A matrix-algebraic algorithm for the Riemannian logarithm on the Stiefel manifold under the canonical metric

Zimmermann, Ralf

Published in:
SIAM Journal on Matrix Analysis and Applications

DOI:
10.1137/16M1074485

Publication date:
2017

Document version:
Accepted manuscript

Citation for pulished version (APA):
Zimmermann, R. (2017). A matrix-algebraic algorithm for the Riemannian logarithm on the Stiefel manifold under the canonical metric. *SIAM Journal on Matrix Analysis and Applications*, *38*(2), 322-342. https://doi.org/10.1137/16M1074485

A MATRIX-ALGEBRAIC ALGORITHM FOR THE RIEMANNIAN LOGARITHM ON THE STIEFEL MANIFOLD UNDER THE CANONICAL METRIC

RALF ZIMMERMANN*

Abstract. We derive a numerical algorithm for evaluating the Riemannian logarithm on the Stiefel manifold with respect to the canonical metric. In contrast to the optimization-based approach known from the literature, we work from a purely matrix-algebraic perspective. Moreover, we prove that the algorithm converges locally and exhibits a linear rate of convergence.

Key words. Stiefel manifold, Riemannian logarithm, Riemannian exponential, Dynkin series, Goldberg series, Baker-Campbell-Hausdorff series

AMS subject classifications. 15A16, 15B10, 15B57, 33B30, 33F05, 53-04, 65F60

1. Introduction.

Consider an arbitrary Riemannian manifold $\mathcal{M}$. Geodesics on $\mathcal{M}$ are locally shortest curves that are parametrized by the arc length. Because they satisfy an initial value problem, they are uniquely determined by specifying a starting point $p_0 \in \mathcal{M}$ and a starting velocity $\Delta \in T_{p_0}\mathcal{M}$ from the tangent space at p_0. Geodesics give rise to the *Riemannian exponential function* that maps a tangent vector $\Delta \in T_{p_0}\mathcal{M}$ to the endpoint $\mathcal{C}(1)$ of a geodesic path $\mathcal{C} : [0,1] \to \mathcal{M}$ starting at $\mathcal{C}(0) = p_0 \in \mathcal{M}$ with velocity $\Delta = \dot{\mathcal{C}}(0) \in T_{p_0}\mathcal{M}$. It thus depends on the base point p_0 and is denoted by

$$Exp_{p_0} : T_{p_0}\mathcal{M} \to \mathcal{M}, Exp_{p_0}(\Delta) := \mathcal{C}(1). \tag{1}$$

The Riemannian exponential is a local diffeomorphism, [13, §5]. This means that it is locally invertible and that its inverse, called the *Riemannian logarithm* is also differentiable. Moreover, the exponential is radially isometric, i.e., the Riemannian distance between the starting point p_0 and the endpoint $p_1 := Exp_{p_0}(\Delta)$ on $\mathcal{M}$ is the same as the length of the velocity vector Δ of the geodesic $t \mapsto Exp_{p_0}(t\Delta)$ when measured on the tangent space $T_{p_0}\mathcal{M}$, [13, Lem. 5.10 & Cor. 6.11]. In this way, the exponential mapping gives a local parametrization from the (flat, Euclidean) tangent space to the (possibly curved) manifold. This is also referred to as to representing the manifold in *normal coordinates* [12, §III.8].

The Riemannian exponential and logarithm are important both from the theoretical perspective as well as in practical applications. The latter fact holds true in particular, when $\mathcal{M}$ is a *matrix manifold* [2]. Examples range from data analysis and signal processing [7, 17, 3, 18] over computer vision [4, 14] to adaptive model reduction and subspace interpolation [5] and, more generally speaking, optimization techniques on manifolds [6, 1, 2]. This list is far from being exhaustive.

Original contribution. In the work at hand, we present a matrix-algebraic derivation of an algorithm for computing the Riemannian logarithm on the *Stiefel manifold*. The matrix-algebraic perspective allows us to prove local linear convergence. The approach is based on an iterative inversion of the closed formula for the associated Riemannian exponential that has been derived in [6, §2.4.2]. Our main tools are Dynkin's explicit Baker-Campbell-Hausdorff formula [19] and Goldberg's exponential

*Department of Mathematics and Computer Science, University of Southern Denmark (SDU) Odense, (zimmermann@imada.sdu.dk)

series [8], both of which represent a solution Z to the matrix equation

$$\exp_m(Z(X,Y)) = \exp_m(X)\exp_m(Y)\,(\Leftrightarrow Z(\log_m(V),\log_M(W)) = \log_m(VW)),$$

where $V=\exp_m(X), W=\exp_m(Y)$ and $\exp_m, \log_m$ are the standard matrix exponential and matrix logarithm [11, §10, §11]. As an aside, we improve Thompson's norm bound from [20] on $\|Z(X,Y)\|$ for the Goldberg series by a factor of 2, where $\|\cdot\|$ is any submultiplicative matrix norm.

The Stiefel log algorithm can be implemented in $\mathcal{O}(10)$ lines of (commented) MATLAB [15] code, which we include in Appendix C.1.

Comparison with previous work. To the best of our knowledge, up to now, the only algorithm for evaluating the Stiefel logarithm appeared in Q. Rentmeesters' thesis [18, Alg. 4, p. 91]. This algorithm is based on a Riemannian optimization problem. It turns out that this approach and the ansatz that is pursued here, though very different in their course of action, lead to essentially the same numerical scheme. Rentmeesters observes numerically a linear rate of convergence [18, p.83, p.100]. Proving linear convergence for [18, Alg. 4, p. 91] would require estimates on the Hessian, see [18, §5.2.1], [2, Thm. 4.5.6]. In contrast, the derivation presented here uses only elementary matrix algebra and the convergence proof given here formally avoids the requirements of computing/estimating step sizes, gradients and Hessians that are inherent to analyzing the convergence of optimization approaches. In fact, the convergence proof applies to [18, Alg. 4, p. 91] and yields the linear convergence of this optimization approach when using a fixed *unit step size*, but only on a sufficiently small domain. The thesis [18] was published under a two-years access embargo and the fundamentals of the work at hand were developed independently before [18] was accessible.

Transition to the complex case. The basic geometric concepts of the Stiefel manifold, the algorithm for the Riemannian log mapping developed here and its convergence proof carry over to complex matrices, where orthogonal matrices have to be replaced with unitary matrices and skew-symmetric matrices with skew-Hermitian matrices and so forth, see also [6, §2.1]. The thus adjusted log mapping algorithm was also confirmed numerically to work in the complex case.

Organization. Background information on the Stiefel manifold are reviewed in Section 2. The new derivation for the Stiefel log algorithm is in Section 3, convergence analysis is performed in Section 4, experimental results are in Section 5, and the conclusions follow in Section 6.

Notational specifics. The $(p\times p)$-*identity matrix* is denoted by $I_p\in\mathbb{R}^{p\times p}$. If the dimension is clear, we will simply write I. The $(p\times p)$-*orthogonal group*, i.e., the set of all square orthogonal matrices is denoted by

$$O_{p\times p} = \{\Phi\in\mathbb{R}^{p\times p} | \Phi^T\Phi = \Phi\Phi^T = I_p\}.$$

The standard matrix exponential and matrix logarithm are denoted by

$$\exp_m(X) := \sum_{j=0}^{\infty}\frac{X^j}{j!}, \quad \log_m(I+X) := \sum_{j=1}^{\infty}(-1)^{j+1}\frac{X^j}{j}.$$

We use the symbols Exp^{St}, Log^{St} for the Riemannian counterparts on the Stiefel manifold.

When we employ the qr-decomposition of a rectangular matrix $A\in\mathbb{R}^{n\times p}$, we implicitly assume that $n\geq p$ and refer to the 'economy size' qr-decomposition $A = QR$, with $Q\in\mathbb{R}^{n\times p}$, $R\in\mathbb{R}^{p\times p}$.

2. The Stiefel manifold in numerical representation. This section reviews the essential aspects of the numerical treatment of Stiefel manifolds, where we rely heavily on the excellent references [2, 6]. The *Stiefel manifold* is the compact homogeneous matrix manifold of all column-orthogonal rectangular matrices

$$St(n,p) := \{U \in \mathbb{R}^{n\times p} | \quad U^T U = I_p\}.$$

The *tangent space* $T_U St(n,p)$ at a point $U \in St(n,p)$ can be thought of as the space of velocity vectors of differentiable curves on $St(n,p)$ passing through U:

$$T_U St(n,p) = \{\dot{\mathcal{C}}(t_o) | \mathcal{C} : (t_0 - \epsilon, t_0 + \epsilon) \to St(n,p), \mathcal{C}(t_0) = U\}.$$

For any matrix representative $U \in St(n,p)$, the tangent space of $St(n,p)$ at U is represented by

$$T_U St(n,p) = \left\{\Delta \in \mathbb{R}^{n\times p} | \quad U^T\Delta = -\Delta^T U\right\} \subset \mathbb{R}^{n\times p}.$$

Every tangent vector $\Delta \in T_U St(n,p)$ may be written as

$$\Delta = UA + (I - UU^T)T, \quad A \in \mathbb{R}^{p\times p} \text{ skew}, \quad T \in \mathbb{R}^{n\times p} \text{ arbitrary.} \tag{2}$$

The dimension of both $T_U St(n,p)$ and $St(n,p)$ is $np - \frac{1}{2}p(p+1)$.

Each tangent space carries an inner product $\langle \Delta, \tilde{\Delta}\rangle_U = tr\left(\Delta^T(I - \frac{1}{2}UU^T)\tilde{\Delta}\right)$ with corresponding norm $\|\Delta\|_U = \sqrt{\langle \Delta, \Delta\rangle_U}$. This is called the *canonical metric* on $T_U St(n,p)$. It is derived from the quotient space representation $St(n,p) = O_{n\times n}/O_{(n-p)\times(n-p)}$ that identifies two square orthogonal matrices in $O_{n\times n}$ as the same point on $St(n,p)$, if their first p columns coincide [6, §2.4]. Endowing each tangent space with this metric (that varies differentiably in U) turns $St(n,p)$ into a *Riemannian manifold*.

We now turn to the Riemannian exponential (1) but for $\mathcal{M} = St(n,p)$. An efficient algorithm for evaluating the Stiefel exponential was derived in [6, §2.4.2]. The algorithm starts with decomposing an input tangent vector $\Delta \in T_U St(n,p)$ into its horizontal and vertical components with respect to the base point U,

$$\Delta = UU^T\Delta + (I - UU^T)\Delta \overset{(\text{qr of } (I-UU^T)\Delta)}{=} UA + Q_E R_E.$$

Because Δ is tangent, $A \in \mathbb{R}^{p\times p}$ is skew. Then the matrix exponential is invoked to compute

$$\begin{pmatrix} M \\ N_E \end{pmatrix} := \exp_m\left(\begin{pmatrix} A & -R_E^T \\ R_E & 0 \end{pmatrix}\right)\begin{pmatrix} I_p \\ 0 \end{pmatrix}. \tag{3}$$

The final output is[1]

$$\tilde{U} := Exp_U^{St}(\Delta) = UM + Q_E N_E \in St(n,p). \tag{4}$$

(A MATLAB function for the Stiefel exponential is in the supplement in Section S4.) The matrix exponential in (3) is related with the solution of the initial value problem that defines a geodesic on $St(n,p)$, see [6, §2.4.2] for details. It turns out that the main obstacle in computing the inverse of the Stiefel exponential and thus the Stiefel logarithm is inverting (3), i.e. finding A, R_E given M, N_E, compare to [18, eq. (5.21)].

[1] The index in Q_E, R_E, N_E is used to emphasize that these matrices stem from the Stiefel exponential as opposed to the closely related matrices Q, R, N that will appear in the procedure for the Stiefel logarithm.

3. Derivation of the Stiefel log algorithm. Let $U, \tilde{U} \in St(n,p)$ and assume that $\tilde{U}$ is contained in a neighborhood $\mathcal{D}$ of U such that Exp_U^{St} is a diffeomorphism from a neighborhood of $0 \in T_U St(n,p)$ onto $\mathcal{D}$. The central objective is to find $\Delta \in T_U St(n,p)$ such that $Exp_U^{St}(\Delta) = \tilde{U}$.

Because of (4), we know that $\tilde{U}$ allows for a representation $\tilde{U} = UM + Q_E N_E$. Hence, we have to determine the unknown matrices $M, N_E \in \mathbb{R}^{p\times p}$, $Q_E \in \mathbb{R}^{n\times p}$, which feature the following properties: $Q_E^T U = 0$ and $M^T M + N_E^T N_E = I_p$. (Note that by (3), M and N_E are the left upper and lower $p\times p$ blocks of a $2p\times 2p$ orthogonal matrix.) We directly obtain

$$M = U^T\tilde{U}, \quad Q_E N_E = (I - UU^T)\tilde{U}.$$

We compute candidates for Q_E, N_E via a qr-decomposition

$$QN \overset{qr}{=} (I - UU^T)\tilde{U}, \quad Q \in St(n,p).$$

The set of all orthogonal matrices with M, N as an upper diagonal and lower off-diagonal block is parametrized via

$$\left\{ \begin{pmatrix} M & X \\ N & Y \end{pmatrix} \mid \quad \begin{pmatrix} X \\ Y \end{pmatrix} = \begin{pmatrix} X_0 \\ Y_0 \end{pmatrix}\Phi, \quad \Phi \in O_{p\times p} \right\},$$

where $(X_0^T, Y_0^T)^T$ is a specific orthogonal completion, computed, say, via the Gram-Schmidt process.

Thus, the objective is reduced to solving the following nonlinear matrix equation

$$0 = \begin{pmatrix} 0 & I_p \end{pmatrix} \log_m \left(\begin{pmatrix} M & X_0 \\ N & Y_0 \end{pmatrix} \begin{pmatrix} I_p & 0 \\ 0 & \Phi \end{pmatrix} \right) \begin{pmatrix} 0 \\ I_p \end{pmatrix}, \quad \Phi \in O_{p\times p}. \tag{5}$$

Writing $\log_m \left(\begin{pmatrix} M & X_0 \\ N & Y_0 \end{pmatrix} \begin{pmatrix} I_p & 0 \\ 0 & \Phi \end{pmatrix} \right) = \begin{pmatrix} A & -B^T \\ B & C \end{pmatrix}$, this means finding a rotation Φ such that $C = 0$.

The first result is that solving (5) indeed leads to the Riemannian logarithm on the Stiefel manifold.

THEOREM 1. *Let $U, \tilde{U} \in St(n,p)$ and assume that $\tilde{U}$ is contained in a neighborhood $\mathcal{D}$ of U such that Exp_U^{St} is a diffeomorphism from a neighborhood of $0 \in T_U St(n,p)$ onto $\mathcal{D}$.*

Let M, Q_E, N_E, Q, N, X_0, Y_0 as introduced in the above setting. Suppose that $\Phi \in O_{p\times p}$ solves (5), i.e.,

$$\log_m \left(\begin{pmatrix} M & X_0\Phi \\ N & Y_0\Phi \end{pmatrix} \right) = \begin{pmatrix} A & -B^T \\ B & 0 \end{pmatrix}.$$

Define $\Delta := UA + QB \in T_U St(n,p)$. Then $Exp_U^{St}(\Delta) = \tilde{U}$, i.e., $\Delta = Log_U^{St}(\tilde{U})$.

Proof. By construction, it holds $QN = (I - UU^T)\tilde{U}$ and hence

$$U^T Q = 0, \quad (I - UU^T)Q = Q, \quad QQ^T\tilde{U} = QQ^T(I - UU^T)\tilde{U} = (I - UU^T)\tilde{U}. \tag{6}$$

Now, we apply the Stiefel exponential (4) to $\Delta = UA + QB$. This gives $(I - UU^T)\Delta = QB$ and

$$Q_E R_E \overset{qr}{=} QB \Leftrightarrow R_E = \Psi B, \text{ where } \Psi := (Q_E^T Q) \in O_{p\times p}.[2]$$

[2]The matrices Q_E and Q differ by an orthogonal rotation but span the same subspace.

With $U^T\Delta = A$, we obtain

$$\begin{aligned}\begin{pmatrix} M \\ N_E \end{pmatrix} &:= \exp_m\left(\begin{pmatrix} A & -R_E^T \\ R_E & 0 \end{pmatrix}\right)\begin{pmatrix} I_p \\ 0 \end{pmatrix} \\ &= \begin{pmatrix} I & 0 \\ 0 & \Psi \end{pmatrix}\exp_m\left(\begin{pmatrix} A & -B^T \\ B & 0 \end{pmatrix}\right)\begin{pmatrix} I & 0 \\ 0 & \Psi^T \end{pmatrix}\begin{pmatrix} I_p \\ 0 \end{pmatrix} \\ &= \begin{pmatrix} I & 0 \\ 0 & \Psi \end{pmatrix}\begin{pmatrix} M & X_0\Phi \\ N & Y_0\Phi \end{pmatrix}\begin{pmatrix} I_p \\ 0 \end{pmatrix} = \begin{pmatrix} M \\ \Psi N \end{pmatrix}.\end{aligned}$$

Keeping in mind that $Q_E\Psi = Q_E Q_E^T Q = Q$, this leads to an output of

$$Exp_U^{St}(\Delta) = UM + Q_E N_E = UM + Q_E\Psi N = UM + QN = \tilde{U}.$$

Thus, Δ is a valid tangent vector in $T_U St(n,p)$ such that $Exp_U^{St}(\Delta) = \tilde{U} \in St(n,p)$. From abstract differential geometry, we know that $Log_U^{St}(\tilde{U}) \in T_U St(n,p)$ is the unique tangent with $Exp_U^{St}(Log_U^{St}(\tilde{U})) = \tilde{U}$. We arrive at the claim

$$\Delta = Log_U^{St}(\tilde{U}).$$ ∎

Having established Theorem 1, we now focus on solving (5). Let

$$\begin{aligned} V_0 &:= \begin{pmatrix} M & X_0 \\ N & Y_0 \end{pmatrix}, \quad \log_m(V_0) := \begin{pmatrix} A_0 & -B_0^T \\ B_0 & C_0 \end{pmatrix}, \\ W_0 &:= \begin{pmatrix} I_p & 0 \\ 0 & \Phi_0 \end{pmatrix}, \quad \log_m(W_0) = \begin{pmatrix} 0 & 0 \\ 0 & \log_m(\Phi_0) \end{pmatrix}. \end{aligned}$$

Up to terms of first order, it holds $\log_m(V_0W_0) = \log_m(V_0) + \log_m(W_0)$. Hence, the choice

$$\Phi_0 = \exp_m(-C_0)$$

gives an approximate solution to (5). We define

$$V_1 := \begin{pmatrix} M & X_0 \\ N & Y_0 \end{pmatrix}\begin{pmatrix} I_p & 0 \\ 0 & \Phi_0 \end{pmatrix}, \quad \log_m(V_1) := \begin{pmatrix} A_1 & -B_1^T \\ B_1 & C_1 \end{pmatrix} \tag{9}$$

and iterate. This is the essential idea of Algorithm 1 for the Riemannian logarithm.[3]

In Section 4 we make use of the Baker-Campbell-Hausdorff formula [19, §1.3, p. 22] that corrects for the misfit in the approximative matrix relation $\log_m(VW) \approx \log_m(V) + \log_m(W)$ for two non-commuting matrices V, W in order to show that the above procedure leads to

$$\|C_{k+1}\|_2 < \alpha\|C_k\|_2$$

for all $k \in \mathbb{N}_0$ and a constant $\alpha < 1$ and is thus convergent.

Since the Riemannian exponential is a local diffeomorphism, we have to postulate a suitable bound on the distance between the input matrices U and $\tilde{U}$. Suppose that $\|U - \tilde{U}\|_2 < \epsilon$. Recalling the definitions $M = U^T\tilde{U}$ and $(I - UU^T)\tilde{U} = QN$, this gives the following bounds for the horizontal and the vertical component of $U - \tilde{U}$ with respect to the subspace spanned by U:

$$\|UU^T(U - \tilde{U})\|_2 = \|I_p - M\|_2 < \epsilon, \quad \|(I - UU^T)(U - \tilde{U})\|_2 = \|QN\|_2 = \|N\|_2 < \epsilon.$$

[3]This is the same algorithm as [18, Alg. 4, p. 91] that Rentmeesters obtains from his geometrical perspective when a fixed unit step length is employed and when [18, §5.3] is taken into account.

However, it turns out that for the convergence proof, estimates on the norms of X_0, Y_0 and $Y_0 - I_p$ are also required. By the CS-decomposition of orthonormal matrices [9, Thm 2.6.3, p. 78], the diagonal blocks M and Y_0 share the same singular values and so do the off-diagonal blocks N, X_0. Hence, $\|N\|_2 = \|X_0\|_2 < \epsilon$. Let $D_1 \Sigma R_1^T$ be the SVD of M and $D_2 \Sigma R_2^T$ be the SVD of Y_0. An estimate for the singular values of M can be obtained as follows:

$$\epsilon^2 > \|N\|_2^2 = \|N^T N\|_2 = \|I - M^T M\|_2 = \|I - \Sigma^2\|_2 = \max_{\sigma_k}(1 - \sigma_k^2), \tag{10}$$

where we have used that $\sigma_1 = \|M\|_2 \leq 1$. Now, we replace the Y_0 that has been obtained via, say, Gram-Schmidt by $Y_0 R_2 D_2^T = D_2 \Sigma D_2^T$ (and, correspondingly, X_0 by $X_0 R_2 D_2^T$). Essentially, this is the orthogonal Procrustes method, [9, §12.4.1, p.601], applied to $\min_{\Psi \in O_{p\times p}} \|I - Y_0 \Psi\|_2$.This operation preserves the orthogonality of $V_0 = \begin{pmatrix} M & X_0 \\ N & Y_0 \end{pmatrix}$, but the new Y_0 is symmetric with eigenvalue decomposition $Y_0 = D_2 \Sigma D_2^T$. This gives

$$\|Y_0 - I_p\|_2 = \|\Sigma - I_p\|_2 = \max_{\sigma_k} |1 - \sigma_k| < \max_{\sigma_k}(1 - \sigma_k^2) < \epsilon^2.$$

In summary, if $\|U - \tilde{U}\|_2 < \epsilon$ and if we start the iterations indicated by (9) with the Procrustes orthogonal completion X_0, Y_0 rather than the standard Gram-Schmidt process, we obtain Algorithm 1 with the starting conditions

$$\|I_p - M\|_2 < \epsilon, \quad \|N\|_2 = \|X_0\|_2 < \epsilon, \quad \|Y_0 - I_p\|_2 < \epsilon^2. \tag{11}$$

Algorithm 1 Stiefel logarithm, iterative procedure

Input: base point $U \in St(n,p)$ and $\tilde{U} \in St(n,p)$ 'close' to base point, $\tau > 0$ convergence threshold

1: $M := U^T \tilde{U} \in \mathbb{R}^{p \times p}$
2: $QN := \tilde{U} - UM \in \mathbb{R}^{n \times p}$ {(thin) qr-decomp. of normal component of $\tilde{U}$}
3: $V_0 := \begin{pmatrix} M & X_0 \\ N & Y_0 \end{pmatrix} \in O_{2p \times 2p}$ {orthogonal completion and Procrustes}
4: **for** $k = 0, 1, 2, \dots$ **do**
5:     $\begin{pmatrix} A_k & -B_k^T \\ B_k & C_k \end{pmatrix} := \log_m(V_k)$ {matrix log, A_k, C_k skew}
6:     **if** $\|C_k\|_2 \leq \tau$ **then**
7:         break
8:     **end if**
9:     $\Phi_k = \exp_m(-C_k)$ {matrix exp, Φ_k orthogonal}
10:     $V_{k+1} := V_k W_k$, where $W_k := \begin{pmatrix} I_p & 0 \\ 0 & \Phi_k \end{pmatrix}$ {update}
11: **end for**

Output: $\Delta := Log_U^{St}(\tilde{U}) = UA_k + QB_k \in T_U St(n,p)$

Computational costs. W.l.o.g. suppose that $n \geq p$. In fact the most important case in practical applications is $n \gg p$. Because of the matrix product in step 1 and the qr-decomposition in step 2 of Algorithm 1, the preparatory steps 1–3 require $\mathcal{O}(np^2)$ FLOPS. The dominating costs in the iterative loop, steps 5–10, are the evaluation of the matrix logarithm for a $2p$-by-$2p$ orthogonal matrix and the matrix exponential for a p-by-p skew-symmetric matrix in every iteration, both of which can be achieved efficiently via the Schur decomposition. The costs are $\mathcal{O}(p^3)$, see [9, Alg. 7.5.2].

A MATLAB function for Algorithm 1 is in Appendix C.1.

4. Convergence proof.

In this section, we establish the convergence of Algorithm 1 under suitable conditions. We state the main result as Theorem 2; the proof is subdivided into the auxiliary results Lemma 3, and Lemma 4 as well as Lemma 5 that appears in Appendix A. An essential requirement is that the point $\tilde{U} \in St(n,p)$ that is to be mapped to the tangent space $T_U St(n,p)$ is sufficiently close to the base point $U \in St(n,p)$ in the sense that $\|U - \tilde{U}\|_2 < \epsilon$. Throughout, we will make extensive use of Dynkin's explicit BCH formula [19, §1.3, p. 22].

THEOREM 2. *Let $U, \tilde{U} \in St(n,p)$. Assume that $\|U - \tilde{U}\|_2 < \epsilon$. Let $(V_k)_k$ be the sequence of orthogonal matrices generated by Algorithm 1.*

If $\epsilon < 0.0912$, then Algorithm 1 converges to a limit matrix $V_\infty := \lim_{k\to\infty} V_k$ such that

$$\log_m(V_\infty) := \begin{pmatrix} A_\infty & -B_\infty^T \\ B_\infty & C_\infty \end{pmatrix} = \begin{pmatrix} A_\infty & -B_\infty^T \\ B_\infty & 0 \end{pmatrix}.$$

Given a numerical convergence threshold $\tau > 0$, see Algorithm 1, line 7, the algorithm requires at most $k = \lceil \frac{\log(\|C_0\|_2) - \log(\tau)}{\log(2)} \rceil - 1$ iteration steps to meet the convergence criterion under the above conditions.

REMARK 1. Algorithm 1 generates a sequence of orthonormal matrices

$$V_{k+1} = V_k W_k = V_0 (W_0 W_1 \ldots W_k) = V_0 \left(\begin{pmatrix} I_p & 0 \\ 0 & \Phi_0 \end{pmatrix} \cdots \begin{pmatrix} I_p & 0 \\ 0 & \Phi_k \end{pmatrix} \right) \in O_{2p\times 2p}. \tag{12}$$

The proof of Theorem 2 will show that $\lim_{k\to\infty} W_k = I_{2p}$, see (23). Therefore, the sequence of orthogonal products $\Phi_0 \ldots \Phi_k$ converges to a limit Φ_∞ for $k \to \infty$. The limit Φ_∞ solves (5). However, it is not required to actually form Φ_∞.

In pursuit of the proof of Theorem 2, we first show that if the norm of the matrix logarithm of the orthogonal matrix V_k produced by Algorithm 1 at iteration k is sufficiently small, then the norm of the lower p-by-p diagonal block of the matrix logarithm of the next iterate V_{k+1} is strictly decreasing by a constant factor.

LEMMA 3. *Let $U, \tilde{U} \in St(n,p)$. Let $(V_k)_k \subset O_{2p\times 2p}$ be the sequence of orthogonal matrices generated by Algorithm 1. Suppose that at stage k, it holds*

$$\|\log_m(V_k)\|_2 := \left\| \begin{pmatrix} A_k & -B_k^T \\ B_k & C_k \end{pmatrix} \right\|_2 < \frac{1}{2}(\sqrt{5} - 1). \tag{13}$$

Then $\log_m(V_{k+1}) := \begin{pmatrix} A_{k+1} & -B_{k+1}^T \\ B_{k+1} & C_{k+1} \end{pmatrix}$ features a lower $(p \times p)$-diagonal block of norm

$$\|C_{k+1}\|_2 < \alpha \|C_k\|_2, \quad 0 < \alpha < \frac{1}{2}.$$

Proof. Given $V_k = \begin{pmatrix} M & X_k \\ N & Y_k \end{pmatrix} = \exp_m \left(\begin{pmatrix} A_k & -B_k^T \\ B_k & C_k \end{pmatrix} \right)$, Algorithm 1 computes the next iterate V_{k+1} via

$$V_{k+1} := V_k W_k,$$

where $W_k := \begin{pmatrix} I_p & 0 \\ 0 & \exp_m(-C_k) \end{pmatrix}$. For brevity, we introduce the notation $L_V := \log_m(V)$ for the matrix logarithm. Recall that $[V,W] := VW - WV$ denotes the commutator or Lie-bracket of the matrices V, W. From Dynkin's formula for the Baker-Campbell-Hausdorff series, see [19, §1.3, p. 22], we obtain

$$\begin{aligned} L_{V_{k+1}} &= \log_m(V_k W_k) \\ &= L_{V_k} + L_{W_k} + \frac{1}{2}[L_{V_k}, L_{W_k}] \\ &\quad + \frac{1}{12}\left([L_{V_k},[L_{V_k},L_{W_k}]] + [L_{W_k},[L_{W_k},L_{V_k}]]\right) \\ &\quad - \frac{1}{24}\Big[L_{W_k},[L_{V_k},[L_{V_k},L_{W_k}]]\Big] + \sum_{l=5}^{\infty} z_l(L_{V_k}, L_{W_k}), \end{aligned}$$

where $\sum_{l=5}^{\infty} z_l(L_{V_k}, L_{W_k}) =: h.o.t.(5)$ are the terms of fifth order and higher in the series. In the case at hand, it holds

$$\begin{aligned} L_{V_k} + L_{W_k} &= \begin{pmatrix} A_k & -B_k^T \\ B_k & C_k \end{pmatrix} + \begin{pmatrix} 0 & 0 \\ 0 & -C_k \end{pmatrix} = \begin{pmatrix} A_k & -B_k^T \\ B_k & 0 \end{pmatrix}, \\ [L_{V_k}, L_{W_k}] &= \begin{pmatrix} 0 & B_k^T C_k \\ C_k B_k & 0 \end{pmatrix}. \end{aligned}$$

(Note that the basic idea in designing Algorithm 1 was exactly to choose W_k such that the lower diagonal block in the BCH-series cancels in the first order terms.)

The third and fourth order terms are

$$\begin{aligned} &\frac{1}{12}\begin{pmatrix} -2B_k^T C_k B_k & A_k B_k^T C_k - 2B_k^T C_k^2 \\ 2C_k^2 B_k - C_k B_k A_k & B_k B_k^T C_k + C_k B_k B_k^T \end{pmatrix}, \text{ and} \\ &\frac{1}{24}\begin{pmatrix} 0 & -B_k^T C_k^3 + A_k B_k^T C_k^2 \\ -C_k^3 B_k + C_k^2 B_k A_k & B_k B_k^T C_k^2 - C_k^2 B_k B_k^T \end{pmatrix}, \text{ respectively.} \end{aligned}$$

Therefore, the series expansion for the lower diagonal block in $\log_m(V_{k+1})$ starts with the terms of third order:

$$\|C_{k+1}\|_2 = \|\frac{1}{12}(B_k B_k^T C_k + C_k B_k B_k^T) - \frac{1}{24}(B_k B_k^T C_k^2 - C_k^2 B_k B_k^T) + h.o.t.(5)\|_2$$

$$\leq \left(\frac{1}{6}\|B_k\|_2^2 + \frac{1}{12}\|B_k\|_2^2\|C_k\|_2 + \frac{\|h.o.t.(5)\|_2}{\|C_k\|_2}\right)\|C_k\|_2. \tag{17a}$$

We tackle the higher order terms via Lemma 5 from the appendix. The lemma applies because $\|C_k\|_2 = \|L_{W_k}\|_2 \leq \|L_{V_k}\|_2 < \frac{1}{2}(\sqrt{5}-1) < 1$. In this setting, it gives

$$\|h.o.t.(5)\|_2 \leq \sum_{l=5}^{\infty} \|z_l(L_{V_k}, L_{W_k})\|_2 < \sum_{l=5}^{\infty} \|L_{V_k}\|^{l-1}\|L_{W_k}\|_2,$$

since each of the "letters" L_{V_k}, L_{W_k} appears at least once in every "word" that contributes to $z_k(L_{V_k}, L_{W_k})$, see Appendix A and [20, 16, 21].

Writing $s := \|L_{V_k}\|_2$ and substituting in (17a) leads to

$$\|C_{k+1}\|_2 < \left(\frac{1}{6}s^2 + \frac{1}{12}s^3 + \sum_{l=4}^{\infty} s^l\right)\|C_k\|_2 =: \alpha\|C_k\|_2. \tag{18}$$

The proof is complete, if we can show that $\alpha < 1$. Note that $\sum_{l=4}^{\infty} s^l = \frac{1}{1-s} - 1 - s - s^2 - s^3$. As a consequence

$$\alpha < 1 \Leftrightarrow \frac{s^2}{1-s} - \frac{5}{6}s^2 - \frac{11}{12}s^3 < 1.$$

An obvious bound on the size of s is obtained via observing that $\frac{s^2}{1-s} < 1$, if $s < \frac{1}{2}(\sqrt{5}-1) \approx 0.618$. The corresponding α is $0.4653 < \frac{1}{2}$. A sharper bound can be obtained via solving the associated quartic equation. This shows that the inequality even holds for all $s < 0.7111$. ∎

In order to make use of Lemma 3, we establish conditions such that $\|\log_m(V_k)\|_2 < \frac{1}{2}(\sqrt{5}-1)$ holds throughout the iterations of Algorithm 1.

This is the goal of the the next lemma. It relies on the auxiliary results Proposition 6, Proposition 7 and Lemma 8 from Appendix B. Proposition 6 shows that $\|\exp_m(C) - I\|_2 < \|C\|_2$ for C skew-symmetric; Proposition 7 establishes a bound in the opposite direction: if V is orthogonal such that $\|V - I\|_2 < r$, then $\|\log_m(V)\|_2 < r\sqrt{1-\frac{r^2}{4}}/(1-\frac{r^2}{2})$. Finally, Lemma 8 shows that $\|V_0 - I\|_2 < 2\epsilon$ for the first iterate V_0 of Algorithm 1, provided that $\|U - \tilde{U}\|_2 < \epsilon$.

LEMMA 4. *Let $U, \tilde{U} \in St(n,p)$ with $\|U - \tilde{U}\|_2 < \epsilon$. Let $(V_k)_k \subset O_{2p\times 2p}$ be the sequence of orthogonal matrices generated by Algorithm 1, where $V_k = \begin{pmatrix} M & X_k \\ N & Y_k \end{pmatrix}$. Let $\tilde{\epsilon} = 2\epsilon\frac{\sqrt{1-\epsilon^2}}{1-2\epsilon^2}$ and $\hat{\epsilon} := (e^{2\tilde{\epsilon}} - 1) + \epsilon + \epsilon^2$. If $0 < \epsilon$ is small enough such that $\hat{\epsilon}\frac{\sqrt{1-\frac{\hat{\epsilon}^2}{4}}}{1-\frac{\hat{\epsilon}^2}{2}} < \frac{1}{2}(\sqrt{5}-1)$, then*

$$\|\log_m(V_k)\|_2 = \|\begin{pmatrix} A_k & -B_k^T \\ B_k & C_k \end{pmatrix}\|_2 < \frac{1}{2}(\sqrt{5}-1) \text{ for all } k.$$

Proof. Let $\delta_0 := \frac{1}{2}(\sqrt{5}-1)$. By Lemma 8 from Appendix B, it holds

$$\|\log_m(V_0)\|_2 < 2\epsilon\frac{\sqrt{1-\epsilon^2}}{1-2\epsilon^2} = \tilde{\epsilon} \quad (< \delta_0).$$

In particular, $\tilde{\epsilon} > \|\begin{pmatrix} A_0 & -B_0^T \\ B_0 & C_0 \end{pmatrix}\|_2 \geq \|C_0\|_2$. By Algorithm 1, $\Phi_0 = \exp_m(-C_0)$, where Φ_0 is orthogonal. By Proposition 6 from Appendix B

$$\|\Phi_0 - I\|_2 \leq \|C_0\|_2 < \tilde{\epsilon}.$$

Writing $V_1 = I + (V_1 - I) =: I + E_1$, this leads to the estimate

$$\begin{aligned}
\|E_1\|_2 &= \|\begin{pmatrix} M - I & X_0\Phi_0 \\ N & Y_0\Phi_0 - I \end{pmatrix}\|_2 \\
&= \|\begin{pmatrix} M - I & 0 \\ 0 & Y_0(\Phi_0 - I) \end{pmatrix} + \begin{pmatrix} 0 & 0 \\ 0 & Y_0 - I \end{pmatrix} + \begin{pmatrix} 0 & X_0\Phi_0 \\ N & 0 \end{pmatrix}\|_2 \\
&\leq \max\{\|M - I\|_2, \|Y_0(\Phi_0 - I)\|_2\} + \|Y_0 - I\|_2 + \max\{\|N\|_2, \|X_0\Phi_0\|_2\} \\
&\leq \max\{\epsilon, \|Y_0\|_2\|(\Phi_0 - I)\|_2\} + \epsilon^2 + \epsilon \leq \tilde{\epsilon} + \epsilon^2 + \epsilon \\
&\leq (e^{2\tilde{\epsilon}} - 1) + \epsilon + \epsilon^2 = \hat{\epsilon},
\end{aligned}$$

where we have used (11) and the fact that $\|Y_0\|_2 \leq 1$, see (33a), (33b). By Lemma 8, $\|\log_m(V_1)\|_2 < \hat{\epsilon}\sqrt{1-\frac{\hat{\epsilon}^2}{4}}/(1-\frac{\hat{\epsilon}^2}{2}) < \delta_0$. Thus, the claim holds for $k = 0, 1$.

Lemma 3 applies to $\|\log_m(V_0)\|_2$ and leads to $\|C_1\|_2 < \frac{1}{2}\|C_0\|_2 < \frac{\tilde{\epsilon}}{2}$ for the lower diagonal block C_1 of the next iterate $\log_m(V_1)$. Therefore, by using Proposition 6 once more, we see that

$$\|\Phi_1 - I\|_2 \leq \|C_1\|_2 < \frac{\tilde{\epsilon}}{2}.$$

By induction, we obtain $V_k = I + (V_k - I) =: I + E_k$ with

$$\begin{aligned}
\|E_k\|_2 &= \| \begin{pmatrix} M & X_0\hat{\Phi}_{k-1} \\ N & Y_0\hat{\Phi}_{k-1} \end{pmatrix} - I\|_2 \\
&= \| \begin{pmatrix} M-I & 0 \\ 0 & Y_0(\hat{\Phi}_{k-1} - I) \end{pmatrix} + \begin{pmatrix} 0 & 0 \\ 0 & Y_0 - I \end{pmatrix} + \begin{pmatrix} 0 & X_0\hat{\Phi}_{k-1} \\ N & 0 \end{pmatrix} \|_2 \\
&\leq \max\{\|M-I\|_2, \|Y_0(\hat{\Phi}_{k-1} - I)\|_2\} + \|Y_0 - I\|_2 + \max\{\|N\|_2, \|X_0\hat{\Phi}_{k-1}\|_2\} \\
&\leq \max\{\epsilon, \|Y_0\|_2\|\hat{\Phi}_{k-1} - I\|_2\} + \epsilon^2 + \epsilon. \qquad \text{(20a)}
\end{aligned}$$

where $\hat{\Phi}_{k-1} = \Phi_0 \cdots \Phi_{k-1}$.

We can estimate $\|\hat{\Phi}_{k-1} - I\|_2$ as follows: By the induction hypothesis, we assume that we have checked that $\|\log_m(V_j)\|_2 < \delta_0$ for $j = 0, \ldots, k-1$. Hence, Lemma 3 ensures that $\|C_j\|_2 < \frac{1}{2}\|C_{j-1}\|_2 < \ldots < \frac{1}{2^j}\|C_0\|_2 < \frac{\tilde{\epsilon}}{2^j}$ for the lower diagonal block of $\log_m(V_j)$, $j = 0, \ldots, k-1$. As above, this gives $\|\Phi_j - I\|_2 \leq \|C_j\|_2 < \frac{\tilde{\epsilon}}{2^j}$. We thus may write $\Phi_j = I + (\Phi_j - I) =: I + \Gamma_j$ with $\|\Gamma_j\|_2 =: g_j < \frac{\tilde{\epsilon}}{2^j}$. This gives

$$\|\hat{\Phi}_{k-1} - I\|_2 = \|(I + \Gamma_1)\cdots(I + \Gamma_{k-1}) - I\|_2 \leq (1 + g_1)\cdots(1 + g_{k-1}) - 1. \qquad (21)$$

It holds

$$\ln\left(\prod_{j=0}^{k-1}(1 + g_j)\right) = \sum_{j=0}^{k-1}\ln(1 + g_j) \leq \sum_{j=0}^{k-1} g_j \leq \sum_{j=0}^{\infty}\frac{\tilde{\epsilon}}{2^j} = 2\tilde{\epsilon}.$$

Using this estimate in (21) gives

$$\|\hat{\Phi}_{k-1} - I\|_2 < e^{2\tilde{\epsilon}} - 1$$

and we finally arrive at

$$\|E_k\|_2 \leq (e^{2\tilde{\epsilon}} - 1) + \epsilon^2 + \epsilon = \hat{\epsilon}.$$

Recalling (20a), we have $V_k = I + E_k$ with $\|E_k\|_2 < \hat{\epsilon}$ at every iteration k. By Lemma 8, $\|\log_m(V_k)\|_2 < \hat{\epsilon}\frac{\sqrt{1-\frac{\hat{\epsilon}^2}{4}}}{1-\frac{\hat{\epsilon}^2}{2}}$ and we see that the postulate on the size of ϵ is such that $\|\log_m(V_k)\|_2 < \delta_0$. Thus Lemma 3 indeed applies at iteration k, which closes the induction. ∎

Remark: The inequality $\hat{\epsilon}\frac{\sqrt{1-\frac{\hat{\epsilon}^2}{4}}}{1-\frac{\hat{\epsilon}^2}{2}} < \delta_0$ holds precisely for $\hat{\epsilon} < \sqrt{2}\left(1 - \frac{1}{\sqrt{1+\delta_0^2}}\right)^{\frac{1}{2}} =: \hat{\epsilon}_0$. A further calculations shows that if $\epsilon < 0.0912$, then $\hat{\epsilon} = (e^{2\tilde{\epsilon}} - 1) + \epsilon^2 + \epsilon < \hat{\epsilon}_0$, i.e., the conditions of Lemma 4 hold, for all $\epsilon < 0.0912$.

With the tools established above at hand, we are now in a position to prove Theorem 2.

Proof (Theorem 2). Let $(V_k)_{k\in\mathbb{N}_0}$ be the sequence of orthogonal matrices generated by Algorithm 1. By Lemma 3 and Lemma 4, it holds

$$\|\log_m V_k\|_2 := \|\begin{pmatrix} A_k & -B_k^T \\ B_k & C_k \end{pmatrix}\|_2 < \frac{1}{2}(\sqrt{5}-1), \quad \|C_{k+1}\|_2 < \alpha^{k+1}\|C_0\|_2 \tag{22}$$

for all $k \geq 0$, where $0 < \alpha < \frac{1}{2}$. From this equation and the continuity of the matrix exponential, we obtain

$$\lim_{k\to\infty} W_k = \lim_{k\to\infty} \begin{pmatrix} I_p & 0 \\ 0 & \exp_m(-C_k) \end{pmatrix} = \begin{pmatrix} I_p & 0 \\ 0 & I_p \end{pmatrix}. \tag{23}$$

The convergence result is now an immediate consequence of Algorithm 1, step 10. The upper bound on the iteration count required for numerical convergence is also obvious from (22). ∎

5. Examples and experimental results. In this section, we discuss a special case that can be treated analytically. Following, we present numerical results on the performance of Algorithm 1.

5.1. A special case. Here, we consider the special situation, where the two points $U, \tilde{U} \in St(n,p)$ are such that their columns span the same subspace.[4] Hence, there exists an orthogonal matrix $M \in O_{p\times p}$ such that $\tilde{U} = UM = UM + (I - UU^T)0$. In this case, Algorithm 1 produces the initial matrices $V_0 = \begin{pmatrix} M & 0 \\ 0 & Y_0 \end{pmatrix}$ and $\Phi_0 = \exp_m(-\log_m(Y_0)) = Y_0^{-1}$. Note that the corresponding $W_0 = \begin{pmatrix} I_p & 0 \\ 0 & Y_0^{-1} \end{pmatrix}$ commutes with V_0. Thus, we have the reduced BCH formula $\log_m(V_0W_0) = \log_m(V_0) + \log_m(W_0) = \begin{pmatrix} \log_m(M) & 0 \\ 0 & 0 \end{pmatrix}$, i.e., Algorithm 1 converges after a single iteration and gives

$$Log_U^{St}(UM) = U\log_m(M). \tag{24}$$

(Of course, it is also straight forward to show this directly without invoking Algorithm 1.) Let $\sigma(M) = \{e^{i\varphi_1}, \dots, e^{i\varphi_p}\}$ be the spectrum of $M \in O_{p\times p}$ and suppose that M is such that none of its eigenvalues is on the negative real axis, i.e., $\varphi_j \in (-\pi,\pi)$. Then, the maximal Riemannian distance between two points U and UM is bounded by

$$\begin{aligned} dist(U,UM) &= \sqrt{\langle U\log_m(M), U\log_m(M)\rangle_U} \\ &= \left(\frac{1}{2}tr(\log_m(M)^T\log_m(M))\right)^{\frac{1}{2}} = \left(\frac{1}{2}\sum_{j=1}^{p}\varphi_j^2\right)^{\frac{1}{2}}. \end{aligned}$$

As a consequence

$$dist(U,UM) < \begin{cases} \sqrt{\frac{p}{2}}\pi, & p \text{ even}, \\ \sqrt{\frac{p-1}{2}}\pi, & p \text{ odd}. \end{cases}$$

The latter fact holds, because the eigenvalues of M come in complex conjugate pairs. Hence, if p is odd, there is at least one real eigenvalue $\lambda_j = e^{i\varphi_j}$ and because $\varphi_j \in (-\pi,\pi)$, there is at least one zero argument $\varphi_j = 0$. Related is [6, eq. (2.15)].

[4] We may alternatively express this by saying that $[U] := \text{colspan}(U)$ and $[\tilde{U}] := \text{colspan}(\tilde{U})$ are the same points on the Grassmann manifold $[U] = [\tilde{U}] \in Gr(n,p)$.

5.2. Numerical performance. First, we try to mimic the experiments featured in [18, §5.4]. Fig. 5.5 (lower left) of the aforementioned reference shows the average iteration count when applying the optimization-based Stiefel logarithm to matrices within a Riemannian annulus of inner radius 0.4π and outer radius 0.44π around $(I_p, 0)^T \in St(n,p)$ for dimensions of $n = 10, p = 2$. Convergence is detected, if $\|C_k\|_F < \tau = 10^{-7}$, where C_k is the same as in Algorithm 1. ([18, Alg. 4, p. 91] uses $\tau^2 < 10^{-14}$). Since [18, §5.4] does not list the precise input data, we create comparable data randomly. To this end, we fix an arbitrary point $U \in St(10,2)$ and create artificially but randomly another point $\tilde{U} \in St(10,2)$ such that the Riemannian distance from U to $\tilde{U}$ is exactly 0.44π. For full comparability, we replace the 2-norm in Algorithm 1, line 7 with the Frobenius norm. We average over 1000 random experiments and arrive at an average iteration count of $\bar{k} = 7.83$. A MATLAB script that performs the required computations is available in Section S2. When the distance of U and $\tilde{U}$ is lowered to 0.4π, the average iteration count drops to a value of $\bar{k} = 6.92$.

As a second experiment, we now return to the 2-norm and lower the convergence threshold to $\|C_k\|_2 < \tau = 10^{-13}$ in the convergence criterion of Algorithm 1. We create randomly points $U, \tilde{U} \in St(n,p)$ that are also a Riemannian distance of 0.44π away from each other, where we consider various different dimensions (n,p), see Table 1. We apply Algorithm 1 to compute $\Delta = Log_U^{St}(\tilde{U})$.

Table 1
Convergence of Algorithm 1 for random data to an accuracy of $\|C_k\|_2 \leq 10^{-13}$.

(n,p)	$\text{dist}\left(U,\tilde{U}\right)$	$\|U-\tilde{U}\|_2$	iters.	$\|\Delta - Log_U^{St}(\tilde{U})\|_2$	time
(10,2)	0.44π	1.0179	16	8.7903e-15	0.01s
(10,2)	0.89π	1.7117	95	4.1934e-13	0.06s
(1,000, 200)	0.44π	0.1616	5	1.5119e-14	0.7s
(1,000, 200)	0.89π	0.3256	7	1.7272e-14	0.8s
(1,000, 900)	0.44π	0.1234	4	9.6999e-14	16.1s
(1,000, 900)	0.89π	0.2491	5	7.9052e-14	21.0s
(100,000, 500)	0.44π	0.0875	4	5.9857e-14	13.1s
(100,000, 500)	0.89π	0.1768	5	6.1041e-14	14.0s

Figure 1 shows the associated convergence histories. The associated computation times[5] are listed in Table 1. As can be seen from the figure and the table, Algorithm 1 converges slowest (in terms of the iteration count) in the case of $St(10,2)$. Note that in this case, the constant $\|U - \tilde{U}\|_2$ that played a major role in the convergence analysis of Algorithm 1 is largest. Moreover, we observe that the algorithm converges in all test cases even though in only one of the experiments the theoretical convergence guarantee $\|U_0 - \tilde{U}\|_2 < 0.09$ is satisfied, so that the theoretical bound derived here can probably be improved. Table 1 suggests that the impact of the size of $\|U - \tilde{U}\|_2$ on the iteration count is more direct than that of the actual Riemannian distance.

We repeat the exercise with random data $U, \tilde{U} \in St(n,p)$ that are a distance of 0.89π apart, which is the lower bound for the injectivity radius on the Stiefel manifold given in [18, eq. (5.14)]. In the case of $St(10,2)$, we hit a random matrix pair $U, \tilde{U}$, where the associated value$\|U - \tilde{U}\|_2$ is so large that the conditions of Theorem 2 *and* Lemma 3, Lemma 4 do not hold. In fact, we have $\|\log_m(V_0)\|_2 = 3.141$ for the

[5] as measured on a Dell desktop computer endowed with six processors of type Intel(R) Core(TM) i7-3770 CPU@3.40GHz

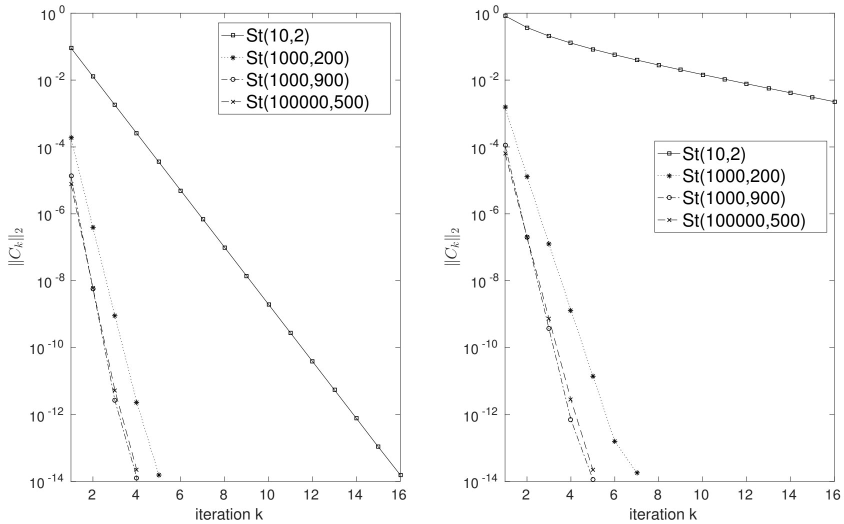

FIG. 1. *Convergence of Algorithm 1 for random data $U, \tilde{U} \in St(n,p)$ for various n and p. Convergence accuracy is set to $\|C_k\|_2 \leq 10^{-13}$. Left: convergence graphs for $dist(U,\tilde{U}) = 0.44\pi$; right: for $dist(U,\tilde{U}) = 0.89\pi$.*

starting point of Algorithm 1 in this case, which is close to π. Yet, the algorithm converges, but very slowly so, see Table 1, second row and Figure 1, right side. In all of the other cases, Algorithm 1 converges in well under ten iterations, even for the larger test cases.

A MATLAB script that performs the required computations is available in Section S2.

5.3. Dependence of the convergence on the Riemannian and the Euclidean distance. In this section, we examine the convergence of Algorithm 1 depending on the Riemannian distance $dist(U,\tilde{U})$ and the distance $\|U - \tilde{U}\|_2$ in the Euclidean operator-2-norm. To this end, we create a random point $U \in St(n,p)$ with MATLAB by computing the thin qr-decomposition of an $(n \times p)$ matrix with entries sampled uniformly from $(0,1)$. Likewise, we create a random tangent vector $\Delta \in T_U St(n,p)$ by chosing randomly a skew-symmetric matrix $A = \tilde{A} - \tilde{A}^T \in \mathbb{R}^{p\times p}$ and a matrix $T \in \mathbb{R}^{n\times p}$, where the entries of $\tilde{A}$ and T are again uniformly sampled from $(0,1)$, and setting $\tilde{\Delta} = UA + (I - UU^T)T$. We normalize $\tilde{\Delta}$ according to the canonical metric $\Delta = \frac{\tilde{\Delta}}{\sqrt{\langle \tilde{\Delta}, \tilde{\Delta}\rangle_U}}$, see Section 2. In this way, we obtain for every $t \in [0,\pi)$ a point $\tilde{U} = U(t)$ that is a Riemannian distance of $dist(U, U(t))) = \|t\Delta\|_U = t$ away from U.

We discretize the interval $[0.1, 0.9\pi)$ by 100 equidistant points $\{x_k | k = 1, \dots, 100\}$ and compute

- the number of iterations until convergence when computing $\log_U^{St}(U(t_k))$ with Algorithm 1 for $k = 1, \dots, 100$.
- the distance in spectral norm $\|U - U(t_k)\|_2$, $k = 1, \dots, 100$.
- the norm of the matrix logarithm of the first iterate $\|\log_m(V_0)\|_2$ from Algorithm 1, step 3.

The results are displayed in Figures 2, 3 and 4 for dimensions of $St(10,000,400)$, $St(100,10)$ and $St(4,2)$, respectively. In all cases, the convergence threshold was set to $\|C_l\|_2 < \tau = 10^{-13}$. The algorithm converged in all cases, where $\|\log_m(V_0)\|_2 < \pi$ and produced a tangent vector $\Delta(t_k) := \log^{St}_U(U(t_k))$ of accuracy $\|\Delta(t_k) - t_k\Delta\|_2 < 10^{-13}$. A MATLAB script that performs the required computations is available in Section S3. In the case of $St(4,2)$, the algorithm starts to fail for $t_k \approx \frac{\pi}{2}$, where $\|\log_m(V_0)\|_2$ jumps

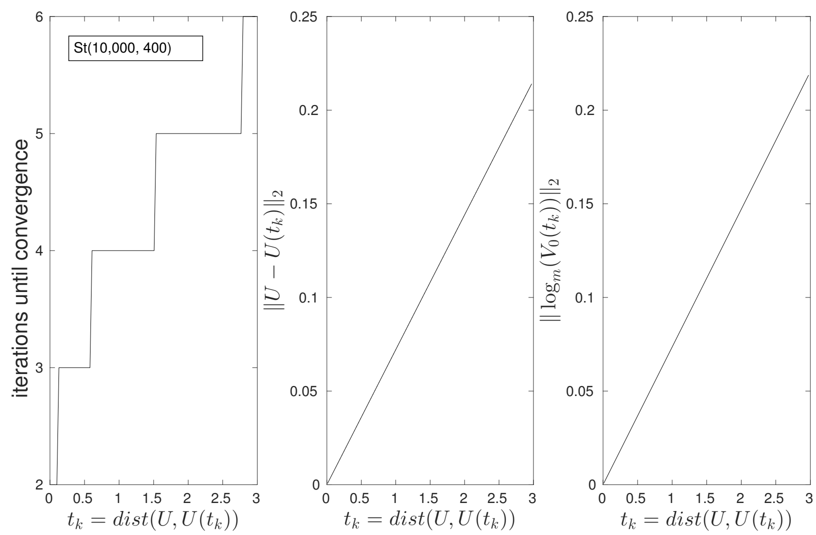

FIG. 2. *Convergence of Algorithm 1 for $U, \tilde{U} = U(t_k) = Exp^{St}_U(t_k\Delta) \in St(n,p)$, where Δ is a random tangent vector of canonical norm 1 and $n = 10,000$, $p = 400$. Convergence accuracy is set to $\|C_k\|_2 \leq 10^{-13}$. Left: number of iterations until convergence vs. $dist(U,\tilde{U})$; middle: $\|U - \tilde{U}\|_2$ vs. $dist(U,\tilde{U})$; right: $\|\log_m(V_0)\|_2$ vs. $dist(U,\tilde{U})$.*

to a value of π. This indicates that V_0 features (up to numerical errors) an eigenvalue $\lambda = -1$ so that the standard principal matrix logarithm is no longer well-defined. In all the experiments that were conducted, this behavior was observed only for small values of $p < 8$, while there was never produced a random data set where Algorithm 1 failed for $t < 0.9\pi$ and $p > 10$. The figures suggest that for small column-numbers p, the ratio between the Riemannian distance $dist(U,\tilde{U})$ and the spectral distance $\|U - \tilde{U}\|_2$ is smaller than in higher dimensions. Moreover, for smaller p, it seems to be more likely to hit a random tangent direction along which Algorithm 1 fails early than for higher p. This may partly be explained by the star-shaped nature of the domain of injectivity of the Riemannian exponential, [13, Lemma 5.7], and the richer variety of directions in higher dimensions.

From these observations, it is tempting to conjecture that Algorithm 1 will converge, whenever $\|\log_m(V_0)\|_2 < \pi$. However, these results are based on a limited notion of randomness and a more thorough examination of the numerical behavior of Algorithm 1 is required to obtained conclusive results, which is beyond the scope of this work. Note that the domain of convergence of Algorithm 1 is related to the injectivity radius of $St(n,p)$ but it does not have to be the same. In Section S1 from the supplement, we state an explicit example in $St(4,2)$, where Algorithm 1 produces a first iterate V_0 with $\lambda = -1$ for an input pair $U, \tilde{U} \in St(4,2)$ with $dist(U,\tilde{U}) = \frac{\pi}{2}$,

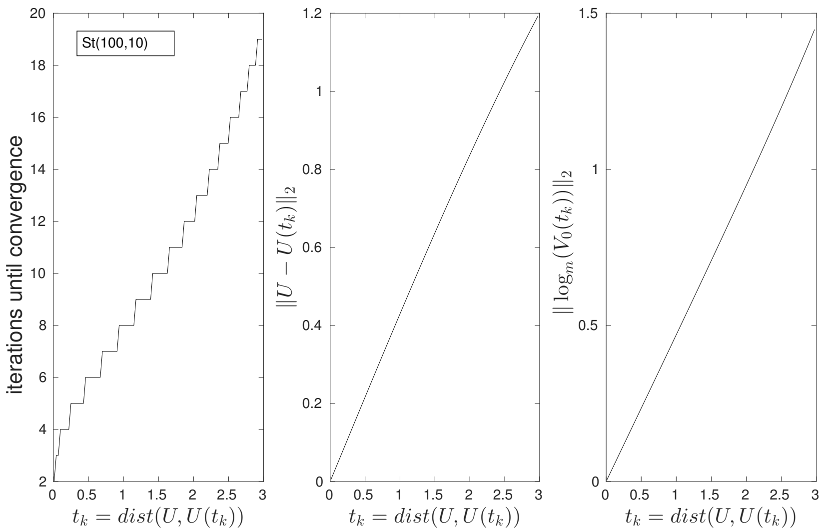

FIG. 3. *Same as Figure 2, but for $n = 100$, $p = 10$.*

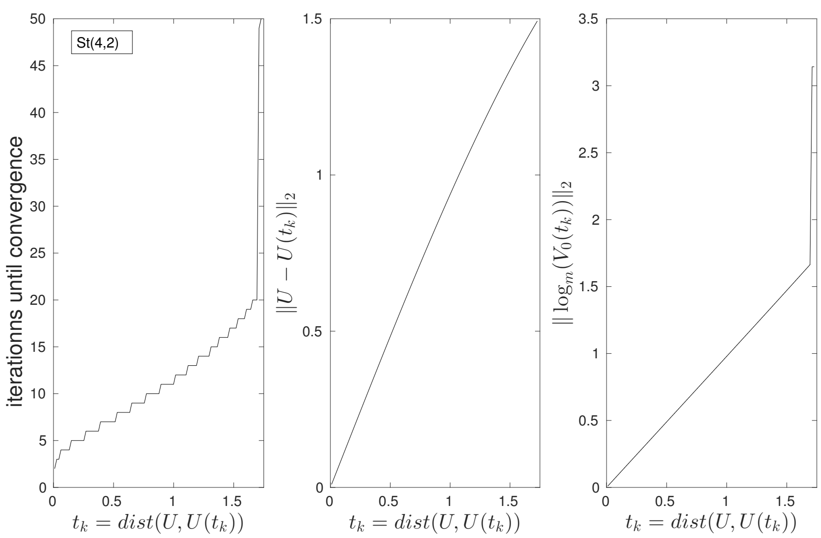

FIG. 4. *Same as Figure 2, but for $n = 4$, $p = 2$.*

while the injectivity radius is estimated to be $\approx 0.71\pi$ in [18, §5]. An analytical investigation in $St(4, 2)$ might be possible and may shed more light on the precise value of the Stiefel manifold's injectivity radius.

6. Conclusions and outlook. We have presented a matrix-algebraic derivation of an algorithm for evaluating the Riemannian logarithm $Log_U^{St}(\tilde{U})$ on the Stiefel manifold. In contrast to [18, Alg. 4, p. 91], the construction here is not based on an optimization procedure but on an iterative solution to a non-linear matrix equation. Yet, it turns out that both approaches lead to essentially the same numerical scheme. More precisely, our Algorithm 1 coincides with [18, Alg. 4, p. 91], when a unit step size is employed in the optimization scheme associated with the latter method. Apart from its comparatively simplicity, a key benefit is that our matrix-algebraic approach allows for a convergence analysis that does not require estimates on gradients nor Hessians and we are able to prove that the convergence rate of Algorithm 1 is at least linear. This, in turn, proves the local linear convergence of [18, Alg. 4, p. 91] when using a unit step size. The algorithm shows a very promising performance in numerical experiments, even when the dimensions n, p become large.

So far, we have carried out a theoretical *local* convergence analysis. Open questions to be tackled in the future include estimates on how large the convergence domain of Algorithm 1 is in terms of the Riemannian distance of the input points $dist(U, \tilde{U})$. This is related with the question of determining the injectivity radius of the Stiefel manifold. Estimates on the injectivity radius are featured in [18, §5.2.1].

Appendix A. A sharper majorizing series for Goldberg's Exponential series. As an alternative to Dynkin's BCH formula of nested commutators, Goldberg has shown in [8] that the solution to the exponential equation

$$\exp_m(X)\exp_m(Y) = \exp_m(Z)$$

can be written as a formal series

$$Z = X + Y + \sum_{k=2}^{\infty} z_k(X,Y), \quad z_k(X,Y) = \sum_{w,|w|=k} g_w w. \tag{26}$$

Each term $z_k(X,Y)$ in (26) is the sum over all *words* of length k in the alphabet $\{X, Y\}$. For example, $YXYX^2$ and X^2YXY^2 are such words of length 5 and 6 and thus contributing to $z_5(X,Y)$ and $z_6(X,Y)$, respectively. The coefficients are rational numbers $g_w \in \mathbb{Q}$, called Goldberg coefficients.

Thompson [20] has shown that the series converges provided that $\|X\|, \|Y\| \leq \mu < 1$ for any submultiplicative norm $\|\cdot\|$. More precisely, his result is that $\|z_k(X,Y)\| = \|\sum_{w,|w|=k} g_w w\| \leq 2\mu^k$ for $k \geq 2$, see also [16, eq. 2]. In the next lemma, we improve this bound by cutting the factor 2.

Lemma 5. *Let $\|X\|, \|Y\| \leq \mu < 1$. The Goldberg series is majorized by*

$$\|Z\| < \|X\| + \|Y\| + \sum_{k=2}^{\infty} \mu^k.$$

Proof. One ingredient of Thompson's proof is the following basic estimate on binomial terms:

$$m \binom{m-1}{\lfloor \frac{m}{2} \rfloor} \geq 2^{m-1}. \tag{27}$$

Here, $\lfloor x \rfloor$ denotes the largest integer smaller or equal to x. Thompson's argument is that $2^{m-1} = (1+1)^{m-1} = \sum_{l=0}^{m-1} \binom{m-1}{l}$ and that $\binom{m-1}{\lfloor \frac{m}{2} \rfloor}$ is the largest out of the m terms in the binomial sum. (It appears twice, if $m-1$ is odd.) In the following, we prefer to write this term with using the ceil-operator as $\binom{m-1}{\lfloor \frac{m}{2} \rfloor} = \binom{m-1}{\lceil \frac{m-1}{2} \rceil}$, because in this way, the same index $m-1$ appears in the upper and lower entry of the binomial coefficient.

For larger m, the inequality (27) can in fact be improved by a factor of 2:

$$\text{Claim: } m\binom{m-1}{\lceil \frac{m-1}{2} \rceil} > 2^m \text{ for all } m \geq 7. \tag{28}$$

For $m = 7$, we have $7\binom{7-1}{\lceil \frac{7-1}{2} \rceil} = 7 \cdot 20 = 140 > 128 = 2^7$; for $m = 8$, the inequality evaluates to $280 > 256 = 2^8$. To prove the claim, we proceed by induction.

Case 1: "m even". In this case, $\lceil \frac{m}{2} \rceil = \frac{m}{2} = \lceil \frac{m-1}{2} \rceil$ and

$$(m+1)\binom{m}{\lceil \frac{m}{2} \rceil} = (m+1)\left(\binom{m-1}{\frac{m}{2}-1} + \binom{m-1}{\frac{m}{2}}\right)$$

$$= 2(m+1)\binom{m-1}{\lceil \frac{m-1}{2} \rceil} > 2(m+1)\frac{2^m}{m} > 2^{m+1}, \tag{29a}$$

where we have used the symmetry in the Pascal triangle ($m-1$ is odd) and the induction hypothesis to arrive at (29a).

Case 2: "m odd". In this case, $\lceil \frac{m}{2} \rceil = \frac{m+1}{2}$ and

$$(m+1)\binom{m}{\lceil \frac{m}{2} \rceil} = (m+1)\left(\binom{m-1}{\frac{m+1}{2}-1} + \binom{m-1}{\frac{m+1}{2}}\right)$$

$$= (m+1)\left(\binom{m-1}{\lceil \frac{m-1}{2} \rceil} + \binom{m-1}{\lceil \frac{m}{2} \rceil}\right). \tag{30a}$$

Note that $\binom{m-1}{\lceil \frac{m}{2} \rceil}$ is the second-to-largest term in the binomial expansion of $(1+1)^{m-1}$. Moreover, since $m-1$ is even, the relation to the largest term is

$$\binom{m-1}{\lceil \frac{m}{2} \rceil} = \frac{m-1}{m+1}\binom{m-1}{\lceil \frac{m-1}{2} \rceil}.$$

Substituting in (30a) and applying the induction hypothesis gives

$$(m+1)\binom{m}{\lceil \frac{m}{2} \rceil} > (m+1)\left(\frac{2^m}{m} + \frac{m-1}{m+1}\frac{2^m}{m}\right) = \left(\frac{m+1}{m} + \frac{m-1}{m}\right)2^m = 2^{m+1}.$$

Using (28) rather than (27) in Thompson's original proof leads to the improved bound of $\|z_k(X;Y)\| \leq \mu^k$ for $k \geq 7$.

We tackle the terms involving words of lengths $k = 2, 3, \ldots, 6$ manually. The reference [21] lists explicit expressions of the summands in the Goldberg BCH series up to z_8. The first three of them read

$$z_2(X,Y) = \frac{1}{2}(XY - YX) \Rightarrow \|z_2(X,Y)\| \leq \frac{2}{2}\mu^2.\checkmark$$

$$z_3(X,Y) = \frac{1}{12}\big(X^2Y - 2XYX + XY^2 + YX^2 - 2YXY + Y^2X\big)$$

$$\Rightarrow \|z_3(X,Y)\| \leq \frac{8}{12}\mu^3.\checkmark$$

$$z_4(X,Y) = \frac{1}{24}\big(X^2Y^2 - 2XYXY + 2YXYX - Y^2X^2\big) \Rightarrow \|z_4(X,Y)\| \leq \frac{6}{24}\mu^4.\checkmark$$

The expressions for $z_5(X,Y)$ and $z_6(X,Y)$ are too cumbersome to be restated here. However, for our purposes, a very rough counting argument is sufficient: The expression for $z_5(X,Y)$ features 30 length-5 words with non-zero Goldberg coefficient and the largest Goldberg coefficient is $\frac{1}{30}$. Hence, $\|z_5(X,Y)\| = \|\sum_{w,|w|=5} g_w w\| < \frac{30}{30}\mu^5$.✓ (A more careful consideration reveals $\|z_5(X,Y)\| \leq \frac{176}{720}\mu^5$.)

The expression for $z_6(X,Y)$ features 28 length-6 words with non-zero Goldberg coefficient and the largest Goldberg coefficient is $\frac{1}{60}$. Hence, $\|z_6(X,Y)\| = \|\sum_{w,|w|=6} g_w w\| \leq \frac{28}{60}\mu^6$.✓ ∎

Appendix B. Norm bound for $\log_m(V_0)$.

Proposition 6. *Let $C \in \mathbb{R}^{p\times p}$ be skew-symmetric with $\|C\|_2 < \pi$. Then*

$$\|\exp_m(C) - I\|_2 < \|C\|_2.$$

Proof. Since C is skew-symmetric, it features an EVD $C = Q\Lambda Q^H$ with $\Lambda = diag(\lambda_1,\ldots,\lambda_p) = diag(i\varphi_1,\ldots,i\varphi_p)$, where $\varphi \in (-\pi,\pi)$ and $\max_j |i\varphi_j| = \|C\|_2$. Therefore, $\exp_m(C) = Q\exp_m(\Lambda)Q^H$ with $\exp_m(\Lambda) = diag(e^{i\varphi_1},\ldots,e^{i\varphi_p})$ and

$$\|\exp_m(C) - I\|_2 = \max_j |e^{i\varphi_j} - 1| < \max_j |\varphi_j| = \|C\|_2.$$

(The latter estimate may also be deduced from Figure 5.) ∎

Proposition 7. *Let $V \in O_{n\times n}$ be such that $\|V - I\|_2 < r < 1$. Then*

$$\|\log_m(V)\|_2 < r\frac{\sqrt{1-\frac{r^2}{4}}}{1-\frac{r^2}{2}}.$$

Proof. Let $E = V - I$. The matrices V and E share the same (orthonormal) basis of eigenvectors Q and the spectrum of V is precisely the spectrum of E shifted by $+1$. By assumption, $r > \|E\|_2 = \max_{\mu\in\sigma(E)} |\mu|$. Hence, the eigenvalues $\lambda \in \sigma(V)$ are complex numbers of modulus one of the form $\lambda = e^{i\alpha} = 1 + \mu$, with $|\mu| < r$. Thus, λ lies on the unit circle but within a ball of radius r around $1 \in \mathbb{C}$, see Figure 5. The maximal angle α for such a λ is bounded by the slope of the line that starts in $0 \in \mathbb{C}$ and crosses the points of intersection of the two circles $\{|z| < 1\}$ and $\{|z-1| < r\}$. The intersection points are $(x_s, \pm y_s) = \left(1 - \frac{r^2}{2}, \pm r\sqrt{1-\frac{r^2}{4}}\right)$. Therefore

$$|\alpha| < \arctan\left(\frac{y_s}{x_s}\right) = \arctan\left(\frac{r\sqrt{1-\frac{r^2}{4}}}{1-\frac{r^2}{2}}\right) < r\frac{\sqrt{1-\frac{r^2}{4}}}{1-\frac{r^2}{2}}.$$

As a consequence,

$$\|\log_m(V)\|_2 = \|Q\log_m(\Lambda)Q^H\|_2 = \max_{\lambda\in\sigma(V)} |\ln(\lambda)| = \max_{\lambda=e^{i\alpha}\in\sigma(V)} |i\alpha| < r\frac{\sqrt{1-\frac{r^2}{4}}}{1-\frac{r^2}{2}}.$$ ∎

Lemma 8. *Let $U, \tilde{U} \in St(n,p)$ with $\|U - \tilde{U}\|_2 < \epsilon$. Let M, N, X_0, Y_0 and $V_0 := \begin{pmatrix} M & X_0 \\ N & Y_0 \end{pmatrix} \in O_{2p\times 2p}$ be as constructed in the first steps of Algorithm 1.*

Then

$$\|\log_m(V_0)\|_2 < 2\epsilon\frac{\sqrt{1-\epsilon^2}}{1-2\epsilon^2}. \tag{32}$$

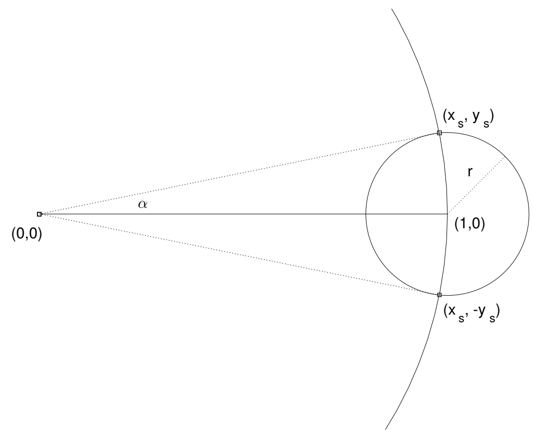

FIG. 5. *Geometrical illustration of Proposition 7 in the complex plane.*

Proof. Because V_0 is orthogonal,

$$1 = \|V_0\|_2 \geq \nu(V_0) = \max_{\|w\|_2=1} |w^H \begin{pmatrix} M & X_0 \\ N & Y_0 \end{pmatrix} w| \tag{33a}$$

$$\geq \max_{\|v\|_2=1} |(0, v^H) \begin{pmatrix} M & X_0 \\ N & Y_0 \end{pmatrix} \begin{pmatrix} 0 \\ v \end{pmatrix} | = \|Y_0\|_2, \tag{33b}$$

where $\nu(V_0)$ denotes the numerical radius of V_0, see [10, eq. 1.21, p. 21]. Likewise, $\|M\|_2 \leq 1$ so that the singular values of M and Y_0 range between 0 and 1. Moreover, by the Procrustes preprocessing outlined at the end of Section 3,

$$\|N\|_2 = \|X_0\|_2 < \epsilon, \quad \|M - I\|_2 < \epsilon, \quad \|Y_0 - I_p\|_2 < \epsilon^2,$$

see (11). Combining these facts, we obtain $V = I + (V - I) = I + E$, where

$$\|E\|_2 = \| \begin{pmatrix} M - I & X_0 \\ N & Y_0 - I \end{pmatrix} \|_2 \leq \| \begin{pmatrix} M - I & 0 \\ 0 & Y_0 - I \end{pmatrix} + \begin{pmatrix} 0 & X_0 \\ N & 0 \end{pmatrix} \|_2$$
$$\leq \max\{\|M - I\|_2, \|Y_0 - I\|_2\} + \max\{\|N\|_2, \|X_0\|_2\} < \max\{\epsilon, \epsilon^2\} + \epsilon = 2\epsilon.$$

Applying Proposition 7 to $V = I + E$ proves the claim. ∎

Appendix C. MATLAB code.

C.1. Algorithm 1.

```
%
function [Delta, k, conv_hist, norm_logV0] = ...
                                Stiefel_Log_supp(U0, U1, tau)
%------------------------------------------------------------
%@author: Ralf Zimmermann, IMADA, SDU Odense
%
% Input arguments
%  U0, U1 : points on St(n,p)
```

```
%     tau : convergence threshold
% Output arguments
%   Delta : Log^{St}_U0(U1),
%           i.e. tangent vector such that Exp^St_U0(Delta) = U1
%       k : iteration count upon convergence
% supplementary output
%  conv_hist : convergence history
% norm_logV0 : norm of matrix log of first iterate V0
%------------------------------------------------------------
% get dimensions
[n,p] = size(U0);
% store convergence history
conv_hist = [0];

% step 1
M = U0'*U1;
% step 2
[Q,N] = qr(U1 - U0*M,0);   % thin qr of normal component of U1
% step 3
[V, ~] = qr([M;N]);                   % orthogonal completion

% "Procrustes preprocessing"
[D,S,R]      = svd(V(p+1:2*p,p+1:2*p));
V(:,p+1:2*p) = V(:,p+1:2*p)*(R*D');
V            = [[M;N], V(:,p+1:2*p)];  %          |M  X0|
                                       % now, V = |N  Y0|
% just for the record
norm_logV0 = norm(logm(V),2);

% step 4: FOR-Loop
for k = 1:10000
    % step 5
    [LV, exitflag] = logm(V);
                                  % standard matrix logarithm
                                  %             |Ak  -Bk'|
                                  % now, LV =   |Bk   Ck |
    C = LV(p+1:2*p, p+1:2*p);     % lower (pxp)-diagonal block
    % steps 6 - 8: convergence check
    normC = norm(C, 2);
    conv_hist(k) = normC;
    if normC<tau;
        disp(['Stiefel log converged after ', num2str(k),...
              ' iterations.']);
        break;
    end
    % step 9
    Phi = expm(-C);              % standard matrix exponential
    % step 10
    V(:,p+1:2*p) = V(:,p+1:2*p)*Phi;   % update last p columns
end
```

```
% prepare output                         |A  -B'|
% upon convergence, we have  logm(V) =   |B   0 | = LV
%     A = LV(1:p,1:p);     B = LV(p+1:2*p, 1:p)
% Delta = U0*A+Q*B
Delta = U0*LV(1:p,1:p) + Q*LV(p+1:2*p, 1:p);
return;
end
```

Note: The performance of this method may be enhanced by computing $\exp_m$, $\log_m$ via a Schur decomposition.

REFERENCES

[1] P.-A. Absil, R. Mahony, and R. Sepulchre, *Riemannian geometry of Grassmann manifolds with a view on algorithmic computation*, Acta Applicandae Mathematica, 80 (2004), pp. 199–220, doi:10.1023/B:ACAP.0000013855.14971.91, http://dx.doi.org/10.1023/B%3AACAP.0000013855.14971.91.

[2] P.-A. Absil, R. Mahony, and R. Sepulchre, *Optimization Algorithms on Matrix Manifolds*, Princeton University Press, Princeton, New Jersey, 2008, http://press.princeton.edu/titles/8586.html.

[3] L. Balzano, R. Nowak, and B. Recht, *Online identification and tracking of subspaces from highly incomplete information*, in Proceedings of Allerton, September 2010. Available at http://arxiv.org/abs/1006.4046.

[4] E. Begelfor and M. Werman, *Affine invariance revisited*, 2012 IEEE Conference on Computer Vision and Pattern Recognition, 2 (2006), pp. 2087–2094, doi:http://doi.ieeecomputersociety.org/10.1109/CVPR.2006.50.

[5] P. Benner, S. Gugercin, and K. Willcox, *A survey of projection-based model reduction methods for parametric dynamical systems*, SIAM Review, 57 (2015), pp. 483–531, doi:10.1137/130932715, http://dx.doi.org/10.1137/130932715, arXiv:http://dx.doi.org/10.1137/130932715.

[6] A. Edelman, T. A. Arias, and S. T. Smith, *The geometry of algorithms with orthogonality constraints*, SIAM Journal on Matrix Analysis and Application, 20 (1999), pp. 303–353, doi:10.1137/S0895479895290954, http://dx.doi.org/10.1137/S0895479895290954.

[7] K. Gallivan, A. Srivastava, X. Liu, and P. Van Dooren, *Efficient algorithms for inferences on Grassmann manifolds*, in Statistical Signal Processing, 2003 IEEE Workshop on, 2003, pp. 315–318, doi:10.1109/SSP.2003.1289408.

[8] K. Goldberg, *The formal power series for* $\log e^x e^y$, Duke Math. J., 23 (1956), pp. 13–21, doi:10.1215/S0012-7094-56-02302-X, http://dx.doi.org/10.1215/S0012-7094-56-02302-X.

[9] G. H. Golub and C. F. Van Loan, *Matrix Computations*, The John Hopkins University Press, Baltimore – London, 3 ed., 1996.

[10] A. Greenbaum, *Iterative Methods for solving linear systems*, vol. 17 of Frontiers in applied mathematics, SIAM Society for Industrial and Applied Mathematics, Philadelphia, 1997.

[11] N. J. Higham, *Functions of Matrices: Theory and Computation*, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2008.

[12] S. Kobayashi and K. Nomizu, *Foundations of Differential Geometry*, vol. I of Interscience Tracts in Pure and Applied Mathematics no. 15, John Wiley & Sons, New York – London – Sidney, 1963.

[13] J. Lee, *Riemannian Manifolds: an Introduction to Curvature*, Springer Verlag, New York – Berlin – Heidelberg, 1997.

[14] Y. Man Lui, *Advances in matrix manifolds for computer vision*, Image and Vision Computing, 30 (2012), pp. 380–388, doi:http://dx.doi.org/10.1016/j.imavis.2011.08.002, http://www.sciencedirect.com/science/article/pii/S0262885611000692.

[15] MATLAB, *version 7.10.0 (R2010a)*, The MathWorks Inc., Natick, Massachusetts, 2010.

[16] M. Newman, S. Wasin, and R. C. Thompson, *Convergence domains for the Campbell-Baker-Hausdorff formula*, Linear and Multilinear Algebra, 24 (1989), pp. 301–310.

[17] I. U. Rahman, I. Drori, V. C. Stodden, D. L. Donoho, and P. Schröder, *Multiscale representations for manifold-valued data*, SIAM J. Mult. Model. Simul., 4 (2005), pp. 1201–1232.

[18] Q. Rentmeesters, *Algorithms for data fitting on some common homogeneous spaces*, PhD thesis, Université Catholique de Louvain, Louvain, Belgium, July 2013/2015, http://hdl.

handle.net/2078.1/132587.

[19] W. Rossmann, *Lie Groups: An Introduction Through Linear Groups*, Oxford graduate texts in mathematics, Oxford University Press, 2006, https://books.google.de/books?id=bAjulQ65W-UC.

[20] R. C. Thompson, *Convergence proof for Goldberg's exponential series*, Linear Algebra and its Applications, 121 (1989), pp. 3–7.

[21] A. Van-Brunt and M. Visser, *Simplifying the Reinsch algorithm for the Baker-Campbell-Hausdorff series.* arXiv:1501.05034, 2015, http://arxiv.org/abs/1501.05034.