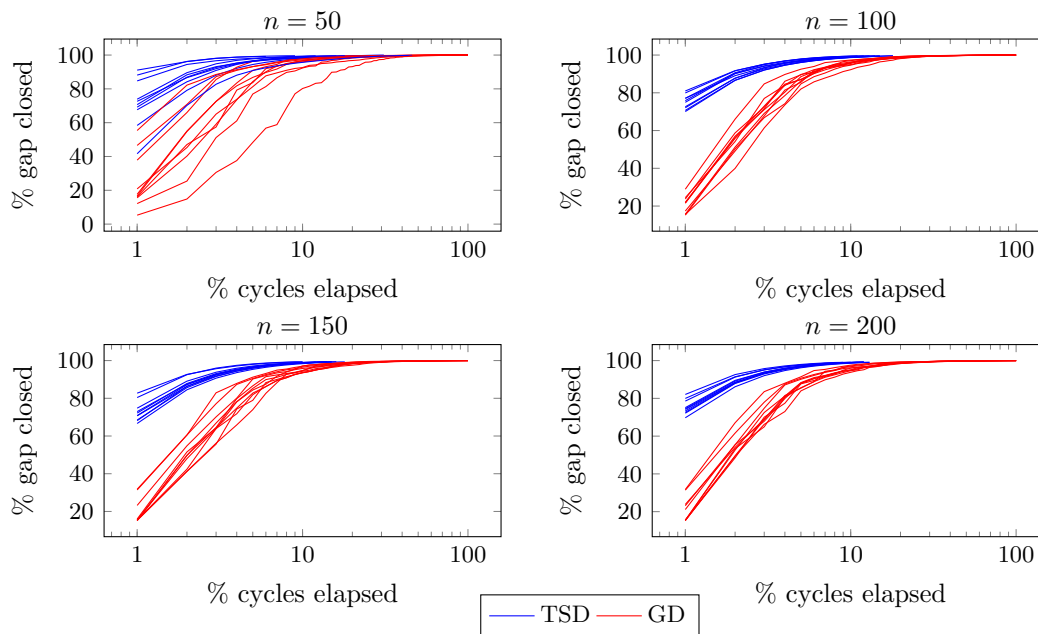


Figure 1: Results for TSD (blue) vs GD (red). The horizontal axis shows the cycles elapsed as a percentage of the largest number of cycles for that instance, while the vertical axis shows the gap closed, as a percentage of the best objective value found across both algorithms.

Acknowledgements

This research is, in part, funded by Award N660011824020 from the DARPA Lagrange Program and NSF Award 1740707.

References

- [1] P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, Princeton, NJ, USA, 2007. ISBN 0691132984, 9780691132983.
- [2] A. Beck and L. Tetruashvili. On the convergence of block coordinate descent type methods. *SIAM Journal on Optimization*, 23(4):2037–2060, 2013. doi: 10.1137/120887679.
- [3] S. Chen, S. Ma, A. M.-C. So, and T. Zhang. Proximal gradient method for nonsmooth optimization over the stiefel manifold. Technical report, November 2018. <https://arxiv.org/abs/1811.00980>.
- [4] F. Deutsch. The angle between subspaces of a hilbert space. In *Approximation theory, wavelets and applications*, pages 107–130. Springer, 1995.
- [5] A. Edelman, T. A. Arias, and S. T. Smith. The geometry of algorithms with orthogonality constraints. *SIAM Journal on Matrix Analysis and Applications*, 20(2):303–353, 1998. doi: 10.1137/S0895479895290954.
- [6] O. P. Ferreira and P. R. Oliveira. Subgradient algorithm on riemannian manifolds. *Journal of Optimization Theory and Applications*, 97(1):93–104, 1998. doi: 10.1023/A:1022675100677.

- [7] O. P. Ferreira, M. S. Louzeiro, and L. F. Prudente. Gradient method for optimization on riemannian manifolds with lower bounded curvature. *SIAM Journal on Optimization*, 29(4): 2517–2541, 2019. doi: 10.1137/18M1180633.
- [8] R. Frongillo and M. D. Reid. Convergence analysis of prediction markets via randomized subspace descent. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 3034–3042. Curran Associates, Inc., 2015.
- [9] B. Gao, X. Liu, X. Chen, and Y.-x. Yuan. A new first-order algorithmic framework for optimization problems with orthogonality constraints. *SIAM Journal on Optimization*, 28(1): 302–332, 2018. doi: 10.1137/16M1098759.
- [10] G. H. Golub and C. F. Van Loan. *Matrix computations*. The Johns Hopkins University Press, fourth edition edition, 2013. ISBN 9781421407944.
- [11] W. Huang and K. Wei. Riemannian proximal gradient methods. Technical report, September 2019. <https://arxiv.org/abs/1909.06065>.
- [12] J. M. Lee. *Introduction to Smooth Manifolds*, volume 218 of *Graduate Texts in Mathematics*. Springer-Verlag New York, 2nd edition, 2012. ISBN 978-1-4419-9982-5.
- [13] J. M. Lee. *Introduction to Riemannian Manifolds*, volume 176 of *Graduate Texts in Mathematics*. Springer International Publishing, 2nd edition, 2018. ISBN 978-3-319-91755-9.
- [14] D. Leventhal and A. S. Lewis. Randomized methods for linear constraints: convergence rates and conditioning. *Mathematics of Operations Research*, 35(3):641–654, 2010.
- [15] A. S. Lewis and J. Malick. Alternating projections on manifolds. *Mathematics of Operations Research*, 33(1):216–234, 2008.
- [16] C. Liu and N. Boumal. Simple algorithms for optimization on riemannian manifolds with constraints. *Applied Mathematics & Optimization*, Mar 2019. ISSN 1432-0606. doi: 10.1007/s00245-019-09564-3.
- [17] Y. Liu, F. Shang, J. Cheng, H. Cheng, and L. Jiao. Accelerated first-order methods for geodesically convex optimization on riemannian manifolds. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 4868–4877. Curran Associates, Inc., 2017.
- [18] Y. Nesterov. Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM Journal on Optimization*, 22(2):341–362, 2012. doi: 10.1137/100802001.
- [19] P. Richtárik and M. Takáč. Iteration complexity of randomized block-coordinate descent methods for minimizing a composite function. *Math. Program.*, 144(1-2):1–38, Apr. 2014. ISSN 0025-5610. doi: 10.1007/s10107-012-0614-z.
- [20] H. Sato, H. Kasai, and B. Mishra. Riemannian stochastic variance reduced gradient algorithm with retraction and vector transport. *SIAM Journal on Optimization*, 29(2):1444–1472, 2019. doi: 10.1137/17M1116787.

- [21] U. Shalit and G. Chechik. Coordinate-descent for learning orthogonal matrices through givens rotations. In E. P. Xing and T. Jebara, editors, *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 548–556, Beijing, China, 22–24 Jun 2014. PMLR.
- [22] N. K. Vishnoi. Geodesic convex optimization: Differentiation on manifolds, geodesics, and convexity. Technical report, June 2018. <https://arxiv.org/abs/1806.06373>.
- [23] Y. Xu. Alternating proximal gradient method for sparse nonnegative tucker decomposition. *Mathematical Programming Computation*, 7(1):39–70, Mar 2015. doi: 10.1007/s12532-014-0074-y.
- [24] H. Zhang and S. Sra. An estimate sequence for geodesically convex optimization. In S. Bubeck, V. Perchet, and P. Rigollet, editors, *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pages 1703–1723. PMLR, 06–09 Jul 2018.
- [25] H. Zhang, S. J. Reddi, and S. Sra. Riemannian svrg: Fast stochastic optimization on riemannian manifolds. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 4592–4600. Curran Associates, Inc., 2016.
- [26] J. Zhang, S. Ma, and S. Zhang. Primal-dual optimization algorithms over riemannian manifolds: an iteration complexity analysis. *Mathematical Programming*, Aug 2019. ISSN 1436-4646. doi: 10.1007/s10107-019-01418-8.

A Auxiliary Results

Proposition A.1. Let $\{a_t, b_t\}_{t \in \mathbb{N}}$ be integrable sequences of random variables adapted to a filtration $\mathcal{F}_1 \subseteq \mathcal{F}_2 \subseteq \dots \subseteq \mathcal{F}$ of some probability space $(\Omega, \mathcal{F}, \mathbb{P})$, which almost surely satisfy for all $t \in \mathbb{N}$

$$\begin{aligned} 0 &\leq b_t \leq a_t - \mathbb{E}[a_{t+1} \mid \mathcal{F}_t] \\ a_t &\geq a^* \in \mathbb{R}. \end{aligned}$$

Then $b_t \rightarrow 0$ almost surely.

Proof of Proposition A.1. First, observe that by taking expectation, and using the law of total expectation, for each $t \in \mathbb{N}$, we have $\mathbb{E}[b_t] \leq \mathbb{E}[a_t] - \mathbb{E}[a_{t+1}]$. Since the right hand side telescopes, we have

$$\sum_{t \in T} \mathbb{E}[b_t] \leq \mathbb{E}[a_1] - \mathbb{E}[a_{T+1}] \leq \mathbb{E}[a_1] - a^*,$$

where the second inequality follows since $a_{t+1} \geq a^*$ almost surely.

Now fix some $\epsilon > 0$. By Markov's inequality, $\mathbb{E}[b_t]/\epsilon \geq \mathbb{P}[b_t \geq \epsilon]$, therefore substituting this into the previous inequality and taking $T \rightarrow \infty$ gives

$$\sum_{t=1}^{\infty} \mathbb{P}[b_t \geq \epsilon] \leq \frac{\mathbb{E}[a_1] - a^*}{\epsilon} < \infty.$$

By the Borel-Cantelli lemma, we know that $\mathbb{P}[b_t \geq \epsilon \text{ infinitely often}] = 0$.

In fact, this holds for any $\epsilon > 0$. Let $\{\epsilon_n\}_{n \in \mathbb{N}}$ be a sequence such that $\epsilon_n \rightarrow 0$. Then by the union bound

$$\mathbb{P}[\exists n \text{ s.t. } b_t \geq \epsilon_n \text{ infinitely often}] \leq \sum_{n=1}^{\infty} \mathbb{P}[b_t \geq \epsilon_n \text{ infinitely often}] = 0.$$

Finally, we observe that $\mathbb{P}[b_t \rightarrow 0] = 1 - \mathbb{P}[\exists n \text{ s.t. } b_t \geq \epsilon_n \text{ infinitely often}] = 1$. \square

Lemma A.2. Let V, W be two finite-dimensional vector spaces of equal dimension. Let $U : V \rightarrow W$ be an invertible unitary map (so $UU^{-1} = \text{id}_W$, $U^{-1}U = \text{id}_V$). Let $P : V \rightarrow V$ be a projection (i.e., $P = P^2$) onto a subspace S of V . Then the projection onto the subspace $US = \{Uv : v \in S\} \subset W$ is UPU^{-1} .

Proof of Lemma A.2. Clearly, $(UPU^{-1})(UPU^{-1}) = UP^2U^{-1} = UPU^{-1}$ so it is a projection operator. It remains to check that the image is US . Clearly, we have $\text{Im}(UPU^{-1}) \subseteq US$. Now take some $v \in S$. We want to show that there exists some $w \in W$ such that $UPU^{-1}w = Uv$. We can easily take $w = Uv$, which gives $UPU^{-1}w = UPU^{-1}(Uv) = UPv = Uv$. Thus, $\text{Im}(UPU^{-1}) \supseteq US$. \square

Proposition A.3. Let $U, U' \subseteq \mathbb{R}^n$ be subspaces of dimension d in some Hilbert space with projection operators $P_U, P_{U'}$. If u_1, \dots, u_d and u'_1, \dots, u'_d are orthonormal bases for U and U' respectively satisfying $\min_{i=1, \dots, d} \langle u_i, u'_i \rangle \geq \beta$ then $\|P_U - P_{U'}\| \leq d\sqrt{1 - \beta^2}$.

Proof of Proposition A.3. Observe by orthonormality of the bases that

$$P_U(\cdot) = \sum_{i=1}^d u_i \langle u_i, \cdot \rangle, \quad P_{U'}(\cdot) = \sum_{i=1}^d u'_i \langle u'_i, \cdot \rangle,$$

so

$$\|P_U - P_{U'}\|^2 = \left\| \sum_{i=1}^d (u_i \langle u_i, \cdot \rangle - u'_i \langle u'_i, \cdot \rangle) \right\|^2 \leq \sum_{i=1}^d \|u_i \langle u_i, \cdot \rangle - u'_i \langle u'_i, \cdot \rangle\|^2.$$

It suffices to bound $\|u_i \langle u_i, \cdot \rangle - u'_i \langle u'_i, \cdot \rangle\|$ for $1 \leq i \leq d$. To simplify notation, let $u = u_i$ and $v = u'_i$. We can write $v = \alpha u + \sqrt{1 - \alpha^2} w$ where $\alpha = \langle u, v \rangle$ and $w = \frac{v - \alpha u}{\sqrt{1 - \alpha^2}}$ is a normal vector orthogonal to u . Then we can write

$$u \langle u, \cdot \rangle - v \langle v, \cdot \rangle = (1 - \alpha^2) u \langle u, \cdot \rangle - \alpha \sqrt{1 - \alpha^2} (u \langle w, \cdot \rangle + w \langle u, \cdot \rangle) - (1 - \alpha^2) w \langle w, \cdot \rangle.$$

Thus, $u \langle u, \cdot \rangle - v \langle v, \cdot \rangle$ essentially acts on the 2-dimensional subspace $\text{Span}(\{u, w\})$ of the Hilbert space. A basis for this subspace is $\{u, w\}$, and the matrix representation of $u \langle u, \cdot \rangle - v \langle v, \cdot \rangle$ on this subspace is

$$\begin{bmatrix} 1 - \alpha^2 & -\alpha \sqrt{1 - \alpha^2} \\ -\alpha \sqrt{1 - \alpha^2} & -(1 - \alpha^2) \end{bmatrix}.$$

Hence, the eigenvalues are readily computed to be $\pm \sqrt{1 - \alpha^2}$, which is bounded above by $\sqrt{1 - \beta^2}$. Thus, $\|u \langle u, \cdot \rangle - v \langle v, \cdot \rangle\| \leq \sqrt{1 - \beta^2}$, and

$$\|P_U - P_{U'}\|^2 \leq \sum_{i=1}^d \|u_i \langle u_i, \cdot \rangle - u'_i \langle u'_i, \cdot \rangle\|^2 \leq d \sqrt{1 - \beta^2}.$$

□