

The leapfrog algorithm as nonlinear Gauss–Seidel

Marco Sutti & Bart Vandereycken

October 23, 2020

Abstract

Several applications in optimization, image and signal processing deal with data that belong to the Stiefel manifold $\text{St}(n, p)$, that is, the set of $n \times p$ matrices with orthonormal columns. Some applications, like the Karcher mean, require evaluating the geodesic distance between two arbitrary points on $\text{St}(n, p)$. This can be done by explicitly constructing the geodesic connecting these two points. An existing method for finding geodesics is the leapfrog algorithm of J. L. Noakes. This algorithm is related to the Gauss–Seidel method, a classical iterative method for solving a linear system of equations that can be extended to nonlinear systems. We propose a convergence proof of leapfrog as a nonlinear Gauss–Seidel method. Our discussion is limited to the case of the Stiefel manifold, however it may be generalized to other embedded submanifolds. We discuss other aspects of leapfrog and present some numerical experiments.

Keywords Riemannian manifolds geodesics Stiefel manifold nonlinear Gauss–Seidel

1 Introduction

The object of study in this paper is the compact Stiefel manifold, i.e., the set of orthonormal n -by- p matrices

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

Here, we are concerned with computing the minimal distance between two points on the Stiefel manifold. To define a distance on a manifold, one has to introduce the concept of geodesics. A geodesic $\gamma: [0, t] \rightarrow \mathcal{M}$ is a curve with zero acceleration, which generalizes the notion of straight lines in Euclidean space to a Riemannian manifold [1]. Geodesics allow us to introduce the *Riemannian exponential* $\text{Exp}_x: T_x \mathcal{M} \rightarrow \mathcal{M}$ that maps a tangent vector $\xi = \dot{\gamma}(0) \in T_x \mathcal{M}$ to the geodesic endpoint $\gamma(1) = y$ such that $\text{Exp}_x(\xi) = y$. The Riemannian exponential is a local diffeomorphism, i.e., it is locally invertible and its inverse is called the *Riemannian logarithm* of y at x satisfying $\text{Log}_x(y) = \xi$. The *injectivity radius* at a point p of a Riemannian manifold \mathcal{M} is the largest radius for which the exponential map Exp_x is a diffeomorphism from the tangent space to the manifold. The global injectivity radius of a manifold is the infimum of all the injectivity radii over all points of the manifold. Given two points x and y on a manifold \mathcal{M} , if the Riemannian distance $d(x, y)$ is smaller than $\text{inj}(\mathcal{M})$, then there exists a unique geodesic from x to y . For the Stiefel manifold, the injectivity radius is at least 0.89π ; see [19, Eq. (5.13)].

The distance on the Stiefel manifold is involved in numerous fields of applications, among which, computer vision [5, 22, 24, 25], statistics [20], reduced-order models [2, 4]. Given two points on the Stiefel manifold, our goal is to compute the length of the minimal geodesic connecting them. For some manifolds, there are explicit formulas available for computing the geodesic distance. For the Stiefel manifold there is no closed-form solution known. In general, the problem of finding the geodesic distance given two points on a Riemannian manifold is related to the Riemannian logarithm. The problem of computing the Riemannian logarithm on the Stiefel manifold has already been tackled by several authors, who proposed some numerical algorithms. Rentmeesters [19] and Zimmermann [26, 27] proposed

a similar algorithm which is only locally convergent and depends upon the definition of the (standard) matrix logarithm function.

Another method for finding geodesics is the leapfrog algorithm introduced by J. L. Noakes [17]. This method has global convergence properties, but it slows down when the solution is approached [14, p. 2796]. Moreover, Noakes realized that his leapfrog algorithm was related to the Gauss–Seidel method [17, p. 39]. The link between leapfrog and nonlinear Gauss–Seidel was not further investigated, since there is no trace of this idea being developed in the other related papers [13, 14]. In this paper, we will prove convergence of leapfrog as a nonlinear block Gauss–Seidel method. Even though our focus will be on $\text{St}(n, p)$, most of our discussion may be generalized to other embedded submanifolds.

A Riemannian metric has to be specified in order to turn $\text{St}(n, p)$ into a Riemannian manifold, and in general different choices are possible. In this paper, we consider the non-Euclidean *canonical metric* inherited by $\text{St}(n, p)$ from its definition as a quotient space of the orthogonal group [6, Eq. (2.39)]

$$g_c(\xi, \xi) = \text{trace}(\xi^\top (I - \frac{1}{2}YY^\top) \xi). \quad (1)$$

Tangent vectors to the Stiefel manifold may be expressed in the form

$$\xi = Y_0\Omega + Y_{0\perp}K, \quad \text{with } \Omega \in \mathcal{S}_{\text{skew}}(p), \quad K \in \mathbb{R}^{(n-p) \times p},$$

with $\mathcal{S}_{\text{skew}}(p)$ the vector space of p -by- p skew-symmetric matrices.

An explicit formula for a geodesic with initial acceleration the tangent vector ξ and base point Y_0 has been provided by Ross Lippert [6, Eq. (2.42)]

$$Y(t) = Q \exp\left(\begin{bmatrix} \Omega & -K^\top \\ K & O_{n-p} \end{bmatrix} t\right) \begin{bmatrix} I_p \\ O_{(n-p) \times p} \end{bmatrix}, \quad (2)$$

with $Q = [Y_0 \ Y_{0\perp}]$ and $Y_{0\perp}$ being any matrix whose range is $(\text{span}(Y_0))^\perp$.

Given two points Y_0, Y_1 on $\text{St}(n, p)$ that are sufficiently close to each other, finding the minimal geodesic distance between them is equivalent to finding the tangent vector $\xi^* \in T_{Y_0}\text{St}(n, p)$ with the shortest possible length such that $\text{Exp}_{Y_0}(\xi^*) = Y_1$. The solution to this problem is equivalent to the Riemannian logarithm of Y_1 with base point Y_0 , i.e., $\xi^* = \text{Log}_{Y_0}(Y_1)$. Given the endpoints Y_0 and Y_1 , we do not know what the matrices Ω and K in (2) are. So the problem becomes: Find the matrices Ω and K such that the explicit formula (2) gives the endpoint Y_1 .

2 Leapfrog algorithm

When $X, Y \in \mathcal{M}$ are sufficiently close, their connecting geodesic can be found by applying Newton's method to (2) such that $Y(1) = Y$ with $Y(0) = X$. This is more generally known as single shooting¹. However, when X and Y are far apart, it is well-known that single shooting will have difficulty finding the connecting geodesic. The main idea behind the leapfrog algorithm of Noakes [17] is to exploit the success of single shooting by subdividing the global problem into several local problems, where intermediate points $X_i \in \mathcal{M}$ are introduced between X and Y , for which the endpoint geodesic problem can be solved again by single shooting. The algorithm then iteratively updates a piecewise geodesic to obtain a globally smooth geodesic between X and Y . This idea is actually not new and goes back as early as 1963 by Milnor [16, III.§16]. It also resembles the better known multiple shooting method for boundary value problems but it is different.

¹In this context, there is no need to solve an ordinary differential equation as in a normal shooting method, because we have the solution (2). Hence, it is actually Newton's method, but we keep the shooting terminology because it is typical for boundary value problems.

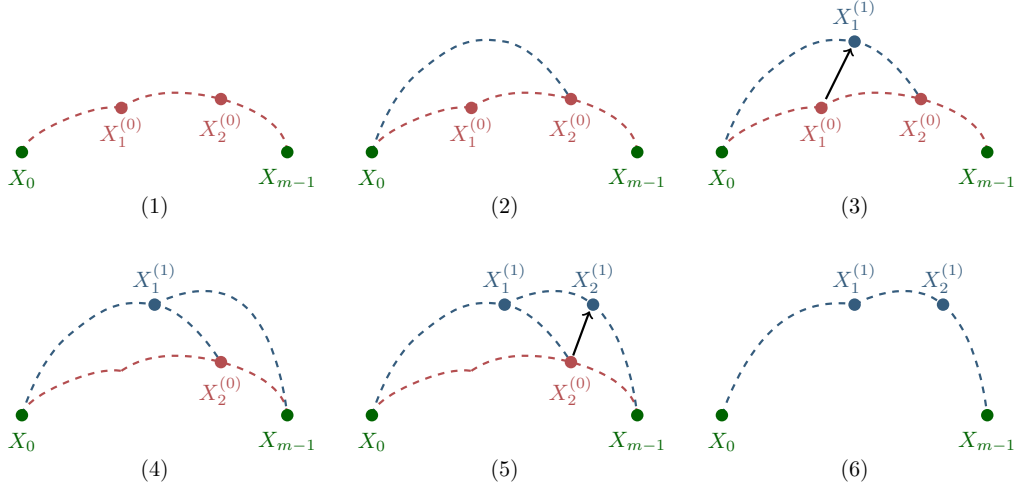


Figure 1: Illustration of one full iteration of the leapfrog scheme for some non-Euclidean metric (the lengths for the Euclidean metric clearly increase during iteration).

2.1 Formal description of the algorithm

In this section we describe the leapfrog algorithm by following the presentation in [17]. Let \mathcal{M} be a C^∞ path-connected Riemannian manifold. Consider a *piecewise* (or *broken*) geodesic ω_X joining X_0 to X_{m-1} , having $m - 1$ geodesic segments. Assuming X_i and X_{i+1} are sufficiently close to each other, ω_X is uniquely identified by the m -tuple $X = (X_0, X_1, \dots, X_{m-1}) \in \mathcal{M}^m$, where X_i are the junctions of the geodesic segments. The leapfrog algorithm now proceeds as follows: For $i = 1, \dots, m - 2$, each X_i is mapped onto the minimal geodesic joining X_{i-1} and X_{i+1} . This achieves the largest possible decrease in length while keeping other variables fixed. Though there are several choices to do this, leapfrog maps X_i onto the midpoint of the geodesic joining X_{i-1} and X_{i+1} . By iterating this procedure, the algorithm generates a sequence $\Omega = \{\omega_{X^{(k)}} : [0, 1] \rightarrow \mathcal{M} : k = 0, 1, \dots\}$ of broken geodesics whose lengths are decreasing. Figure 1 illustrates one iteration of the leapfrog algorithm. It is clear that leapfrog generates a sequence of broken geodesics $\omega_{X^{(k)}}$ that are defined from $X^{(k)}$. In addition, the length of $\omega_{X^{(k)}}$ is non-increasing in k since at each step two neighboring geodesics get replaced by one global geodesic connecting their endpoints.

2.2 Known results

Let \mathcal{Y} be the set of all tuples $X = (X_0, X_1, \dots, X_{m-1}) \in \mathcal{M}^m$ satisfying $d(X_{i-1}, X_i) \leq \delta$ for all $i = 1, 2, \dots, m - 2$. In [17, §2], δ is related to the notion of Lebesgue number of an open cover. Here, we can assume that δ is equal to $\frac{1}{2} \text{inj}(\mathcal{M})$, where inj is the injectivity radius (see Sect. 1). Let $\mathcal{F} : \mathcal{Y} \rightarrow \mathcal{Y}$ represent one full leapfrog iteration and let X^* be the limit of any convergent subsequence of $S = \{\mathcal{F}^k(X^{(0)}) : k \geq 1\}$ with $X^{(0)} \in \mathcal{Y}$. By compactness, [17] shows that at least one convergent subsequence of S exists and that the limit of this subsequence are points that lie on a global geodesic connecting the endpoints X_0 and X_{m-1} . The following result is stated in [14, Theorem 5.2].

Theorem 2.1. *S has a unique accumulation point.*

The theorem guarantees convergence of the iterates $X^{(k)} = \mathcal{F}(X^{(k-1)})$ with $X^{(0)} \in \mathcal{Y}$. From [17, Lemma 3.2] we also know that leapfrog will converge to a uniformly distributed m -tuple $X^* = (X_0, X_1^*, \dots, X_{m-2}^*, X_{m-1})$, i.e., $d(X_i^*, X_{i+1}^*)$ are all equal, for $i = 0, \dots, m - 2$. In other words, at convergence, the geodesic segments connecting the junction points will all have the same length.

An apparent drawback in the current theory is that it lacks a classical convergence proof as a fixed point iteration method, although leapfrog can be easily recognized as such. In the next section, we will provide the details of how to analyze leapfrog as a nonlinear block Gauss–Seidel method.

3 Convergence of leapfrog as nonlinear Gauss–Seidel

Let $\mathcal{M} = \text{St}(n, p)$ with the Riemannian distance function d . The starting point is to realize that leapfrog solves the optimization problem

$$\min_{X_1, \dots, X_{m-2} \in \text{St}(n, p)} F(X_1, \dots, X_{m-2}) \quad \text{with} \quad F(X_1, \dots, X_{m-2}) = \sum_{i=1}^{m-1} d^2(X_{i-1}, X_i),$$

by cyclically minimizing over each variable X_i for $i = 1, 2, \dots, m-2$. Specifically, at the k th iteration, leapfrog updates $X_i^{(k-1)}$ by the minimizer of the problem

$$\begin{aligned} & \min_{X_i \in \text{St}(n, p)} F(X_1^{(k)}, \dots, X_{i-1}^{(k)}, X_i, X_{i+1}^{(k-1)}, \dots, X_{m-2}^{(k-1)}) \\ &= \min_{X_i \in \text{St}(n, p)} d^2(X_{i-1}^{(k)}, X_i) + d^2(X_i, X_{i+1}^{(k-1)}) + \text{constant}. \end{aligned} \tag{3}$$

Since d is the Riemannian distance, this problem coincides with the definition of the Karcher mean² between the two points $X_{i-1}^{(k)}$ and $X_{i+1}^{(k-1)}$; see [12, Eq. (1.1)]. For the compact Stiefel manifold, a Karcher mean always exists, but it does not need to be unique [19, p. 37]. However, a sufficient condition for uniqueness is $d(X_{i-1}^{(k)}, X_{i+1}^{(k-1)}) < \text{inj}(\text{St}(n, p))$, where inj is the injectivity radius (see Sect. 1). This is true if all X_i are close enough (we will make this more precise later). In that case, the unique solution that solves (3) is the midpoint of the minimizing geodesic between $X_{i-1}^{(k)}$ and $X_{i+1}^{(k-1)}$. Leapfrog now proceeds to update the X_i in a Gauss–Seidel fashion where the most recent $X_{i-1}^{(k)}$ is used to update $X_i^{(k-1)}$. This kind of optimization scheme is known as *block coordinate descent method* of Gauss–Seidel type [18].

3.1 Nonlinear block Gauss–Seidel method

Let us first consider the case of Gauss–Seidel in \mathbb{R}^n . Let the variable $x \in \mathbb{R}^n$ be partitioned as $x = (x_1, x_2, \dots, x_m)$, where $x_i \in \mathbb{R}^{q_i}$ and $\sum_i q_i = n$, and group correspondingly the components of $\tilde{F}: D \subset \mathbb{R}^n \rightarrow \mathbb{R}^n$ into mappings $\tilde{F}_i: \mathbb{R}^n \rightarrow \mathbb{R}^{q_i}$, $i = 1, \dots, m$. The minimizers of the function $\tilde{F}(x)$ satisfy the first-order optimality condition $\nabla \tilde{F}(x) = 0$. Let us define $\mathcal{G}_i = \nabla \tilde{F}_i$, $i = 1, \dots, m$. If we interpret the linear Gauss–Seidel iteration in terms of obtaining $x_i^{(k)}$ as the solution of the i th equation of the system with the other $m-1$ block variables held fixed, then we may immediately consider the same prescription for nonlinear equations [18, p. 219]. Then solving

$$\mathcal{G}_i(x_1^{(k)}, \dots, x_{i-1}^{(k)}, y, x_{i+1}^{(k-1)}, \dots, x_m^{(k-1)}) = 0 \tag{4}$$

for y and defining $x_i^{(k)} = y$ describes a nonlinear block Gauss–Seidel process in which a complete iteration requires the solution of m nonlinear systems of dimensions q_i , $i = 1, \dots, m$; see [18, p. 225]. The convergence theory in [18] applies only to functions whose domain of definition is Euclidean space \mathbb{R}^n . It cannot be applied to functions which are defined on manifolds, such as the Riemannian distance d that is only defined on a subset of \mathbb{R}^n , namely, the embedded submanifold. For this reason, in the next section we will introduce a smooth extension of the Riemannian distance function that can also be evaluated for points that do not belong to the manifold.

²Also known as Riemannian center of gravity or Riemannian center of mass; see [15, Theorem 9.1, p. 109]. A numerical experiment involving the Karcher mean on the Stiefel manifold is discussed in Sect. 5.




Figure 2: The extended distance function.

3.2 Extended objective function

As we have seen above, leapfrog solves in an alternating way the problem

$$\min_{X_1, \dots, X_{m-2} \in \text{St}(n, p)} F(X_1, \dots, X_{m-2}) = \sum_{i=1}^{m-1} d^2(X_{i-1}, X_i),$$

where X_0 and X_{m-1} are the fixed endpoints. This objective function F is only defined on the manifold $\text{St}(n, p)$. In this section, we will identify an *extended objective function* \tilde{F} that is defined on $\mathbb{R}^{n \times p}$ for which the standard nonlinear block Gauss–Seidel method produces the same iterates as the leapfrog algorithm. The key result of this section is stated in Prop. 3.1. This will allow us to analyze the convergence of leapfrog using standard results for nonlinear Gauss–Seidel.

We claim the extended cost function can be chosen as

$$\min_{X_1, \dots, X_{m-2} \in \mathbb{R}^{n \times p}} \tilde{F}(X_1, \dots, X_{m-2}) = \sum_{i=1}^{m-1} \tilde{d}^2(X_{i-1}, X_i),$$

with *extended distance function*

$$\tilde{d}^2(\tilde{X}, \tilde{Y}) = \begin{cases} d^2(\text{P}_{\text{St}} \tilde{X}, \text{P}_{\text{St}} \tilde{Y}) + \|\tilde{X} - \text{P}_{\text{St}} \tilde{X}\|_F^2 + \|\tilde{Y} - \text{P}_{\text{St}} \tilde{Y}\|_F^2, & \text{if } \sigma_p(\tilde{X}) > 0 \text{ and } \sigma_p(\tilde{Y}) > 0; \\ +\infty, & \text{otherwise,} \end{cases} \quad (5)$$

where P_{St} denotes the orthogonal projector onto the Stiefel manifold.

The condition $\sigma_p(\tilde{X}) > 0$ is equivalent to the existence of a unique best approximation of \tilde{X} in $\text{St}(n, p)$. In other words, $\text{P}_{\text{St}} \tilde{X}$ is well defined. Concretely, we can define the projector $\text{P}_{\text{St}}: \mathbb{R}^{n \times p} \rightarrow \text{St}(n, p)$ by $\text{P}_{\text{St}}(Z) = Z(Z^\top Z)^{-1/2}$, that is, the orthogonal factor of the polar decomposition of Z [1, p. 58]. Figure 2 illustrates the extended distance function $\tilde{d}^2(\tilde{X}, \tilde{Y})$.

3.3 Leapfrog as nonlinear Gauss–Seidel

In order to show that nonlinear Gauss–Seidel applied to \tilde{F} is equivalent to leapfrog for F , we need a few lemmas. The first one addresses the problem of how close the points on $\text{St}(n, p)$ need to be so that their connecting geodesic is unique.

Lemma 3.1. *Let $X, Y \in \text{St}(n, p)$ such that $d(X, Y) \leq \delta_g$, with $\delta_g = 0.89\pi$. Then there exists a unique length-minimizing geodesic between X and Y . As a consequence, also the Karcher mean between X and Y exists and is uniquely defined.*

Proof. By definition of injectivity radius, if $d(X, Y) < \text{inj}(\text{St}(n, p))$, then there is only one length-minimizing geodesic between X and Y . From [19, Eq. (5.13)], we know that the injectivity radius is lower bounded by 0.89π . \square

Remark 3.1. We can compare the Riemannian and Euclidean distances between X and $Y \in \text{St}(n, p)$ asymptotically in the following way³. From the expansion of the canonical distance in (19), it is clear that

$$d(X, Y) \leq \|X - Y\|_F + O(\|X - Y\|_F^2) \quad \text{for} \quad \|X - Y\|_F \rightarrow 0.$$

By neglecting $O(\|X - Y\|_F^2)$, we thus have $d(X, Y) \lesssim \|X - Y\|_F$. In particular, $\|X - Y\|_F \leq \delta_g$ implies $d(X, Y) \lesssim \delta_g$.

Let $X_{i-1}, X_{i+1} \in \text{St}(n, p)$. Denote

$$F_i(Y) = d^2(X_{i-1}, Y) + d^2(Y, X_{i+1}), \quad \tilde{F}_i(\tilde{Y}) = \tilde{d}^2(X_{i-1}, \tilde{Y}) + \tilde{d}^2(\tilde{Y}, X_{i+1}),$$

where X_{i-1}, X_{i+1} are constant and hidden in the notation.

Lemma 3.2. With the notation from above assume that $d(X_{i-1}, X_{i+1}) \leq \delta_g$, then the i th substep of leapfrog produces the same solution Y^* as the minimization of \tilde{F}_i

$$\arg \min_{Y \in \text{St}(n, p)} F_i(Y) = \arg \min_{\tilde{Y} \in \mathbb{R}^{n \times p}} \tilde{F}_i(\tilde{Y}) = Y^*,$$

with Y^* the Karcher mean on $\text{St}(n, p)$ of X_{i-1} and X_{i+1} .

Proof. Since $d(X_{i-1}, X_{i+1}) \leq \delta_g$, Lemma 3.1 gives that the minimizer of F_i on $\text{St}(n, p)$ is unique and equals the Karcher mean Y^* . To show that it also equals the minimizer of \tilde{F}_i on $\mathbb{R}^{n \times p}$, take any $\tilde{Y} \in \mathbb{R}^{n \times p}$. If $\sigma_k(\tilde{Y}) > 0$, then we can write

$$\tilde{Y} = Y + \Delta, \quad Y = P_{\text{St}} \tilde{Y} \in \text{St}(n, p).$$

Using that Y^* is the minimizer of F_i on $\text{St}(n, p)$, we thus get

$$\tilde{F}_i(\tilde{Y}) = d^2(X_{i-1}, Y) + d^2(Y, X_{i+1}) + 2\|\Delta\|_F^2 \geq F_i(Y) \geq F_i(Y^*).$$

The same inequality holds trivially if $\sigma_k(\tilde{Y}) = 0$ since then $\tilde{F}_i(\tilde{Y}) = +\infty$. Finally, since $\tilde{F}_i(Y^*) = F_i(Y^*)$, we obtain that \tilde{F}_i is also uniquely minimized by Y^* . \square

Lemma 3.3. Suppose that for all iterations $k = 0, 1, \dots$, the iterates of leapfrog satisfy

$$d(X_{i-1}^{(k)}, X_{i+1}^{(k-1)}) \leq \delta_g,$$

for all $i = 1, 2, \dots, m-2$. Then, the leapfrog algorithm started in $X^{(0)}$ generates the same iterates as the nonlinear Gauss–Seidel algorithm started in $X^{(0)}$ and applied to

$$\min_{X_1, \dots, X_{m-2} \in \mathbb{R}^{n \times p}} \tilde{F}(X_1, \dots, X_{m-2}).$$

Proof. By induction. Suppose true until substep $i-1$ of iteration k . Then, leapfrog computes the new iterate as

$$X_i^{(k)} = \arg \min_{Y \in \text{St}(n, p)} d^2(X_{i-1}^{(k)}, Y) + d^2(Y, X_{i+1}^{(k-1)}).$$

The uniqueness of the minimizer follows from Lemma 3.1 and $d(X_{i-1}^{(k)}, X_{i+1}^{(k-1)}) \leq \delta_g$. Likewise, nonlinear Gauss–Seidel computes

$$\tilde{X}_i^{(k)} = \arg \min_{\tilde{Y} \in \mathbb{R}^{n \times p}} \tilde{F}(X_1^{(k)}, \dots, X_{i-1}^{(k)}, \tilde{Y}, X_{i+1}^{(k-1)}, \dots, X_{m-2}^{(k-1)}),$$

³For the Riemannian distance d_e based on the embedded metric, it is easy to see that $\|X - Y\|_F \leq d_e(X, Y)$ since the Euclidean length of a geodesic on $\text{St}(n, p)$ is always larger than that of a straight line.

and the uniqueness of the minimizer follows from our reasoning below. Both minimization problems are the same as minimizing F_i and \tilde{F}_i from Lemma 3.2 but with $X_{i-1}^{(k)}$ and $X_{i+1}^{(k-1)}$ taking the roles of X_{i-1} and X_{i+1} , respectively. By Lemma 3.2, the minimizers of both problems are the same and hence $X_i^{(k)} = \tilde{X}_i^{(k)}$. The above reasoning can also be applied to the base case $k = i = 1$ since $X_0^{(1)} = X_0^{(0)}$. Hence, we have proven the result. \square

If the initial points are close enough, the iterates in leapfrog stay close.

Lemma 3.4. *Let $X^{(0)} \in \text{St}(n, p)^m$ be such that $d(X_{i-1}^{(0)}, X_i^{(0)}) \leq \frac{1}{2}\delta_g$ for all $1 \leq i \leq m-1$. Then, leapfrog started at $X^{(0)}$ is well defined and all its iterates $X^{(k)}$ satisfy for all $1 \leq i \leq m-2$ and $k \geq 1$*

$$d(X_{i-1}^{(k)}, X_i^{(k)}) = d(X_i^{(k)}, X_{i+1}^{(k-1)}) \leq \frac{1}{2}\delta_g. \quad (6)$$

Proof. By induction. Suppose true for all substeps i until iteration $k-1$ and until substep $i-1$ of iteration k . This implies in particular

$$d(X_{i-1}^{(k)}, X_i^{(k-1)}) \leq \frac{1}{2}\delta_g, \quad d(X_i^{(k-1)}, X_{i+1}^{(k-1)}) \leq \frac{1}{2}\delta_g.$$

By triangle inequality for the Riemannian distance,

$$d(X_{i-1}^{(k)}, X_{i+1}^{(k-1)}) \leq d(X_{i-1}^{(k)}, X_i^{(k-1)}) + d(X_i^{(k-1)}, X_{i+1}^{(k-1)}) \leq \delta_g,$$

Lemma 3.1 gives that the leapfrog iteration is well defined and produces the unique minimizer

$$X_i^{(k)} = \arg \min_{Y \in \text{St}(n, p)} d^2(X_{i-1}^{(k)}, Y) + d^2(Y, X_{i+1}^{(k-1)}).$$

We thus have

$$d^2(X_{i-1}^{(k)}, X_i^{(k)}) + d^2(X_i^{(k)}, X_{i+1}^{(k-1)}) \leq d^2(X_{i-1}^{(k)}, X_i^{(k-1)}) + d^2(X_i^{(k-1)}, X_{i+1}^{(k-1)}) \leq \frac{1}{2}\delta_g^2.$$

Since $X_i^{(k)}$ is the midpoint of the geodesic connecting $X_{i-1}^{(k)}$ to $X_{i+1}^{(k-1)}$, we also have

$$d(X_{i-1}^{(k)}, X_i^{(k)}) = d(X_i^{(k)}, X_{i+1}^{(k-1)}).$$

Combining these two results proves (6) until substep i at iteration k . Since $X_0^{(k+1)} = X_0^{(k)} = X_0^{(0)}$, the case for substep $i = 1$ and iteration $k+1$ satisfies the same reasoning as above. The same is true for the base case $i = k = 1$, which ends the proof. \square

Hence, combining Lemmas 3.3 and 3.4, we get our desired result:

Proposition 3.1. *Let $X^{(0)} \in \text{St}(n, p)^m$ be such that $d(X_{i-1}^{(0)}, X_i^{(0)}) \leq \frac{1}{2}\delta_g$ for all $1 \leq i \leq m$. Then the leapfrog algorithm applied to F is equivalent to the nonlinear Gauss–Seidel method applied to \tilde{F} .*

We can now proceed and analyze the convergence of this nonlinear Gauss–Seidel method using standard theory.

3.4 First-order optimality

From Prop. 3.1, we know that at iteration $k \geq 1$ and for subinterval $i \in \{1, \dots, m-2\}$, leapfrog solves the following unconstrained optimization problem

$$\min_{X_i \in \mathbb{R}^{n \times p}} \tilde{F}_i^k(X_i),$$

where the objective function is defined as

$$\tilde{F}_i^k(Y) = \tilde{d}^2(X_{i-1}^{(k)}, Y) + \tilde{d}^2(Y, X_{i+1}^{(k-1)}).$$

Recall that $X_{i-1}^{(k)}, X_{i+1}^{(k-1)} \in \text{St}(n, p)$ are the neighboring points of X_i and that $X_{i-1}^{(k)}$ was previously updated and that $X_{i+1}^{(k-1)}$ will be updated next.

Let us define

$$\mathcal{G}_i(Y) = \nabla_Y \tilde{F}_i^k(Y) = \nabla_Y \tilde{d}^2(X_{i-1}^{(k)}, Y) + \nabla_Y \tilde{d}^2(Y, X_{i+1}^{(k-1)}).$$

At the minimizer X_i , the gradient of \tilde{F}_i^k vanishes, i.e., $\mathcal{G}_i(X_i) = 0$. Likewise, if we take all the minimizers $X = (X_1, \dots, X_{m-2})$ together, they will satisfy

$$\begin{cases} \mathcal{G}_1(X) = \nabla_{X_1} \tilde{d}^2(X_0, X_1) + \nabla_{X_1} \tilde{d}^2(X_1, X_2) = 0, \\ \mathcal{G}_2(X) = \nabla_{X_2} \tilde{d}^2(X_1, X_2) + \nabla_{X_2} \tilde{d}^2(X_2, X_3) = 0, \\ \vdots \\ \mathcal{G}_{m-2}(X) = \nabla_{X_{m-2}} \tilde{d}^2(X_{m-3}, X_{m-2}) + \nabla_{X_{m-2}} \tilde{d}^2(X_{m-2}, X_{m-1}) = 0. \end{cases}$$

This can be written compactly as $\mathcal{G}(X) = 0$, where \mathcal{G} is defined component-wise $\mathcal{G}_i: \mathbb{R}^{n \times p} \rightarrow \mathbb{R}^{n \times p}$, for $i = 1, \dots, m-2$.

3.5 Known results on local convergence

Assuming convergence to the limit point $X_1^*, X_2^*, \dots, X_{m-2}^*$, the asymptotic convergence rate is determined by the spectral radius of a certain blockwise partitioning of the Hessian of \tilde{F} at this limit point.

Theorem 3.5 (Nonlinear block Gauss–Seidel theorem). *Let $\mathcal{G}: \mathcal{D} \subset \mathbb{R}^{(m-2)np} \rightarrow \mathbb{R}^{(m-2)np}$ be continuously differentiable in an open neighborhood $\mathcal{B}_0 \subset \mathcal{D}$ of a point $X^* \in \mathcal{D}$ for which $\mathcal{G}(X^*) = 0$. Consider the decomposition of $\mathcal{G}' = D - L - U$ into its block diagonal, strictly lower-, and strictly upper-triangular parts, and suppose that $D(X^*)$ is nonsingular and $\rho(M^{\text{BGS}}(X^*)) < 1$, where $M^{\text{BGS}} = (D - L)^{-1}U$. Then there exists an open ball $\mathcal{B} = \mathcal{B}(X^*, \delta)$ in \mathcal{B}_0 such that, for any $X^{(0)} \in \mathcal{B}$, there is a unique sequence $\{X^{(k)}\} \subset \mathcal{B}$ which satisfies the nonlinear Gauss–Seidel prescription. Moreover, $\lim_{k \rightarrow \infty} X^{(k)} = X^*$ and for any $X^{(0)} \in \mathcal{B}_0$, the convergence rate in the form $\limsup_{k \rightarrow \infty} \sqrt[k]{\|X^{(k)} - X^*\|}$ is upper bounded by $\rho(M^{\text{BGS}}(X^*))$.*

Proof. As a direct extension of [18, Theorem 10.3.5]. □

This theorem shows the need for the Hessian of \tilde{F} (i.e., \mathcal{G}') and its block $D - L - U$ decomposition. As we shall see, our matrix \mathcal{G}' is given by the sum of two matrices $\mathcal{G}' = A + E$, where A is symmetric block tridiagonal and positive definite, and E can be regarded as a perturbation matrix. Since it is very difficult to compute the spectral radius of M^{BGS} with this perturbation E , we will not use Theorem 3.5 directly. Instead, we will use the *Householder–John theorem* [10, Corollary 3.42], which states that if \mathcal{G}' is positive definite, then the M^{BGS} from Theorem 3.5 satisfies $\rho(M^{\text{BGS}}) < 1$. In other words, (linear) block Gauss–Seidel for a symmetric and positive definite \mathcal{G}' always converges monotonically in the energy norm [10, Theorem 3.53]. Therefore, we only need to restrict the perturbation E such that the whole matrix \mathcal{G}' is symmetric and positive definite. In order to do that, we will also use a block version of the Gershgorin circle theorem [7, Theorem 2].

3.6 Local convergence

As required in Theorem 3.5, we compute the Hessian as the Jacobian matrix $\mathcal{G}'(X)$, a square matrix of size $(m-2)np$.

$$\mathcal{G}' = \begin{bmatrix} \nabla_{X_1}^2 \tilde{d}^2(X_0, X_1) + \nabla_{X_1}^2 \tilde{d}^2(X_1, X_2) & \nabla_{X_2} \nabla_{X_1} \tilde{d}^2(X_1, X_2) & & \\ \nabla_{X_1} \nabla_{X_2} \tilde{d}^2(X_1, X_2) & \nabla_{X_2}^2 \tilde{d}^2(X_1, X_2) + \nabla_{X_2}^2 \tilde{d}^2(X_2, X_3) & \nabla_{X_3} \nabla_{X_2} \tilde{d}^2(X_2, X_3) & \\ & \ddots & \ddots & \\ & & \nabla_{X_{m-2}} \nabla_{X_{m-3}} \tilde{d}^2(X_{m-3}, X_{m-2}) & \ddots \end{bmatrix}.$$

By symmetry of the Hessian, we can write this compactly as

$$\mathcal{G}' = \begin{bmatrix} D_{10} + D_{12} & L_{12}^\top & & \\ L_{12} & D_{21} + D_{23} & L_{23}^\top & \\ & \ddots & \ddots & \\ & & L_{m-3, m-2} & D_{m-2, m-3} + D_{m-2, m-1} \end{bmatrix},$$

where

$$L_{ij} = \nabla_{X_i} \nabla_{X_j} \tilde{d}^2(X_i, X_j) \quad \text{and} \quad D_{ij} = \nabla_{X_i}^2 \tilde{d}^2(X_i, X_j)$$

denote the mixed and double derivatives⁴.

We now turn to the computation of these derivatives L_{ij} and D_{ij} . To that end, the following lemma is convenient since it writes $\tilde{d}^2(X_i, X_j)$ as an expansion that does not explicitly use the geodesic distance.

Lemma 3.6. *Let $\tilde{X}, \tilde{Y} \in \mathbb{R}^{n \times p}$ such that $\sigma_p(\tilde{X}) > 0$ and $\sigma_p(\tilde{Y}) > 0$, then*

$$\begin{aligned} \tilde{d}^2(\tilde{X}, \tilde{Y}) &= \|\text{P}_{\text{St}} \tilde{X} - \text{P}_{\text{St}} \tilde{Y}\|_F^2 - \frac{1}{2} \|I_p - (\text{P}_{\text{St}} \tilde{X})^\top \text{P}_{\text{St}} \tilde{Y}\|_F^2 \\ &\quad + \|\tilde{X} - \text{P}_{\text{St}} \tilde{X}\|_F^2 + \|\tilde{Y} - \text{P}_{\text{St}} \tilde{Y}\|_F^2 + O(\|\text{P}_{\text{St}} \tilde{X} - \text{P}_{\text{St}} \tilde{Y}\|_F^4). \end{aligned} \quad (7)$$

Proof. See App. A. □

In the following, denote $\delta_{ij} = \|X_i - X_j\|_2$ for any $X_i, X_j \in \text{St}(n, p)$.

Lemma 3.7. *Let $X_i \in \text{St}(n, p)$. Then*

$$D_{ij} = 2I_{np} + \frac{1}{2} (X_i^\top \otimes X_i) \Pi_{p,n} - \frac{1}{2} (I_p \otimes X_i X_i^\top) + \Delta_{ij}, \quad (8)$$

$$L_{ij} = -2I_{np} + \frac{1}{2} (X_i^\top \otimes X_i) \Pi_{p,n} + \frac{3}{2} (I_p \otimes X_i X_i^\top) + \Lambda_{ij}, \quad (9)$$

with $\|\Delta_{ij}\|_2 \leq 14\delta_{ij} + 10\delta_{ij}^2$ and $\|\Lambda_{ij}\|_2 \leq \frac{11}{2}\delta_{ij} + 10\delta_{ij}^2 + 4\delta_{ij}^3$. Here, $\Pi_{p,n}$ is the vec-permutation matrix defined as the permutation matrix that satisfies $\text{vec}(X) = \Pi_{n,p} \text{vec}(X^\top)$; see, e.g., [11, Eq. (5)].

Proof. See App. B. □

Our aim is to diagonalize \mathcal{G}' . We will do this in a few steps. First, observe that \mathcal{G}' remains block-tridiagonal if it is transformed using a compatible block diagonal matrix $\mathcal{Q} = \text{diag}\{Q_1, Q_2, \dots, Q_{m-2}\}$:

$$\mathcal{Q}^\top \mathcal{G}' \mathcal{Q} = \begin{bmatrix} Q_1^\top (D_{10} + D_{12}) Q_1 & Q_1^\top L_{12}^\top Q_2 & & \\ Q_2^\top L_{12} Q_1 & Q_2^\top (D_{21} + D_{23}) Q_2 & Q_2^\top L_{23}^\top Q_3 & \\ & \ddots & \ddots & \\ & & Q_{m-2}^\top L_{m-3, m-2} Q_{m-1} & D_{m-2, m-3} + D_{m-2, m-1} \end{bmatrix}, \quad (10)$$

Here, the $Q_1, \dots, Q_{m-2} \in \mathbb{R}^{np \times np}$ can be any orthogonal matrices. The lemma below shows us how to choose these matrices so that we obtain diagonal blocks in $\mathcal{Q}^\top \mathcal{G}' \mathcal{Q}$, up to first order in δ_{ij} .

⁴Observe that $L_{ij} = L_{ji}^\top$ by equality of mixed derivatives but in general $D_{ij} \neq D_{ji}^\top$ since only the variable corresponding to the first index is derived.

Lemma 3.8. Let $X_i^\perp \in \mathbb{R}^{n \times (n-p)}$ be such that $X_i^\top X_i^\perp = O_{p \times (n-p)}$ and $(X_i^\perp)^\top X_i^\perp = I_{(n-p)}$. Define the orthogonal matrices

$$\bar{Q}_i = [I_p \otimes X_i \quad I_p \otimes X_i^\perp],$$

and similarly for \bar{Q}_j . Then, there exists an orthogonal matrix \hat{Q} , only depending on n and p , such that $Q_i = \bar{Q}_i \hat{Q}$ and $Q_j = \bar{Q}_j \hat{Q}$ satisfy

$$\|Q_i^\top D_{ij} Q_i - D\|_2 \leq C_D^{(ij)}, \quad D = \text{diag} \{I_{p(p-1)/2}, 2I_{np-p(p-1)/2}\}, \quad (11)$$

$$\|Q_j^\top L_{ij} Q_i - L\|_2 \leq C_L^{(ij)}, \quad L = \text{diag} \{-I_{p(p-1)/2}, -2I_{(n-p)p}, O_{p(p+1)/2}\}, \quad (12)$$

where $C_D^{(ij)} = 14\delta_{ij} + 10\delta_{ij}^2$ and $C_L^{(ij)} = \frac{15}{2}\delta_{ij} + \frac{31}{2}\delta_{ij}^2 + 14\delta_{ij}^3 + 4\delta_{ij}^4$.

Proof. See App. C. □

The matrix \hat{Q} above is related to the diagonalization of the vec-permutation matrix $\Pi_{p,p}$; see (25) in App. C for its definition. It is therefore also independent of X_i . This is a crucial property to obtain the following result.

Lemma 3.9. Define $\delta = \max_{0 \leq i \leq m-2} \delta_{i,i+1}$ and assume $\delta \leq 1$. Then the minimal eigenvalue of \mathcal{G}' is bounded by

$$\lambda_{\min}(\mathcal{G}') \geq 2 - 2 \cos \frac{\pi}{m-1} - 43\delta - 90\delta^2.$$

As a consequence, \mathcal{G}' is symmetric and positive definite when

$$\delta < \frac{1}{180} \left(\sqrt{2569 - 720 \cos \frac{\pi}{m-1}} - 43 \right).$$

Proof. From Lemma 3.8, recall the diagonal matrices D and L , and the orthogonal matrices Q_1, \dots, Q_{m-2} . Define $\mathcal{Q} = \text{diag}\{Q_1, Q_2, \dots, Q_{m-2}\}$. Substituting the nonzero blocks in (10) by

$$Q_i^\top (D_{i,i-1} + D_{i,i+1}) Q_i = 2D + E_{ii}, \quad Q_{i+1}^\top L_{i,i+1} Q_i = L + E_{i,i+1},$$

we can write $\mathcal{Q}^\top \mathcal{G}' \mathcal{Q}$ as

$$\mathcal{Q}^\top \mathcal{G}' \mathcal{Q} = \begin{bmatrix} 2D & L & & \\ L & 2D & L & \\ & \ddots & \ddots & \ddots \end{bmatrix} + \begin{bmatrix} E_{11} & E_{12}^\top & & \\ E_{12} & E_{22} & E_{23}^\top & \\ & \ddots & \ddots & \ddots \end{bmatrix} =: A + E. \quad (13)$$

Eq. (13) is an approximate tridiagonalization of the matrix \mathcal{G}' . Observe that the symmetric matrices A and E have compatible block partitioning. Furthermore, from Lemma 3.8, we get immediately that

$$\|E_{ii}\|_2 \leq 28\delta + 20\delta^2 =: C_D, \quad \|E_{i,i+1}\|_2 \leq \frac{15}{2}\delta + \frac{31}{2}\delta^2 + 14\delta^3 + 4\delta^4 =: C_L.$$

We will regard $\mathcal{Q}^\top \mathcal{G}' \mathcal{Q}$ as an $O(\delta)$ perturbation of A . Using properties of Kronecker products, we can write

$$A = 2I_{m-2} \otimes D + M \otimes L, \quad M = \begin{bmatrix} 0 & 1 & & \\ 1 & \ddots & \ddots & \\ & \ddots & \ddots & 1 \\ & & 1 & 0 \end{bmatrix} \in \mathbb{R}^{(m-2) \times (m-2)}. \quad (14)$$

Thanks to the Kronecker structure in (14) and the diagonal matrices D and L , the eigenvalues of A are easily determined as

$$\lambda_{jk} = 2d_j + \mu_k \ell_j, \quad j = 1, \dots, np, \quad k = 1, \dots, m-2,$$

where d_j and ℓ_j are the diagonal entries of D and L , respectively, and μ_k are the eigenvalues of the Toeplitz matrix M . Using [8, Eq. (2.7)], we find

$$\mu_k = -2 \cos \frac{k\pi}{m-1}, \quad k = 1, \dots, m-2.$$

Together with (11) and (12), this allows us to determine that the minimal value among all λ_{jk} corresponds to $j = 1$ and $k = m-2$. We thus obtain

$$\lambda_{\min}(A) = 2 - 2 \cos \frac{\pi}{m-1} > 0 \quad \text{for all } m \geq 2.$$

By Weyl's inequality [21, Corollary 4.9], $\lambda_{\min}(\mathcal{G}') = \lambda_{\min}(A + E) > 0$ is guaranteed if $\|E\|_2 < \lambda_{\min}(A)$. To bound $\|E\|_2$, we use a block version of the Gershgorin circle theorem (see [7, Theorem 2] and also [23, Remark 1.13.2]). Applied to the symmetric block tridiagonal matrix E , it guarantees that its eigenvalues are included in the union of intervals

$$\bigcup_{i=1}^{m-2} \bigcup_{k=1}^{np} [\varepsilon_k^{(i)} - R_i, \varepsilon_k^{(i)} + R_i], \quad R_i = \|E_{i-1,i}\|_2 + \|E_{i,i+1}^\top\|_2 \leq 2C_L,$$

where $\varepsilon_k^{(i)}$ is the k th eigenvalue of E_{ii} . These eigenvalues $\varepsilon_k^{(i)}$ are all bounded in magnitude by C_D . Hence $\|E\|_2 \leq C_D + 2C_L = 43\delta + 51\delta^2 + 28\delta^3 + 8\delta^4$. Since $\delta < 1$, it is easily verified that $\|E\|_2 \leq 43\delta + 90\delta^2$ and thus the matrix \mathcal{G}' remains positive definite if $43\delta + 90\delta^2 < \lambda_{\min}(A)$, i.e.,

$$\delta < \frac{1}{180} \left(\sqrt{2569 - 720 \cos \frac{\pi}{m-1}} - 43 \right).$$

□

All put together, we have the final result of local convergence.

Theorem 3.10. *If the leapfrog algorithm is started with δ satisfying the condition of Lemma 3.9, then it converges to the unique length-minimizing geodesic connecting X_0 and X_{m-1} , provided that the initial intermediate points are sufficiently close to that geodesic.*

Proof. We use [10, Corollary 3.42] which states that if \mathcal{G}' is positive definite and can be split into the sum of an arbitrary positive definite matrix and an arbitrary symmetric matrix, then the scalar Gauss–Seidel converges, i.e., $\rho(M^{\text{BGS}}) < 1$, and the convergence is monotone with respect to the energy norm $\|\cdot\|_{\mathcal{G}'}$. By [10, Theorem 3.53], we know that this theorem remains valid for any block version.

Now, the splitting (13) has exactly the form prescribed by [10, Corollary 3.42], because A is positive definite and E is symmetric. By Lemma 3.9, we know that \mathcal{G}' remains positive definite if $\delta < \frac{1}{180} \left(\sqrt{2569 - 720 \cos \frac{\pi}{m-1}} - 43 \right)$. Under these conditions, the leapfrog algorithm converges as a block Gauss–Seidel method to the length-minimizing geodesic connecting X_0 and X_{m-1} . □

4 Some observations and open problems

For m large, Lemma 3.9 gives that \mathcal{G}' is positive definite when $\delta \lesssim \pi^2/43m^2$. Let $d_0 = \|X_0 - X_{m-1}\|_2$ be the distance between the two endpoints. Then by equidistant partitioning of the intermediate points, one has $\delta \simeq d_0/m$. To guarantee a positive definite \mathcal{G}' , we would then need $d_0/m \lesssim \pi^2/43m^2$ which implies $m \lesssim 0.23/d_0$.

This result is unsatisfactory, since it would have been desirable to guarantee positive definiteness of $\mathcal{Q}^\top \mathcal{G}' \mathcal{Q} = A + E$ with orthogonal \mathcal{Q} by increasing the number of points m given a fixed d_0 . Unfortunately, we cannot guarantee this with our proof. The problem is that $\|E\|_2 = O(\delta)$ whereas




Figure 3: Eigenvalue perturbations – not at the limiting geodesic.




Figure 4: Eigenvalue perturbations – at the limiting geodesic.

$\lambda_{\min}(A) = O(1/m^2)$, which lead to our condition that m needed to be smaller than some fixed fraction of the original distance d_0 . If $\|E\|_2 = O(\delta^2)$, then there would be no condition on m since $\delta^2 \simeq d_0^2/m^2 \lesssim 1/m^2$ is sufficient to guarantee $\lambda_{\min}(A + E) > 0$. However, it would still not be satisfactory since the perturbation does not lead to an improvement with increasing m , for which one probably needs $\|E\|_2 = O(\delta^3)$. As we show below, there is strong numerical indication that with our choice of extended distance function this is not the case.

Numerical experiments reported in Figure 3 suggest that the minimal eigenvalues of \mathcal{G}' and A differ by $O(\delta^2)$, whereas our perturbation analysis only showed $\|E\|_2 = O(\delta)$. It is however not trivial to prove this result. Indeed, up to first order, we can study the eigenvalues of the symmetric matrix $A + E$ by using the derivative formula [21, Theorem 2.3]

$$\lambda_{\min}(A + E) = \lambda_{\min}(A) + v_{\min}^{\top} E v_{\min} + O(\|E\|^2), \quad (15)$$

where $\lambda_{\min}(A)$ is assumed to be isolated (as it is the case) and v_{\min} is its associated eigenvector. One possibility to improve on our bounds, at least asymptotically, would be to prove that $|v_{\min}^{\top} E v_{\min}| = O(\delta^3)$. However, in the same figure, $|v_{\min}^{\top} E v_{\min}|$ seems to be again $O(\delta^2)$. In addition, all these conclusions remain true in the limiting geodesic.

Another problem with the matrix A and \mathcal{G}' is that it has a bad spectral gap γ (i.e., the difference of smallest and second smallest eigenvalue) when m grows. Numerical observations suggest that the spectral gap might be $O(1/m)$ which complicates non-asymptotic bounds.

As a last remark, one could resort to a more general theory for the convergence of nonlinear block Gauss–Seidel for a quasi-convex objective function [9], which requires quasi-convexity for each X_i alone. Looking at the Hessian \mathcal{G}' where all X_j except X_i are constant, the only block that is left in the matrix \mathcal{G}' is the diagonal one, namely $D_{i,i-1} + D_{i,i+1}$. Using Lemma 3.8, we immediately get the eigenvalues of this block. Now, for $C_D^{(ij)} < 1$ in (11) we get strong convexity in X_i alone. One problem with this approach is that the feasible set has to be a Cartesian product of convex subsets of $\mathbb{R}^{n \times p}$. Moreover, the result in [9] only guarantees subsequence convergence, and there is no rate of convergence or contraction rate for the whole sequence. Hence the convergence behavior could also be slower than linear.

5 Numerical experiments and applications

As a concrete example to demonstrate the leapfrog algorithm, let us consider the Stiefel manifold $\text{St}(12, 3)$. We fix one point $X = [I_p \ O_{(n-p) \times p}]^{\top}$, while the other point Y is placed at the geodesic

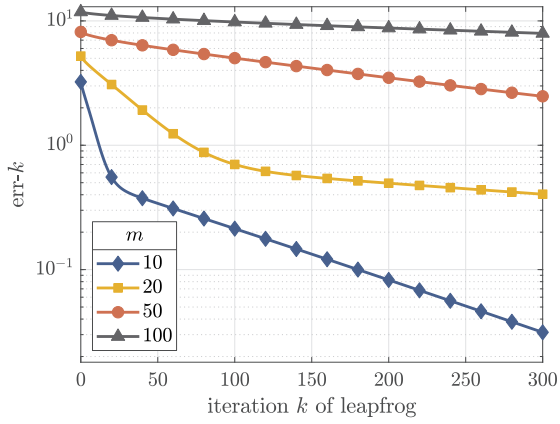


Figure 5: Convergence behavior of $\text{err-}k$ for increasing values of m .

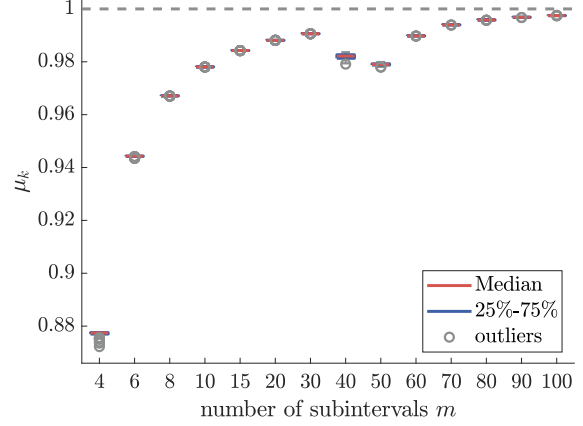


Figure 6: Boxplot of $\max_k \{\mu_k^{(i)}\}$ for increasing values of m .

distance $L^* = 0.95\pi$ from X . This is done by creating a tangent vector to $\text{St}(12, 3)$ at X of length L^* , and then mapping it to $\text{St}(12, 3)$ via the Riemannian exponential (2). For this choice of L^* , single shooting will not work (recall that the injectivity radius on $\text{St}(n, p)$ is at least 0.89π). We want to recover this geodesic distance using the leapfrog algorithm and study its convergence.

For each value of $m \in \{10, 20, 50, 100\}$, we construct an initial guess $X^{(0)}$ by placing $m - 2$ intermediate points randomly along the linear segment connecting X and Y in the embedding space, and projecting them to the Stiefel manifold. We then apply leapfrog for 300 iterations and monitor the convergence behavior of

$$\text{err-}k = \|X^{(k)} - X^*\|_F,$$

where X^* is the solution of leapfrog (i.e., a uniformly distributed tuple corresponding to the global geodesic that was constructed above), and $X^{(k)}$ is the approximate solution at iteration k of leapfrog. This is illustrated in Figure 5, from which it is clear that for large m leapfrog always converges albeit very slowly.

Next, we apply leapfrog for 50 iterations and for each $m \in \{4, 6, 8, 10, \dots, 100\}$ we repeat this experiment for 100 random initializations of the initial guess $X^{(0)}$. For each experiment i , we define the error reduction rate⁵ as

$$\mu_k^{(i)} = \frac{\text{err-}(k+1)}{\text{err-}k}, \quad \text{for } k = 0, 1, \dots, 49, \quad i = 1, \dots, 100,$$

and we compute the worst and the median reduction rates across all the experiments, $\max_{i,k} \{\mu_k^{(i)}\}$ and $\text{med}_i \max_k \{\mu_k^{(i)}\}$. Since during the first iterations leapfrog is faster, we also compute the convergence factor given by $\max_i \{\mu_0^{(i)}\}$.

From Table 1, we see that the convergence of leapfrog deteriorates as m increases but it remains strictly smaller than 1. For small values of m , $\max_i \{\mu_0^{(i)}\}$ and $\max_{i,k} \{\mu_k^{(i)}\}$ are significantly different, whereas for large values of m , they are quite similar. The same conclusion can be reached from Figure 6 where boxplots show the dispersion and skewness in the $\mu_k^{(i)}$. Clearly, the convergence factors become very concentrated for large m .

A Proof of Lemma 3.6

The expansion (7) is simple to obtain once the Riemannian distance is related to the Euclidean one.

⁵In the limit $k \rightarrow \infty$, this gives the asymptotic Q-rate of convergence of the sequence.

Table 1: Values of $\max_i \{\mu_0^{(i)}\}$, $\max_{i,k} \{\mu_k^{(i)}\}$ and $\text{med}_i \max_k \{\mu_k^{(i)}\}$ versus number of points m , for the experiment described in Sect. 5.

m	4	6	8	10	15	20	30
$\max_i \{\mu_0^{(i)}\}$	0.5577	0.7058	0.7829	0.8296	0.8604	0.8824	0.8980
$\max_{i,k} \{\mu_k^{(i)}\}$	0.8776	0.9443	0.9671	0.9781	0.9843	0.9881	0.9906
$\text{med}_i \max_k \{\mu_k^{(i)}\}$	0.8774	0.9443	0.9671	0.9781	0.9843	0.9881	0.9906
m	40	50	60	70	80	90	100
$\max_i \{\mu_0^{(i)}\}$	0.5577	0.7058	0.7829	0.8296	0.8604	0.8824	0.8980
$\max_{i,k} \{\mu_k^{(i)}\}$	0.8776	0.9443	0.9671	0.9781	0.9843	0.9881	0.9906
$\text{med}_i \max_k \{\mu_k^{(i)}\}$	0.8774	0.9443	0.9671	0.9781	0.9843	0.9881	0.9906

Proof of Lemma 3.6. Take $X, Y \in \text{St}(n, p)$ sufficiently close so that we can define the Riemannian logarithm $\xi = \text{Log}_X(Y)$ (see Remark 3.1). By definition of the Riemannian distance d for the canonical metric g , we have

$$d^2(X, Y) = \|\xi\|^2 = g(\xi, \xi).$$

Writing a tangent vector as $\xi = X\Omega + X_\perp K \in T_X \text{St}(n, p)$ (see Sect. 1) and using (1), we can evaluate g as

$$g(\xi, \xi) = \text{trace}(\xi^\top (I_n - \frac{1}{2} X X^\top) \xi) = \frac{1}{2} \|\Omega\|_F^2 + \|K\|_F^2 = \|\xi\|_F^2 - \frac{1}{2} \|\Omega\|_F^2.$$

Using $\Omega = X^\top \xi$, we also have

$$d^2(X, Y) = \|\xi\|_F^2 - \frac{1}{2} \|X^\top \xi\|_F^2. \quad (16)$$

Since ξ is the initial velocity vector of the geodesic connecting X to Y , it follows that

$$\xi = Y - X + O(\|\xi\|_F^2). \quad (17)$$

This can be seen by expanding the matrix exponential in the expression (2) of the geodesic:

$$\begin{aligned} Y &= [X \ X_\perp] \left(I_n + \begin{bmatrix} X^\top \xi & -\xi^\top X_\perp \\ X_\perp^\top \xi & O_{n-p} \end{bmatrix} + O(\|\xi\|_F^2) \right) \begin{bmatrix} I_p \\ O_{(n-p) \times p} \end{bmatrix} \\ &= X + [X \ X_\perp] [X \ X_\perp]^\top \xi + O(\|\xi\|_F^2). \end{aligned}$$

We obtain (17) using the fact that $[X \ X_\perp]$ is an orthogonal matrix. In addition, [3, Lemma 4.2.1, p. 59] shows that

$$\|\xi\|_F^2 = \|X - Y\|_F^2 + O(\|X - Y\|_F^4). \quad (18)$$

Then inserting the equations (17) and (18) into (16) leads to

$$d^2(X, Y) = \|X - Y\|_F^2 - \frac{1}{2} \|X^\top (X - Y)\|_F^2 + O(\|X - Y\|_F^4). \quad (19)$$

Using this in (5), one obtains (7). □

B Proof of Lemma 3.7

The aim is to compute $L_{ij} = \nabla_{X_i} \nabla_{X_j} \tilde{d}^2(X_i, X_j)$ and $D_{ij} = \nabla_{X_i}^2 \tilde{d}^2(X_i, X_j)$, where $X_j \in \text{St}(n, p)$. Let us simplify notation and define

$$\begin{aligned} \tilde{d}^2(X, Y) &= \|\text{P}_{\text{St}} X - \text{P}_{\text{St}} Y\|_F^2 - \frac{1}{2} \|I_p - (\text{P}_{\text{St}} X)^\top \text{P}_{\text{St}} Y\|_F^2 \\ &\quad + \|X - \text{P}_{\text{St}} X\|_F^2 + \|Y - \text{P}_{\text{St}} Y\|_F^2 + O(\|\text{P}_{\text{St}} X - \text{P}_{\text{St}} Y\|_F^4). \end{aligned}$$

Clearly, $L_{ij} = \nabla_X \nabla_Y \tilde{d}^2(X, Y)$ and $D_{ij} = \nabla_X^2 \tilde{d}^2(X, Y)$ with $X = X_i$ and $Y = X_j$. Recall from Sect. 3.2 that we can specify the projector on the Stiefel manifold as $\text{P}_{\text{St}}(Y) = Y(Y^\top Y)^{-1/2}$, that is, the orthogonal factor of the polar decomposition of Y .

Proof of Lemma 3.7. To compute the gradient and the Hessian of $\tilde{d}^2(\tilde{X}, \tilde{Y})$, consider the perturbation $\tilde{X} = X + E$, with $X \in \text{St}(n, p)$, $\|E\|_F$ small, and expand the previous expression.

First, for a symmetric matrix A , one can easily show by diagonalizing that

$$(I + A)^{-1/2} = I - \frac{1}{2}A + \frac{3}{8}A^2 + O(\|A\|^3), \quad \|A\| \rightarrow 0,$$

from which we can obtain the expansion for the perturbed projector

$$\begin{aligned} \text{P}_{\text{St}} \tilde{X} &= \tilde{X}(\tilde{X}^\top \tilde{X})^{-1/2} \\ &= X + E - \frac{1}{2}XX^\top E - \frac{1}{2}XE^\top X - \frac{1}{2}XE^\top E - \frac{1}{2}EX^\top E - \frac{1}{2}EE^\top X \\ &\quad + \frac{3}{8}X(X^\top E)^2 + \frac{3}{8}X(E^\top X)^2 + \frac{3}{8}XX^\top EE^\top X + \frac{3}{8}XE^\top XX^\top E + O(\|E\|_F^3). \end{aligned} \quad (20)$$

After substituting the expansion (20) for $\text{P}_{\text{St}}(\tilde{X})$ in $\tilde{d}^2(X, Y)$ and isolating first- and second-order terms in E , we find the expressions for the gradient and the Hessian. Here, only the final results are reported.

The gradient with respect to X is

$$\nabla_X \tilde{d}^2(X, \tilde{Y}) = -(I_n - \frac{1}{2}XX^\top) \text{P}_{\text{St}} \tilde{Y} + \frac{1}{2}X(\text{P}_{\text{St}} \tilde{Y})^\top X - (I_n - XX^\top) \text{P}_{\text{St}} \tilde{Y} (\text{P}_{\text{St}} \tilde{Y})^\top X, \quad (21)$$

and the gradient with respect to Y is

$$\nabla_Y \tilde{d}^2(\tilde{X}, Y) = -(I_n - \frac{1}{2}YY^\top) \text{P}_{\text{St}} \tilde{X} + \frac{1}{2}Y(\text{P}_{\text{St}} \tilde{X})^\top Y - (I_n - YY^\top) \text{P}_{\text{St}} \tilde{X} (\text{P}_{\text{St}} \tilde{X})^\top Y.$$

The Hessian matrix with respect to X is

$$\begin{aligned} \nabla_X^2 \tilde{d}^2(X, Y) &= \text{Sym} \left[Y^\top X \otimes I_n + (Y^\top \otimes X) \Pi_{p,n} + I_p \otimes YX^\top \right] \\ &\quad - \frac{3}{4} \text{Sym} \left[(Y^\top XX^\top \otimes X) \Pi_{p,n} + (X^\top \otimes XY^\top X) \Pi_{p,n} + I_p \otimes XX^\top YX^\top + Y^\top X \otimes XX^\top \right] \\ &\quad + 2 \text{Sym} \left[(X^\top YY^\top \otimes X) \Pi_{p,n} + I_p \otimes XX^\top YY^\top - (X^\top YY^\top XX^\top \otimes X) \Pi_{p,n} \right] \\ &\quad + (X^\top \otimes X) \Pi_{p,n} + I_p \otimes XX^\top - I_p \otimes YY^\top + X^\top YY^\top X \otimes I_n \\ &\quad - I_p \otimes XX^\top YY^\top XX^\top - X^\top YY^\top X \otimes XX^\top, \end{aligned}$$

where $\text{Sym}(A) = (A + A^\top)/2$. In order to simplify $\nabla_X^2 \tilde{d}^2(X, Y)$, we will take $Y = X + \Delta$ with $\|\Delta\| \rightarrow 0$. After some algebraic manipulations, we obtain⁶

$$\begin{aligned} \nabla_X^2 \tilde{d}^2(X, X + \Delta) &= 2I_{np} + \frac{1}{2}(X^\top \otimes X) \Pi_{p,n} - \frac{1}{2}(I_p \otimes XX^\top) + 3 \text{Sym} \left(X^\top \Delta \otimes I_n + (\Delta^\top \otimes X) \Pi_{p,n} \right) \\ &\quad + \text{Sym}(I_p \otimes \Delta X^\top) - \frac{11}{4} \text{Sym} \left((\Delta^\top XX^\top \otimes X) \Pi_{p,n} + \Delta^\top X \otimes XX^\top \right) \\ &\quad - \frac{3}{4} \text{Sym} \left((X^\top \otimes X \Delta^\top X) \Pi_{p,n} + I_p \otimes XX^\top \Delta X^\top \right) \\ &\quad + 2 \text{Sym} \left((X^\top \Delta \Delta^\top \otimes X) \Pi_{p,n} + I_p \otimes XX^\top \Delta \Delta^\top - (X^\top \Delta \Delta^\top XX^\top \otimes X) \Pi_{p,n} \right) \\ &\quad - I_p \otimes \Delta \Delta^\top + X^\top \Delta \Delta^\top X \otimes I_n - I_p \otimes XX^\top \Delta \Delta^\top XX^\top - X^\top \Delta \Delta^\top X \otimes XX^\top. \end{aligned}$$

⁶We stress that $\nabla_X^2 \tilde{d}^2$ denotes the derivative with respect to the first argument of \tilde{d}^2 .

Observe that every term on the right-hand side above can be bounded by at most a second power of $\|\Delta\|_2$ since $\|\text{Sym}(A)\|_2 \leq \|A\|_2$, $\|A \otimes B\|_2 = \|A\|_2 \|B\|_2$ and $X \in \text{St}(n, p)$. Hence, we obtain after some manipulation that

$$\|\nabla_X^2 \tilde{d}^2(X, X + \Delta) - \nabla_X^2 \tilde{d}^2(X, X)\|_2 \leq 14\|\Delta\|_2 + 10\|\Delta\|_2^2.$$

Writing the result with $X = X_i$ and $X + \Delta = X_j$, we recover the expression (8) for the Hessian $D_{ij} = \nabla_{X_i}^2 \tilde{d}^2(X_i, X_j)$.

Next, for the term L_{ij} , to obtain the gradient of (21) with respect to X we can expand $P_{\text{St}} \tilde{X}$ at first order in E

$$P_{\text{St}} \tilde{X} = \tilde{X}(\tilde{X}^\top \tilde{X})^{-1/2} = X + E - \frac{1}{2}XX^\top E - \frac{1}{2}XE^\top X + O(\|E\|_F^2).$$

After some manipulations, we arrive at the mixed term

$$\begin{aligned} \nabla_X \nabla_Y \tilde{d}^2(X, Y) = & -I_{np} + \frac{1}{2}(I_p \otimes YY^\top) - \frac{1}{4}(I_p \otimes YY^\top XX^\top) - \frac{1}{4}(X^\top \otimes YY^\top X) \Pi_{p,n} \\ & + \frac{1}{2}(I_p \otimes XX^\top) + \frac{1}{2}(X^\top \otimes X) \Pi_{p,n} + \frac{1}{2}(Y^\top \otimes Y) \Pi_{p,n} - \frac{1}{4}(Y^\top XX^\top \otimes Y) \Pi_{p,n} \\ & - \frac{1}{4}(Y^\top X \otimes YX^\top) + (Y^\top \otimes YY^\top X) \Pi_{p,n} - (Y^\top XX^\top \otimes YY^\top X) \Pi_{p,n} \\ & - Y^\top X \otimes YY^\top XX^\top + Y^\top X \otimes YY^\top - (Y^\top \otimes X) \Pi_{p,n} + (Y^\top XX^\top \otimes X) \Pi_{p,n} \\ & + Y^\top X \otimes XX^\top - Y^\top X \otimes I_n. \end{aligned}$$

We observe that the other mixed term is $\nabla_Y \nabla_X \tilde{d}^2(X, Y) = (\nabla_X \nabla_Y \tilde{d}^2(Y, X))^\top$.

As above, in order to bound the spectrum of $\nabla_X \nabla_Y \tilde{d}^2(X, Y)$, we expand it with $Y = X + \Delta$ with $\|\Delta\| \rightarrow 0$. After some algebraic manipulations, we obtain

$$\begin{aligned} \nabla_X \nabla_Y \tilde{d}^2(X, X + \Delta) = & -2I_{np} + \frac{1}{2}(X^\top \otimes X) \Pi_{p,n} + \frac{3}{2}(I_p \otimes XX^\top) - \frac{1}{4}(X^\top \otimes X \Delta^\top X) \Pi_{p,n} + \frac{1}{2}(\Delta^\top \otimes X) \Pi_{p,n} \\ & - \frac{1}{4}(\Delta^\top XX^\top \otimes X) \Pi_{p,n} + \frac{3}{4}(\Delta^\top X \otimes XX^\top) - \frac{5}{4}(I_p \otimes X \Delta^\top XX^\top) + \frac{3}{2}(I_p \otimes X \Delta^\top) \\ & - \Delta^\top X \otimes I_n - \frac{5}{4}(\Delta^\top XX^\top \otimes \Delta) \Pi_{p,n} - \frac{5}{4}(I_p \otimes \Delta \Delta^\top XX^\top) + \frac{3}{2}(I_p \otimes \Delta \Delta^\top) \\ & + \Delta^\top X \otimes X \Delta^\top + \frac{3}{2}(\Delta^\top \otimes \Delta) \Pi_{p,n} - (\Delta^\top XX^\top \otimes X \Delta^\top X) \Pi_{p,n} - \frac{1}{4}(X^\top \otimes \Delta \Delta^\top X) \Pi_{p,n} \\ & + (\Delta^\top \otimes X \Delta^\top X) \Pi_{p,n} - \frac{1}{4}(\Delta^\top X \otimes \Delta X^\top) - \Delta^\top X \otimes X \Delta^\top XX^\top + (\Delta^\top \otimes \Delta \Delta^\top X) \Pi_{p,n} \\ & + \Delta^\top X \otimes \Delta \Delta^\top - \Delta^\top X \otimes \Delta \Delta^\top XX^\top - (\Delta^\top XX^\top \otimes \Delta \Delta^\top X) \Pi_{p,n}. \end{aligned}$$

Observe that every term on the right-hand side above can be bounded by at most a third power of $\|\Delta\|_2$. Hence, we obtain that

$$\|\nabla_X \nabla_Y \tilde{d}^2(X, X + \Delta) - \nabla_X \nabla_Y \tilde{d}^2(X, X)\|_2 \leq \frac{11}{2}\|\Delta\|_2 + 10\|\Delta\|_2^2 + 4\|\Delta\|_2^3.$$

Writing the result with $X = X_i$ and $Y = X_j$, we recover the expression (9) for the gradient L_{ij} . \square

C Proof of Lemma 3.8

We first start with $i = j$, which corresponds to $\Delta_{ij} = \Lambda_{ij} = 0$ in Lemma 3.7.

Lemma C.1. Define the orthogonal matrix $\bar{Q}_i = [I_p \otimes X_i \quad I_p \otimes X_i^\perp]$. Then there exists an orthogonal matrix \hat{Q} , only depending on n and p , such that $Q_i = \bar{Q}_i \hat{Q}$ satisfies

$$Q_i^\top D_{ii} Q_i = D = \text{diag} \{ I_{p(p-1)/2}, 2I_{np-p(p-1)/2} \}, \quad (22)$$

and

$$Q_i^\top L_{ii} Q_i = L = \text{diag} \{ -I_{p(p-1)/2}, -2I_{(n-p)p}, O_{p(p+1)/2} \}. \quad (23)$$

Proof. By properties of the so-called vec-permutation matrices (see [11, Eq. (5), (6), (23)]), there exists a permutation matrix $\Pi_{p,n} \in \mathbb{R}^{np \times np}$ that satisfies

$$\Pi_{p,n}(X_i^\top \otimes X_i) \Pi_{p,n} = X_i \otimes X_i^\top, \quad \Pi_{p,n}^{-1} = \Pi_{p,n}^\top.$$

This shows that $(X_i^\top \otimes X_i) \Pi_{p,n} = \Pi_{p,n}^\top (X_i \otimes X_i^\top)$ is symmetric. Furthermore,

$$((X_i^\top \otimes X_i) \Pi_{p,n})^2 = (X_i^\top \otimes X_i) \Pi_{p,n} \Pi_{p,n}^\top (X_i \otimes X_i^\top) = I_p \otimes X_i X_i^\top.$$

Denoting the symmetric matrix $S_i = (X_i^\top \otimes X_i) \Pi_{p,n}$, we can therefore write using Lemma 3.7 that

$$D_{ii} = 2I_{np} + \frac{1}{2}S_i - \frac{1}{2}S_i^2, \quad L_{ii} = -2I_{np} + \frac{1}{2}S_i + \frac{3}{2}S_i^2. \quad (24)$$

It thus suffices to diagonalize S_i . To this end, define the orthogonal matrix

$$\bar{Q}_i = [I_p \otimes X_i \quad I_p \otimes X_i^\perp] \in O(np).$$

Direct calculation shows that

$$\bar{Q}_i^\top S_i \bar{Q}_i = \text{diag} \left\{ (X_i^\top \otimes I_p) \Pi_{p,n} (I_p \otimes X_i), O_{(n-p)p} \right\} = \text{diag} \{ \Pi_{p,p}, O_{(n-p)p} \} =: \hat{\Pi},$$

where we used that $\Pi_{p,n}(I_p \otimes X_i) \Pi_{p,p} = X_i \otimes I_p$, with $\Pi_{p,p} \in \mathbb{R}^{p^2 \times p^2}$ another vec-permutation matrix that is also symmetric (see [11, Eq. (6), (15)]). The matrix $\hat{\Pi}$ above therefore has the spectral decomposition

$$\hat{\Pi} = \hat{Q} \hat{\Lambda} \hat{Q}^\top, \quad \hat{\Lambda} = \text{diag} \{ -I_{p(p-1)/2}, O_{(n-p)p}, I_{p(p+1)/2} \} \quad (25)$$

for some orthogonal matrix \hat{Q} that indeed does not depend on X_i , as claimed. By defining the orthogonal matrix $Q_i = \bar{Q}_i \hat{Q}$, we have thus shown that $Q_i^\top S_i Q_i = \hat{\Lambda}$ and by (24) also that

$$Q_i^\top D_{ii} Q_i = 2I_{np} + \frac{1}{2}\hat{\Lambda} - \frac{1}{2}\hat{\Lambda}^2, \quad Q_i^\top L_{ii} Q_i = -2I_{np} + \frac{1}{2}\hat{\Lambda} + \frac{3}{2}\hat{\Lambda}^2.$$

It is straightforward to verify that these matrices can be written as the claimed matrices D and L . □

Lemma 3.8 is now proven as a perturbation of the case above.

Proof of Lemma 3.8. From Lemma 3.7, we know that $L_{ij} = L_{ii} + \Lambda_{ij}$. Lemma C.1 therefore gives

$$Q_j^\top L_{ij} Q_i = (Q_j - Q_i)^\top L_{ij} Q_i + Q_i^\top L_{ij} Q_i = (Q_j - Q_i)^\top L_{ij} Q_i + L + Q_i^\top \Lambda_{ij} Q_i.$$

Taking norms and recalling that $\delta_{ij} = \|Q_j - Q_i\|_2$, we obtain

$$\|Q_j^\top L_{ij} Q_i - L\|_2 \leq \delta_{ij} (\|L_{ii}\|_2 + \|\Lambda_{ij}\|_2) + \|\Lambda_{ij}\|_2.$$

Since $\|L_{ii}\|_2 = \|L\|_2 \leq 2$ by Lemma C.1, this shows (12). The bound (11) is similarly proven. □

References

- [1] Absil, P.-A. and Mahony, R. and Sepulchre, R. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, Princeton, NJ, 2008.
- [2] D. Amsallem and C. Farhat. An Online Method for Interpolating Linear Parametric Reduced-Order Models. *SIAM Journal on Scientific Computing*, 33(5):2169–2198, 2011.
- [3] M. Belkin. *Problems of Learning on Manifolds*. PhD thesis, The University of Chicago, 2003.

- [4] P. Benner, S. Gugercin, and K. Willcox. A Survey of Projection-Based Model Reduction Methods for Parametric Dynamical Systems. *SIAM Review*, 57(4):483–531, 2015.
- [5] H. E. Cetingul and R. Vidal. Intrinsic mean shift for clustering on Stiefel and Grassmann manifolds. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1896–1902, 2009.
- [6] 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.
- [7] D. G. Feingold and R. S. Varga. Block diagonally dominant matrices and generalizations of the Gerschgorin circle theorem. *Pacific J. Math.*, 12:1241–1250, 1962.
- [8] M. J. C. Gover. The Eigenproblem of a Tridiagonal 2-Toeplitz Matrix. *Linear Algebra and its Applications*, 197-198:63 – 78, 1994.
- [9] L. Grippo and M. Sciandrone. On the convergence of the block nonlinear Gauss–Seidel method under convex constraints. *Operations Research Letters*, 26(3):127 – 136, 2000.
- [10] W. Hackbusch. *Iterative Solution of Large Sparse Systems of Equations*. Springer, 2016.
- [11] H. V. Henderson and S. R. Searle. The vec-permutation matrix, the vec operator and Kronecker products: a review. *Linear and Multilinear Algebra*, 9(4):271–288, 1981.
- [12] H. Karcher. Riemannian center of mass and mollifier smoothing. *Comm. Pure Appl. Math.*, 30(5):509–541, September 1977.
- [13] C. Y. Kaya and J. L. Noakes. Geodesics and an optimal control algorithm. *Proceedings of the 36th IEEE Conference on Decision and Control (CDC)*, pages 4918–4919, 1997.
- [14] C. Y. Kaya and J. L. Noakes. Leapfrog for Optimal Control. *SIAM Journal on Numerical Analysis*, 46(6):2795–2817, 2008.
- [15] S. Kobayashi and K. Nomizu. *Foundations of Differential Geometry*. Wiley, 1969.
- [16] J. W. Milnor. *Morse Theory*, volume 51 of *Annals of Math. Studies*. Princeton University Press, 1963.
- [17] J. L. Noakes. A global algorithm for geodesics. *Journal of the Australian Mathematical Society. Series A. Pure Mathematics and Statistics*, 65(1):37–50, 1998.
- [18] J. Ortega and W. Rheinboldt. *Iterative Solution of Nonlinear Equations in Several Variables*. Society for Industrial and Applied Mathematics, 2000.
- [19] Q. Rentmeesters. *Algorithms for data fitting on some common homogeneous spaces*. PhD thesis, Université catholique de Louvain, Louvain, Belgium, 2013.
- [20] A. Srivastava and E. P. Klassen. *Functional and Shape Data Analysis*. Springer series in Statistics. Springer, 2016.
- [21] Stewart, G. W. and Sun, Ji-guang. *Matrix Perturbation Theory*. Academic Press, 1990.
- [22] Ganesh Sundaramoorthi, Andrea Mennucci, Stefano Soatto, and Anthony Yezzi. A New Geometric Metric in the Space of Curves, and Applications to Tracking Deforming Objects by Prediction and Filtering. *SIAM Journal on Imaging Sciences*, 4(1):109–145, 2011.
- [23] C. Tretter. *Spectral Theory of Block Operator Matrices and Applications*. Imperial College Press, 2008.
- [24] P. Turaga, A. Veeraraghavan, A. Srivastava, and R. Chellappa. Statistical Computations on Grassmann and Stiefel Manifolds for Image and Video-Based Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(11):2273–2286, 2011.
- [25] M. Yin, J. Gao, and Y. Guo. Nonlinear low-rank representation on Stiefel manifolds. *Electronics Letters*, 51(10):749–751, 2015.
- [26] R. Zimmermann. A Matrix-Algebraic Algorithm for the Riemannian Logarithm on the Stiefel Manifold under the Canonical Metric. *SIAM Journal on Matrix Analysis and Applications*, 38(2):322–342, 2017.
- [27] Ralf Zimmermann and Kristian Debrabant. Parametric Model Reduction via Interpolating Orthonormal Bases. In Florin Adrian Radu, Kundan Kumar, Inga Berre, Jan Martin Nordbotten, and Iuliu Sorin Pop, editors, *Numerical Mathematics and Advanced Applications ENUMATH 2017*, pages 683–691, Cham, 2019. Springer International Publishing.