

RIEMANNIAN OPTIMIZATION ON THE SYMPLECTIC STIEFEL MANIFOLD*

BIN GAO[†], NGUYEN THANH SON[‡], P.-A. ABSIL[†], AND TATJANA STYKEL[§]

Abstract. The symplectic Stiefel manifold, denoted by $\text{Sp}(2p, 2n)$, is the set of linear symplectic maps between the standard symplectic spaces \mathbb{R}^{2p} and \mathbb{R}^{2n} . When $p = n$, it reduces to the well-known set of $2n \times 2n$ symplectic matrices. Optimization problems on $\text{Sp}(2p, 2n)$ find applications in various areas, such as optics, quantum physics, numerical linear algebra and model order reduction of dynamical systems. The purpose of this paper is to propose and analyze gradient-descent methods on $\text{Sp}(2p, 2n)$, where the notion of gradient stems from a Riemannian metric. We consider a novel Riemannian metric on $\text{Sp}(2p, 2n)$ akin to the canonical metric of the (standard) Stiefel manifold. In order to perform a feasible step along the antigradient, we develop two types of search strategies: one is based on quasi-geodesic curves, and the other one on the symplectic Cayley transform. The resulting optimization algorithms are proved to converge globally to critical points of the objective function. Numerical experiments illustrate the efficiency of the proposed methods.

Key words. Riemannian optimization, symplectic Stiefel manifold, quasi-geodesic, Cayley transform

AMS subject classifications. 65K05, 70G45, 90C48

1. Introduction. We consider the following optimization problem with symplectic constraints:

$$(1.1) \quad \begin{aligned} \min_{X \in \mathbb{R}^{2n \times 2p}} \quad & f(X) \\ \text{s. t.} \quad & X^\top J_{2n} X = J_{2p}, \end{aligned}$$

where $p \leq n$, $J_{2m} = \begin{bmatrix} 0 & I_m \\ -I_m & 0 \end{bmatrix}$, and I_m is the $m \times m$ identity matrix for any positive integer m . When there is no confusion, we omit the subscript of J_{2m} and I_m for simplicity. We assume that the objective function $f : \mathbb{R}^{2n \times 2p} \rightarrow \mathbb{R}$ is continuously differentiable. The feasible set of problem (1.1) is known as the *symplectic Stiefel manifold* [35, 4]

$$\text{Sp}(2p, 2n) := \{X \in \mathbb{R}^{2n \times 2p} : X^\top J_{2n} X = J_{2p}\}.$$

Whereas the term usually refers to the $p = n$ case, we call a matrix X *symplectic* whenever $X \in \text{Sp}(2p, 2n)$, as [33] does. When $p = n$, the symplectic Stiefel manifold $\text{Sp}(2p, 2n)$ becomes a matrix Lie group, termed the *symplectic group* and denoted by $\text{Sp}(2n)$.

Note that the name “symplectic Stiefel manifold” is also found in the literature [26] to denote the set of orthogonal frames in the quaternion n -space. This definition differs fundamentally from $\text{Sp}(2p, 2n)$ since the latter is noncompact, as we will see in section 3.

Optimization with symplectic constraints appears in various fields. In the case $p = n$, i.e., the symplectic group, most applications are found in physics. In the study of optical systems, such as human eyeballs [24, 21], the problem of averaging optical transference matrices can be formulated as an optimization problem with symplectic constraints. Another application can be found in accelerator design and charged-particle beam transport [18], where symplectic constraints are required for characterizing beam dynamics. A recent application is also found in the optimal control of quantum symplectic gates [39], where one has to find a (square) symplectic matrix such that the distance from this sought matrix to a

*Submitted to the editors June 26, 2020.

Funding: This work was supported by the Fonds de la Recherche Scientifique – FNRS and the Fonds Wetenschappelijk Onderzoek – Vlaanderen under EOS Project no. 30468160.

[†]ICTEAM Institute, UCLouvain, Louvain-la-Neuve, Belgium (gaobin@lsec.cc.ac.cn, pa.absil@uclouvain.be).

[‡]ICTEAM Institute, UCLouvain, Louvain-la-Neuve, Belgium; Department of Mathematics and Informatics, Thai Nguyen University of Sciences, Thai Nguyen, Vietnam (thanh.son.nguyen@uclouvain.be).

[§]Institute of Mathematics, University of Augsburg, Augsburg, Germany (tatjana.stykel@math.uni-augsburg.de).

given symplectic matrix is minimal. It is indeed an optimization problem on $\text{Sp}(2n)$, where n corresponds to the number of quantum observables.

Furthermore, several problems in scientific computing require solving (1.1) for $p < n$. For instance, [38] states that there exists a symplectic matrix that diagonalizes a given symmetric positive definite matrix, which is popularly termed the symplectic eigenvalue problem. In many cases, one is interested in computing only a few extreme eigenvalues. In [7], it is established that the sum of the p smallest symplectic eigenvalues is equal to the minimal value of an optimization problem on $\text{Sp}(2p, 2n)$. We also mention the projection-based symplectic model order reduction problem in systems and control theory [33, 3, 14]. This problem requires to reduce the order ($p \ll n$) of a Hamiltonian system, and at the same time, preserve the Hamiltonian structure. This can only be done via finding so-called symplectic projection matrices, and it is formulated as the problem (1.1) where the objective function describes the projection error.

The symplectic constraints make the problem (1.1) out of the reach of several optimization techniques: the feasible set is nonconvex and, in contrast with the (standard) Stiefel manifold [8, Theorem 10.2], the projection onto the symplectic Stiefel manifold does not admit a known closed-form expression. It appears that all the existing methods that explicitly address (1.1) restrict either to a specific objective function or to the case $p = n$ (symplectic group). Optimality conditions of the Broyden function [13] over quadratic matrix Lie groups, which include the symplectic group, were studied in [30]. In [40], the critical landscape topology for optimization on the symplectic group was investigated, where the corresponding cost function is the Frobenius distance from a target symplectic transformation. In [20], Fiori studied the geodesics of the symplectic group under a pseudo-Riemannian metric and proposed a geodesic-based method for computing the empirical mean of a set of symplectic matrices. Follow-up work can be found in [21, 36]. More recently, Birtea et al. [10] studied the first and second order optimality conditions for optimization problems on the symplectic group. Their proposed method computes the steepest-descent direction with respect to a left-invariant metric, and adopts the symplectic Cayley transform [30, 17] as a retraction to preserve the symplectic group constraint.

In this paper, we propose Riemannian gradient methods for optimization problems on $\text{Sp}(2p, 2n)$. To this end, we first prove that $\text{Sp}(2p, 2n)$ is a closed unbounded embedded submanifold of the Euclidean space $\mathbb{R}^{2n \times 2p}$. Then, leveraging two explicit characterizations of its tangent space, we develop a class of novel canonical-like Riemannian metrics, with respect to which we obtain expressions for the normal space, the tangent and normal projections, and the Riemannian gradient.

We propose two strategies to select a search curve on $\text{Sp}(2p, 2n)$ along a given tangent direction. One is based on a quasi-geodesic curve, and it needs to compute two matrix exponentials. The other is based on the symplectic Cayley transform, which requires to solve a $2n \times 2n$ linear system. By exploiting the low-rank structure of the tangent vectors, we construct a numerically efficient update for the symplectic Cayley transform. In addition, we find that the Cayley transform can be interpreted as an instance of the trapezoidal rule for solving ODEs on quadratic Lie groups.

We develop and analyze Riemannian gradient algorithms that combine each of the two curve selection strategies with a non-monotone line search scheme. We prove that the accumulation points of the sequences of iterates produced by the proposed algorithms are critical points of (1.1).

Note that all these results subsume the case of the symplectic group ($p = n$). Along the way, we extend the convergence analysis of general non-monotone Riemannian gradient methods to the case of retractions that are not globally defined; see Theorem 5.6.

Numerical experiments investigate the impact of various algorithmic choices—curve se-

lection strategy, Riemannian metric, line search scheme—on the convergence and feasibility of the iterates. Moreover, tests on various instances of (1.1)—nearest symplectic matrix problem, minimization of the Brockett cost function, and symplectic eigenvalue problem—illustrate the efficiency of the proposed algorithms.

The paper is organized as follows. In section 3, we study the geometric structure of $\text{Sp}(2p, 2n)$. The Riemannian geometry of $\text{Sp}(2p, 2n)$ endowed with the canonical-like metric is investigated in section 4. We construct two different curve selection strategies and propose a Riemannian gradient framework with a non-monotone line search scheme in section 5. Numerical results on several problems are reported in section 6. Conclusions are drawn in section 7.

2. Notation. The Euclidean inner product of two matrices $X, Y \in \mathbb{R}^{2n \times 2p}$ is denoted by $\langle X, Y \rangle = \text{tr}(X^\top Y)$, where $\text{tr}(\cdot)$ denotes the trace of the matrix argument. The Frobenius norm of X is denoted by $\|X\|_F := \sqrt{\langle X, X \rangle}$. Given $A \in \mathbb{R}^{m \times m}$, e^A or $\exp(A)$ represents the matrix exponential of A . Moreover,

$$\text{sym}(A) := \frac{1}{2}(A + A^\top) \quad \text{and} \quad \text{skew}(A) := \frac{1}{2}(A - A^\top)$$

stand for the symmetric part and the skew-symmetric part of A , respectively, and $\det(A)$ denotes the determinant of A . We let $\text{diag}(v) \in \mathbb{R}^{m \times m}$ denote the diagonal matrix with the components of $v \in \mathbb{R}^m$ on the diagonal. We use $\text{span}(A)$ to express the subspace spanned by the columns of A . Furthermore, $\mathcal{S}_{\text{sym}}(n)$ and $\mathcal{S}_{\text{skew}}(n)$ denote the sets of all symmetric and skew-symmetric $n \times n$ matrices, respectively. Let \mathcal{E}_1 and \mathcal{E}_2 be two finite-dimensional vector spaces over \mathbb{R} . The Fréchet derivative of a map $F : \mathcal{E}_1 \rightarrow \mathcal{E}_2$ at $X \in \mathcal{E}_1$ is the linear operator

$$DF(X) : \mathcal{E}_1 \rightarrow \mathcal{E}_2 : Z \mapsto DF(X)[Z]$$

satisfying $F(X + Z) = F(X) + DF(X)[Z] + o(\|Z\|)$. The rank of F at $X \in \mathcal{E}_1$, denoted by $\text{rank}(F)(X)$, is the dimension of the range of $DF(X)$. The domain of F is denoted by $\text{dom}(F)$.

3. The symplectic Stiefel manifold. In this section, we investigate the submanifold structure of $\text{Sp}(2p, 2n)$.

The matrix J defined in section 1 satisfies the following properties

$$J^\top = -J, \quad J^\top J = I, \quad J^2 = -I, \quad J^{-1} = J^\top,$$

which imply that J is skew-symmetric and orthogonal. Table 1 collects the notation and definition of several matrix manifolds that appear in this paper.

TABLE 1
Notation for matrix manifolds

SPACE	SYMBOL	ELEMENT
Orthogonal group	$\mathcal{O}(n)$	$Q \in \mathbb{R}^{n \times n} : Q^\top Q = I_n$
Stiefel manifold	$\text{St}(p, n)$	$V \in \mathbb{R}^{n \times p} : V^\top V = I_p$
Symplectic group	$\text{Sp}(2n)$	$U \in \mathbb{R}^{2n \times 2n} : U^\top J_{2n} U = J_{2n}$
Symplectic Stiefel manifold	$\text{Sp}(2p, 2n)$	$X \in \mathbb{R}^{2n \times 2p} : X^\top J_{2n} X = J_{2p}$

First we show that $\text{Sp}(2p, 2n)$ is an embedded submanifold of the Euclidean space $\mathbb{R}^{2n \times 2p}$.

PROPOSITION 3.1. *The symplectic Stiefel manifold $\text{Sp}(2p, 2n)$ is a closed embedded submanifold of the Euclidean space $\mathbb{R}^{2n \times 2p}$. Moreover, it has dimension $4np - p(2p - 1)$.*

Proof. Consider the map

$$(3.1) \quad F : \mathbb{R}^{2n \times 2p} \rightarrow \mathcal{S}_{\text{skew}}(2p) : X \mapsto X^\top JX - J.$$

We have

$$(3.2) \quad \text{Sp}(2p, 2n) = F^{-1}(0),$$

which implies that $\text{Sp}(2p, 2n)$ is closed since it is the inverse image of the closed set $\{0\}$ under the continuous map F .

Next, we prove that the rank of F is $p(2p - 1)$ at every point of $\text{Sp}(2p, 2n)$. Let $X \in \text{Sp}(2p, 2n)$. It is sufficient to show that $DF(X)$ is a surjection, i.e., for all $\bar{Z} \in \mathcal{S}_{\text{skew}}(2p)$, there exists $Z \in \mathbb{R}^{2n \times 2p}$ such that $DF(X)[Z] = \bar{Z}$. We have

$$(3.3) \quad DF(X)[Z] = X^\top JZ + Z^\top JX$$

for all $Z \in \mathbb{R}^{2n \times 2p}$. Let $\bar{Z} \in \mathcal{S}_{\text{skew}}(2p)$. Then substituting $Z = \frac{1}{2}XJ^\top \bar{Z}$ into the above equation, we obtain $DF(X)[\frac{1}{2}XJ^\top \bar{Z}] = \frac{1}{2}X^\top JXJ^\top \bar{Z} + \frac{1}{2}\bar{Z}^\top JX^\top JX = \frac{1}{2}\bar{Z} - \frac{1}{2}\bar{Z}^\top = \bar{Z}$. Hence F has full rank, namely, $\text{rank}(F)(X) = \dim(\mathcal{S}_{\text{skew}}(2p)) = p(2p - 1)$. Using (3.2) and the submersion theorem [1, Proposition 3.3.3], it follows that $\text{Sp}(2p, 2n)$ is a closed embedded submanifold of $\mathbb{R}^{2n \times 2p}$. Its dimension is $\dim(\text{Sp}(2p, 2n)) = \dim(F^{-1}(0)) = \dim(\mathbb{R}^{2n \times 2p}) - \dim(\mathcal{S}_{\text{skew}}(2p)) = 4np - p(2p - 1)$. \square

Observe that the dimension of the symplectic Stiefel manifold $\text{Sp}(2p, 2n)$ is larger than the dimension of Stiefel manifold $\text{St}(2p, 2n)$, which is equal to $4np - p(2p + 1)$. Another essential difference between $\text{Sp}(2p, 2n)$ and $\text{St}(2p, 2n)$ is that the symplectic Stiefel manifold is unbounded, hence noncompact. We show it for the simplest case, $p = n = 1$. For $X = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \in \mathbb{R}^{2 \times 2}$, we readily obtain that $X \in \text{Sp}(2)$ if and only if $ad - bc = 1$. Hence,

$$\text{Sp}(2) = \left\{ \begin{bmatrix} a & b \\ c & d \end{bmatrix} \in \mathbb{R}^{2 \times 2} : ad - bc = 1 \right\}$$

and it has dimension 3. In particular, the matrix $\begin{bmatrix} a & 0 \\ 0 & 1/a \end{bmatrix}$ is symplectic for all $a \in \mathbb{R} \setminus \{0\}$, which implies that $\text{Sp}(2)$ is unbounded. On the other hand, the orthogonal group

$$\mathcal{O}(2) = \left\{ \begin{bmatrix} \sin \theta & \cos \theta \\ \cos \theta & -\sin \theta \end{bmatrix} : \theta \in [0, 2\pi] \right\}$$

has dimension 1 and is compact.

We now set the scene for the descriptions of the tangent space that will come in Proposition 3.3. Given $X \in \text{Sp}(2p, 2n)$, we let $X_\perp \in \mathbb{R}^{2n \times (2n-2p)}$ be a full rank matrix such that $\text{span}(X_\perp)$ is the orthogonal complement of $\text{span}(X)$, and we let

$$(3.4) \quad E := [XJ \quad JX_\perp] \in \mathbb{R}^{2n \times 2n}.$$

Note that X_\perp is not assumed to be an orthonormal matrix. The next lemma gathers basic linear algebra results that will be useful later on.

LEMMA 3.2. *The matrix $E = [XJ \quad JX_\perp]$ defined in (3.4) has the following properties.*

- (i) E is invertible;
- (ii) $E^\top J E = \begin{bmatrix} J & 0 \\ 0 & X_\perp^\top J X_\perp \end{bmatrix}$ and $X_\perp^\top J X_\perp$ is invertible;

$$(iii) E^{-1} = \begin{bmatrix} X^\top J^\top \\ (X_\perp^\top J X_\perp)^{-1} X_\perp^\top \end{bmatrix};$$

(iv) Every matrix $Z \in \mathbb{R}^{2n \times 2p}$ can be represented as $Z = E \begin{bmatrix} W \\ K \end{bmatrix}$, i.e.,

$$(3.5) \quad Z = X J W + J X_\perp K,$$

where $W \in \mathbb{R}^{2p \times 2p}$ and $K \in \mathbb{R}^{(2n-2p) \times 2p}$. Moreover, we have

$$(3.6) \quad W = X^\top J^\top Z, \quad K = (X_\perp^\top J X_\perp)^{-1} X_\perp^\top Z.$$

Proof. (i) Suppose $E \begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = 0$. Multiplying this equation from the left by $X^\top J$ yields $y_1 = 0$. The equation thus reduces to $J X_\perp y_2 = 0$, which yields $y_2 = 0$ since $J X_\perp$ has full column rank. This shows that E has full rank.

(ii) By using $X^\top J X = J$ and $X^\top X_\perp = 0$, we have $E^\top J E = \begin{bmatrix} J & 0 \\ 0 & X_\perp^\top J X_\perp \end{bmatrix}$. From (i), we know that E is invertible, hence $E^\top J E$ is invertible, and so is $X_\perp^\top J X_\perp$.

(iii) From (ii), we have $E^{-1} = \begin{bmatrix} J & 0 \\ 0 & X_\perp^\top J X_\perp \end{bmatrix}^{-1} E^\top J$, and the result follows.

(iv) The first claim follows from the invertibility of E . Using (iii), we have

$$\begin{bmatrix} W \\ K \end{bmatrix} = E^{-1} Z = \begin{bmatrix} X^\top J^\top Z \\ (X_\perp^\top J X_\perp)^{-1} X_\perp^\top Z \end{bmatrix}. \quad \square$$

Given $X \in \text{Sp}(2p, 2n)$, there are infinitely many possible choices of X_\perp . The choice of X_\perp affects E in (3.4) and K in the decomposition (3.5) of Z . However, it does not affect $J X_\perp K$. In fact, it follows from (iii) in Lemma 3.2 that

$$(3.7) \quad I = E E^{-1} = X J X^\top J^\top + J X_\perp (X_\perp^\top J X_\perp)^{-1} X_\perp^\top,$$

which further implies that $J X_\perp K = J X_\perp (X_\perp^\top J X_\perp)^{-1} X_\perp^\top Z = (I - X J X^\top J^\top) Z$, where we used the expression of K in (3.6).

The tangent space of the symplectic Stiefel manifold at $X \in \text{Sp}(2p, 2n)$, denoted by $T_X \text{Sp}(2p, 2n)$, is defined by

$$T_X \text{Sp}(2p, 2n) := \{\gamma'(0) : \gamma(t) \text{ is a smooth curve in } \text{Sp}(2p, 2n) \text{ with } \gamma(0) = X\}.$$

The next result gives an implicit form and two explicit forms of $T_X \text{Sp}(2p, 2n)$.

PROPOSITION 3.3. *Given $X \in \text{Sp}(2p, 2n)$, the tangent space of $\text{Sp}(2p, 2n)$ at X admits the following expressions*

$$(3.8a) \quad T_X \text{Sp}(2p, 2n) = \{Z \in \mathbb{R}^{2n \times 2p} : Z^\top J X + X^\top J Z = 0\}$$

$$(3.8b) \quad = \{X J W + J X_\perp K : W \in \mathcal{S}_{\text{sym}}(2p), K \in \mathbb{R}^{(2n-2p) \times 2p}\}$$

$$(3.8c) \quad = \{S J X : S \in \mathcal{S}_{\text{sym}}(2n)\}.$$

Proof. Let F be as in (3.1). According to (3.3), we observe that the right-hand side of (3.8a) is the null space of $DF(X)$, namely, $\{Z \in \mathbb{R}^{2n \times 2p} : DF(X)[Z] = 0\}$. It follows from [1, (3.19)] that it coincides with $T_X \text{Sp}(2p, 2n)$.

Using (3.5), relation (3.8a) is equivalent to $W = W^\top$, which yields (3.8b).

Finally, we prove that the form (3.8a) is equivalent to (3.8c). We readily obtain

$$(3.9) \quad \{S J X : S \in \mathcal{S}_{\text{sym}}(2n)\} \subset \{Z \in \mathbb{R}^{2n \times 2p} : Z^\top J X + X^\top J Z = 0\}.$$

Since both sets are linear subspaces of $\mathbb{R}^{2n \times 2p}$, it remains to show that they have the same dimension in order to conclude that they are equal. The dimension of the right-hand set is the

dimension of $\mathrm{T}_X \mathrm{Sp}(2p, 2n)$, which is the dimension of $\mathrm{Sp}(2p, 2n)$, i.e., $4np - p(2p - 1)$; see Proposition 3.1. As for the dimension of the left-hand set, choose X_\perp as above and observe that $P := JE \begin{bmatrix} J^\top & 0 \\ 0 & -I \end{bmatrix} = [JX \ X_\perp] \in \mathbb{R}^{2n \times 2n}$ is invertible. Let $B := P^\top SP = \begin{bmatrix} B_{11} & B_{21}^\top \\ B_{21} & B_{22} \end{bmatrix}$. Then we have

$$\begin{aligned} \dim\{SJX : S \in \mathcal{S}_{\mathrm{sym}}(2n)\} &= \dim\{P^\top SJX : S \in \mathcal{S}_{\mathrm{sym}}(2n)\} \\ &= \dim\{P^\top SP \begin{bmatrix} I_{2p} \\ 0 \end{bmatrix} : S \in \mathcal{S}_{\mathrm{sym}}(2n)\} = \dim\{B \begin{bmatrix} I_{2p} \\ 0 \end{bmatrix} : B \in \mathcal{S}_{\mathrm{sym}}(2n)\} \\ &= \dim\left\{\begin{bmatrix} B_{11} \\ B_{21} \end{bmatrix} : B_{11} \in \mathcal{S}_{\mathrm{sym}}(2p), B_{21} \in \mathbb{R}^{(2n-2p) \times 2p}\right\} \\ &= p(2p+1) + (2n-2p)2p = 4np - p(2p-1), \end{aligned}$$

yielding the equality of dimensions that concludes the proof. The first equality follows from $\dim(PS) = \dim(S)$ for any subspace $\mathcal{S} \subset \mathbb{R}^{2n \times 2p}$ and invertible matrix P . The second equality follows from $JX = P \begin{bmatrix} I_{2p} \\ 0 \end{bmatrix}$. The third equality is due to $P^\top \mathcal{S}_{\mathrm{sym}}(2n)P = \mathcal{S}_{\mathrm{sym}}(2n)$, and the last equality comes from the fact $\dim(\mathcal{S}_{\mathrm{sym}}(2p)) = p(2p+1)$. \square

In fact, it follows from the proof above that the derivation of (3.8c) also works for $\{JSX : S \in \mathcal{S}_{\mathrm{sym}}(2n)\}$, namely, $\mathrm{T}_X \mathrm{Sp}(2p, 2n) = \{JSX : S \in \mathcal{S}_{\mathrm{sym}}(2n)\}$, but we will restrict to (3.8c) in the following.

Unlike (3.8b), when $p < n$, (3.8c) is an over-parameterization, i.e., $\dim(\mathcal{S}_{\mathrm{sym}}(2n)) > \dim(\mathrm{T}_X \mathrm{Sp}(2p, 2n))$. However, S can be chosen with a low-rank structure. Indeed, letting $P := JE \begin{bmatrix} J^\top & 0 \\ 0 & -I \end{bmatrix} = [JX \ X_\perp]$ and $B := P^\top SP$, we have that $SJX = P^{-\top} B P^{-1} JX = P^{-\top} B \begin{bmatrix} I_{2p} \\ 0 \end{bmatrix}$, which shows that only the first $2p$ columns of the symmetric matrix B have an impact on SJX . Hence there is no restriction on SJX , $S \in \mathcal{S}_{\mathrm{sym}}(2n)$, if we restrict B to the form $B = M \begin{bmatrix} I_{2p} & 0 \\ 0 & 0 \end{bmatrix} M^\top$, with $M \in \mathbb{R}^{2n \times 2p}$, or equivalently, if we restrict S to the form

$$(3.10) \quad S = P^{-\top} B P^{-1} = L(XJ)^\top + XJL^\top = \begin{bmatrix} L & XJ \end{bmatrix} \begin{bmatrix} (XJ)^\top \\ L^\top \end{bmatrix},$$

where $L = P^{-\top} M \in \mathbb{R}^{2n \times 2p}$ and the second equality follows from Lemma 3.2. Note that such S has rank at most $4p$. This low-rank structure will have a crucial impact on the computational cost of the Cayley retraction introduced in subsection 5.2.

4. The canonical-like metric. Given an objective function $f : \mathrm{Sp}(2p, 2n) \rightarrow \mathbb{R}$, finding the steepest-descent direction at a point $X \in \mathrm{Sp}(2p, 2n)$ amounts to finding $Z \in \mathrm{T}_X \mathrm{Sp}(2p, 2n)$ subject to $\|Z\|_X = 1$ that minimizes $\mathrm{D}f(X)[Z]$. In order to define $\|\cdot\|_X$, it is customary to endow $\mathrm{T}_X \mathrm{Sp}(2p, 2n)$ with an inner product $\langle \cdot, \cdot \rangle_X$ that depends smoothly on X ; $\|\cdot\|_X$ is then the norm induced by the inner product. Such an inner product is termed a Riemannian metric—*metric* for short—and turns $\mathrm{Sp}(2p, 2n)$ into a Riemannian manifold. In this section, we propose a class of metrics on $\mathrm{Sp}(2p, 2n)$ that are inspired from the canonical metric on the Stiefel manifold. Then we work out formulas for the normal space and gradient, and we propose a class of curves that we term quasi-geodesics.

For the Stiefel manifold $\mathrm{St}(p, n)$, any tangent vector at $V \in \mathrm{St}(p, n)$ has a unique expression $\Delta = VA + V_\perp B$, where V_\perp satisfies $V_\perp^\top V = 0$ and $V_\perp^\top V_\perp = I$, $A \in \mathcal{S}_{\mathrm{skew}}(p)$ and $B \in \mathbb{R}^{(n-p) \times p}$, see [19]. The canonical metric [19, (2.39)] on $\mathrm{St}(p, n)$ is defined as

$$g_c(\Delta_1, \Delta_2) = \mathrm{tr} \left(\Delta_1^\top \left(I - \frac{1}{2} V V^\top \right) \Delta_2 \right) = \frac{1}{2} \mathrm{tr}(A_1^\top A_2) + \mathrm{tr}(B_1^\top B_2)$$

for $\Delta_i = VA_i + V_\perp B_i$, $i = 1, 2$.

By using the tangent vector representation in (3.8b), we develop a similar metric for the symplectic Stiefel manifold $\text{Sp}(2p, 2n)$. We choose the inner product on $T_X \text{Sp}(2p, 2n)$ to be

$$(4.1) \quad g_{\rho, X_{\perp}}(Z_1, Z_2) \equiv \langle Z_1, Z_2 \rangle_X := \frac{1}{\rho} \text{tr}(W_1^{\top} W_2) + \text{tr}(K_1^{\top} K_2),$$

where $\rho > 0$ is a parameter and $W_i \in \mathcal{S}_{\text{sym}}(2p)$ and $K_i \in \mathbb{R}^{(2n-2p) \times 2p}$ are obtained from Z_i as in (3.6), i.e., $\begin{bmatrix} W_i \\ K_i \end{bmatrix} = E^{-1} Z_i$ for $i = 1, 2$. Hence, we have

$$\begin{aligned} g_{\rho, X_{\perp}}(Z_1, Z_2) &= \text{tr} \left(\begin{bmatrix} W_1 \\ K_1 \end{bmatrix}^{\top} \begin{bmatrix} \frac{1}{\rho} I & 0 \\ 0 & I \end{bmatrix} \begin{bmatrix} W_2 \\ K_2 \end{bmatrix} \right) = \text{tr} \left((E^{-1} Z_1)^{\top} \begin{bmatrix} \frac{1}{\rho} I & 0 \\ 0 & I \end{bmatrix} E^{-1} Z_2 \right) \\ &= \text{tr} \left(Z_1^{\top} E^{-\top} \begin{bmatrix} \frac{1}{\rho} I & 0 \\ 0 & I \end{bmatrix} E^{-1} Z_2 \right) = \text{tr}(Z_1^{\top} B_X Z_2), \end{aligned}$$

where

$$(4.2) \quad B_X := E^{-\top} \begin{bmatrix} \frac{1}{\rho} I & 0 \\ 0 & I \end{bmatrix} E^{-1} = \frac{1}{\rho} J X X^{\top} J^{\top} - X_{\perp} (X_{\perp}^{\top} J X_{\perp})^{-2} X_{\perp}^{\top}.$$

The last expression of B_X follows from (iii) in Lemma 3.2. In view of its definition, B_X is positive definite; this confirms that (4.1) is a bona-fide inner product. The expressions of B_X also confirm that $g_{\rho, X_{\perp}}$ depend on ρ and X_{\perp} . There is however an invariance: $g_{\rho, X_{\perp} Q} = g_{\rho, X_{\perp}}$ for all $Q \in \mathcal{O}(2n - 2p)$.

In order to make $g_{\rho, X_{\perp}}$ a bona-fide Riemannian metric, it remains to choose the “ X_{\perp} map” $\text{Sp}(2p, 2n) \ni X \mapsto X_{\perp}$ in such a way that B_X smoothly depends on X . Choosing a smooth $X \mapsto X_{\perp}$ map would be sufficient, but it is unknown whether such a smooth map globally exists. However, if orthonormalization conditions are imposed on X_{\perp} such that the X_{\perp} -term of B_X can be rewritten as a smooth expression of $X \in \text{Sp}(2p, 2n)$, then $g_{\rho, X_{\perp}}$ becomes a bona-fide Riemannian metric, which we term *canonical-like metric*. We will consider the following two such orthonormalization conditions on X_{\perp} , which are readily seen to be achievable since the set of all admissible X_{\perp} matrices has the form $\{X_{\perp} M : M \in \mathbb{R}^{(2n-2p) \times (2n-2p)} \text{ is invertible}\}$:

- (I) X_{\perp} is orthonormal.
- (II) $X_{\perp} (X_{\perp}^{\top} J X_{\perp})^{-1}$ is orthonormal.

The announced smooth expressions of B_X are given next.

PROPOSITION 4.1. *Under the orthonormalization condition (I), the X_{\perp} -term of B_X in (4.2), namely, $-X_{\perp} (X_{\perp}^{\top} J X_{\perp})^{-2} X_{\perp}^{\top}$, is equal to $-(J X J X^{\top} J^{\top} - J)^2$. Under the orthonormalization condition (II), it is equal to $I - X (X^{\top} X)^{-1} X^{\top}$.*

Proof. For the first claim, since X_{\perp} is restricted to be orthonormal, it follows that

$$X_{\perp} (X_{\perp}^{\top} J X_{\perp})^{-2} X_{\perp}^{\top} = X_{\perp} (X_{\perp}^{\top} J X_{\perp})^{-1} X_{\perp}^{\top} X_{\perp} (X_{\perp}^{\top} J X_{\perp})^{-1} X_{\perp}^{\top} = (J X J X^{\top} J^{\top} - J)^2.$$

The last equality is due to the fact $X_{\perp} (X_{\perp}^{\top} J X_{\perp})^{-1} X_{\perp}^{\top} = J X J X^{\top} J^{\top} - J$ which is readily derived from (3.7). For the second claim, since $X_{\perp} (X_{\perp}^{\top} J X_{\perp})^{-1}$ is now restricted to be orthonormal, we have that $(X_{\perp}^{\top} J X_{\perp})^{-\top} X_{\perp}^{\top} X_{\perp} (X_{\perp}^{\top} J X_{\perp})^{-1} = I$, which yields $(X_{\perp}^{\top} J X_{\perp})^2 = -X_{\perp}^{\top} X_{\perp}$. Consequently, we obtain that

$$X_{\perp} (X_{\perp}^{\top} J X_{\perp})^{-2} X_{\perp}^{\top} = -X_{\perp} (X_{\perp}^{\top} X_{\perp})^{-1} X_{\perp}^{\top} = X (X^{\top} X)^{-1} X^{\top} - I,$$

where the last equality can be checked by observing that multiplying each side by the invertible matrix $\begin{bmatrix} X & X_{\perp} \end{bmatrix}$ yields $-\begin{bmatrix} 0 & X_{\perp} \end{bmatrix}$. \square

We will use g_ρ to denote g_{ρ, X_\perp} if there is no confusion.

In the particular case $p = n$, where $\text{Sp}(2p, 2n)$ reduces to the symplectic group $\text{Sp}(2n)$, the K -term in (3.8b) disappears. Further choose $\rho = 1$. Then, for all $U \in \text{Sp}(2n)$, we have $B_U = JUU^\top J^\top$. Since $U^{-1} = JU^\top J^\top$, we also have $B_U = U^{-\top}U^{-1}$, and the canonical-like metric reduces to $\langle Z_1, Z_2 \rangle_U := \text{tr}(Z_1^\top B_U Z_2) = \langle U^{-1}Z_1, U^{-1}Z_2 \rangle$, which is the left-invariant metric used in [36, 10].

4.1. Normal space and projections. The symmetric matrix B_X in (4.2) is positive definite if and only if $X \in \mathbb{R}_*^{2n \times 2p} := \{X \in \mathbb{R}^{2n \times 2p} : \det(X^\top JX) \neq 0\}$. Hence the Riemannian metric g_ρ defined in (4.1), extended to $\mathbb{R}_*^{2n \times 2p}$, remains a Riemannian metric, which we also denote by g_ρ . In this section, we give an expression for the normal space of $\text{Sp}(2p, 2n)$ viewed as a submanifold of $(\mathbb{R}_*^{2n \times 2p}, g_\rho)$, and we obtain expression for the projections onto the tangent and normal spaces. This will be instrumental in the expression of the gradient in subsection 4.2.

The normal space at $X \in \text{Sp}(2p, 2n)$ with respect to g_ρ is defined as

$$(\text{T}_X \text{Sp}(2p, 2n))^\perp := \{N \in \mathbb{R}^{2n \times 2p} : g_\rho(N, Z) = 0 \text{ for all } Z \in \text{T}_X \text{Sp}(2p, 2n)\},$$

where we have used the fact that $\text{T}_X \mathbb{R}_*^{2n \times 2p} \simeq \text{T}_X \mathbb{R}^{2n \times 2p} \simeq \mathbb{R}^{2n \times 2p}$.

PROPOSITION 4.2. *Given $X \in \text{Sp}(2p, 2n)$, we have*

$$(4.3) \quad (\text{T}_X \text{Sp}(2p, 2n))^\perp = \{XJ\Omega : \Omega \in \mathcal{S}_{\text{skew}}(2p)\}.$$

Proof. Using Lemma 3.2 and (3.8b), we have that $N \in \mathbb{R}^{2n \times 2p}$ and $Z \in \text{T}_X \text{Sp}(2p, 2n)$ if and only if

$$\begin{aligned} N &= XJ\Omega + JX_\perp K_N, & \Omega &\in \mathbb{R}^{2p \times 2p}, & K_N &\in \mathbb{R}^{(2n-2p) \times 2p}, \\ Z &= XJW + JX_\perp K_Z, & W &\in \mathcal{S}_{\text{sym}}(2p), & K_Z &\in \mathbb{R}^{(2n-2p) \times 2p}. \end{aligned}$$

The normal space condition, $g_\rho(N, Z) = 0$ for all $Z \in \text{T}_X \text{Sp}(2p, 2n)$, is thus equivalent to

$$\frac{1}{\rho} \text{tr}(\Omega^\top W) + \text{tr}(K_N^\top K_Z) = 0 \quad \text{for all } W \in \mathcal{S}_{\text{sym}}(2p) \text{ and } K_Z \in \mathbb{R}^{(2n-2p) \times 2p}.$$

In turn, this equation is equivalent to $\Omega \in \mathcal{S}_{\text{skew}}(2p)$ (since $\mathcal{S}_{\text{skew}}(2p)$ is the orthogonal complement of $\mathcal{S}_{\text{sym}}(2p)$ with respect to the Euclidean inner product) and $K_N = 0$. \square

Every $Y \in \mathbb{R}^{2n \times 2p}$ admits a decomposition $Y = \mathcal{P}_X(Y) + \mathcal{P}_X^\perp(Y)$, where \mathcal{P}_X and \mathcal{P}_X^\perp denote the orthogonal (in the sense of g_ρ) projections onto the tangent and normal space, respectively. Next, we derive explicit expressions of \mathcal{P}_X —in the forms (3.8b) and (3.8c)—and \mathcal{P}_X^\perp .

PROPOSITION 4.3. *Given $X \in \text{Sp}(2p, 2n)$ and $Y \in \mathbb{R}^{2n \times 2p}$, the orthogonal projection onto $\text{T}_X \text{Sp}(2p, 2n)$ and $(\text{T}_X \text{Sp}(2p, 2n))^\perp$ of Y with respect to the metric g_ρ has the following expressions*

$$(4.4a) \quad \mathcal{P}_X(Y) = XJ \text{sym}(X^\top J^\top Y) + (I - XJX^\top J^\top)Y$$

$$(4.4b) \quad = S_{X,Y} JX,$$

$$(4.5) \quad \mathcal{P}_X^\perp(Y) = XJ \text{skew}(X^\top J^\top Y),$$

where $S_{X,Y} = G_X Y (XJ)^\top + XJ(G_X Y)^\top$ and $G_X = I - \frac{1}{2}XJX^\top J^\top$.

Proof. First we prove (4.4a) and (4.5). For all $Y \in \mathbb{R}^{2n \times 2p}$, according to (3.8b) and Proposition 4.2, we have

$$\begin{aligned}\mathcal{P}_X(Y) &= XJW_Y + JX_\perp K_Y, \\ \mathcal{P}_X^\perp(Y) &= XJ\Omega_Y,\end{aligned}$$

where $W_Y \in \mathcal{S}_{\text{sym}}(2p)$, $K_Y \in \mathbb{R}^{(2n-2p) \times 2p}$ and $\Omega_Y \in \mathcal{S}_{\text{skew}}(2p)$. We thus have

$$(4.6) \quad Y = \mathcal{P}_X(Y) + \mathcal{P}_X^\perp(Y) = XJW_Y + JX_\perp K_Y + XJ\Omega_Y.$$

Multiplying (4.6) from the left by $X^\top J^\top$ yields $X^\top J^\top Y = W_Y + \Omega_Y$, hence $W_Y = \text{sym}(X^\top J^\top Y)$ and $\Omega_Y = \text{skew}(X^\top J^\top Y)$. Replacing these expressions in (4.6) yields $JX_\perp K_Y = Y - XJW_Y - XJ\Omega_Y = (I - XJX^\top J^\top)Y$. We have thus obtained (4.4a) and (4.5).

It remains to prove (4.4b). A fairly straightforward development confirms that (4.4b) is equal to $Y - XJ \text{skew}(X^\top J^\top Y) = Y - \mathcal{P}_X^\perp(Y) = \mathcal{P}_X(Y)$. Instead we give a constructive proof which explains how we obtained the expression of $S_{X,Y}$. We thus seek $S_{X,Y}$ symmetric such that $S_{X,Y}JX = \mathcal{P}_X(Y)$. In view of (3.10) and Lemma 3.2, we can restrict our search to $S_{X,Y} = L(XJ)^\top + (XJ)L^\top$ with $L = XJW_L + JX_\perp K_L$, and the task thus reduces to finding $W_L \in \mathbb{R}^{2p \times 2p}$ and $K_L \in \mathbb{R}^{(2n-2p) \times 2p}$ such that

$$(4.7) \quad 2XJ \text{sym}(W_L) + JX_\perp K_L = \mathcal{P}_X(Y) = Y - \mathcal{P}_X^\perp(Y) = Y - XJ \text{skew}(X^\top J^\top Y),$$

where the last equality comes from (4.5). Multiplying (4.7) from the left by $X^\top J^\top$ yields $2 \text{sym}(W_L) = \text{sym}(X^\top J^\top Y)$, a solution of which is $W_L = \frac{1}{2}X^\top J^\top Y$. It further follows from (4.7) that $JX_\perp K_L = (I - XJX^\top J^\top)Y$. Replacing the obtained expressions in the above-mentioned decomposition of L yields $L = XJW_L + JX_\perp K_L = \frac{1}{2}XJX^\top J^\top Y + (I - XJX^\top J^\top)Y = (I - \frac{1}{2}XJX^\top J^\top)Y$. Replacing this expression in (3.10) yields the sought expression of $S_{X,Y}$. \square

Note that the projections do not depend on ρ nor X_\perp , though the metric g_ρ does. Nevertheless, by using (3.7), we obtain

$$(4.8) \quad \mathcal{P}_X(Y) = XJ \text{sym}(X^\top J^\top Y) + JX_\perp (X_\perp^\top JX_\perp)^{-1} X_\perp^\top Y.$$

Therefore, the projection $\mathcal{P}_X(Y)$ in (4.4a) can alternatively be computed by involving X_\perp .

For later reference, we record the following result which follows from (4.4b) and the fact that $\mathcal{P}_X(Z) = Z$ when $Z \in \text{T}_X \text{Sp}(2p, 2n)$. It shows how to write a tangent vector in the form (3.8c).

COROLLARY 4.4. *If $X \in \text{Sp}(2p, 2n)$ and $Z \in \text{T}_X \text{Sp}(2p, 2n)$, then $Z = S_{X,Z}JX$ with $S_{X,Z}$ as in Proposition 4.3.*

4.2. Riemannian gradient. All is now in place to specify how, given the Euclidean gradient $\nabla \bar{f}(X)$ where \bar{f} is any smooth extension of $f : \text{Sp}(2p, 2n) \rightarrow \mathbb{R}$ around X in $\mathbb{R}^{2n \times 2p}$, one can compute the steepest-descent direction of f at X in the sense of the metric g_ρ given in (4.1), as announced at the beginning of section 4.

It is well known (see, e.g., [1, §3.6]) that the steepest-descent direction is along the opposite of the Riemannian gradient. Denoted by $\text{grad}f(X)$, the Riemannian gradient of f with respect to a metric g is defined to be the unique element of $\text{T}_X \text{Sp}(2p, 2n)$ satisfying the condition

$$(4.9) \quad g(\text{grad}f(X), Z) = D\bar{f}(X)[Z] = \langle \nabla \bar{f}(X), Z \rangle \quad \text{for all } Z \in \text{T}_X \text{Sp}(2p, 2n),$$

where the last expression is the standard Euclidean inner product. (The second equality follows from the definition of Fréchet derivative in $\mathbb{R}^{2n \times 2p}$.)

PROPOSITION 4.5. *The forms (3.8b) and (3.8c) of the Riemannian gradient of a function $f : \text{Sp}(2p, 2n) \rightarrow \mathbb{R}$ with respect to the metric g_ρ defined in (4.1) are*

$$(4.10a) \quad \text{grad}_\rho f(X) = \rho X J \text{sym}(J^\top X^\top \nabla \bar{f}(X)) + J X_\perp X_\perp^\top J^\top \nabla \bar{f}(X)$$

$$(4.10b) \quad = S_X J X,$$

where $\nabla \bar{f}(X)$ is the Euclidean gradient of any smooth extension \bar{f} of f around X in $\mathbb{R}^{2n \times 2p}$, and $S_X = (H_X \nabla \bar{f}(X)) (X J)^\top + X J (H_X \nabla \bar{f}(X))^\top$ with $H_X = \frac{\rho}{2} X X^\top + J X_\perp X_\perp^\top J^\top$.

Proof. According to the definitions of the Riemannian gradient (4.9) and g_ρ , for all $Z \in \mathbb{R}^{2n \times 2p}$, it follows that $\langle B_X \text{grad}_\rho \bar{f}(X), Z \rangle = g_\rho(\text{grad}_\rho \bar{f}(X), Z) = D\bar{f}(X)[Z] = \langle \nabla \bar{f}(X), Z \rangle$. Hence, $\text{grad}_\rho \bar{f}(X) = B_X^{-1} \nabla \bar{f}(X)$. The expression of B_X in (4.2) yields $\text{grad}_\rho \bar{f}(X) = B_X^{-1} \nabla \bar{f}(X) = E \begin{bmatrix} \rho I & 0 \\ 0 & I \end{bmatrix} E^\top \nabla \bar{f}(X) = (\rho X X^\top + J X_\perp X_\perp^\top J^\top) \nabla \bar{f}(X)$. Owing to [1, (3.37)], it holds that $\text{grad}_\rho f(X) = \mathcal{P}_X(\text{grad}_\rho \bar{f}(X))$. By fairly straightforward developments, the expression (4.8) of \mathcal{P}_X yields (4.10a) while the expression (4.4b) of \mathcal{P}_X yields (4.10b). \square

Just as the Riemannian metric g_{ρ, X_\perp} in (4.1), the gradient $\text{grad}_\rho f(X)$ depends on ρ and X_\perp . The dependence on X_\perp is through the factor $J X_\perp X_\perp^\top J^\top$. However, we have seen that when we impose the orthonormalization condition (I) or (II), the metric g_{ρ, X_\perp} no longer depends on the remaining freedom in the choice of X_\perp . Hence the same can be said of $\text{grad}_\rho f(X)$. Specifically:

- (I) X_\perp is orthonormal. Then $X_\perp X_\perp^\top$ is the symmetric projection matrix onto the orthogonal complement of $\text{span}(X)$, i.e.,

$$X_\perp X_\perp^\top = I - X(X^\top X)^{-1} X^\top;$$

- (II) $X_\perp (X_\perp^\top J X_\perp)^{-1}$ is orthonormal. Then the factor

$$\begin{aligned} J X_\perp X_\perp^\top J^\top &= J X_\perp (X_\perp^\top J X_\perp)^{-\top} X_\perp^\top X_\perp (X_\perp^\top J X_\perp)^{-1} X_\perp^\top J^\top \\ &= (I - X J X^\top J^\top) (I - X J X^\top J^\top)^\top, \end{aligned}$$

where the last equality follows from (3.7).

In practice, we compute the gradient by means of above two formulations. Note that we do not need to compute X_\perp .

Even when a restriction on X_\perp makes $\text{grad}_\rho f(X)$ independent of the choice of X_\perp , it is still impacted by the choice of ρ . The influence of ρ on the algorithmic performance will be investigated in subsection 6.2.

As we have seen, in the particular case $p = n$ and $\rho = 1$, the matrix X_\perp disappears and the canonical-like metric g_ρ reduces to the left-invariant metric used in [36, 10]. In that particular case, the Riemannian gradient (4.10a) is thus also the same as that developed in [36, 10].

4.3. Quasi-geodesics. Searching along geodesics can be viewed as the most ‘‘geometrically natural’’ way to minimize a function on a manifold. In this section, we observe that it is difficult, perhaps impossible, to obtain a closed-form expression of the geodesics of $\text{Sp}(2p, 2n)$ with respect to g_ρ in (4.1). Instead, we develop an approximation, termed quasi-geodesics, defined by imposing that its *Euclidean* second derivative (in lieu of the Riemannian second derivative in $\mathbb{R}^{2n \times 2p}$ endowed with g_ρ) belongs to the normal space (4.3).

The geodesics on $(\text{Sp}(2p, 2n), g_\rho)$ are the smooth curves on $\text{Sp}(2p, 2n)$ with zero Riemannian acceleration. Since $(\text{Sp}(2p, 2n), g_\rho)$ is a Riemannian submanifold of $(\mathbb{R}^{2n \times 2p}, g_\rho)$,

it follows that the Riemannian acceleration of a curve on $(\mathrm{Sp}(2p, 2n), g_\rho)$ is the orthogonal projection (in the sense of g_ρ) onto $T_X\mathrm{Sp}(2p, 2n)$ of its Riemannian acceleration as a curve in $(\mathbb{R}^{2n \times 2p}, g_\rho)$; see, e.g., [1, Prop. 5.3.2 and (5.23)]. In other words, the geodesics on $(\mathrm{Sp}(2p, 2n), g_\rho)$ are the curves $Y(t)$ in $\mathrm{Sp}(2p, 2n)$ that satisfy

$$(4.11) \quad \frac{D}{dt}\dot{Y}(t) + Y(t)J\Omega = 0 \quad \text{with } \Omega \in \mathcal{S}_{\text{skew}}(2p),$$

where we have used the characterization (4.3) of the normal space, and where $\frac{D}{dt}$ denotes the Riemannian covariant derivative [1, (5.23)] on $(\mathbb{R}^{2n \times 2p}, g_\rho)$.

Finding a closed-form solution of (4.11) is elusive. Instead, we replace $\frac{D}{dt}\dot{Y}$ in (4.11) by the classical (Euclidean) second derivative \ddot{Y} in $\mathbb{R}^{2n \times 2p}$, yielding

$$(4.12) \quad \ddot{Y}(t) + Y(t)J\Omega = 0 \quad \text{with } \Omega \in \mathcal{S}_{\text{skew}}(2p),$$

which we term *quasi-geodesic* equation.

The purpose of the rest of this section is to obtain a closed-form solution of (4.12).

For ease of notation, we write Y instead of $Y(t)$ if there is no confusion. Since the curve $Y(t)$ remains on $\mathrm{Sp}(2p, 2n)$, we have $Y(t)^\top JY(t) = J$ for all t . Differentiating this equation twice with respect to t , we have

$$(4.13) \quad Y^\top J\dot{Y} + \dot{Y}^\top JY = 0, \quad Y^\top J\ddot{Y} + 2\dot{Y}^\top J\dot{Y} + \ddot{Y}^\top JY = 0.$$

Substituting \ddot{Y} of (4.12) into the second equation of (4.13), we arrive at $\Omega = -\dot{Y}^\top J\dot{Y}$. Therefore, (4.12) can equivalently be written as

$$(4.14) \quad \ddot{Y} - YJ(\dot{Y}^\top J\dot{Y}) = 0.$$

Following an approach similar to the one used in [19] to obtain the geodesics on the Stiefel manifold, we rewrite the differential equation (4.14) by involving three types of first integrals. Let

$$C = Y^\top JY, \quad W = Y^\top J\dot{Y}, \quad S = \dot{Y}^\top J\dot{Y}.$$

According to (4.14), it is straightforward to verify that

$$\dot{C} = W - W^\top, \quad \dot{W} = S + CJS, \quad \dot{S} = SJW - W^\top JS.$$

Since $Y \in \mathrm{Sp}(2p, 2n)$, we have that $C = J$ and therefore $\dot{C} = 0$. This in turn implies that $W = W^\top$ and $\dot{W} = 0$. As a consequence, we get $W = W(0)$. Thus, integrals of the above motion are

$$(4.15) \quad C(t) = J, \quad W(t) = W(0), \quad S(t) = e^{(JW)^\top t} S(0) e^{JWt}.$$

Since W is symmetric, the matrix JW is a Hamiltonian matrix. We will make use of the following classical result stating that the exponential of every Hamiltonian matrix is symplectic.

PROPOSITION 4.6. *For every symmetric matrix $W \in \mathbb{R}^{2p \times 2p}$, the matrix exponential e^{JW} is symplectic.*

Proof. A proof can be found, e.g., in [31, Corollary 3.1] and [16, Proposition 2.15]. Briefly, letting $U(t) = (e^{tJW})^\top J e^{tJW}$, it follows from $\frac{d}{dt}e^{At} = Ae^{At} = e^{At}A$ that $\dot{U}(t) = (e^{tJW})^\top ((JW)^\top J + J(JW))e^{tJW} = (e^{tJW})^\top (W^\top - W)e^{tJW} = 0$ for all t . Hence $U(t)$ is constant, and since $U(0) = J$, the claim follows. \square

Next, we reorganize (4.15) as an integrable equation to derive the closed-form expression of quasi-geodesics.

PROPOSITION 4.7. *Let $Y(t)$ on $\mathrm{Sp}(2p, 2n)$ satisfy the quasi-geodesic equation (4.14) with the initial conditions $Y(0) = X \in \mathrm{Sp}(2p, 2n)$ and $\dot{Y}(0) = Z \in \mathrm{T}_X \mathrm{Sp}(2p, 2n)$. Then*

$$(4.16) \quad Y(t) = Y_X^{\mathrm{qgeo}}(t; Z) := [X, Z] \exp \left(t \begin{bmatrix} -JW & JZ^\top JZ \\ I_{2p} & -JW \end{bmatrix} \right) \begin{bmatrix} I_{2p} \\ 0 \end{bmatrix} e^{JWt},$$

where $W = X^\top JZ$.

Proof. Using the notation and results introduced above, we have

$$\begin{aligned} \frac{d}{dt} (Y(t)e^{-JWt}) &= \dot{Y}(t)e^{-JWt} - Y(t)e^{-JWt} JW, \\ \frac{d}{dt} (\dot{Y}(t)e^{-JWt}) &= \ddot{Y}(t)e^{-JWt} - \dot{Y}(t)e^{-JWt} JW \\ &= Y(t)JS(t)e^{-JWt} - \dot{Y}(t)e^{-JWt} JW \\ &= Y(t)Je^{(JW)^\top t} S(0)e^{JWt}e^{-JWt} - \dot{Y}(t)e^{-JWt} JW \\ &= Y(t)e^{-JWt} JS(0) - \dot{Y}(t)e^{-JWt} JW, \end{aligned}$$

where the last equality uses the identity $Je^{(JW)^\top t} = e^{-JWt} J$ follows from the fact (Proposition 4.6) that e^{JWt} is symplectic. Stacking the above equations together yields

$$\frac{d}{dt} \begin{bmatrix} Y(t)e^{-JWt} \\ \dot{Y}(t)e^{-JWt} \end{bmatrix} = \begin{bmatrix} Y(t)e^{-JWt} \\ \dot{Y}(t)e^{-JWt} \end{bmatrix} \begin{bmatrix} -JW & JS(0) \\ I_{2p} & -JW \end{bmatrix},$$

which implies that

$$Y(t) = \begin{bmatrix} Y(0) \\ \dot{Y}(0) \end{bmatrix} \exp \left(t \begin{bmatrix} -JW & JS(0) \\ I_{2p} & -JW \end{bmatrix} \right) \begin{bmatrix} I_{2p} \\ 0 \end{bmatrix} e^{JWt}.$$

Since $Y(0) = X$, $\dot{Y}(0) = Z$, $W = W(0) = X^\top JZ$, $S(0) = Z^\top JZ$, we arrive at (4.16). \square

Note that if a tangent vector is given by $Z = XJW_Z + JX_\perp K_Z$ with $W_Z \in \mathcal{S}_{\mathrm{sym}}(2p)$ and $K_Z \in \mathbb{R}^{(2n-2p) \times 2p}$, then $W = -W_Z$ in (4.16).

In the case of the symplectic group ($p = n$), quasi-geodesics can be computed directly by solving $W = Y^\top J\dot{Y} = W(0)$, which leads to $Y_U^{\mathrm{qgeo}}(t; Z) = Ue^{-JWt}$, where $Y(0) = U \in \mathrm{Sp}(2n)$ and $W = U^\top JZ$. It turns out that the quasi-geodesic curves for $p = n$ coincide with the geodesic curves corresponding to the indefinite Khvedelidze–Mladenov metric $\langle Z_1, Z_2 \rangle_U := \mathrm{tr}(U^{-1}Z_1U^{-1}Z_2)$; see [20, Theorem 2.4].

5. Optimization on the symplectic Stiefel manifold. With the geometric tools of section 3 and section 4 at hand, almost everything is in place to apply to problem (1.1) first-order Riemannian optimization methods such as those proposed and analyzed in [1, 34, 29, 28, 11, 27]. The remaining task is to choose a mechanism which, given a search direction in the tangent space $\mathrm{T}_X \mathrm{Sp}(2p, 2n)$, returns a curve on $\mathrm{Sp}(2p, 2n)$ along that direction. In Riemannian optimization, it is customary to provide such a mechanism by choosing a retraction. In subsection 5.1 and subsection 5.2, we propose two retractions on $\mathrm{Sp}(2p, 2n)$. Then, in subsection 5.3, we work out a non-monotone gradient method for problem (1.1).

First we briefly recall the concept of (first-order) retraction on a manifold \mathcal{M} ; see [2] or [1, §4.1] for details. Let $\mathrm{T}\mathcal{M} := \bigcup_{X \in \mathcal{M}} \mathrm{T}_X \mathcal{M}$ be the tangent bundle to \mathcal{M} . A smooth mapping $\mathcal{R} : \mathrm{T}\mathcal{M} \rightarrow \mathcal{M}$ is called a *retraction* if the following properties hold for all $X \in \mathcal{M}$:

- 1) $\mathcal{R}_X(0_X) = X$, where 0_X is the origin of $\mathrm{T}_X \mathcal{M}$;
- 2) $\frac{d}{dt} \mathcal{R}_X(tZ)|_{t=0} = Z$ for all $Z \in \mathrm{T}_X \mathcal{M}$,

where \mathcal{R}_X denotes the restriction of \mathcal{R} to $T_X\mathcal{M}$. Then, given a search direction Z at a point X , the above-mentioned mechanism simply returns the curve $t \mapsto \mathcal{R}_X(tZ)$.

In most analyses available in the literature, the retraction \mathcal{R} is assumed to be globally defined, i.e., defined everywhere on the tangent bundle $T\mathcal{M}$. In [11], however, \mathcal{R}_X is only required to be defined locally, in a closed ball of radius $\varrho(X) > 0$ centered at $0_X \in T_X\mathcal{M}$, provided that $\inf_k \varrho(X^k) > 0$ where the X^k 's denote the iterates of the considered method. Henceforth, unless otherwise stated, we only assume that, for all $X \in \mathcal{M}$, \mathcal{R}_X is defined in a neighborhood of 0_X in $T_X\mathcal{M}$.

5.1. Quasi-geodesic curve. In view of the results of subsection 4.3, we define $\mathcal{R}^{\text{qgeo}}$ by

$$(5.1) \quad \mathcal{R}_X^{\text{qgeo}}(Z) := Y_X^{\text{qgeo}}(1; Z) = [X, Z] \exp \left(\begin{bmatrix} -JW & JZ^\top JZ \\ I_{2p} & -JW \end{bmatrix} \right) \begin{bmatrix} I_{2p} \\ 0 \end{bmatrix} e^{JW},$$

where $Z \in T_X\text{Sp}(2p, 2n)$ and $W = X^\top JZ$. The concept is illustrated in Figure 1. We prove next that $\mathcal{R}^{\text{qgeo}}$ is a well-defined retraction on $\text{Sp}(2p, 2n)$.




FIG. 1. Quasi-geodesic curve

LEMMA 5.1. *The map $\mathcal{R}^{\text{qgeo}} : T\text{Sp}(2p, 2n) \rightarrow \text{Sp}(2p, 2n)$ defined in (5.1) is a globally defined retraction.*

Proof. In view of (5.1), $\mathcal{R}^{\text{qgeo}}$ is well defined and smooth on $T\text{Sp}(2p, 2n)$. From the power series definition of the matrix exponential, we obtain that $Y_X^{\text{qgeo}}(1; 0_X) = X$. It also follows from this definition (or from the Baker–Campbell–Hausdorff formula) that $\frac{d}{dt} \exp(A(t))|_{t=0} = \exp(A(0))\dot{A}(0)$ when $A(0)$ and $\dot{A}(0)$ commute. This property can be exploited along with the product rule to deduce that $\frac{d}{dt} Y_X^{\text{qgeo}}(1; tZ)|_{t=0} = Z$. \square

Numerically, computing the exponential will dominate the complexity when p is relatively large; see Figure 3 for timing experiments.

5.2. Symplectic Cayley transform. In this section, we present another retraction, based on the Cayley transform. It follows naturally from the Cayley transform on quadratic Lie groups [23, Lemma 8.7] and the Cayley retraction on the Stiefel manifold [37, (7)], with the crucial help of the tangent vector representation given in Corollary 4.4.

Given $Q \in \mathbb{R}^{n \times n}$, and considering the quadratic Lie group $\mathcal{G}_Q := \{X \in \mathbb{R}^{n \times n} : X^\top Q X = Q\}$ and its Lie algebra $\mathfrak{g}_Q := \{A \in \mathbb{R}^{n \times n} : QA + A^\top Q = 0\}$, the Cayley transform is given by [23, Lemma 8.7]

$$(5.2) \quad \text{cay} : \mathfrak{g}_Q \rightarrow \mathcal{G}_Q : A \mapsto \text{cay}(A) := (I - A)^{-1}(I + A),$$

which is well defined whenever $I - A$ is invertible. As for the Cayley retraction on the Stiefel manifold, it is defined, in view of by [37, (7)], by $\mathcal{R}_X(Z) = \text{cay}(A_{X,Z})X$ where $A_{X,Z} = (I - \frac{1}{2}XX^\top)ZX^\top - XZ^\top(I - \frac{1}{2}XX^\top)$. This inspires the following definition.

DEFINITION 5.2. *The Cayley retraction on the symplectic Stiefel manifold $\text{Sp}(2p, 2n)$ is defined, for $X \in \text{Sp}(2p, 2n)$ and $Z \in \text{T}_X \text{Sp}(2p, 2n)$, by*

$$(5.3) \quad \mathcal{R}_X^{\text{cay}}(Z) := \left(I - \frac{1}{2} S_{X,Z} J \right)^{-1} \left(I + \frac{1}{2} S_{X,Z} J \right) X,$$

where $S_{X,Z}$ is as in Proposition 4.3, i.e., $S_{X,Z} = G_X Z (XJ)^\top + XJ (G_X Z)^\top$ and $G_X = I - \frac{1}{2} X J X^\top J^\top$. It is well defined whenever $I - \frac{1}{2} S_{X,Z} J$ is invertible.

In other words, the selected curve along $Z \in \text{T}_X \text{Sp}(2p, 2n)$ is

$$(5.4) \quad Y_X^{\text{cay}}(t; S) := \left(I - \frac{t}{2} S J \right)^{-1} \left(I + \frac{t}{2} S J \right) X = \text{cay} \left(\frac{t}{2} S J \right) X,$$

where $S = S_{X,Z}$ as defined above.

We confirm right away that \mathcal{R}^{cay} is indeed a retraction.

PROPOSITION 5.3. *The map $\mathcal{R}^{\text{cay}} : \text{TSp}(2p, 2n) \rightarrow \text{Sp}(2p, 2n)$ in (5.3) is a retraction.*

Proof. When $Z = 0$, we have $S_{X,Z} = 0$ and we obtain $\mathcal{R}_X^{\text{cay}}(Z) = X$, which is the first defining property of retractions. For the second property, we have $\frac{d}{dt} \mathcal{R}_X^{\text{cay}}(tZ)|_{t=0} = \frac{d}{dt} \text{cay} \left(\frac{t}{2} S_{X,Z} J \right) X|_{t=0} = D \text{cay} \left(\frac{t}{2} S_{X,Z} J \right) \left[\frac{1}{2} S_{X,Z} J \right] X|_{t=0} = S_{X,Z} J X = Z$, where the last two equalities come from $D \text{cay}(A)[\dot{A}] = 2(I - A)^{-1} \dot{A} (I + A)^{-1}$ (see [23, Lemma 8.8]) and Corollary 4.4. \square

Incidentally, we point out that the Cayley transform (5.2) can be interpreted as the trapezoidal rule for solving ODEs on quadratic Lie groups. According to [23, IV. (6.3)], notice that $\text{T}_X \mathcal{G}_Q = \{AX : A \in \mathfrak{g}_Q\}$. Hence, the following defines a differential equation on \mathcal{G}_Q :

$$(5.5) \quad \dot{Y}(t) = AY(t), \quad Y(0) = X \in \mathcal{G}_Q,$$

where $A \in \mathfrak{g}_Q$. To solve it numerically, we adopt one step of the trapezoidal rule over $[0, t]$,

$$(5.6) \quad Y(t) = X + \frac{t}{2} (AX + AY(t)).$$

Its solution is given by $Y(t) = \text{cay} \left(\frac{t}{2} A \right) X$ whenever $I - \frac{t}{2} A$ is invertible. Since $\text{cay} \left(\frac{t}{2} A \right) X \in \mathcal{G}_Q$, it follows that $Y(t) \in \mathcal{G}_Q$. This means that the trapezoidal rule for (5.5) is indeed achieved by the Cayley transform (5.2) and remains on \mathcal{G}_Q . In particular, letting $A = S_{X,Z} J \in \mathfrak{g}_{\text{Sp}(2n)}$, $Z = AX \in \text{T}_X \text{Sp}(2p, 2n)$ and $Y = \mathcal{R}_X^{\text{cay}}(Z)$, the Cayley retraction (5.3) can be exactly recovered by the same trapezoidal rule (5.6); see Figure 2 for an illustration.




FIG. 2. *Trapezoidal rule on the symplectic Stiefel manifold*

In the case of the symplectic group ($p = n$), it can be shown that \mathcal{R}^{cay} reduces to the retraction proposed in [10, Proposition 2].

The retraction $\mathcal{R}_X^{\text{cay}}(tZ)$ in (5.3) is not globally defined. It is defined if and only if $I - \frac{t}{2}S_{X,Z}J$ is invertible, i.e., $\frac{2}{t}$ is not an eigenvalue of the Hamiltonian matrix $S_{X,Z}J$. Since the eigenvalues of Hamiltonian matrices come in opposite pairs, we have that $Y_X^{\text{cay}}(t; S)$ exists for all $t \geq 0$ if and only if SJ has no real eigenvalue. This situation contrasts with the Cayley retraction on the (standard) Stiefel manifold, which is everywhere defined due to the fact that $I - A$ is invertible for all skew-symmetric A .

The fact that \mathcal{R}^{cay} is not globally defined does not preclude us from applying the convergence and complexity results of [11]. However, we have to ensure that \mathcal{R}_X is defined in a closed ball of radius $\varrho(X) > 0$ centered at $0_X \in \mathbb{T}_X \mathcal{M}$, with $\inf_k \varrho(X^k) > 0$ where the X^k 's denote the iterates of the considered method. An assumption that guarantees this condition is that (i) the objective function f in (1.1) has compact sublevel sets and (ii) the considered optimization scheme guarantees that $f(X^k) \leq f(X^0)$ for all k . Indeed, in that case, since $\{X^k\}_{k=0,1,\dots}$ remains in a compact subset of $\text{Sp}(2p, 2n)$, it is possible to find ρ such that, for all k , if $\|Z\|_{X^k} \leq \rho$, then the spectral radius of $\frac{1}{2}S_{X^k,Z}J$ is strictly smaller than one, making $I - \frac{1}{2}S_{X^k,Z}J$ invertible.

The main computational cost of $\mathcal{R}_X^{\text{cay}}(Z)$ (Definition 5.2) is the symplectic Cayley transform, which requires solving linear systems with the $2n \times 2n$ system matrix $(I - \frac{1}{2}S_{X,Z}J)$. In general, this has complexity $O(n^3)$. However, as we now show, the low-rank structure of $S_{X,Z}$ can be exploited to reduce the linear system matrix to the size $4p \times 4p$, thereby inducing a considerable reduction of computational cost when $p \ll n$. The development parallels the one given in [37, Lemma 4(1)] for the (standard) Stiefel manifold.

PROPOSITION 5.4. *Let $S = LR^\top + RL^\top = UV^\top$, where $L, R \in \mathbb{R}^{2n \times 2p}$ and $U = [L \ R] \in \mathbb{R}^{2n \times 4p}$, $V = [R \ L] \in \mathbb{R}^{2n \times 4p}$. If $I + \frac{t}{2}V^\top J^\top U$ is invertible, then (5.4) admits the expression*

$$(5.7) \quad Y_X^{\text{cay}}(t; S) = X + tU \left(I + \frac{t}{2}V^\top J^\top U \right)^{-1} V^\top JX.$$

In particular, if we choose $L = -H_X \nabla \bar{f}(X)$ and $R = XJ$, then we get $S = -S_X$ with S_X given in the gradient formula (Proposition 4.5), and we have that $\mathcal{R}_X^{\text{cay}}(-t\text{grad}_\rho f(X))$ is given by

$$(5.8) \quad Y_X^{\text{cay}}(t; -S_X) = X + t[-P_f \ XJ] \left(I + \frac{t}{2} \begin{bmatrix} E_\rho & J^\top \\ P_f^\top J^\top P_f & -E_\rho^\top \end{bmatrix} \right)^{-1} \begin{bmatrix} I \\ -E_\rho^\top J \end{bmatrix},$$

where H_X is defined in Proposition 4.5, $P_f := H_X \nabla \bar{f}(X)$, and $E_\rho := \frac{\rho}{2}X^\top \nabla \bar{f}(X)$.

Proof. By using Sherman–Morrison–Woodbury (SMW) formula [22, (2.1.4)]:

$$(A + \bar{U}\bar{V}^\top)^{-1} = A^{-1} - A^{-1}\bar{U} \left(I + \bar{V}^\top A^{-1}\bar{U} \right)^{-1} \bar{V}^\top A^{-1}$$

with $A = I$, $\bar{U} = -\frac{t}{2}U$ and $\bar{V} = J^\top V$, it follows that $(I - \frac{t}{2}SJ)^{-1} = (I - \frac{t}{2}UV^\top J)^{-1} = (A + \bar{U}\bar{V}^\top)^{-1} = I - \frac{t}{2}U \left(I + \frac{t}{2}V^\top J^\top U \right)^{-1} V^\top J^\top$. In view of (5.4), it turns out that

$$\begin{aligned} Y_X^{\text{cay}}(t; S) &= \text{cay} \left(\frac{t}{2}SJ \right) X = \left(I - \frac{t}{2}SJ \right)^{-1} \left(I + \frac{t}{2}SJ \right) X \\ &= X + \frac{t}{2}U \left(I + \left(I + \frac{t}{2}V^\top J^\top U \right)^{-1} \left(I - \frac{t}{2}V^\top J^\top U \right) \right) V^\top JX \\ &= X + tU \left(I + \frac{t}{2}V^\top J^\top U \right)^{-1} V^\top JX, \end{aligned}$$

which completes the proof of (5.7). Substituting $L = -H_X \nabla \bar{f}(X)$ and $R = XJ$ into $Y_X^{\text{cay}}(t; S)$, it is straightforward to arrive at (5.8). \square

5.3. A non-monotone line search scheme on manifolds. Now, since we can compute a gradient $\text{grad}f(X)$ and a retraction $\mathcal{R}_X(Z)$, various first-order Riemannian optimization methods can be applied to problem (1.1). Here, we adopt an approach called non-monotone line search [41]. It was extended to the Stiefel manifold in [37] and to general Riemannian manifolds in [29] and [27, §3.3]. The non-monotone approach has been observed to work well on various Riemannian optimization problems. In this section, we first present and analyze the non-monotone line search algorithm on general Riemannian manifolds endowed with a retraction that need not be globally defined. Then we apply the algorithm and its analysis to the case of the symplectic Stiefel manifold $\text{Sp}(2p, 2n)$.

Given $\beta \in (0, 1)$, a backtracking parameter $\delta \in (0, 1)$, a search direction Z^k and a trial step size $\gamma_k > 0$, the non-monotone line search procedure proposed in [27] uses the step size $t_k = \gamma_k \delta^h$, where h is the smallest integer such that

$$(5.9) \quad f(\mathcal{R}_{X^k}(t_k Z^k)) \leq c_k + \beta t_k \langle \text{grad}f(X^k), Z^k \rangle_{X^k},$$

and the next iterate is given by $X^{k+1} = \mathcal{R}_{X^k}(t_k Z^k)$. The scalar c_k in (5.9) is a convex combination of c_{k-1} and $f(X^k)$. Specifically, we set $q_0 = 1, c_0 = f(X^0)$,

$$(5.10) \quad \begin{aligned} q_k &= \alpha q_{k-1} + 1, \\ c_k &= \frac{\alpha q_{k-1}}{q_k} c_{k-1} + \frac{1}{q_k} f(X^k) \end{aligned}$$

with a parameter $\alpha \in [0, 1]$. When $\alpha = 0$, it follows that $q_k = 1$ and $c_k = f(X^k)$, and the non-monotone condition (5.9) reduces to the standard Armijo backtracking line search

$$(5.11) \quad f(\mathcal{R}_{X^k}(t_k Z^k)) \leq f(X^k) + \beta t_k \langle \text{grad}f(X^k), Z^k \rangle_{X^k}.$$

The corresponding Riemannian gradient method is presented in Algorithm 1. The search direction is chosen as the Riemannian antigradient in line 4. On the other hand, there is no restriction on the choice of the trial step size γ_k in line 5. One possible strategy is the Barzilai–Borwein (BB) method [5], which often accelerates the convergence of gradient methods when the search space is a Euclidean space. This method was extended to general Riemannian manifolds in [29]. We implement and compare several choices of γ_k ; see subsection 6.2 for details.

Algorithm 1: Riemannian gradient method with non-monotone line search

- 1 **Input:** $X^0 \in \mathcal{M}$.
 - 2 **Require:** Continuously differentiable function $f : \mathcal{M} \rightarrow \mathbb{R}$; retraction \mathcal{R} on \mathcal{M} defined on $\text{dom}(\mathcal{R})$; $\beta, \delta \in (0, 1)$, $\alpha \in [0, 1]$, $0 < \gamma_{\min} < \gamma_{\max}$, $c_0 = f(X^0)$, $q_0 = 1$, $\gamma_0 = f(X^0)$.
 - 3 **for** $k = 0, 1, 2, \dots$ **do**
 - 4 Choose $Z^k = -\text{grad}f(X^k)$.
 - 5 Choose a trial step size $\gamma_k > 0$ and set $\gamma_k = \max(\gamma_{\min}, \min(\gamma_k, \gamma_{\max}))$. Find the smallest integer h such that $\gamma_k \delta^h Z^k \in \text{dom}(\mathcal{R})$ and the non-monotone condition (5.9) hold. Set $t_k = \gamma_k \delta^h$.
 - 6 Set $X^{k+1} = \mathcal{R}_{X^k}(t_k Z^k)$.
 - 7 Update q_{k+1} and c_{k+1} by (5.10).
 - 8 **Output:** Sequence of iterates $\{X^k\}$.
-

Next we prove the convergence of Algorithm 1 on a general Riemannian manifold \mathcal{M} . Note that, in terms of convergence analysis, the only relevant difference between Algorithm 1

and [27, Algorithm 1] is that the latter assumes that the retraction \mathcal{R} is globally defined, whereas we only assume that, for all $X \in \mathcal{M}$, \mathcal{R}_X is defined in a neighborhood of 0_X in $\mathbb{T}_X\mathcal{M}$. In other words, we only assume that, for every $X \in \mathcal{M}$, there exists a ball $\mathcal{B}_{0_X}^r := \{Z \in \mathbb{T}_X\mathcal{M} : \|Z\|_X < r\} \subseteq \text{dom}(\mathcal{R})$. Hence, the convergence result in [27, Theorem 3.3] does not directly apply to our case.

First, we show that Algorithm 1 does not abort.

PROPOSITION 5.5. *Algorithm 1 generates an infinite sequence of iterates.*

Proof. Let X^k be the current iterate. In view of $Z^k = -\text{grad}f(X^k)$, by applying [41, Lemma 1.1] to the Riemannian case, it yields $f(X^k) \leq c_k$. Since \mathcal{R} is locally defined and $\delta, \beta \in (0, 1)$, it follows that

$$\begin{aligned} \lim_{h \rightarrow +\infty} \frac{f(\mathcal{R}_{X^k}(\gamma_k \delta^h Z^k)) - c_k}{\gamma_k \delta^h} &\leq \lim_{h \rightarrow +\infty} \frac{f(\mathcal{R}_{X^k}(\gamma_k \delta^h Z^k)) - f(X^k)}{\gamma_k \delta^h} \\ &= \langle \text{grad}f(X^k), Z^k \rangle_{X^k} < \beta \langle \text{grad}f(X^k), Z^k \rangle_{X^k}. \end{aligned}$$

It implies that there exists $\bar{h} \in \mathbb{N}$ such that $t_k = \gamma_k \delta^{\bar{h}} \in (0, \frac{r}{\|Z^k\|_{X^k}})$ and the non-monotone condition (5.9) hold. It means that $t_k Z^k \in \mathcal{B}_{0_{X^k}}^r \subseteq \text{dom}(\mathcal{R})$, and hence it is accepted. \square

Next, we give the proof of the convergence for Algorithm 1.

THEOREM 5.6. *Let $\{X^k\}$ be an infinite sequence of iterates generated by Algorithm 1. Then every accumulation point X^* of $\{X^k\}$ such that $0_{X^*} \in \mathbb{T}_{X^*}\mathcal{M}$ in the interior of $\text{dom}(\mathcal{R})$, is a critical point of f , i.e., $\text{grad}f(X^*) = 0$.*

Proof. We adapt the proof strategy of [27, Theorem 3.3] (that invokes [1, Theorem 4.3.1], which itself is a generalization to manifolds of the proof of [6, Proposition 1.2.1]) in order to handle the locally defined retraction. The adaptation consists in defining and making use of the neighborhood $\mathcal{N}_{0_{X^*}}$.

Since X^* is an accumulation point of $\{X^k\}$, there exists a subsequence $\{X^k\}_{k \in \mathcal{K}}$ that converges to it. In view of (5.9) and (5.10), it holds that $q_{k+1} = 1 + \alpha q_k = 1 + \alpha + \alpha^2 q_{k-1} = \dots = \sum_{i=0}^{k+1} \alpha^i \leq k + 2$. Moreover, it follows that

$$\begin{aligned} c_{k+1} - c_k &= \frac{\alpha q_k c_k + f(X^{k+1})}{q_{k+1}} - c_k = \frac{\alpha q_k c_k + f(X^{k+1})}{q_{k+1}} - \frac{\alpha q_k + 1}{q_{k+1}} c_k \\ &= \frac{f(X^{k+1}) - c_k}{q_{k+1}} \leq -\frac{\beta t_k \|\text{grad}f(X^k)\|_{X^k}^2}{q_{k+1}} < 0. \end{aligned}$$

Hence, $\{c_k\}$ is monotonically decreasing. Since $f(X^k) \leq c_k$ (in the proof of Proposition 5.5) and $\{f(X^k)\}_{k \in \mathcal{K}} \rightarrow f(X^*)$, it readily follows that $c_\infty := \lim_{k \rightarrow +\infty} c_k > -\infty$. Summing above inequalities and using $q_{k+1} \leq k + 2$, it turns out

$$\sum_{k=0}^{\infty} \frac{\beta t_k \|\text{grad}f(X^k)\|_{X^k}^2}{k+2} \leq \sum_{k=0}^{\infty} \frac{\beta t_k \|\text{grad}f(X^k)\|_{X^k}^2}{q_{k+1}} \leq \sum_{k=0}^{\infty} (c_k - c_{k+1}) = c_0 - c_\infty < \infty.$$

This inequality implies that

$$(5.12) \quad \lim_{k \rightarrow +\infty} t_k \|\text{grad}f(X^k)\|_{X^k}^2 = 0, \quad k \in \mathcal{K}.$$

For the sake of contradiction, suppose that X^* is not a critical point of f , i.e., that $\text{grad}f(X^*) \neq 0$. It readily follows from (5.12) that $\{t_k\}_{k \in \mathcal{K}} \rightarrow 0$. By the construction

of Algorithm 1, the step size t_k has the form of $t_k = \gamma_k \delta^{h_k}$ with $\gamma_k \geq \gamma_{\min} > 0$. Since 0_{X^*} is in the interior of $\text{dom}(\mathcal{R})$, there exists a neighborhood $\mathcal{N}_{0_{X^*}} := \{(X, Z) \in \text{TM} : \text{dist}^2(X, X^*) + \|Z\|_X^2 < r\} \subseteq \text{dom}(\mathcal{R})$. Due to $\{t_k\}_{k \in \mathcal{K}} \rightarrow 0$, it follows that there exists $\bar{k} \in \mathcal{K}$ and $h_k \in \mathbb{N}$ with $h_k \geq 1$ such that for all $k > \bar{k}$, $t_k = \gamma_k \delta^{h_k}$ and $(X^k, \frac{t_k}{\delta} Z^k) \in \mathcal{N}_{0_{X^*}}$. Since h_k is the smallest integer such that the non-monotone condition (5.9) holds, it follows that $\frac{t_k}{\delta} Z^k$ does not satisfy (5.9), i.e.,

$$(5.13) \quad f(X^k) - f\left(\mathcal{R}_{X^k}\left(\frac{t_k}{\delta} Z^k\right)\right) \leq c_k - f\left(\mathcal{R}_{X^k}\left(\frac{t_k}{\delta} Z^k\right)\right) < \beta \frac{t_k}{\delta} \|\text{grad}f(X^k)\|_{X^k}^2.$$

The mean value theorem for (5.13) ensures that there exists $\bar{t}_k \in [0, \frac{t_k}{\delta}]$ such that

$$D(f \circ \mathcal{R}_{X^k})(\bar{t}_k Z^k)[Z^k] < \beta \|\text{grad}f(X^k)\|_{X^k}^2, \quad \text{for all } k \in \mathcal{K}, k > \bar{k}.$$

We now take the limit in the above inequality as $k \rightarrow \infty$ over \mathcal{K} . Using the fact that $D(f \circ \mathcal{R}_{X^*})(0_{X^*}) = Df(X^*)$ in view of the defining properties of a retraction, we obtain $\|\text{grad}f(X^*)\|_{X^*}^2 \leq \beta \|\text{grad}f(X^*)\|_{X^*}^2$. Since $\beta < 1$, this is a contradiction with the supposition that $\text{grad}f(X^*) \neq 0$. \square

Algorithm 1 applies to $\mathcal{M} = \text{Sp}(2p, 2n)$ as follows. Pick a constant $\rho > 0$ and one of the orthonormalization conditions (I) or (II) for X_\perp in order to make (4.1) a bona-fide Riemannian metric. In line 4, the gradient is then as stated in Proposition 4.5. Finally, choose the retraction \mathcal{R} as either the quasi-geodesic curve (5.1) or the symplectic Cayley transform (5.3). The initialization techniques for $X^0 \in \text{Sp}(2p, 2n)$ and the stopping criterion will be discussed in subsection 6.1.

COROLLARY 5.7. *Apply Algorithm 1 to $\text{Sp}(2p, 2n)$ as specified in the previous paragraph. Let $\{X^k\}$ be an infinite sequence of iterates generated by Algorithm 1. Then every accumulation point X^* of $\{X^k\}$ is a critical point of f , i.e., $\text{grad}f(X^*) = 0$.*

Proof. First, the quasi-geodesic retraction is globally defined. Hence, Theorem 5.6 can be directly applied.

Next, we consider the Cayley-based algorithm. In order to conclude by invoking Theorem 5.6, it is sufficient to show that, for all $X \in \text{Sp}(2p, 2n)$, $0_X \in \text{T}_X \text{Sp}(2p, 2n)$ is in the interior of $\text{dom}(\mathcal{R}^{\text{cay}})$. In view of the properties of retractions, we have that $0_X \in \text{dom}(\mathcal{R}^{\text{cay}})$, and we finish the proof by showing that all points of $\text{dom}(\mathcal{R}^{\text{cay}})$ are in its interior, namely, $\text{dom}(\mathcal{R}^{\text{cay}})$ is open. To this end, let $F : \text{TSp}(2p, 2n) \rightarrow \mathbb{R} : (X, Z) \mapsto \det(I - \frac{1}{2} S_{X,Z})$, where $S_{X,Z}$ is as in Definition 5.2. Since F is continuous and $\{0\}$ is closed, it follows that $F^{-1}(0) = \{(X, Z) \in \text{TSp}(2p, 2n) : F(X, Z) = 0\}$ is a closed set. Its complement in the tangent bundle is thus open, and it is also the domain of \mathcal{R}^{cay} . \square

6. Numerical experiments. In this section, we report the numerical performance of Algorithm 1. Both methods based on the quasi-geodesics (5.1) and the symplectic Cayley transform (5.3) are evaluated. We first introduce implementation details in subsection 6.1. To determine default settings, we investigate the parameters of our proposed algorithms in subsection 6.2. Finally, the efficiency of Algorithm 1 is assessed by solving several different problems. The experiments are performed on a workstation with two Intel(R) Xeon(R) Processors Silver 4110 (at 2.10GHz \times 8, 12M Cache) and 384GB of RAM running MATLAB R2018a under Ubuntu 18.10. The code that produced the results is available from <https://github.com/opt-gaobin/spopt>.

6.1. Implementation details. As we mentioned in subsection 4.2, we propose two strategies to compute the Riemannian gradient $\text{grad}_\rho f(X)$. Both algorithms with type (I)

and type (II) perform well in our preliminary experiments. In this section, Algorithm 1 with the symplectic Cayley transform is denoted by ‘‘Sp-Cayley’’, and its instances using type (I) and (II) for the gradient are represented as ‘‘Sp-Cayley-I’’ and ‘‘Sp-Cayley-II’’, respectively. Similarly, quasi-geodesic algorithms are denoted by ‘‘Quasi-geodesics’’, ‘‘Quasi-geodesics-I’’ and ‘‘Quasi-geodesics-II’’.

We adopt formula (4.10b), i.e., $H_X \nabla \bar{f}(X)(XJ)^\top JX + XJ(H_X \nabla \bar{f}(X))^\top JX$, to assemble $\text{grad}_\rho f(X)$ for all the algorithms. In order to save flops and obtain a good feasibility (see Figure 3), we choose (5.8) to compute the Cayley retraction. Note that we keep the calculation of $(XJ)^\top JX$ in the first term of $\text{grad}_\rho f(X)$, although it is trivial that $(XJ)^\top JX = I$ for $X \in \text{Sp}(2p, 2n)$. This is due to our observation that the feasibility of Quasi-geodesics gradually degrades when we omit this calculation.

At the beginning of Algorithm 1, we need a feasible point $X^0 \in \text{Sp}(2p, 2n)$ to start the iteration. The easiest way to generate a symplectic matrix is to choose the ‘‘identity’’ matrix in $\text{Sp}(2p, 2n)$, namely, $I^0 = \begin{bmatrix} I_p & 0 & 0 & 0 \\ 0 & 0 & I_p & 0 \end{bmatrix}^\top$. Moreover, by using Proposition 4.6, we suggest the following strategies to generate an initial point:

- 1) $X^0 = I^0$;
- 2) $X^0 = I^0 e^{J(W+W^\top)}$, where W is randomly generated by $W = \text{randn}(2*p, 2*p)$;
- 3) X^0 is assembled by the first p columns and $(n+1)$ -th to $(n+p)$ -th columns of $e^{J(W+W^\top)}$, where W is randomly generated by $W = \text{randn}(2*n, 2*n)$.

The matrix exponential is computed by the function `expm` in MATLAB. Unless otherwise specified, we choose strategy 2) as our initialization.

For a stopping criterion, we check the following two conditions:

$$(6.1) \quad \|\text{grad}_\rho f(X^k)\|_F \leq \epsilon,$$

$$(6.2) \quad \frac{\|X^k - X^{k+1}\|_F}{\sqrt{2n}} < \epsilon_x \quad \text{and} \quad \frac{|f(X^k) - f(X^{k+1})|}{|f(X^k)| + 1} < \epsilon_f$$

with given tolerances $\epsilon, \epsilon_x, \epsilon_f > 0$. We terminate the algorithms once one of the criteria (6.1)-(6.2) or a maximum iteration number `MaxIter` is reached. The default tolerance parameters are chosen as $\epsilon = 10^{-5}$, $\epsilon_x = 10^{-5}$, $\epsilon_f = 10^{-8}$ and `MaxIter` = 1000. For parameters to control the non-monotone line search, we follow the choices in the code `OptM`¹ [37], specifically, $\beta = 10^{-4}$, $\delta = 0.1$, $\alpha = 0.85$, as our default settings. In addition, we choose $\gamma_{\min} = 10^{-15}$, $\gamma_{\max} = 10^{15}$, and a trial step size γ_k as in (6.4).

6.2. Default settings of the algorithms. In this section, we study the performance and robustness of our algorithms by choosing different parameters. All the comparisons and results are based on a test problem, called the nearest symplectic matrix problem, which aims to calculate the nearest symplectic matrix to a target matrix $A \in \mathbb{R}^{2n \times 2p}$ with respect to the Frobenius norm, i.e.,

$$(6.3) \quad \min_{X \in \text{Sp}(2p, 2n)} \|X - A\|_F^2.$$

The special case of this problem on the symplectic group (i.e., $p = n$) was studied in [40]. In our experiments, A is randomly generated by $A = \text{randn}(2*n, 2*p)$, then it is scaled using $A = A / \text{norm}(A)$.

6.2.1. A comparison between quasi-geodesics and symplectic Cayley transform. In Algorithm 1, we can choose between two different retractions: quasi-geodesic and Cayley. In

¹ Available from <https://github.com/wenstone/OptM>.

order to investigate the numerical performance of the two alternatives, we solve the nearest symplectic matrix problem (6.3) of size 2000×400 and choose the parameter for the metric g_ρ to be $\rho = 1$. We stop our algorithms only when $\text{MaxIter} = 120$ is reached. The evolution of the norm of the gradient and the feasibility violation $\|X^\top JX - J\|_F$ for both algorithms is shown in Figure 3. It illustrates that Sp-Cayley performs better than Quasi-geodesics in terms of efficiency and feasibility. Therefore, we choose Sp-Cayley as our default algorithm and the following experiments will focus on Sp-Cayley.

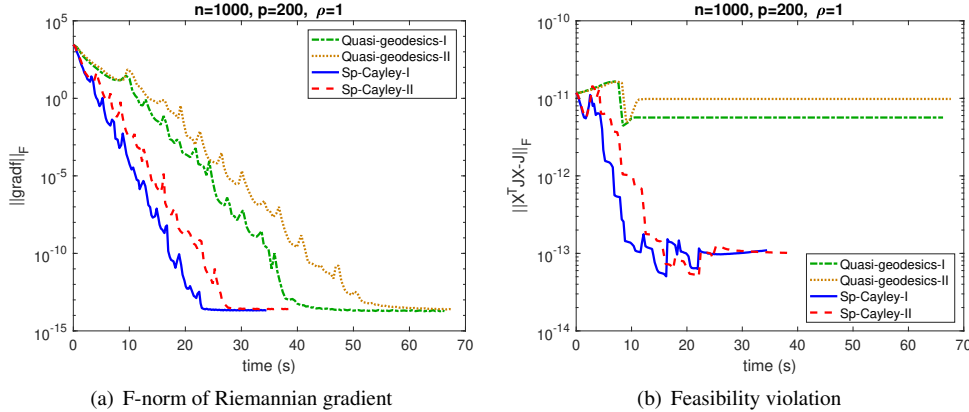


FIG. 3. A comparison of Algorithm 1 based on the quasi-geodesic curve and the symplectic Cayley transform

6.2.2. A comparison of different metrics. In this section, we compare the performance of Sp-Cayley with different metrics. Namely, we compare Sp-Cayley-I and Sp-Cayley-II for a set of parameters $\rho = 2^l$ with l chosen from $\{-3, -2, -1, 0, 1, 2, 3\}$. We run both algorithms 100 times on randomly generated nearest symplectic matrix problems of size 2000×40 . Note that for each instance, Sp-Cayley-I and Sp-Cayley-II use the same initial guess. In order to get an average performance, we stop algorithms only when $\|\text{grad}_\rho f(X^k)\|_F \leq 10^{-4}$. A summary of numerical results is reported in Figure 4. It displays average iteration numbers and the feasibility violation for the two algorithms with different ρ . We can learn from the figures that:

- The value ρ^* at which Sp-Cayley-I and Sp-Cayley-II achieve the best performance is approximately equal to $1/2$ and 1 , respectively. The difference of ρ^* is due to the different choice of X_\perp , which also has an effect on the metric g_ρ . We have observed that ρ^* may vary for different objective functions. This indicates that tuning ρ for the problem class of interest may significantly improve the performance of Algorithm 1.
- Sp-Cayley-I has a lower average iteration number than Sp-Cayley-II when $\rho \leq 1$. Over all the variants considered in Figure 4, Sp-Cayley-I with $\rho = 1/2$ has the best performance.
- Both Sp-Cayley-I and Sp-Cayley-II show a loss of feasibility when ρ becomes large. A possible reason is that the non-normalized second term of $H_X = JX_\perp X_\perp^\top J^\top + \frac{\rho}{2} X X^\top$ in the Riemannian gradient becomes dominant.

According to the above observations, we choose $\rho = 1/2$ for Sp-Cayley-I and $\rho = 1$ for Sp-Cayley-II as our default settings.

6.2.3. A comparison of different line search schemes. The non-monotone line search strategy (subsection 5.3) depends on several parameters. The purpose of this section is to

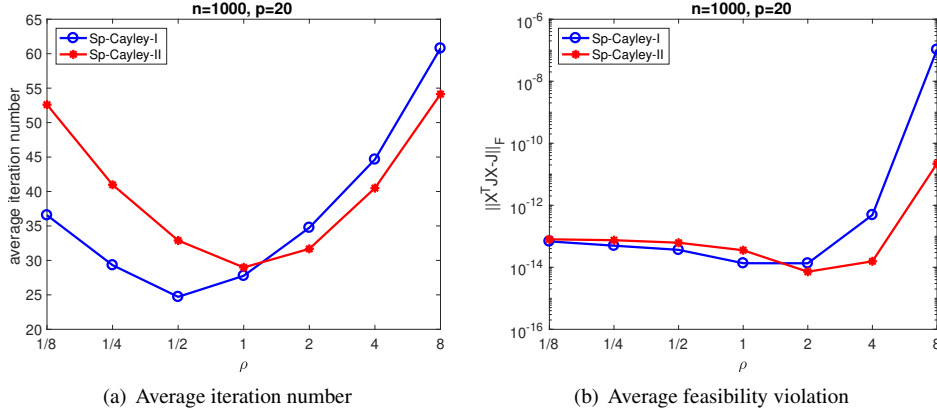


FIG. 4. A comparison of Sp-Cayley-I and Sp-Cayley-II with different parameter ρ

investigate which among those parameters have a significant impact on the performance of Algorithm 1.

First, we consider the choice of the trial step size γ_k in Algorithm 1. In our case, the ambient space is Euclidean. Thus, we can use the BB method proposed in [37] and define

$$\gamma_k^{\text{BB1}} := \frac{\langle S^{k-1}, S^{k-1} \rangle}{|\langle S^{k-1}, Y^{k-1} \rangle|}, \quad \gamma_k^{\text{BB2}} := \frac{|\langle S^{k-1}, Y^{k-1} \rangle|}{\langle Y^{k-1}, Y^{k-1} \rangle},$$

where $S^{k-1} = X^k - X^{k-1}$ and $Y^{k-1} = \text{grad}_\rho f(X^k) - \text{grad}_\rho f(X^{k-1})$. Note that this differs from the Riemannian BB method in [29] since it adopts the Euclidean inner product $\langle \cdot, \cdot \rangle$ rather than g_ρ . This choice is cheaper in flops and we have observed that it speeds up the algorithm. Owing to its efficiency, we further adopt the *alternating* BB strategy [15] and choose the trial step size as

$$(6.4) \quad \gamma_k^{\text{ABB}} := \begin{cases} \gamma_k^{\text{BB1}}, & \text{for odd } k, \\ \gamma_k^{\text{BB2}}, & \text{for even } k. \end{cases}$$

We next compare γ_k^{BB1} , γ_k^{BB2} , γ_k^{ABB} , and the step size

$$\gamma_k^{\text{M}} := 2 \left| \frac{f(X^k) - f(X^{k-1})}{\text{D}f(X)[Z]} \right|$$

proposed in [32, (3.60)], where γ_k^{M} is also used in the line search function in Manopt² [12]. In this test, we opt for the monotone line search ($\alpha = 0$) adopted with tolerances $\epsilon = 10^{-10}$, $\epsilon_x = 10^{-10}$, $\epsilon_f = 10^{-14}$. Figure 5 reveals that the BB strategies greatly improve the performance of the Riemannian gradient method, and outperform the classical initial step size γ^{M} in iteration number and function value decreasing. We have obtained similar results, omitted here, for Sp-Cayley-II. Since γ_k^{ABB} is the most efficient choice in this experiment, we employ it as our default setting henceforth.

Next we investigate the impact of the parameter α , which controls the degree of non-monotonicity. If $\alpha = 0$, then condition (5.9) reduces to the usual monotone condition (5.11). Here we scale the problem as $A=2 * A / \text{norm}(A)$ because we found that it reveals better the

²A MATLAB toolbox for optimization on manifolds (available from <https://www.manopt.org/>).

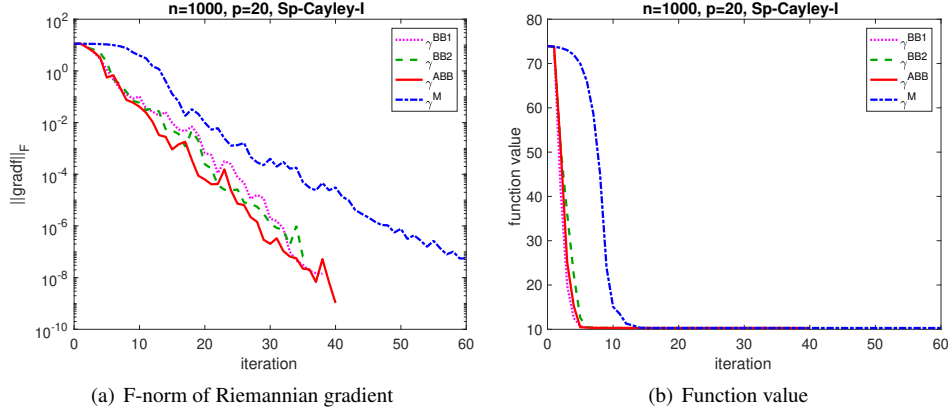


FIG. 5. A comparison of *Sp-Cayley-I* with different initial step size γ_k in the monotone line search scheme

advantage that the non-monotone approach can have. We test *Sp-Cayley-II* with $\alpha = 0, 0.85$, i.e., the monotone and non-monotone schemes. The results are shown in Figure 6. The purpose of the non-monotone strategy is to make the line search condition (5.9) more prone than the monotone strategy to accept the trial step size γ_k , here γ_k^{ABB} . We see that this results in a faster convergence in this experiment. Since the non-monotone condition (5.9) works well in our problem, thus we choose it as a default setting henceforth.

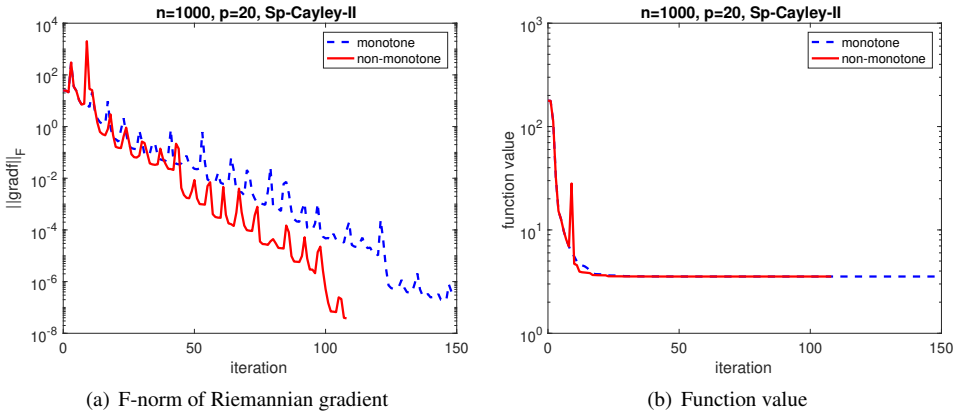


FIG. 6. A comparison of *Sp-Cayley-II* with the monotone and non-monotone line search schemes

6.3. Nearest symplectic matrix problem. In this section, we still focus on the nearest symplectic matrix problem (6.3). We first compare the algorithms *Sp-Cayley-I* and *Sp-Cayley-II* on an open matrix dataset: SuiteSparse Matrix Collection³. Due to the different size and scale of data matrices, we choose the column number p from the set $\{5, 10, 20, 40, 80\}$, and the target matrix $A \in \mathbb{R}^{2n \times 2p}$ is generated by the first $2p$ columns of an original data matrix. In order to obtain a comparable error, we normalize all matrices as $A/\|A\|_{\max}$, where $\|A\|_{\max} := \max_{i,j} |A_{ij}|$. Numerical results are presented in Table 2 for representative problem instances. Here, “fval” represents the function value, “gradf”, “feasi”, “iter”, and “time”

³Available from <https://sparse.tamu.edu/>.

stand for $\|\text{grad}_\rho f\|_F$, $\|X^\top JX - J\|_F$, the number of iterations, and the wall-clock time in seconds, respectively. From the table, we find that both algorithms perform well on most of the instances, and Sp-Cayley-I performs better than Sp-Cayley-II with fewer iteration number and fewer running time. In the largest problem “2cubes_sphere”, both methods converge and obtain comparable results for the function value and gradient error. In addition, Sp-Cayley-II diverges on the problem “msc23052” with $p = 10$, while Sp-Cayley-I converges. Therefore, we conclude that for the nearest symplectic matrix problem, Sp-Cayley-I is more robust and efficient than Sp-Cayley-II.

TABLE 2
Numerical results for the nearest symplectic matrix problem

p	Sp-Cayley-I					Sp-Cayley-II				
	fval	gradf	feasi	iter	time	fval	gradf	feasi	iter	time
2cubes_sphere, $2n = 101492$										
5	3.995e+00	8.38e-05	1.39e-14	24	2.11	3.995e+00	1.40e-04	4.53e-15	23	2.18
10	7.331e+00	2.23e-04	1.95e-14	27	4.14	7.331e+00	5.28e-04	9.13e-15	28	4.60
20	1.602e+01	2.19e-04	4.76e-14	30	8.81	1.602e+01	6.72e-04	5.26e-14	36	12.53
40	3.423e+01	2.51e-03	7.08e-14	32	23.60	3.423e+01	7.20e-04	4.58e-14	40	31.79
80	1.056e+02	7.98e-04	4.05e-13	39	55.97	1.056e+02	1.50e-03	2.81e-10	41	71.11
cvxbqp1, $2n = 50000$										
5	2.649e+00	2.03e-04	1.00e-14	22	0.78	2.649e+00	1.49e-03	3.15e-15	27	1.01
10	6.230e+00	2.16e-04	1.65e-14	26	1.54	6.230e+00	6.35e-04	6.35e-15	36	2.30
20	1.289e+01	2.86e-04	3.00e-14	26	2.79	1.289e+01	6.85e-04	1.08e-14	38	4.31
40	2.542e+01	1.33e-03	5.74e-14	28	8.19	2.542e+01	6.75e-04	3.72e-14	36	11.78
80	5.156e+01	6.38e-04	1.08e-13	26	19.81	5.156e+01	1.47e-03	7.33e-14	36	30.05
msc23052, $2n = 23052$										
5	6.204e+00	1.23e-04	2.47e-04	92	1.30	6.204e+00	7.61e-04	5.26e-06	66	1.04
10	1.361e+01	5.87e-04	7.06e-08	67	1.66	4.556e-01	2.72e-02	4.24e+00	637	17.69
20	2.732e+01	9.82e-04	1.48e-04	80	3.70	2.730e+01	1.14e-02	8.69e-03	112	5.91
40	5.721e+01	2.18e-03	1.74e-06	70	6.78	5.721e+01	4.36e-03	1.98e-03	90	9.90
80	1.302e+02	9.96e-03	8.89e-03	98	27.87	1.302e+02	1.60e-02	1.00e-02	88	28.35
Na5, $2n = 5832$										
5	2.170e+00	3.34e-04	1.29e-14	34	0.16	2.170e+00	1.05e-03	3.79e-15	35	0.17
10	4.378e+00	1.53e-04	3.62e-14	35	0.25	4.378e+00	4.34e-03	9.72e-15	41	0.29
20	9.033e+00	3.54e-04	6.24e-14	34	0.44	9.033e+00	1.46e-03	1.41e-14	37	0.53
40	1.828e+01	7.82e-04	8.58e-14	40	1.09	1.828e+01	1.75e-03	2.81e-14	44	1.24
80	3.731e+01	3.42e-04	3.32e-13	48	3.00	3.731e+01	6.74e-04	5.79e-14	56	3.92

We next compare our algorithms on randomly generated datasets. Given a set of samples $\{X_1, X_2, \dots, X_N\}$ with $X_i \in \text{Sp}(2p, 2n)$ for $i = 1, \dots, N$, the extrinsic mean problem [9, Section 3] on $\text{Sp}(2p, 2n)$ is defined as

$$(6.5) \quad \min_{X \in \text{Sp}(2p, 2n)} \frac{1}{N} \sum_{i=1}^N \|X - X_i\|_F^2.$$

In view of [9, Section 3], the solutions of (6.5) are those of (6.3) with $A = \frac{1}{N} \sum_{i=1}^N X_i$. This allows us to reuse the code that addressed (6.3). We test Sp-Cayley-I and Sp-Cayley-II for solving the problem (6.5) with three different random sample sets: (i) $N = 100, n = p = 2$; (ii) $N = 100, n = p = 10$; (iii) $N = 1000, n = 1000, p = 20$. In each dataset, samples are randomly generated around a selected center $Y^0 \in \text{Sp}(2p, 2n)$. Specifically, we choose $X_i = Y^0 e^{J(W_i + W_i^\top)}$, where $W_i = 0.1 * \text{randn}(2 * p, 2 * p)$. The initial point X^0 and the center Y^0 are calculated by the strategy 3) (subsection 6.1) for datasets (i)-(ii), and 2) for dataset (iii). In the first two sets (the symplectic group), Sp-Cayley-I and Sp-Cayley-II reduce to the same algorithm due to the same Riemannian gradient. Therefore, we omit the results of Sp-Cayley-II. We run our algorithms twice with different stopping tolerances: default and $\{\epsilon = 10^{-10}, \epsilon_x = 10^{-10}, \epsilon_f = 10^{-14}\}$. The detailed results for these two setting are presented in Table 3. It reveals that our algorithms converge for three different sample sets

with different stopping tolerances. Moreover, we show the initial and final errors of each sample in Figure 7, where X^* denotes the solution obtained by Sp-Cayley. From the figure, we observe that for both dataset (i) and (ii), the sample error greatly decreases with Sp-Cayley.

TABLE 3
Numerical results for the extrinsic mean problem on $\text{Sp}(2p, 2n)$

ϵ	Sp-Cayley-I					Sp-Cayley-II				
	fval	gradf	feasi	iter	time	fval	gradf	feasi	iter	time
dataset (i), $n = 2, p = 2, N = 100$										
1e-05	1.627e+00	3.74e-05	1.07e-14	82	0.01	-	-	-	-	-
1e-10	1.627e+00	9.14e-10	8.84e-15	158	0.01	-	-	-	-	-
dataset (ii), $n = 10, p = 10, N = 100$										
1e-05	3.068e+01	1.13e-04	8.97e-14	158	0.02	-	-	-	-	-
1e-10	3.068e+01	6.91e-09	1.12e-13	316	0.03	-	-	-	-	-
dataset (iii), $n = 1000, p = 20, N = 1000$										
1e-05	1.333e+02	1.54e-04	3.56e-13	154	0.96	1.333e+02	2.54e-04	3.38e-13	178	1.05
1e-10	1.333e+02	3.67e-08	5.32e-13	256	1.33	1.333e+02	2.83e-07	4.63e-13	262	1.50




FIG. 7. A comparison of initial and final errors for each sample

6.4. Minimization of the Brockett cost function. In [13], Brockett investigated least squares matching problems on matrix Lie groups with an objective function

$$(6.6) \quad f(X) := \text{tr}(X^\top AXN - 2BX^\top),$$

where $A, N, B \in \mathbb{R}^{n \times n}$ are given matrices. This function is widely known as the Brockett cost function. Recently, in [30], the results of [13] were extended to P -orthogonal matrices satisfying $X^\top PX = P$ with a given orthogonal matrix $P \in \mathbb{R}^{n \times n}$. For $P = J$, the problem reduces to an optimization problem on the symplectic group $\text{Sp}(2n)$. Such a problem was more recently considered in [10], where several optimization algorithms were proposed.

In this section, we study a well-defined (bounded from below) minimization problem based on the Brockett cost function (6.6) with $N = I$ and $B = 0$. Specifically, we consider the following optimization problem

$$\min_{X \in \text{Sp}(2p, 2n)} \text{tr}(X^\top AX),$$

where $A \in \mathbb{R}^{2n \times 2n}$ is a symmetric positive definite matrix. Such a problem arises, for example, in the symplectic eigenvalue problem [38, 7] which will be investigated in more detail in subsection 6.5. In this test, the matrix $A \in \mathbb{R}^{2n \times 2n}$ is randomly generated as $A = Q\Lambda Q^\top$, where $Q \in \mathbb{R}^{2n \times 2n}$ is an orthogonal matrix computed from a QR factorization $Q = \text{qr}(\text{randn}(2*n, 2*n))$, and $\Lambda \in \mathbb{R}^{2n \times 2n}$ is a diagonal matrix with diagonal elements $\Lambda_{ii} = \lambda^{1-i}$ for $i = 1, 2, \dots, 2n$. The parameter $\lambda \geq 1$ determines the decay of eigenvalues of A . The experiments are divided into two parts. At first, we solve the problem on the symplectic group, i.e., $n = p$, for matrices of different size 20×20 , 80×80 and 160×160 and different parameters $\lambda \in \{1.01, 1.04, 1.07, 1.1\}$. The numerical results are presented in Table 4. From the table, we observe that Sp-Cayley (since Sp-Cayley-I and Sp-Cayley-II reduce to the same method when $n = p$) works well on different problems. In the second part, we move to the symplectic Stiefel manifold and test our algorithms on different problems with parameters $n = 1000, 2000, 3000$, $p = 5, 10, 20, 40, 80$, and λ as above. The corresponding results are also presented in Table 4. It illustrates that for relatively large problems, Sp-Cayley-I and Sp-Cayley-II still perform well, and have the comparable function values and feasibility violations.

TABLE 4
Numerical results in the Brockett cost function minimization

		Sp-Cayley-I					Sp-Cayley-II				
		fval	gradf	feasi	iter	time	fval	gradf	feasi	iter	time
λ	$n = p = 10$										
1.01		3.642e+01	5.73e-06	3.98e-14	12	0.01	-	-	-	-	-
1.04		2.793e+01	5.75e-06	5.19e-14	14	0.01	-	-	-	-	-
1.07		2.168e+01	1.10e-05	2.20e-14	14	0.01	-	-	-	-	-
1.10		1.737e+01	6.26e-06	4.85e-14	19	0.01	-	-	-	-	-
	$n = p = 40$										
1.01		1.094e+02	4.25e-06	4.90e-13	14	0.01	-	-	-	-	-
1.04		4.197e+01	1.61e-04	6.56e-13	25	0.02	-	-	-	-	-
1.07		2.032e+01	4.69e-05	6.15e-13	46	0.04	-	-	-	-	-
1.10		1.222e+01	1.00e-04	5.28e-13	72	0.07	-	-	-	-	-
	$n = p = 80$										
1.01		1.531e+02	2.98e-05	3.30e-12	19	0.07	-	-	-	-	-
1.04		3.256e+01	1.46e-04	3.95e-12	59	0.21	-	-	-	-	-
1.07		1.421e+01	2.06e-04	2.89e-12	154	0.58	-	-	-	-	-
1.10		8.442e+00	1.90e-04	2.39e-12	252	0.94	-	-	-	-	-
p	$n = 1000, \lambda = 1.01$										
5		3.150e-04	2.10e-04	1.20e-14	182	0.90	2.015e-04	1.62e-04	1.50e-14	208	1.05
10		3.631e-04	1.61e-04	3.07e-14	293	2.05	4.028e-04	1.95e-04	4.40e-14	260	1.93
20		5.902e-04	1.85e-04	9.02e-14	362	3.18	5.350e-04	1.81e-04	1.37e-13	280	2.97
40		7.764e-04	1.61e-04	3.82e-13	484	7.06	6.669e-04	1.52e-04	7.34e-13	548	9.77
80		1.037e-03	5.72e-04	1.65e-12	649	24.33	1.094e-03	1.72e-04	3.02e-12	619	21.14
	$n = 2000, \lambda = 1.04$										
10		1.241e-04	1.77e-04	2.95e-14	146	5.15	1.105e-04	1.85e-04	3.08e-14	174	6.22
20		1.748e-04	1.84e-04	2.12e-13	164	7.06	1.671e-04	2.06e-04	1.99e-13	234	10.61
40		2.216e-04	1.77e-04	3.37e-13	220	11.93	1.499e-04	1.36e-04	3.13e-13	270	17.54
80		2.948e-04	1.69e-04	1.91e-12	303	27.72	2.031e-04	1.47e-04	1.75e-12	273	24.48
	$n = 3000, \lambda = 1.10$										
5		1.794e-05	7.72e-04	9.63e-15	87	5.65	3.433e-05	1.70e-04	7.48e-15	74	4.98
10		5.133e-05	1.55e-04	5.01e-14	110	8.19	3.337e-05	2.34e-04	4.91e-14	125	8.65
20		7.352e-05	1.49e-04	1.58e-13	141	12.33	4.253e-05	1.32e-04	1.53e-13	142	12.41
40		1.226e-04	1.80e-04	4.43e-13	176	20.75	1.085e-04	1.91e-04	4.14e-13	134	16.43
80		1.661e-04	1.70e-04	3.96e-12	180	31.62	1.637e-04	1.91e-04	3.48e-12	188	35.02

6.5. The symplectic eigenvalue problem. It was shown in [38] that for every symmetric positive definite matrix $M \in \mathbb{R}^{2n \times 2n}$, there exists $X \in \text{Sp}(2n)$ such that

$$(6.7) \quad X^\top M X = \begin{bmatrix} D & \\ & D \end{bmatrix},$$

where $D = \text{diag}(d_1, \dots, d_n)$ and $0 < d_1 \leq \dots \leq d_n$. These entries are called *symplectic eigenvalues*, and (6.7) is referred to as the *symplectic eigenvalue problem*. Note that the symplectic eigenvalues are uniquely defined, whereas the symplectic transformation X is not unique. It can be shown (see [25, (12)]) that the symplectic eigenvalues of M coincide with the positive (standard) eigenvalues of the matrix $G = iJ^\top M$, where $i = \sqrt{-1}$ is the imaginary unit.

In practice, it can be of interest to compute only a few extreme symplectic eigenvalues. They can be determined by exploiting the following relationship between the $p \leq n$ smallest symplectic eigenvalues of M and a symplectic optimization problem

$$2 \sum_{j=1}^p d_j = \min_{X \in \text{Sp}(2p, 2n)} \text{tr}(X^\top M X)$$

which was first established in [25] and further investigated in [7]. Based on this relation, we aim to compute the smallest symplectic eigenvalue d_1 using Algorithm 1 with $p = 1$. We test this algorithm on different data matrices from the MATLAB matrix gallery: (i) the Lehmer matrix; (ii) the Wilkinson matrix; (iii) the companion matrix of the polynomial whose coefficients are $1, \dots, 2n + 1$; (iv) the central finite difference matrix. Whenever M is not positive definite, which happens to the second and third case, we use $M^\top M$ instead of M to generate the appropriate problem. The parameters in Algorithm 1 are default settings. For a comparison, we also compute the smallest positive eigenvalue of G by using the MATLAB function `eig`. The obtained results are shown in Table 5. We observe that the symplectic eigenvalues computed by Sp-Cayley-I are comparable with that provided by `eig`.

TABLE 5
The smallest symplectic eigenvalues

Model matrix	n	<code>eig</code>	Sp-Cayley-I
Lehmer	50	7.67480301454e-03	7.67480302204e-03
Wilkinson	75	1.53471652403e+01	1.53471650305e+01
Companion	500	5.47240371331e-02	5.47244189951e-02
Centr. Finite Diff.	500	2.23005375485e-05	2.23005375834e-05

7. Conclusion and perspectives. We have developed the ingredients—retraction and Riemannian gradient—that turn general first-order Riemannian optimization methods into concrete numerical algorithms for the optimization problem (1.1) on the symplectic Stiefel manifold $\text{Sp}(2p, 2n)$. The algorithms only need to be provided with functions that evaluate the objective function f and the Euclidean gradient $\nabla \bar{f}$. In order to cover the case of the Cayley retraction, we have extended the convergence analysis of a Riemannian non-monotone gradient descent method to encompass the situation where the retraction is not globally defined. This extended analysis leads to the conclusion that, for the sequences generated by the proposed algorithms, every accumulation point is a stationary point. Numerical experiments demonstrate the efficiency of the proposed algorithms. All the results in this paper apply to the symplectic group as the special case $p = n$.

This paper opens several perspectives for further research. In particular, it is tempting to further exploit the leeway in the choice of the metric and the retraction. Extensions to quotients of other quadratic Lie groups are also worth considering, as well as other applications.

- [1] P.-A. ABSIL, R. MAHONY, AND R. SEPULCHRE, *Optimization Algorithms on Matrix Manifolds*, Princeton University Press, 2008, <https://press.princeton.edu/absil>.
- [2] R. L. ADLER, J.-P. DEDIEU, J. Y. MARGULIES, M. MARTENS, AND M. SHUB, *Newton's method on Riemannian manifolds and a geometric model for the human spine*, IMA J. Numer. Anal., 22 (2002), pp. 359–390, <https://doi.org/10.1093/imanum/22.3.359>.
- [3] B. AFKHAM AND J. HESTHAVEN, *Structure preserving model of parametric Hamiltonian systems*, SIAM J. Sci. Comput., 39 (2017), pp. A2616–A2644, <https://doi.org/10.1137/17M1111991>.
- [4] D. O. A. AJAYI AND A. BANYAGA, *An explicit retraction of symplectic flag manifolds onto complex flag manifolds*, J. Geometry, 104 (2013), pp. 1–9, <https://doi.org/10.1007/s00022-013-0148-4>.
- [5] J. BARZILAI AND J. M. BORWEIN, *Two-point step size gradient methods*, IMA J. Numer. Anal., 8 (1988), pp. 141–148, <https://doi.org/10.1093/imanum/8.1.141>.
- [6] D. P. BERTSEKAS, *Nonlinear Programming*, Athena Scientific, 1995, <http://www.athenasc.com/nonlinbook.html>.
- [7] R. BHATIA AND T. JAIN, *On symplectic eigenvalues of positive definite matrices*, J. Math. Phys., 56 (2015), p. 112201, <https://doi.org/10.1063/1.4935852>.
- [8] A. BHATTACHARYA AND R. BHATTACHARYA, *Nonparametric Inference on Manifolds: With Applications to Shape Spaces*, Cambridge University Press, 2012, <https://doi.org/10.1017/CBO9781139094764>.
- [9] R. BHATTACHARYA AND V. PATRANGENARU, *Large sample theory of intrinsic and extrinsic sample means on manifolds*, Ann. Statist., 31 (2003), pp. 1–29, <https://doi.org/10.1214/aos/1046294456>.
- [10] P. BIRTEA, I. CAȘU, AND D. COMĂNESCU, *Optimization on the real symplectic group*, Monatsh. Math., 191 (2020), pp. 465–485, <https://doi.org/10.1007/s00605-020-01369-9>.
- [11] N. BOUMAL, P.-A. ABSIL, AND C. CARTIS, *Global rates of convergence for nonconvex optimization on manifolds*, IMA J. Numer. Anal., 39 (2018), pp. 1–33, <https://doi.org/10.1093/imanum/drx080>.
- [12] N. BOUMAL, B. MISHRA, P.-A. ABSIL, AND R. SEPULCHRE, *Manopt, a Matlab toolbox for optimization on manifolds*, J. Machine Learn. Res., 15 (2014), pp. 1455–1459, <https://www.manopt.org>.
- [13] R. W. BROCKETT, *Least squares matching problems*, Linear Algebra Appl., 122-124 (1989), pp. 761–777, [https://doi.org/10.1016/0024-3795\(89\)90675-7](https://doi.org/10.1016/0024-3795(89)90675-7).
- [14] P. BUCHFINK, A. BHATT, AND B. HAASDONK, *Symplectic model order reduction with non-orthonormal bases*, Math. Comput. Appl., 24 (2019), <https://doi.org/10.3390/mca24020043>.
- [15] Y.-H. DAI AND R. FLETCHER, *Projected Barzilai–Borwein methods for large-scale box-constrained quadratic programming*, Numer. Math., 100 (2005), pp. 21–47, <https://doi.org/10.1007/s00211-004-0569-y>.
- [16] M. A. DE GOSSON, *Symplectic Geometry and Quantum Mechanics*, vol. 166 of Advances in Partial Differential Equations, Birkhuser Basel, 2006, <https://doi.org/10.1007/3-7643-7575-2>.
- [17] M. A. DE GOSSON AND F. LUEF, *Metaplectic group, symplectic Cayley transform, and fractional Fourier transforms*, J. Math. Anal. Appl., 416 (2014), pp. 947–968, <https://doi.org/10.1016/j.jmaa.2014.03.013>.
- [18] A. DRAFT, F. NERI, G. RANGARAJAN, D. R. DOUGLAS, L. M. HEALY, AND R. D. RYNE, *Lie algebraic treatment of linear and nonlinear beam dynamics*, Ann. Rev. Nuc. Part. Sci., 38 (1988), pp. 455–496, <https://doi.org/10.1146/annurev.ns.38.120188.002323>.
- [19] A. EDELMAN, T. A. ARIAS, AND S. T. SMITH, *The geometry of algorithms with orthogonality constraints*, SIAM J. Matrix Anal. Appl., 20 (1998), pp. 303–353, <https://doi.org/10.1137/S0895479895290954>.
- [20] S. FIORI, *Solving minimal-distance problems over the manifold of real-symplectic matrices*, SIAM J. Matrix Anal. Appl., 32 (2011), pp. 938–968, <https://doi.org/10.1137/100817115>.
- [21] S. FIORI, *A Riemannian steepest descent approach over the inhomogeneous symplectic group: Application to the averaging of linear optical systems*, Appl. Math. Comput., 283 (2016), pp. 251–264, <https://doi.org/10.1016/j.amc.2016.02.018>.
- [22] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, Johns Hopkins University Press, 4th ed., 2013, <https://jhupbooks.press.jhu.edu/title/matrix-computations>.
- [23] E. HAIRER, C. LUBICH, AND G. WANNER, *Geometric Numerical Integration: Structure-Preserving Algorithms for Ordinary Differential Equations*, vol. 31, Springer Science & Business Media, 2 ed., 2006, <https://doi.org/10.1007/3-540-30666-8>.
- [24] W. HARRIS, *The average eye*, Ophthal. Physiol. Opt., 24 (2004), pp. 580–585, <https://doi.org/10.1111/j.1475-1313.2004.00239.x>.
- [25] T. HIROSHIMA, *Additivity and multiplicativity properties of some Gaussian channels for Gaussian inputs*, Phys. Rev. A, 73 (2006), p. 012330, <https://doi.org/10.1103/PhysRevA.73.012330>.
- [26] W.-Y. HSIANG AND J. SU, *On the classification of transitive effective actions on Stiefel manifolds*, Trans. Amer. Math. Soc., 130 (1968), pp. 322–336, <https://doi.org/10.1090/S0002-9947-1968-0221529-1>.
- [27] J. HU, X. LIU, Z.-W. WEN, AND Y.-X. YUAN, *A brief introduction to manifold optimization*, J. Oper. Res. Soc. China, (2020), <https://doi.org/10.1007/s40305-020-00295-9>.
- [28] W. HUANG, P.-A. ABSIL, AND K. GALLIVAN, *A Riemannian BFGS method without differentiated retraction for nonconvex optimization problems*, SIAM J. Optim., 28 (2018), pp. 470–495, <https://doi.org/10.1137/17M1127582>.
- [29] B. IANNAZZO AND M. PORCELLI, *The Riemannian Barzilai–Borwein method with nonmonotone line search*

- and the matrix geometric mean computation, *IMA J. Numer. Anal.*, 38 (2018), pp. 495–517, <https://doi.org/10.1093/imanum/drx015>.
- [30] L. M. MACHADO AND F. S. LEITE, *Optimization on quadratic matrix Lie groups*, (2002), <http://hdl.handle.net/10316/11446>.
- [31] K. R. MEYER AND G. R. HALL, *Linear Hamiltonian Systems*, Springer New York, 1992, pp. 33–71, https://doi.org/10.1007/978-1-4757-4073-8_2.
- [32] J. NOCEDAL AND S. WRIGHT, *Numerical Optimization*, Springer Science & Business Media, 2006, <https://doi.org/10.1007/978-0-387-40065-5>.
- [33] L. PENG AND K. MOHSENI, *Symplectic model reduction of Hamiltonian systems*, *SIAM J. Sci. Comput.*, 38 (2016), pp. A1–A27, <https://doi.org/10.1137/140978922>.
- [34] W. RING AND B. WIRTH, *Optimization methods on Riemannian manifolds and their application to shape space*, *SIAM J. Optim.*, 22 (2012), pp. 596–627, <https://doi.org/10.1137/11082885X>.
- [35] F. SIGRIST AND U. SUTER, *Cross-sections of symplectic Stiefel manifolds*, *Trans. Amer. Math. Soc.*, 184 (1973), pp. 247–259, <https://doi.org/10.1090/S0002-9947-1973-0326728-8>.
- [36] J. WANG, H. SUN, AND S. FIORI, *A Riemannian-steepest-descent approach for optimization on the real symplectic group*, *Math. Meth. Appl. Sci.*, 41 (2018), pp. 4273–4286, <https://doi.org/10.1002/mma.4890>.
- [37] Z. WEN AND W. YIN, *A feasible method for optimization with orthogonality constraints*, *Math. Program.*, 142 (2013), pp. 397–434, <https://doi.org/10.1007/s10107-012-0584-1>.
- [38] J. WILLIAMSON, *On the algebraic problem concerning the normal forms of linear dynamical systems*, *Amer. J. Math.*, 58 (1936), pp. 141–163, <https://doi.org/10.2307/2371062>.
- [39] R. WU, R. CHAKRABARTI, AND H. RABITZ, *Optimal control theory for continuous-variable quantum gates*, *Phys. Rev. A*, 77 (2008), p. 052303, <https://doi.org/10.1103/PhysRevA.77.052303>.
- [40] R.-B. WU, R. CHAKRABARTI, AND H. RABITZ, *Critical landscape topology for optimization on the symplectic group*, *J. Optim. Theory Appl.*, 145 (2010), pp. 387–406, <https://doi.org/10.1007/s10957-009-9641-1>.
- [41] H. ZHANG AND W. W. HAGER, *A nonmonotone line search technique and its application to unconstrained optimization*, *SIAM J. Optim.*, 14 (2004), pp. 1043–1056, <https://doi.org/10.1137/S1052623403428208>.