

RIEMANNIAN PRECONDITIONING*

BAMDEV MISHRA[†] AND RODOLPHE SEPULCHRE[‡]

Abstract. This paper exploits a basic connection between sequential quadratic programming and Riemannian gradient optimization to address the general question of selecting a *metric* in Riemannian optimization, in particular when the Riemannian structure is sought on a *quotient* manifold. The proposed method is shown to be particularly insightful and efficient in quadratic optimization with orthogonality and/or rank constraints, which covers most current applications of Riemannian optimization in matrix manifolds.

Key words. Riemannian optimization, sequential quadratic programming, metric tuning, quotient manifold, Grassmann, low rank, generalized eigenvalue problem, Lyapunov equation

AMS subject classifications. 65K05, 90C30, 90C22, 90C25, 90C26, 90C27, 90C46, 58C05, 49M15

DOI. 10.1137/140970860

1. Introduction. Gradient algorithms are a method of choice for large-scale optimization, but their convergence properties critically depend on the choice of a suitable *metric*. Good metrics can lead to superlinear convergence, whereas bad metrics can lead to very slow convergence. Goodness of the metric depends on its ability to encode second-order information about the optimization problem. For general optimization problems with equality constraints, sequential quadratic programming (SQP) methods provide an efficient selection procedure based on (approximating) the Hessian of a local quadratic approximation of the problem [32, Chapter 18]. This approach is Lagrangian; that is, it lifts the constraint into the cost function. An alternative is to embed the constraint into the search space, leading to unconstrained optimization on a nonlinear search space. Selecting the metric then amounts to equipping the search space with a Riemannian structure [2, 13, 35]. Riemannian optimization has gained popularity in recent years because of the particular nature of the constraints that show up in many matrix or tensor applications, in particular *orthogonality* and *rank* constraints. Such constraints are nonlinear and nonconvex but nevertheless very special [13, 26]. In particular, they are efficiently encoded through matrix factorizations, and their underlying geometry makes the search space sufficiently structured to make the machinery of Riemannian optimization competitive with alternative approaches, including convex relaxations. A current limitation of Riemannian optimization is, however, in the choice of the metric. Previous work has

*Received by the editors May 29, 2014; accepted for publication (in revised form) December 29, 2015; published electronically March 3, 2016. This paper presents research results of the Belgian Network DYSCO (Dynamical Systems, Control, and Optimization), funded by the Interuniversity Attraction Poles Programme, initiated by the Belgian State, Science Policy Office. The scientific responsibility rests with its authors. This work was also partly supported by the Belgian FRFC (Fonds de la Recherche Fondamentale Collective).

<http://www.siam.org/journals/siopt/26-1/97086.html>

[†]Amazon Development Centre India, Bangalore 560055, India (Bamdevm@amazon.com). This author was supported as a research fellow of the Belgian National Fund for Scientific Research (FNRS). Most of the work was done while this author was with the Department of Electrical Engineering and Computer Science, University of Liège, 4000 Liège, Belgium, and was visiting the Department of Engineering (Control Group), University of Cambridge, Cambridge CB2 1PZ, UK.

[‡]Department of Engineering, University of Cambridge, Cambridge CB2 1PZ, UK, and Department of Electrical Engineering and Computer Science, University of Liège, 4000 Liège, Belgium (R.Sepulchre@eng.cam.ac.uk).

mostly focused on the search space, exploiting the differential geometry of the constraint, but disregarding the role of the cost function. This limitation was pointed out early [22] and has been addressed in a number of recent contributions that emphasized the importance of preconditioning [25, 28, 29, 31, 38], but with no general procedure. The simple observation, and the main contribution, of the present paper is that SQP provides a systematic framework for choosing a metric in Riemannian optimization in a way that takes into consideration both the cost function and the constrained search space. This connection seems novel and insightful, especially in the situation where the unconstrained optimization problem is formulated on a quotient manifold, leading to a situation where the Hessian of the Lagrangian is singular in the total space. Most notably, this situation covers optimization problems on the Grassmann manifold and on the manifold of matrices of fixed rank.

This paper advocates that the use of SQP to select a metric in Riemannian optimization is general and connects two rather independent areas of constrained optimization. We focus in particular on the special case of quadratic cost functions with orthogonality and/or rank constraints. This particular situation encompasses a great deal of current successful applications of Riemannian optimization, including the popular generalized eigenvalue problem [3, 13] and linear matrix equation problems [8, 38]. Even in these highly researched problems, we show that SQP methods unify a number of recent disparate results and provide novel metrics. In the eigenvalue problem, where both the cost and the constraints are quadratic, the SQP method suggests a parameterized family of Riemannian metrics that provides novel insight on the role of shifts in the power, inverse, and Rayleigh quotient iteration methods. In the problem of solving linear matrix equations, low-rank matrix factorizations make the cost function quadratic in each of the factors, leading to Riemannian metrics rooted in block-diagonal approximations of the Hessian. In all of the mentioned applications, we stress the complementary, but not equivalent, role of SQP and Riemannian optimization: the SQP method provides a framework for selecting the metric from the (degenerate) Lagrangian in the total space, while the Riemannian framework provides the necessary generalization of unconstrained optimization to quotient manifolds, allowing for rigorous design and convergence analysis of a broad class of quasi-Newton algorithms in optimization problems over *classes of equivalences* of matrices.

We view this approach of selecting a metric from the Lagrangian as a form of *Riemannian preconditioning*. Similar to the notion of preconditioning in the unconstrained case [32, Chapter 5], the chosen Riemannian metrics have a preconditioning effect on optimization algorithms.

This paper does not aim at a comprehensive treatment of the topic, but rather focuses on connections between several classical branches of matrix calculus: matrix factorizations and shifts in numerical linear algebra, Riemannian submersions in differential geometry, and SQP in constrained optimization. The connection between SQP and the Riemannian Newton method on submanifolds is shown in section 2. The general idea of using the Lagrangian to select a metric in Riemannian optimization for constraints with symmetries is presented in section 3. We discuss the choice of the metric depending on the curvature properties of both the cost and the constraint and the interpretation of the Lagrange parameter as a shift. Section 4 develops the specific situation of quadratic cost and orthogonality constraints, revisiting the classical generalized eigenvalue problem. Section 5 further develops the specific situation of quadratic cost and rank constraints, with applications to solving large-scale matrix Lyapunov equations. All numerical illustrations use the MATLAB toolbox Manopt [10].

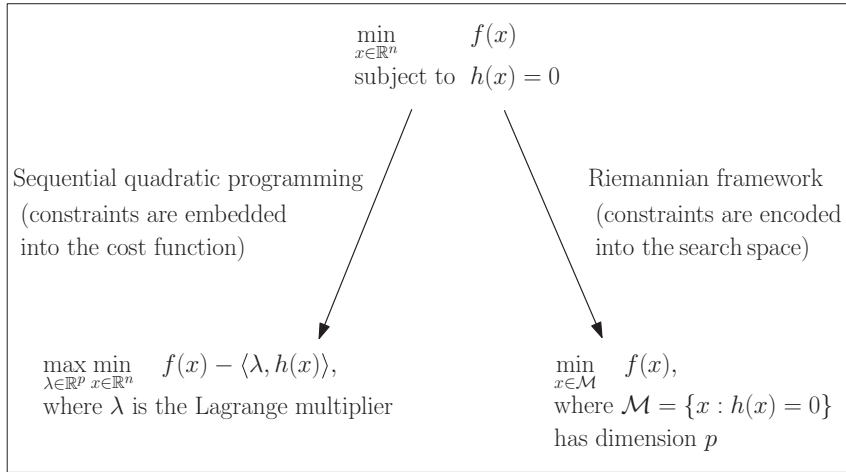


FIG. 1. Two complementary viewpoints on optimization with equality constraints.

2. SQP as Riemannian Newton method. Consider the optimization problem

$$(2.1) \quad \begin{aligned} \min_{x \in \mathbb{R}^n} \quad & f(x) \\ \text{subject to} \quad & h(x) = 0, \end{aligned}$$

where $f : \mathbb{R}^n \rightarrow \mathbb{R}$ and $h : \mathbb{R}^n \rightarrow \mathbb{R}^p$ are smooth functions. To illustrate the connection between SQP and Riemannian optimization, we first consider the simplified situation in which the set $\mathcal{M} := \{x : h(x) = 0\}$ has the structure of an *embedded differentiable submanifold* of \mathbb{R}^n [2, section 3.3], i.e., $h_x(x)$ is full rank everywhere, where $h_x(x)$ is the first-order derivative of h at x and is represented as an $n \times p$ matrix.

We recall two complementary approaches for the problem (2.1), namely the sequential quadratic approach and the Riemannian approach. The schematic view is shown in Figure 1. We aim at connecting these two approaches in order to tune the metric on the search space in such a way that it incorporates *second-order* information of the problem.

2.1. The constrained optimization viewpoint (SQP). SQP is a particularly well known approach for equality constrained nonlinear optimization [32, Chapter 18]. The core idea is to deal with the unconstrained *Lagrangian* $\mathcal{L} : \mathbb{R}^n \times \mathbb{R}^p \rightarrow \mathbb{R} : (x, \lambda) \mapsto \mathcal{L}(x, \lambda)$ with two sets of parameters, defined as

$$(2.2) \quad \mathcal{L}(x, \lambda) = f(x) - \langle \lambda, h(x) \rangle,$$

where $\langle \cdot, \cdot \rangle$ is the Euclidean inner product and $\lambda \in \mathbb{R}^p$ is the *Lagrange multiplier*. With the introduction of λ , the optimization problem (2.1) translates to the problem

$$(2.3) \quad \max_{\lambda \in \mathbb{R}^p} \min_{x \in \mathbb{R}^n} \mathcal{L}(x, \lambda),$$

which leads to a *primal-dual* iterative algorithm in $(x, \lambda) \in \mathbb{R}^n \times \mathbb{R}^p$, but the linearity in λ is exploited to reduce the number of variables. For example, in the neighborhood of a local minimum, the best *least-squares estimate* of the Lagrange multiplier is

$$(2.4) \quad \lambda_x = (h_x(x)^T h_x(x))^{-1} h_x(x)^T f_x(x),$$

TABLE 1
The SQP algorithm.

<p>The SQP algorithm for</p> $\begin{aligned} & \min_{x \in \mathbb{R}^n} && f(x) \\ & \text{subject to} && h(x) = 0. \end{aligned}$ <ol style="list-style-type: none"> 1. Compute the search direction ζ_x^* that is the solution of (2.5). 2. The next iterate x_+ is obtained by projecting $x + s\zeta_x^*$ onto the constrained space, where the step-size s is obtained with backtracking line search. 3. Repeat until convergence.
--

where $h_x(x)$ and $f_x(x)$ are the first-order derivatives of the functions h and f at x , respectively [32, equation (18.21)]. We use the convention that $f_x(x)$ is a column vector of length n and $h_x(x)$ is a full column rank matrix of size $n \times p$. The least-squares estimate of λ corresponds to the least-squares solution of the *optimality condition* $h_x(x)\lambda = f_x(x)$. It should be noted that any other estimate for $h_x(x)\lambda = f_x(x)$ is equally valid, provided that the equality is satisfied at optimality.

Substituting λ with λ_x transforms the primal-dual iteration in (x, λ) to a purely primal iteration in the variable x alone [32, page 539]. Once the Lagrange multiplier is estimated, the SQP optimization approach proceeds by minimizing the quadratic programming problem

$$(2.5) \quad \begin{aligned} & \arg \min_{\zeta_x \in \mathbb{R}^n} && f(x) + \langle f_x(x), \zeta_x \rangle + \frac{1}{2} \langle \zeta_x, D^2 \mathcal{L}(x, \lambda_x)[\zeta_x] \rangle \\ & \text{subject to} && Dh(x)[\zeta_x] = 0 \end{aligned}$$

at each iteration, where x is assumed to be in the constraint set (i.e., it is a feasible point), $f_x(x)$ is the first-order derivative of the cost function f , $D^2 \mathcal{L}(x, \lambda_x)[\zeta_x]$ is the second-order derivative of $\mathcal{L}(x, \lambda_x)$ with respect to x (keeping λ_x fixed) that is applied in the direction $\zeta_x \in \mathbb{R}^n$, $\langle \cdot, \cdot \rangle$ is the Euclidean inner product, and $Dh(x)[\zeta_x]$ is the Euclidean directional derivative of $h(x)$ in the direction $\zeta_x \in \mathbb{R}^n$, i.e., $Dh(x)[\zeta_x] = \lim_{t \rightarrow 0} (h(x + t\zeta_x) - h(x))/t$.

If the quantity $\langle \zeta_x, D^2 \mathcal{L}(x, \lambda_x)[\zeta_x] \rangle$ is strictly positive in the tangent space of constraints, then the problem (2.5) has a *unique solution* [32, section 18.1]. The solution ζ_x^* of (2.5) is interpreted as a *search direction*. Following the search direction with a step-size s , the next iterate x_+ in the SQP algorithm is obtained by projecting $x + s\zeta_x^*$ onto the constrained space to maintain strict feasibility of the iterates, where ζ_x^* is the solution to (2.5). This specific SQP algorithm is also called the *feasibly projected SQP method*, as every iterate x_+ is feasible [5, 39].

The resulting iterative algorithm is shown in Table 1 and has the properties of a Newton algorithm (locally in the neighborhood of a minimum) with favorable convergence properties [39]. The algorithm in Table 1 is the simplified version of the algorithm in [39]. Additionally, the SQP approach is appealing for the simplicity of its formulation.

2.2. Connecting SQP to Riemannian Newton method. An alternative treatment of problem (2.1) is to recast it as an unconstrained optimization problem on a nonlinear search space that encodes the constraint. For special constraints that are sufficiently structured, the framework leads to an efficient computational framework [2, 13, 35]. Particularly, concrete numerical algorithms, both first-order and second-order, are developed by endowing the set $\mathcal{M} := \{x : h(x) = 0\}$ with a *Riemannian*

structure.

The connection between SQP and the Riemannian framework for constraints that are embedded in \mathbb{R}^n (i.e., the manifold \mathcal{M} has the structure of an embedded submanifold) has been studied in [15], [13, section 4.9], and more recently in [5, section 4]. In particular, the following two results are relevant for the present paper.

PROPOSITION 2.1 (see [5, Proposition 4.1]). *If \mathcal{M} is an embedded submanifold of \mathbb{R}^n , then the SQP algorithm in Table 1 is equivalent to the Riemannian Newton method for the choice of the Lagrange multiplier λ in (2.4).*

PROPOSITION 2.2 (from the remark made by Gabay [15, Remark 4.3]). *Assume that \mathcal{M} is an embedded submanifold of \mathbb{R}^n and that $f : \mathcal{M} \rightarrow \mathbb{R}$ is a smooth function with isolated minima on \mathcal{M} . If $x^* \in \mathcal{M}$ is a local minimum of $f : \mathcal{M} \rightarrow \mathbb{R}$ on \mathcal{M} , then the second-order derivative of the Lagrangian along $\zeta_{x^*} \in T_{x^*}\mathcal{M}$ (the tangent space of \mathcal{M} at x^*), i.e., $\langle \zeta_{x^*}, D^2\mathcal{L}_x(x^*, \lambda_{x^*})[\zeta_{x^*}] \rangle$, captures all second-order information of the cost function f on \mathcal{M} .*

Propositions 2.1 and 2.2 show that SQP and the Riemannian Newton method are equivalent for embedded submanifolds. Furthermore, this connection emphasizes the role of the Hessian of the Lagrangian in the total space in effectively capturing second-order information of the problem, locally in the neighborhood of a minimum.

Our goal in the present paper is to build upon this connection in situations where the constrained search space is equipped with the differentiable structure of a *quotient manifold* rather than an *embedded submanifold*. This situation is frequent in applications due to symmetry properties of the optimization problem. The symmetries make the connection between SQP and Riemannian optimization less straightforward because the Lagrangian is *degenerate* at a local minimum. Nevertheless, we aim at showing that the structure of the Lagrangian in the total space is still very insightful in suggesting Riemannian metrics that capture second-order information on the quotient manifold.

3. Metric selection on quotient manifolds. Quotient manifolds are embedded in the manifold \mathcal{M} (that itself is in \mathbb{R}^n) up to an *equivalence relation* \sim . The search space is then the set of equivalence classes. Optima are not isolated in the computational *total space* \mathcal{M} , but they become isolated once we consider the abstract *quotient space* \mathcal{M}/\sim . (A precise characterization of the space \mathcal{M}/\sim is shown in section 3.1.)

Because of their prevalence in applications, quotient manifolds have been a focus of much research in Riemannian optimization [2, 13, 35]. Algorithms are run in the total space, but under appropriate compatibility between the Riemannian structure of the total space and the Riemannian structure of the quotient manifold, they define algorithms in the quotient space. The exposition for quotient manifolds here follows [2, Chapters 3, 5, and 8].

3.1. The Riemannian optimization viewpoint. Consider an equivalence relation \sim in the *total* (computational) space \mathcal{M} . The quotient manifold \mathcal{M}/\sim generated by this equivalence property consists of elements that are *equivalence classes* of the form $[x] = \{y \in \mathcal{M} : y \sim x\}$. Equivalently, if $[x]$ is an element in \mathcal{M}/\sim , then its matrix representation in \mathcal{M} is x . For example, the Grassmann manifold $\text{Gr}(r, n)$, which is the set of r -dimensional subspaces in \mathbb{R}^n , is obtained by the equivalence relationship $\text{Gr}(r, n) = \text{St}(r, n)/\mathcal{O}(r)$. $\text{St}(r, n)$ is the set of matrices of size $n \times r$ with orthonormal columns, and $\mathcal{O}(r)$ is the set of square $r \times r$ *orthogonal matrices*. Each element in the total space $\mathcal{M} := \text{St}(r, n) = \{\mathbf{X} \in \mathbb{R}^{n \times r} : \mathbf{X}^T \mathbf{X} = \mathbf{I}\}$ is characterized by

a matrix $\mathbf{X} \in \mathbb{R}^{n \times r}$ such that $\mathbf{X}^T \mathbf{X} = \mathbf{I}$. And an abstract element on the Grassmann manifold $\text{Gr}(r, n)$ is characterized by the equivalence class $[\mathbf{X}] := \{\mathbf{X}\mathbf{O} : \mathbf{O} \in \mathcal{O}(r)\}$ at $\mathbf{X} \in \text{St}(r, n)$.

Since the manifold \mathcal{M}/\sim is an abstract space, the elements of its tangent space $T_{[x]}(\mathcal{M}/\sim)$ at $[x]$ also call for a matrix representation in the tangent space $T_x\mathcal{M}$ that respects the equivalence relation \sim . Equivalently, the matrix representation of $T_{[x]}(\mathcal{M}/\sim)$ should be restricted to the directions in the tangent space $T_x\mathcal{M}$ on the total space \mathcal{M} at x that do not induce a displacement along the equivalence class $[x]$. This is realized by decomposing $T_x\mathcal{M}$ into complementary subspaces, the *vertical* and *horizontal* subspaces, such that $\mathcal{V}_x \oplus \mathcal{H}_x = T_x\mathcal{M}$. The vertical space \mathcal{V}_x is the tangent space of the equivalence class $[x]$. On the other hand, the horizontal space \mathcal{H}_x , which is any complementary subspace to \mathcal{V}_x in $T_x\mathcal{M}$, provides a valid matrix representation of the abstract tangent space $T_{[x]}(\mathcal{M}/\sim)$ [2, section 3.5.8]. An abstract tangent vector $\xi_{[x]} \in T_{[x]}(\mathcal{M}/\sim)$ at $[x]$ has a unique element in the horizontal space $\xi_x \in \mathcal{H}_x$ that is called its *horizontal lift*. Our specific choice of the horizontal space is the subspace of $T_x\mathcal{M}$ that is the *orthogonal complement* of \mathcal{V}_x in the sense of a Riemannian metric.

A Riemannian metric $g_x : T_x\mathcal{M} \times T_x\mathcal{M} \rightarrow \mathbb{R}$ at $x \in \mathcal{M}$ in the total space defines a Riemannian metric $g_{[x]} : T_{[x]}(\mathcal{M}/\sim) \times T_{[x]}(\mathcal{M}/\sim) \rightarrow \mathbb{R}$, i.e.,

$$(3.1) \quad g_{[x]}(\xi_{[x]}, \eta_{[x]}) := g_x(\xi_x, \eta_x),$$

on the quotient manifold \mathcal{M}/\sim , provided that the expression $g_x(\xi_x, \eta_x)$ does not depend on a specific representation along the equivalence class $[x]$. Here $\xi_{[x]}$ and $\eta_{[x]}$ are tangent vectors in $T_{[x]}(\mathcal{M}/\sim)$, and ξ_x, η_x are their horizontal lifts in \mathcal{H}_x at x . Equivalently, the definition (3.1) is well posed when $g_x(\xi_x, \zeta_x) = g_y(\xi_y, \zeta_y)$ for all $x, y \in [x]$, where $\xi_x, \zeta_x \in \mathcal{H}_x$ and $\xi_y, \zeta_y \in \mathcal{H}_y$ are the horizontal lifts of $\xi_{[x]}, \zeta_{[x]} \in T_{[x]}(\mathcal{M}/\sim)$ along the same equivalence class $[x]$ [2, section 3.6.2]. In other words, the metric g_x is *invariant* along the equivalence class $[x]$. Endowed with this Riemannian metric, the quotient manifold \mathcal{M}/\sim is called a *Riemannian quotient manifold* of \mathcal{M} .

The choice of the metric (3.1), which is invariant along the equivalence class $[x]$, and of the horizontal space \mathcal{H}_x as the orthogonal complement of \mathcal{V}_x , in the sense of the Riemannian metric (3.1), makes the space \mathcal{M}/\sim a *Riemannian submersion*. It allows for a convenient matrix representation of the gradient of a cost function. Figure 2 presents a schematic view of the search space. Consequently, the steepest-descent algorithm on the manifold \mathcal{M} that respects the equivalence property \sim on the space acquires the form shown in Table 2. Convergence of the steepest-descent algorithm in the neighborhood of a local minimum is shown by Absil, Mahony, and Sepulchre [2, Theorems 4.3.1 and 4.5.1]. The main ingredients of Table 2 are the gradient direction computation and the *retraction* mapping.

Riemannian gradient. The horizontal lift of the Riemannian gradient $\text{grad}_{[x]}$ of a cost function, say $f : \mathcal{M} \rightarrow \mathbb{R}$, on the quotient manifold \mathcal{M}/\sim is uniquely represented by the matrix representation, i.e., the

$$(3.2) \quad \text{horizontal lift of } \text{grad}_{[x]}f = \text{grad}_x f,$$

where $\text{grad}_x f$ is the gradient of f on the computational space \mathcal{M} at x . The equality in (3.2) is possible due to the invariance of the cost function along the equivalence class $[x]$, the choice of the Riemannian metric (3.1), and the choice of the horizontal space \mathcal{H}_x as the orthogonal complement of the vertical space \mathcal{V}_x [2, section 3.6.2].




FIG. 2. A schematic view of optimization on constraints with symmetries. The dotted lines represent abstract objects, and the solid lines are their matrix representations. The points x and y in the total (computational) space \mathcal{M} belong to the same equivalence class (shown in solid blue, discernible online only), and they represent a single point $[x] := \{y \in \mathcal{M} : y \sim x\}$ in the quotient space \mathcal{M}/\sim . An algorithm by necessity is implemented in the computation space, but conceptually, the search is on the quotient manifold. With the Riemannian metric g from (3.1), the quotient manifold \mathcal{M}/\sim is submersed into \mathcal{M} . The vertical space \mathcal{V}_x is the linearization of the equivalence class. The horizontal space \mathcal{H}_x is complementary to the vertical space \mathcal{V}_x and provides a matrix representation to the abstract tangent space $T_{[x]}(\mathcal{M}/\sim)$ of the Riemannian quotient manifold. Consequently, tangent vectors on the quotient space are lifted to the horizontal space. Given ξ_x as the horizontal lift, i.e., the matrix representation of a tangent vector $\xi_{[x]}$ belonging to the abstract space $T_{[x]}(\mathcal{M}/\sim)$, R_x maps it onto an element in \mathcal{M} .

TABLE 2
The Riemannian steepest-descent algorithm.

<p>The Riemannian steepest-descent algorithm for</p> $\min_{x \in \mathcal{M}} f(x).$ <ol style="list-style-type: none"> 1. Search direction: compute the negative Riemannian gradient $\xi_x = -\text{grad}_x f$ with respect to the Riemannian metric g_x (3.1). Equivalently, we solve the problem (3.4). 2. Retract with backtracking line search: the next iterate is computed using the retraction (3.5) such that $x_+ = R_x(s\xi_x)$, where the step-size s is obtained with backtracking line search. 3. Repeat until convergence.
--

The gradient $\text{grad}_x f$ on the computational space is computed from its definition: it is the unique element of $T_x \mathcal{M}$ that satisfies [2, equation (3.31)]

$$(3.3) \quad Df(x)[\eta_x] = g_x(\text{grad}_x f, \eta_x) \quad \text{for all } \eta_x \in T_x \mathcal{M},$$

where g_x is the Riemannian metric (3.1) and $Df(x)[\eta_x]$ is the Euclidean directional derivative of f in the direction η_x , i.e., $Df(x)[\eta_x] = \lim_{t \rightarrow 0} (f(x + t\eta_x) - f(x))/t$ [2, section 3.6]. An equivalent way of computing $\text{grad}_x f$ is by solving the *convex quadratic* problem

$$(3.4) \quad \text{grad}_x f = \arg \min_{\zeta_x \in T_x \mathcal{M}} f(x) - \langle f_x(x), \zeta_x \rangle + \frac{1}{2} g_x(\zeta_x, \zeta_x),$$

where $f_x(x)$ is the first-order derivative of the cost function f , $\langle \cdot, \cdot \rangle$ is the Euclidean

inner product, and g_x is the Riemannian metric (3.1) at $x \in \mathcal{M}$. It should be noted that $\langle f_x(x), \zeta_x \rangle = \text{D}f(x)[\zeta_x]$. The equivalence between solutions to (3.3) and (3.4) is established by observing that the condition (3.3) is, in fact, equivalent to the optimality conditions of the convex quadratic problem (3.4).

Retraction. An iterative optimization algorithm involves computing a search direction and then “moving in that direction.” The default option on a Riemannian manifold is to move along geodesics, leading to the definition of the exponential map [2, section 5.4]. Because the calculation of the exponential map can be computationally demanding, it is customary in the context of manifold optimization to relax the constraint of moving along geodesics. The exponential map is then relaxed to a *retraction*, which is any map $R_x : \mathcal{H}_x \rightarrow \mathcal{M}$ that locally approximates the exponential map, up to first order, on the manifold [2, Definition 4.1.1]. A natural update on the manifold is, thus, based on the update formula

$$(3.5) \quad x_+ = R_x(s\xi_x),$$

where $\xi_x \in \mathcal{H}_x$ is a search direction, s is the step-size, and $x_+ \in \mathcal{M}$.

The retraction R_x defines a retraction $R_{[x]}(\xi_{[x]}) := [R_x(\xi_x)]$ on the Riemannian quotient manifold \mathcal{M}/\sim , provided that the equivalence class $[R_x(\xi_x)]$ does not depend on the specific choice of the matrix representations of $[x]$ and $\xi_{[x]}$. Here ξ_x is the horizontal lift of an abstract tangent vector $\xi_{[x]} \in T_{[x]}(\mathcal{M}/\sim)$ in \mathcal{H}_x and $[\cdot]$ is the equivalence class defined earlier in the section. Equivalently, the retraction operation $R_{[x]}(\xi_{[x]}) := [R_x(\xi_x)]$ is well defined on \mathcal{M}/\sim when $R_x(\xi_x)$ and $R_y(\xi_y)$ belong to the same equivalence class, i.e., $[R_x(\xi_x)] = [R_y(\xi_y)]$ for all $x, y \in [x]$.

3.2. Inferring a metric from the Lagrangian. The practical performance of the Riemannian steepest-descent algorithm in Table 2 greatly depends on the choice of the metric (3.1). The dominant trend in Riemannian optimization has been to infer the metric from the geometry of the search space. Symmetry properties of the search space suggest choosing invariant metrics, that is, metrics that are not affected by a symmetry transformation of variables. In many situations, invariance properties of search space single out the choice of a unique metric [2, 13, 30]. However, the metrics that are motivated solely by the search space may not perform well in an optimization setup, as they do not take into consideration the cost function [22].

To address the above issue, we connect the SQP approach in Table 1 to the Riemannian steepest-descent algorithm in Table 2 with a specific metric that is induced by the Lagrangian (2.2). In particular, we extend the ideas of Propositions 2.1 and 2.2 to the case of quotient manifolds. The extension has a twofold objective. First, it provides guidance in choosing metrics on a manifold with symmetries. Second, it provides a framework to extend SQP to constraints with symmetries.

THEOREM 3.1. *Consider an equivalence relation \sim in the space \mathcal{M} . Assume that both \mathcal{M} and \mathcal{M}/\sim have the structure of a Riemannian manifold with the metric \bar{g} on \mathcal{M} and that the function $f : \mathcal{M} \rightarrow \mathbb{R}$ is a smooth function with isolated minima on the quotient manifold \mathcal{M}/\sim . Assume also that \mathcal{M} has the structure of an embedded submanifold in \mathbb{R}^n .*

If $x^ \in \mathcal{M}$ is a local minimum of $f : \mathcal{M} \rightarrow \mathbb{R}$ on \mathcal{M} , then the following hold:*

- (i) $\langle \eta_{x^*}, \text{D}^2\mathcal{L}(x^*, \lambda_{x^*})[\eta_{x^*}] \rangle = 0$ for all $\eta_{x^*} \in \mathcal{V}_{x^*}$, and
- (ii) the quantity $\langle \xi_{x^*}, \text{D}^2\mathcal{L}(x^*, \lambda_{x^*})[\xi_{x^*}] \rangle$ captures all second-order information of the cost function f on \mathcal{M}/\sim for all $\xi_{x^*} \in \mathcal{H}_{x^*}$,

where \mathcal{V}_{x^} is the vertical space, \mathcal{H}_{x^*} is the horizontal space (the subspace of $T_{x^*}\mathcal{M}$ that is complementary to \mathcal{V}_{x^*}), $\langle \cdot, \cdot \rangle$ is the Euclidean inner product, and $\text{D}^2\mathcal{L}(x^*, \lambda_{x^*})[\xi_{x^*}]$*

is the second-order derivative of $\mathcal{L}(x, \lambda_x)$ with respect to x at $x^* \in \mathcal{M}$ applied in the direction $\xi_{x^*} \in \mathcal{H}_{x^*}$ and keeping λ_{x^*} fixed to its least-squares estimate (2.4).

Proof. Statement (i) corresponds to the degeneracy of the Lagrangian at a local minimum, and statement (ii) is an extension of Proposition 2.2 to the case of quotient manifolds.

The Lagrangian $\mathcal{L}(x^*, \lambda)$ at x^* , because both the cost and the constraint terms remain invariant under the equivalence relation \sim , is *constant* along the equivalence class $[x^*] := \{y^* \in \mathcal{M} : y^* \sim x^*\}$. It should be noted that since x^* is a local minimum, the first-order derivative of the Lagrangian with respect to x at x^* is zero, i.e., $\mathcal{L}_x(x^*, \lambda_{x^*}) = 0$. It should be noted that $\mathcal{L}_x(x, \lambda_x)$ is the derivative of $\mathcal{L}(x, \lambda_x)$ with respect to the first argument, i.e., λ_x is kept fixed. Differentiating the Lagrangian $\mathcal{L}(x, \lambda_x)$ twice at x^* along the linearization of the equivalence class $[x^*]$, that is, the vertical space \mathcal{V}_{x^*} , results in the first equality in (i).

To prove (ii) of the theorem, we exploit the Riemannian structure on \mathcal{M}/\sim . The theory of Riemannian submersion [2, Chapter 3] states that the horizontal space \mathcal{H}_x is *orthogonal* to the vertical space \mathcal{V}_x with respect to the metric \bar{g}_x and allows us to compute the Riemannian gradient and Hessian of f on \mathcal{M}/\sim using the *orthogonal projection* of their counterparts in the total space \mathcal{M} . From the computation of the Riemannian gradient $\text{grad}_x f$ and the first-order derivative of the Lagrangian we have the equality [2, equation (3.31)]

$$Df(x)[\xi_x] = \langle \mathcal{L}_x(x, \lambda_x), \xi_x \rangle = \bar{g}_x(\text{grad}_x f, \xi_x) \quad \text{for all } \xi_x \in T_x \mathcal{M}.$$

Taking the Euclidean directional derivative of the above equation at x^* along $\xi_{x^*} \in \mathcal{H}_{x^*}$ with the additional information that $\text{grad}_{x^*} f = 0$ (x^* is a local minimum),

$$(3.6) \quad \langle \xi_{x^*}, D^2 f(x^*)[\xi_{x^*}] \rangle = \langle \xi_{x^*}, D^2 \mathcal{L}(x^*, \lambda_{x^*})[\xi_{x^*}] \rangle = \bar{g}_{x^*}(\xi_{x^*}, D \text{grad}_{x^*} f[\xi_{x^*}]),$$

where the specific operation $D \text{grad}_{x^*} f[\xi_{x^*}]$ should be treated as the *Euclidean differential of a vector field*, the Riemannian gradient $\text{grad}_{x^*} f$, along ξ_{x^*} . In this case, $D^2 \mathcal{L}(x, \lambda_x)[\xi_x]$ is the second-order derivative of $\mathcal{L}(x, \lambda_x)$ with respect to x along ξ_x keeping λ_x fixed.

We define $\Pi_x : \mathbb{R}^n \rightarrow \mathcal{H}_x$ as the orthogonal projection operator in the metric \bar{g}_x , and $\text{Hess}_x f[\xi_x]$ as the Riemannian Hessian in the total space \mathcal{M} applied along the direction $\xi_x \in \mathcal{H}_x$. The horizontal lift of the Riemannian Hessian $\text{Hess}_{[x]} f[\xi_{[x]}]$ on the quotient manifold \mathcal{M}/\sim has the characterization $\Pi_x(\text{Hess}_x f[\xi_x])$ [2, Proposition 5.3.3]. Additionally, the *Koszul formula* relates the operation $D \text{grad}_x f[\xi_x]$ and the Riemannian Hessian in the total space \mathcal{M} [2, Theorem 5.3.1].

At the minimum x^* , second-order information of f on the manifold \mathcal{M}/\sim is captured by the quantity $\bar{g}_{x^*}(\xi_{x^*}, \Pi_{x^*}(\text{Hess}_{x^*} f[\xi_{x^*}]))$ for all $\xi_{x^*} \in \mathcal{H}_{x^*}$. The term simplifies as follows:

$$\begin{aligned} \bar{g}_{x^*}(\xi_{x^*}, \underbrace{\Pi_{x^*}(\text{Hess}_{x^*} f[\xi_{x^*}])}_{\left\{ \begin{array}{l} \text{horizontal lift of the} \\ \text{Hessian on } \mathcal{M}/\sim \end{array} \right\}}) &= \bar{g}_{x^*}(\xi_{x^*}, \underbrace{\text{Hess}_{x^*} f[\xi_{x^*}]}_{\text{the Hessian on } \mathcal{M}}) \quad \text{for } \Pi_{x^*}(\xi_{x^*}) = \xi_{x^*} \\ &= \bar{g}_{x^*}(\xi_{x^*}, D \text{grad}_{x^*} f[\xi_{x^*}]) \quad \text{from the Koszul formula} \\ &\quad \text{and } \text{grad}_{x^*} f = 0 \\ &= \langle \xi_{x^*}, D^2 \mathcal{L}(x^*, \lambda_{x^*})[\xi_{x^*}] \rangle \quad \text{from (3.6)}. \end{aligned}$$

This proves statement (ii). □

Theorem 3.1 states that the underlying symmetries make the Lagrangian *degenerate* in the tangent space $T_x\mathcal{M}$ of the total space \mathcal{M} . The other important observation is that the quantity $\langle \xi_x, D^2\mathcal{L}(x, \lambda_x)[\xi_x] \rangle$ captures second-order information at a local minimum along the horizontal space, where the horizontal space is the subspace of $T_x\mathcal{M}$ that is the orthogonal complement of the vertical space \mathcal{V}_x .

As an immediate consequence of Theorem 3.1, we observe that in the neighborhood of a minimum, a selection of the search direction is given by solving

$$(3.7) \quad \arg \min_{\zeta_x \in \mathcal{H}_x} f(x) + \langle f_x(x), \zeta_x \rangle + \frac{1}{2} \langle \zeta_x, D^2\mathcal{L}(x, \lambda_x)[\zeta_x] \rangle,$$

where \mathcal{H}_x is any predefined horizontal space characterization. Followed by a retraction operation (section 3.1), (3.7) defines locally a steepest-descent algorithm on the quotient manifold \mathcal{M}/\sim , that is, an algorithm that iterates on the classes of equivalences. Here $\langle \cdot, \cdot \rangle$ is the Euclidean inner product and $f_x(x)$ is the first-order derivative of the function f . The scheme with (3.7) has the interpretation of a steepest-descent algorithm on the quotient manifold \mathcal{M}/\sim and can be considered as an extension of the classical SQP approach, that has been defined on embedded constraints, to quotient manifolds. The significance of the result is twofold: first, the computational problem (3.7) is considerably simpler than the computational machinery needed for a steepest-descent or a quasi-Newton algorithm on a general quotient manifold. Second, it decouples the second-order information term, i.e., $D^2\mathcal{L}(x, \lambda_x)$, from the horizontal space characterization \mathcal{H}_x that could be predefined. The expression (3.7) is appealing as a starting point to equip the quotient manifold with a Riemannian structure since it captures second-order information in the neighborhood of a minimum.

3.3. Regularized Lagrangians for Riemannian optimization. Motivated by the expression (3.7), we introduce the family of regularized Lagrangians

$$(3.8) \quad \tilde{\mathcal{L}}(x, \lambda_x) = \omega_1 f(x) - \omega_2 \langle \lambda_x, h(x) \rangle,$$

with the corresponding candidate metrics

$$(3.9) \quad g_x(\xi_x, \eta_x) = \langle \xi_x, D^2\tilde{\mathcal{L}}(x, \lambda_x)[\eta_x] \rangle \\ = \omega_1 \underbrace{\langle \xi_x, D^2f(x)[\eta_x] \rangle}_{\text{cost related}} + \omega_2 \underbrace{\langle \xi_x, D^2c(x, \lambda_x)[\eta_x] \rangle}_{\text{constraint related}}$$

in the total space \mathcal{M} , where $c(x, \lambda_x) = -\langle \lambda_x, h(x) \rangle$. The choice of the regularizing parameters $\omega_1 \in [0, 1]$ and $\omega_2 \in [0, 1]$ is problem dependent, but concrete choices are discussed below and also illustrated in Figure 3. The default choice of $\omega_1 = 1$ and $\omega_2 = 1$ is appealing because it captures the full second-order information. However, it also leads to a degenerate Lagrangian at the minimum. The choices of ω_1 and ω_2 thus result from a trade-off between ensuring the positive definiteness of the metric candidate (3.9) in the tangent space $T_x\mathcal{M}$ while capturing second-order information near minima.

If in addition to being positive definite the metric (3.9) is invariant along the equivalence class $[x]$, and if the horizontal space \mathcal{H}_x is chosen as the orthogonal subspace to the vertical space \mathcal{V}_x , then the manifold \mathcal{M}/\sim has the structure of a Riemannian submersion [2, 13, 35]. This simplifies the computation of the search direction, which then amounts to solving the problem

$$(3.10) \quad \arg \min_{\zeta_x \in T_x\mathcal{M}} f(x) + \langle f_x(x), \zeta_x \rangle + \frac{1}{2} g_x(\zeta_x, \zeta_x),$$

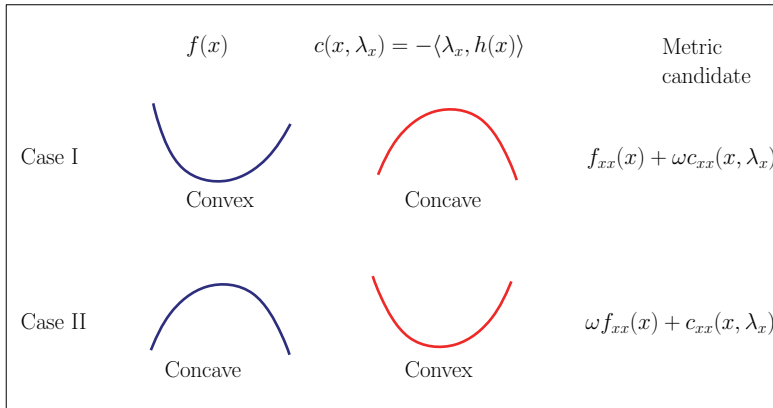


FIG. 3. Choosing metrics for the Riemannian steepest-descent algorithm in Table 2. Shown are two extreme situations in which the regularized Lagrangian (3.8) provides a clear metric candidate locally in the neighborhood of a minimum. $f_{xx}(x)$ is the second-order derivative of $f(x)$, and $c_{xx}(x, \lambda_x)$ is the second-order derivative of $c(x, \lambda_x)$ with respect to x keeping λ_x fixed. Because of local convexity of the Lagrangian (on the tangent space) at a minimum, convex and concave structures of the function f lead to a well-defined family of metrics parameterized by the regularization parameter $\omega \in [0, 1)$. It locally captures second-order information of the problem. The case $\omega = 1$ is discarded to prevent singularity of the metric candidate at a local minimum. Additionally, to extend the metrics away from a minimum, the parameter ω is updated at every iteration, e.g., with the procedure shown in (3.13).

where $\langle \cdot, \cdot \rangle$ is the Euclidean inner product and $f_x(x)$ is the first-order derivative of the function f , and $g_x(\cdot, \cdot)$ is the Riemannian metric (3.9). It should be noted that even though the minimization (3.10) is in the total tangent space $T_x\mathcal{M}$, the solution ζ_x^* of (3.10), by construction, belongs to the chosen horizontal space \mathcal{H}_x that is orthogonal to \mathcal{V}_x .

We further discuss two scenarios below that suggest how to exploit the available structure to construct novel Riemannian metrics. The problem structure can be exploited in more general situations along the same lines.

Case I: Minimizing a strictly convex function. Consider the case when f is a *strictly* convex function. In this case the second-order derivative $f_{xx}(x) \succ 0$ (due to strict convexity assumption) is a good metric candidate. In addition, locally in the neighborhood of a minimum, a family of Riemannian metrics is identified from (3.9) as

$$(3.11) \quad g_x(\xi_x, \eta_x) = \underbrace{\langle \xi_x, D^2 f(x)[\eta_x] \rangle}_{f_{xx} \succ 0 \text{ and dominating}} + \omega \langle \xi_x, D^2 c(x, \lambda_x)[\eta_x] \rangle$$

by selecting $\omega_1 = 1$ and $\omega_2 = \omega$, where $\omega \in [0, 1)$, ξ_x, η_x are tangent vectors in $T_x\mathcal{M}$, $c(x, \lambda_x) = -\langle \lambda_x, h(x) \rangle$, and $D^2 c(x, \lambda_x)[\eta_x]$ is the second-order derivative of $c(x, \lambda_x)$ with respect to x keeping λ_x fixed and that is applied along the direction η_x . The case $\omega = 1$ is discarded to prevent singularity of the metric candidate (3.11) in $T_x\mathcal{M}$ in the neighborhood of a minimum.

Case II: Maximizing a strictly convex function. Consider the problem of maximizing a convex cost function, that is equivalent to minimizing a concave cost function, on a manifold. In this case, $f_{xx}(x, \lambda_x) \prec 0$, and locally in the neighborhood of a minimum, second-order information of $c(x, \lambda_x) = -\langle \lambda_x, h(x) \rangle$ is the source of

convexity. This fact follows from the second-order optimality condition of the optimization problem [32, Chapter 18]. Here the problem structure suggests a family of Riemannian metrics

$$(3.12) \quad g_x(\xi_x, \eta_x) = \omega \langle \xi_x, D^2 f(x)[\eta_x] \rangle + \underbrace{\langle \xi_x, D^2 c(x, \lambda_x)[\eta_x] \rangle}_{c_{xx} \text{ is locally positive definite}}$$

by selecting $\omega_1 = \omega$ and $\omega_2 = 1$ in (3.9), where $\omega \in [0, 1)$, ξ_x, η_x are tangent vectors in $T_x \mathcal{M}$, $c(x) = -\langle \lambda_x, h(x) \rangle$, and $c_{xx}(x, \lambda_x)$ is the second-order derivative of $c(x, \lambda_x)$ with respect to x keeping λ_x fixed. Once again $\omega = 1$ is discarded to prevent singularity of the metric candidate (3.12) in $T_x \mathcal{M}$ in the neighborhood of a minimum.

Globalizing the local metrics. The parameter $\omega \in [0, 1)$ in the metrics (3.11) and (3.12), apart from providing a family of Riemannian metrics, also plays a critical role in the numerical performance of the Riemannian steepest-descent algorithm in Table 2. With $\omega = 0$, the Riemannian metric captures only part of second-order information, and therefore, locally in the neighborhood of a minimum, the Riemannian steepest-descent algorithm may converge poorly, e.g., linearly. On the other hand, with $\omega \rightarrow 1$, the Riemannian metric captures the full second-order information and the Riemannian steepest-descent algorithm is expected to show better convergence. A numerical technique for interpolating between these two extreme scenarios is to vary $\omega = [0, 1)$ at every iteration with an increasing barrier function that tends to 1 as the number of iterations increases. A simple updating technique is $\omega(k) = 1 - 2^{-k-1}$, where k is the iteration number. A strategy to safeguard against a nondescent search direction, e.g., when solving the quadratic programming problem (3.10) with the metric (3.9), is to ignore the updated ω that resulted in a nondescent direction (checking this numerically is straightforward) and restart the procedure of updating ω again.

A different technique is to modify ω as and when required. For example, defining $\delta = 1 - \omega$, we have the strategy where at the k th iteration

$$(3.13) \quad \delta_k = \begin{cases} 0.5\delta_{k-1} & \text{when a descent direction is obtained,} \\ 4\delta_{k-1} & \text{when a nondescent direction is obtained,} \end{cases}$$

with $\delta_0 = 1$. Care is taken to ensure that $\omega \in [0, 1)$ for all iterations.

Safeguards similar to the trust-regions, i.e., by constraining the norm of the search direction, can also be implemented to ensure that the search direction computed with the Riemannian metric remains a locally descent direction [32, section 18.5].

4. Quadratic optimization with orthogonality constraints: Revisiting the generalized eigenvalue problem. Constrained quadratic optimization problems arise naturally in a number of applications, especially while solving linear systems of matrix equations [2, section 2.2]. Also popular are the *orthogonality* constraints in large-scale problems that are imposed to identify relevant smaller dimensional subspaces [13]. Specific optimization problems include the generalized eigenvalue problem [3, 6, 13], the generalized orthogonal Procrustes problem [14], and the joint diagonalization problem in signal processing [36], to name just a few.

For the sake of illustration, we specifically focus on the well-studied generalized eigenvalue problem that computes the smallest eigenvalues and eigenvectors of the matrix $\mathbf{B}^{-1}\mathbf{A}$, where \mathbf{A} is a symmetric matrix of size $n \times n$ and \mathbf{B} is a symmetric positive definite matrix of size $n \times n$ [13, section 4.5], [16, Chapter 8]. This is realized by solving the optimization problem below iteratively, an extensively researched question

in the literature [16, Chapter 8], [4]. In this section we exploit the quadratic nature of the cost function and the constraints to show that a family of Riemannian metrics has a simple characterization. It is also shown that the algorithms that result from the proposed metrics connect to a number of established algorithms, each of which is interpreted as a steepest-descent algorithm with a specific Riemannian metric choice.

The minimal r -eigenspace of $\mathbf{B}^{-1}\mathbf{A}$ is computed iteratively by solving the constrained quadratic optimization problem

$$(4.1) \quad \begin{aligned} \min_{\mathbf{X} \in \mathbb{R}^{n \times r}} \quad & \frac{1}{2} \text{Trace}(\mathbf{X}^T \mathbf{A} \mathbf{X}) \\ \text{subject to} \quad & \mathbf{X}^T \mathbf{B} \mathbf{X} = \mathbf{I}, \end{aligned}$$

where the constraint set of $n \times r$ matrices that satisfy $\mathbf{X}^T \mathbf{B} \mathbf{X} = \mathbf{I}$ is known as the *generalized Stiefel manifold* $\text{St}_{\mathbf{B}}(r, n)$. The constraint enforces orthogonality among vectors in coordinates spanned by $\mathbf{B}^{1/2}$. Specifically, when $\mathbf{B} = \mathbf{I}$, the generalized Stiefel manifold is the popular Stiefel manifold $\text{St}(r, n) := \{\mathbf{X} \in \mathbb{R}^{n \times r} : \mathbf{X}^T \mathbf{X} = \mathbf{I}\}$ [13].

It should be noted that the optimization problem (4.1) is invariant under the transformation $\mathbf{X} \mapsto \mathbf{X}\mathbf{O}$ for all $\mathbf{O} \in \mathcal{O}(r)$, where $\mathcal{O}(r)$ is the set of $r \times r$ orthogonal matrices. As a consequence, the problem (4.1) is an optimization problem on the abstract quotient space $\text{St}_{\mathbf{B}}(r, n)/\mathcal{O}(r)$, also known as the *generalized Grassmann manifold*. For the case $\mathbf{B} = \mathbf{I}$, this again boils down to the well-known Grassmann manifold $\text{Gr}(r, n)$, which is the set of r -dimensional subspaces in \mathbb{R}^n [13]. The optimization problem (4.1) is, therefore, reformulated on the generalized Grassmann quotient manifold, i.e.,

$$(4.2) \quad \begin{aligned} \min_{\mathbf{X} \in \mathbb{R}^{n \times r}} \quad & \frac{1}{2} \text{Trace}(\mathbf{X}^T \mathbf{A} \mathbf{X}) \\ \text{subject to} \quad & [\mathbf{X}] \in \text{St}_{\mathbf{B}}(r, n)/\mathcal{O}(r), \end{aligned}$$

where the optimization is on the set of equivalence classes $[\mathbf{X}] := \{\mathbf{X}\mathbf{O} : \mathbf{O} \in \mathcal{O}(r)\}$ at $\mathbf{X} \in \text{St}_{\mathbf{B}}(r, n)$.

The *conventional* choice of the metric in the Riemannian framework is

$$(4.3) \quad g_x(\eta_x, \xi_x) = \text{Trace}(\eta_x^T \xi_x),$$

where $x = \mathbf{X} \in \text{St}_{\mathbf{B}}(r, n)$ and ξ_x, η_x are vectors in the tangent space of the constraints (the matrix characterization of the tangent space is shown in Table 3). It is the unique metric that is motivated by *invariance* on the orthogonal group, i.e., for the case when $r = n$, it is the unique choice [30]. Hence, it becomes a conventional choice for $r < n$. Because of its simplicity and its geometric consideration, the metric (4.3) is also advocated by Absil, Mahony, and Sepulchre [2] and Edelman, Arias, and Smith [13].

In contrast, the developments in section 3.3 suggest a *family of Riemannian metrics* that take the complete problem structure into account by computing the regularized Lagrangian (3.8) and its second-order derivative. To this end, we have the following matrix representations for (4.1):

$$(4.4) \quad \begin{aligned} & \mathcal{L}(x, \lambda_x) = \text{Trace}(\mathbf{X}^T \mathbf{A} \mathbf{X})/2 - \langle \lambda_x, \mathbf{X}^T \mathbf{B} \mathbf{X} - \mathbf{I} \rangle/2 \\ \Rightarrow \quad & \mathcal{L}_x(x, \lambda_x) = \mathbf{A} \mathbf{X} - \mathbf{B} \mathbf{X} \lambda_x \\ \Rightarrow \quad & \text{D}^2 \mathcal{L}(x, \lambda_x)[\xi_x] = \mathbf{A} \xi_x - \mathbf{B} \xi_x \lambda_x, \end{aligned}$$

where x has the matrix representation $\mathbf{X} \in \text{St}_{\mathbf{B}}(r, n)$, $\langle \cdot, \cdot \rangle$ is the Euclidean inner product, and the least-squares Lagrange multiplier is $\lambda_x = \text{Sym}((\mathbf{X}^T \mathbf{B} \mathbf{B} \mathbf{X})^{-1} (\mathbf{X}^T \mathbf{B} \mathbf{A} \mathbf{X}))$

TABLE 3

Optimization-related ingredients for computing the extreme eigenvalues of $\mathbf{B}^{-1}\mathbf{A}$. The numeric complexity per iteration depends on solving the quadratic programming problem for the search direction computation. In many instances, exploiting sparsity in matrices \mathbf{A} and \mathbf{B} leads to numerically efficient schemes. Here $\text{Sym}(\cdot)$ extracts the symmetric part of a square matrix, i.e., $\text{Sym}(\mathbf{D}) = (\mathbf{D} + \mathbf{D}^T)/2$. A few choices of the regularizing parameters $\omega_1, \omega_2 \in [0, 1]$ for relevant situations are shown in section 4.1.

	$\min_{\mathbf{X} \in \mathbb{R}^{n \times r}}$ subject to $\mathbf{X}^T \mathbf{B} \mathbf{X} = \mathbf{I}$
Matrix representation of an element $x \in \mathcal{M}$	$x = \mathbf{X}$
Computational space \mathcal{M}	$\text{St}_{\mathbf{B}}(r, n) = \{\mathbf{X} \in \mathbb{R}^{n \times r} : \mathbf{X}^T \mathbf{B} \mathbf{X} = \mathbf{I}\}$
Group action	$\mathbf{X} \mathbf{O}$ for all $\mathbf{O} \in \mathcal{O}(r)$ such that $\mathbf{O}^T \mathbf{O} = \mathbf{O} \mathbf{O}^T = \mathbf{I}$
Quotient space	$\text{St}_{\mathbf{B}}(r, n) / \mathcal{O}(r)$
Tangent vectors in $T_x \mathcal{M}$	$\{\xi_x \in \mathbb{R}^{n \times r} : \xi_x^T \mathbf{B} \mathbf{X} + \mathbf{X}^T \mathbf{B} \xi_x = 0\}$
Metric $g_x(\xi_x, \zeta_x)$ for $\xi_x, \zeta_x \in T_x \mathcal{M}$	$g_x(\xi_x, \zeta_x) = \omega_1 \text{Trace}(\zeta_x^T \mathbf{A} \xi_x) - \omega_2 \text{Trace}(\zeta_x^T \mathbf{B} \xi_x \lambda_x)$ or the metrics proposed in section 4.1, where $\lambda_x = \text{Sym}((\mathbf{X}^T \mathbf{B} \mathbf{B} \mathbf{X})^{-1} (\mathbf{X}^T \mathbf{B} \mathbf{A} \mathbf{X}))$
Cost function	$f(x) = \text{Trace}(\mathbf{X}^T \mathbf{A} \mathbf{X}) / 2$
First-order derivative of $f(x)$	$f_x(x) = \mathbf{A} \mathbf{X}$
Search direction	$\arg \min_{\zeta_x \in T_x \mathcal{M}} f(x) + \langle f_x(x), \zeta_x \rangle + \frac{1}{2} g_x(\zeta_x, \zeta_x)$
Retraction $R_x(\xi_x)$ that maps a search direction ξ_x onto \mathcal{M}	$\mathbf{U} \mathbf{V}^T$, where $\mathbf{X} + \xi_x = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T$ is the restriction of thin SVD such that $\mathbf{U}^T \mathbf{B} \mathbf{U} = \mathbf{I}$, $\mathbf{V}^T \mathbf{V} = \mathbf{I}$, and $\mathbf{\Sigma}$ is a positive diagonal matrix.

from (2.4) with the additional symmetry condition from the constraint, where $\text{Sym}(\cdot)$ extracts the symmetric part of a square matrix, i.e., $\text{Sym}(\mathbf{D}) = (\mathbf{D} + \mathbf{D}^T)/2$. $\mathcal{L}_x(x, \lambda_x)$ is the first-order derivative of $\mathcal{L}(x, \lambda_x)$ and $\text{D}^2 \mathcal{L}(x, \lambda_x)[\xi_x]$ is the second-order derivative of $\mathcal{L}(x, \lambda_x)$ applied in the direction ξ_x , both computed while keeping λ_x fixed.

We should say that the least-squares estimate $\lambda_x = \text{Sym}((\mathbf{X}^T \mathbf{B} \mathbf{B} \mathbf{X})^{-1} (\mathbf{X}^T \mathbf{B} \mathbf{A} \mathbf{X}))$ is the solution to the problem $\arg \min_{\lambda \in \mathbb{R}^{r \times r}} \|\mathbf{A} \mathbf{X} - \mathbf{B} \mathbf{X} \lambda\|_{\mathbf{Q}}^2$ such that λ is symmetric, where $\|\mathbf{A} \mathbf{X} - \mathbf{B} \mathbf{X} \lambda\|_{\mathbf{Q}}^2 = \text{Trace}((\mathbf{A} \mathbf{X} - \mathbf{B} \mathbf{X} \lambda)^T \mathbf{Q} (\mathbf{A} \mathbf{X} - \mathbf{B} \mathbf{X} \lambda))$ and $\mathbf{Q} = \mathbf{B} \mathbf{X} (\mathbf{X}^T \mathbf{B} \mathbf{B} \mathbf{X})^{-2} \mathbf{X}^T \mathbf{B}$. A different estimate of λ_x is obtained by solving the problem $\arg \min_{\lambda \in \mathbb{R}^{r \times r}} \|\mathbf{A} \mathbf{X} - \mathbf{B} \mathbf{X} \lambda\|_F^2$ such that λ is symmetric. Both these estimates coincide at a local minimum and can be used in the neighborhood of a minimum.

It is readily checked that the Lagrangian $\mathcal{L}(x, \lambda_x)$ in (4.4) remains unchanged under the action $\mathbf{X} \mapsto \mathbf{X} \mathbf{O}$ for all $\mathbf{O} \in \mathcal{O}(r)$, where $\mathbf{X} \in \text{St}_{\mathbf{B}}(r, n)$. The action $\mathbf{X} \mapsto \mathbf{X} \mathbf{O}$ also leads to the transformation in the least-squares Lagrange multiplier

as $\lambda_x \mapsto \mathbf{O}^T \lambda_x \mathbf{O}$. Subsequently, we have the following proposition for constructing a family of Riemannian metrics for the problem (4.2) on the generalized Grassmann manifold.

PROPOSITION 4.1. *There exists a family of Riemannian metrics*

$$(4.5) \quad g_x(\xi_x, \eta_x) = \omega_1 \underbrace{\langle \xi_x, \mathbf{A} \eta_x \rangle}_{\text{cost related}} - \omega_2 \underbrace{\langle \xi_x, \mathbf{B} \eta_x \lambda_x \rangle}_{\text{constraint related}}$$

on $\text{St}_{\mathbf{B}}(r, n)$ with $\omega_1, \omega_2 \in [0, 1]$, each of which induces a Riemannian metric on the quotient manifold $\text{St}_{\mathbf{B}}(r, n)/\mathcal{O}(r)$ in the neighborhood of the minimum of the quadratic optimization problem (4.2). The parameters ω_1 and ω_2 weight the cost and constraint, respectively. ξ_x and η_x are vectors in the tangent space of the constraints at $x = \mathbf{X}$ such that $\mathbf{X}^T \mathbf{B} \mathbf{X} = \mathbf{I}$ and $\lambda_x = \text{Sym}((\mathbf{X}^T \mathbf{B} \mathbf{B} \mathbf{X})^{-1} (\mathbf{X}^T \mathbf{B} \mathbf{A} \mathbf{X}))$ is the least-squares estimate, where $\text{Sym}(\cdot)$ extracts the symmetric part of a square matrix, i.e., $\text{Sym}(\mathbf{D}) = (\mathbf{D} + \mathbf{D}^T)/2$.

Proof. First, we show that the metrics (4.5) respect the symmetry condition in (3.1). In order to show that the metric does not change along the equivalence class $[x] = [\mathbf{X}] = \{\mathbf{X} \mathbf{O} : \mathbf{O} \in \mathcal{O}(r)\}$ for all $\mathbf{O} \in \mathcal{O}(r)$, it is equivalent, but simplified following [2, Proposition 3.6.1], to show that the metric for tangent vectors ξ_x and η_x does not change under the transformations $\mathbf{X} \mapsto \mathbf{X} \mathbf{O}$, $\eta_x \mapsto \eta_x \mathbf{O}$, and $\xi_x \mapsto \xi_x \mathbf{O}$. These transformations lead to the transformation in the least-squares Lagrange multiplier as $\lambda_x \mapsto \mathbf{O}^T \lambda_x \mathbf{O}$. Finally, a few extra computations show that the metrics (4.5) indeed remain the same under the mentioned transformations.

Second, we show the construction of one particular family of Riemannian metrics. To this end, consider the case when $\mathbf{A} \succ 0$. Consider the parameters $\omega_1 = 1$ and $\omega_2 = \omega$. Restricting $\omega \in [0, 1]$ guarantees that in the neighborhood of a minimum $g_x(\zeta_x, \zeta_x) > 0$ for all ζ_x in the tangent space, satisfying the criterion of positive definiteness.

The proofs for symmetry compatibility and positive definiteness on the tangent space conclude the proof of the proposition. \square

4.1. Metric tuning and shift policies. Due to the quadratic nature of both cost and constraints, the metric (4.5) has the appealing feature of being parameterized by the Lagrangian parameter λ_x . This object is low dimensional when $r \ll n$. It provides an interesting interpretation of various “shift” policies developed in numerical linear algebra for eigenspace computations [16, Chapter 8]. We further particularize the selection of regularization parameters ω_1 and ω_2 in (4.5) when $\mathbf{A} \succ 0$ and when $\mathbf{A} \not\succeq 0$. In both these cases, we propose metrics that connect to a number of classical algorithms for the generalized eigenvalue problem [3, 4, 6].

When $\mathbf{A} \succ 0$. This instance falls under Case I of section 3.3, i.e., with $\omega_1 = 1$ and $\omega_2 = \omega$ in (3.9). Therefore, the family of proposed Riemannian metrics has the structure

$$(4.6) \quad g_x(\xi_x, \zeta_x) = \text{Trace}(\xi_x^T \mathbf{A} \zeta_x) - \omega \text{Trace}(\xi_x^T \mathbf{B} \zeta_x \lambda_x),$$

where ξ_x and ζ_x are vectors in the tangent space of the constraints, the least-squares Lagrange multiplier $\lambda_x = \text{Sym}((\mathbf{X}^T \mathbf{B} \mathbf{B} \mathbf{X})^{-1} (\mathbf{X}^T \mathbf{B} \mathbf{A} \mathbf{X}))$, and $\omega = [0, 1)$.

The metric (4.6) provides two insightful connections to the literature. First, the proposed metric (4.6) with $\omega = 0$ generalizes the well-known *inverse iteration algorithm* for computing the smallest eigenvalues of a symmetric matrix [16, section 8.2.2]. For the case when $\mathbf{B} = \mathbf{I}$, the *negative* Riemannian gradient with the metric (4.6) and $\omega = 0$ is computed as in Table 3 as

$$\left. \begin{array}{l} \arg \min \\ \zeta_x \in \mathbb{R}^{n \times r} \\ \zeta_x^T \mathbf{X} + \mathbf{X}^T \zeta_x = 0 \end{array} \left\langle \mathbf{A}\mathbf{X}, \zeta_x \right\rangle + \frac{1}{2} \text{Trace}(\zeta_x^T \mathbf{A} \zeta_x) \right\} = \mathbf{A}^{-1} \mathbf{X} (\mathbf{X}^T \mathbf{A}^{-1} \mathbf{X})^{-1} - \mathbf{X}.$$

Given an iterate $x = \mathbf{X}$ such that $\mathbf{X}^T \mathbf{X} = \mathbf{I}$, the Riemannian steepest-descent update with *unit step-size* is $x_+ = R_x(\mathbf{A}^{-1} \mathbf{X} (\mathbf{X}^T \mathbf{A}^{-1} \mathbf{X})^{-1} - \mathbf{X})$, which is equivalent to the update $x_+ = \text{qf}(\mathbf{A}^{-1} \mathbf{X})$ on the Grassmann manifold, where $R_x(\cdot)$ is the retraction operator defined in Table 3 and $\text{qf}(\mathbf{A}^{-1} \mathbf{X})$ is the Q-factor of the QR decomposition of $\mathbf{A}^{-1} \mathbf{X}$. This is also the classical inverse iteration update [16, section 8.2.2]. This shows that the inverse iteration has the interpretation of a Riemannian steepest-descent algorithm with the metric (4.11) for $\omega = 0$.

A second insight is obtained for the case when ω is updated with iterations. The Riemannian steepest-descent algorithm with the metric (4.6) generalizes the popular Rayleigh quotient iteration algorithm (see [16, section 8.2.3], [3, Algorithm 4.2], [4]). Given an iterate $x = \mathbf{X}$ such that $\mathbf{X}^T \mathbf{X} = \mathbf{I}$, Absil et al. [3, Algorithm 4.2] propose the *Grassmann-Rayleigh quotient iteration* (GRQI) algorithm for computing the next update x_+ with the steps

$$(4.7) \quad \left. \begin{array}{l} \mathbf{A}\mathbf{Z} - \mathbf{Z}\mathbf{X}^T \mathbf{A}\mathbf{X} = \mathbf{X} \quad (\text{we solve for } \mathbf{Z}) \\ x_+ = \text{qf}(\mathbf{Z}). \end{array} \right\}$$

To show that the Riemannian steepest-descent algorithm with the metric (4.6) generalizes GRQI algorithm (4.7), we consider again the case when $\mathbf{B} = \mathbf{I}$. At each iteration of the Riemannian steepest-descent algorithm with the metric (4.6), we are required to solve the system of linear equations (by looking at the optimality conditions of the quadratic program for computing the search direction) for $\zeta_x \in \mathbb{R}^{n \times r}$ and an auxiliary variable $\mu_x \in \mathbb{R}^{r \times r}$ (μ_x is the Lagrange multiplier corresponding to the constraint $\zeta_x^T \mathbf{X} + \mathbf{X}^T \zeta_x = 0$) of the form

$$(4.8) \quad \left. \begin{array}{l} \arg \min \\ \zeta_x \in \mathbb{R}^{n \times r} \\ \zeta_x^T \mathbf{X} + \mathbf{X}^T \zeta_x = 0 \end{array} \left\langle \mathbf{A}\mathbf{X}, \zeta_x \right\rangle + \frac{1}{2} (\text{Trace}(\zeta_x^T \mathbf{A} \zeta_x) - \omega \text{Trace}(\zeta_x^T \zeta_x \lambda_x)) \right\} \\ \Rightarrow \left\{ \begin{array}{l} \mathbf{A}\zeta_x^* - \omega \zeta_x^* \lambda_x = \mathbf{X} \mu_x^* - \mathbf{A}\mathbf{X}, \\ \mathbf{X}^T \zeta_x^* + \zeta_x^{*T} \mathbf{X} = 0, \end{array} \right.$$

where $\omega \in [0, 1)$, ζ_x^* is the computed search direction, and μ_x^* is the Lagrange multiplier that guarantees that the search direction ζ_x^* belongs to the tangent space of the constraints. It should be noted that the linear system of equations (4.8) can be solved efficiently by exploiting additional sparsity structure in \mathbf{A} [3]. Once the solution for μ_x^* in (4.8) is obtained by eliminating the constraint $\mathbf{X}^T \zeta_x^* + \zeta_x^{*T} \mathbf{X} = 0$, the Riemannian

steepest-descent update x_+ with unit step-size is

$$\begin{aligned}
 & \left. \begin{aligned} \mathbf{A}\zeta_x^* - \omega\zeta_x^*\lambda_x &= \mathbf{X}\mu_x^* - \mathbf{A}\mathbf{X} \\ x_+ &= R_x(\zeta_x^*) \equiv \text{qf}(\mathbf{X} + \zeta_x^*) \end{aligned} \right\} \\
 (4.9) \quad & \Leftrightarrow \left\{ \begin{aligned} \mathbf{A}\mathbf{Z} - \omega\mathbf{Z}\lambda_x &= \mathbf{X}(\mu_x^* - \omega\lambda_x), \\ x_+ &= \text{qf}(\mathbf{Z}), \end{aligned} \right. \\
 & \Leftrightarrow \left\{ \begin{aligned} \mathbf{A}\mathbf{Z} - \omega\mathbf{Z}\mathbf{X}^T\mathbf{A}\mathbf{X} &= \mathbf{X}(\mu_x^* - \omega\lambda_x) \quad (\text{we solve for } \mathbf{Z}), \\ x_+ &= \text{qf}(\mathbf{Z}), \end{aligned} \right.
 \end{aligned}$$

where $\mathbf{Z} := \mathbf{X} + \zeta_x^*$, $R_x(\cdot)$ is the retraction operation defined in Table 3, and $\text{qf}(\mathbf{Z})$ is the Q-factor of the QR decomposition of \mathbf{Z} . It should be noted that $R_x(\zeta_x^*)$ and $\text{qf}(\mathbf{X} + \zeta_x^*)$ define the same element on the Grassmann manifold and hence the equivalence in (4.9). It should also be emphasized that the update (4.9) is similar to (4.7) in the neighborhood of a minimum up to the extra terms ω and $\mu_x^* - \omega\lambda_x$. The similarity of (4.9) and (4.7) is more profound for $r = 1$ and when ω is updated, e.g., using (3.13). This connection also shows a way to extend GRQI beyond the neighborhood of a minimum. Consequently, GRQI update (4.7) proposed by Absil et al. [3] has the interpretation of a Riemannian steepest-descent algorithm with the metric from (4.6).

When $\mathbf{A} \not\prec \mathbf{0}$. Consider first the case when $\mathbf{A} \prec \mathbf{0}$ that falls under Case II of section 3.3, i.e., with $\omega_1 = \omega$ and $\omega_2 = 1$ in (3.9), suggesting (locally) a family of Riemannian metrics that has the form

$$(4.10) \quad g_x(\xi_x, \zeta_x) = \omega \text{Trace}(\xi_x^T \mathbf{A} \zeta_x) - \text{Trace}(\xi_x^T \mathbf{B} \zeta_x \lambda_x),$$

where ξ_x and ζ_x are vectors in the tangent space of the constraints and where $\omega \in [0, 1)$. The expression for the least-squares Lagrange multiplier from (2.4) is $\lambda_x = \text{Sym}((\mathbf{X}^T \mathbf{B} \mathbf{B} \mathbf{X})^{-1} (\mathbf{X}^T \mathbf{B} \mathbf{A} \mathbf{X}))$, where $\text{Sym}(\cdot)$ extracts the symmetric part of a square matrix, i.e., $\text{Sym}(\mathbf{D}) = (\mathbf{D} + \mathbf{D}^T)/2$.

It should be noted that as $-\lambda_x$ is only guaranteed to be positive definite locally in the neighborhood of a local minimum, the metric characterization (4.10) is a Riemannian metric *only* in the neighborhood of a local minimum. In order to extend the metric away from a local minimum, we modify the metric, defined in (4.10), by replacing $-\lambda_x$ with $(\lambda_x^T \lambda_x)^{1/2}$, resulting in the modified metric

$$(4.11) \quad g_x(\xi_x, \zeta_x) = \omega \text{Trace}(\xi_x^T \mathbf{A} \zeta_x) + \text{Trace}(\xi_x^T \mathbf{B} \zeta_x (\lambda_x^T \lambda_x)^{1/2}),$$

where $(\lambda_x^T \lambda_x)^{1/2}$ is the matrix square root of $\lambda_x^T \lambda_x$ that is well defined as long as λ_x is full rank. The full-rank assumption of λ_x is required for the metric, from (4.11), to be a smooth inner product. The modified metric, shown in (4.11), is also a good metric candidate for the case when \mathbf{A} is symmetric indefinite since $(\lambda_x^T \lambda_x)^{1/2}$ is also positive definite in this case.

The proposed metric in (4.11) with $\omega = 0$ generalizes the well-known *power iteration algorithm* for computing the dominant eigenvalues of a matrix [16, section 8.2.1]. To show this consider the case $\mathbf{B} = \mathbf{I}$. Given an iterate $x = \mathbf{X}$ such that $\mathbf{X}^T \mathbf{X} = \mathbf{I}$, the update of the Riemannian steepest-descent algorithm with unit step-size has the characterization (after a few computations) that is equivalent to $x_+ = \text{qf}(\mathbf{X}(\mathbf{I} + \lambda_x(\lambda_x^T \lambda_x)^{-1/2}) - \mathbf{A}\mathbf{X}(\lambda_x^T \lambda_x)^{-1/2})$. Locally, in the neighborhood of a minimum, $\mathbf{I} + \lambda_x(\lambda_x^T \lambda_x)^{-1/2} \approx \mathbf{0}$, and therefore the equivalent update is $x_+ = \text{qf}(\mathbf{A}\mathbf{X})$, which is the classical power iteration update [16, section 8.2.1]. In other words, the




FIG. 4. Benefits of the proposed metric (4.6) for the generalized eigenvalue problem to compute the extreme 5-dimensional subspace (corresponding to the smallest 5 eigenvalues) of the matrix pencil (\mathbf{A}, \mathbf{B}) of size 500×500 . The problem instance is described in section 4.2. Shown are 10 runs of the Riemannian steepest-descent algorithms with random initializations for the problem instance. The distance to the solution is defined as the square root of the sum of squared canonical angles between the current subspace and the dominant 5-dimensional subspace of $\mathbf{B}^{-1}\mathbf{A}$.

power algorithm has the interpretation of a Riemannian steepest-descent algorithm with the metric from (4.11) and with $\omega = 0$. Similarly, the steepest-descent algorithm with a shifted version of the metric (4.11), i.e., for ω updated with iterations, generalizes the GRQI algorithm proposed by Absil et al. [3].

A similar insight still holds when the quadratic cost is generalized to a strictly concave function, i.e., minimizing a concave cost (or maximizing a convex cost) with orthogonality constraints. For the metric with $\omega = 0$, i.e., taking only the constraint-related term, this is the essence of the *generalized power method* proposed by Journée et al. [18].

4.2. A numerical illustration. As a numerical comparison, we consider the example proposed by Manton [22, section 8]. \mathbf{A} is a diagonal matrix of size 500×500 with entries equispaced on the interval $[10, 11]$. \mathbf{B} is chosen as the identity matrix of size 500×500 . In Figure 4, we seek to compute the $r = 5$ smallest eigenvalues of $\mathbf{B}^{-1}\mathbf{A}$. The algorithms compared are the Riemannian steepest-descent algorithms with the Euclidean metric (4.3) and the preconditioned Riemannian metric in (4.6) with the ω -updating procedure (3.13). Both the algorithms are stopped when either the norm of the gradient is below 10^{-8} or when they complete 500 iterations. Distances of the iterates to the solution are plotted for the algorithms. The distance of an iterate \mathbf{X} to the solution \mathbf{X}_{opt} is defined as the square root of the sum of canonical angles between \mathbf{X} and \mathbf{X}_{opt} . In MATLAB it is computed using the command `norm(acos(svd(orth(X)'*orth(Xopt))))`. Figure 4 shows the initial 50 iterations, where we see that tuning the metric to the problem structure leads to improved performance.

5. Quadratic optimization with rank constraints. This class of problems has met with considerable interest in recent years. Applications include collaborative filtering [34], multivariate linear classification [7], dimensionality reduction [12], learning of low-rank distances [20, 24, 27], filter design problems [22], model reduction in

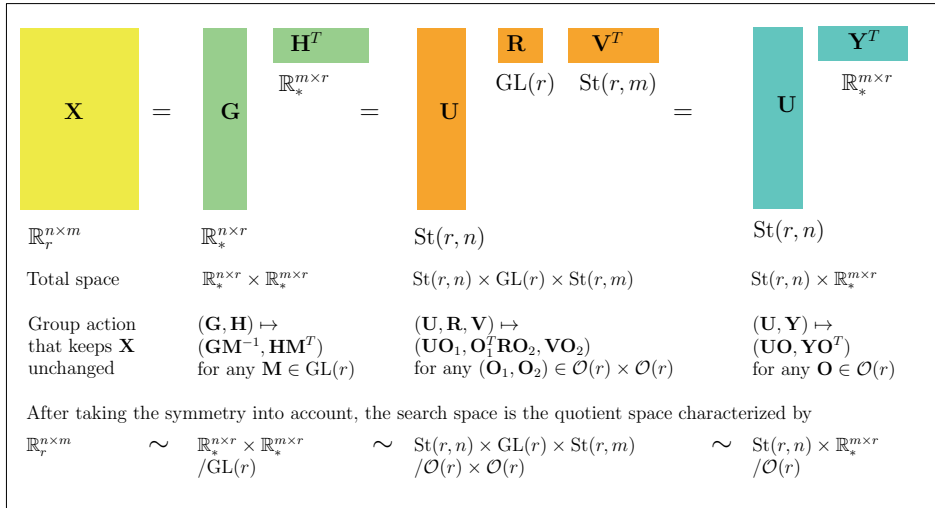


FIG. 5. Fixed-rank matrix factorizations lead to quotient search spaces due to their intrinsic symmetries. The pictures emphasize the situation of interest, i.e., the rank r is small compared to the matrix dimensions. $\text{St}(r, n)$ is the set of $n \times r$ matrices with orthogonal columns, $\mathbb{R}_*^{n \times r}$ is the set of $n \times r$ matrices with full column rank, $\text{GL}(r)$ is the set of $r \times r$ square matrices with nonzero determinant, and $\mathcal{O}(r)$ is the set of $r \times r$ square matrices with orthonormal columns and rows.

dynamical systems [8, 21, 38], sparse principal components analysis [11, 17], computing maximal cut of a graph [11, 17], and low-rank matrix completion [9, 19, 26, 31, 37], to name just a few.

In all those applications, the discussion in section 3.2 allows us to propose novel preconditioned Riemannian metrics. These metrics connect to those proposed by Mishra, Adithya Apuroop, and Sepulchre [25], Mishra and Sepulchre [28], Mishra and Vandereycken [29], and Ngo and Saad [31] for specific optimization problems.

A popular way to characterize the set of fixed-rank matrices is through fixed-rank matrix factorizations. Most matrix factorizations have symmetry properties that make them nonunique. And in many cases the set of rank r of $n \times m$ matrices $\mathbb{R}_r^{n \times m}$ is identified with structured (smooth and differentiable) quotient spaces [1, 23, 26]. Figure 5 shows three different fixed-rank matrix factorizations and the quotient manifold structure of the set $\mathbb{R}_r^{n \times m}$.

To identify Riemannian metrics on the low-rank manifold $\mathbb{R}_r^{n \times m}$, we consider minimization of a convex quadratic cost function. Specifically, we focus on the parameterization $\mathbf{X} = \mathbf{GH}^T$, where $\mathbf{X} \in \mathbb{R}_r^{n \times m}$, $\mathbf{G} \in \mathbb{R}_*^{n \times r}$ (the set of full column rank matrices), and $\mathbf{H} \in \mathbb{R}_*^{m \times r}$. Other fixed-rank matrix factorizations are dealt with similarly.

Consider the optimization problem

$$(5.1) \quad \min_{\mathbf{X} \in \mathbb{R}_r^{n \times m}} \frac{1}{2} \text{Trace}(\mathbf{X}^T \mathbf{A} \mathbf{X} \mathbf{B}) + \text{Trace}(\mathbf{X}^T \mathbf{C}),$$

where $\mathbf{A} \succ 0$ of size $n \times n$, $\mathbf{B} \succ 0$ of size $m \times m$, and $\mathbf{C} \in \mathbb{R}^{n \times m}$. Positive definiteness of \mathbf{A} and \mathbf{B} implies that the cost function is bounded from below and is *convex* in \mathbf{X} . Invoking the low-rank parameterization $\mathbf{X} = \mathbf{GH}^T$, shown in Figure 5, the problem

(5.1) translates to

$$\begin{aligned} & \min_{(\mathbf{G}, \mathbf{H}) \in \mathbb{R}^{n \times r} \times \mathbb{R}^{m \times r}} \frac{1}{2} \text{Trace}(\mathbf{H}\mathbf{G}^T \mathbf{A}\mathbf{G}\mathbf{H}^T \mathbf{B}) + \text{Trace}(\mathbf{H}\mathbf{G}^T \mathbf{C}) \\ & \text{subject to} \quad [(\mathbf{G}, \mathbf{H})] \in \mathbb{R}_*^{n \times r} \times \mathbb{R}_*^{m \times r} / \text{GL}(r), \end{aligned}$$

where the equivalence class $[(\mathbf{G}, \mathbf{H})] := \{(\mathbf{G}\mathbf{M}^{-1}, \mathbf{H}\mathbf{M}^T) : \mathbf{M} \in \text{GL}(r)\}$ and $\text{GL}(r)$ is the set of $r \times r$ square matrices of nonzero determinant.

A conventional way to handle this symmetry in the Riemannian framework is endowing the set $\mathbb{R}_*^{n \times r}$ with the *natural metric* [2, section 3.6.4]. Since the computational space \mathcal{M} is the product space $\mathbb{R}_*^{n \times r} \times \mathbb{R}_*^{m \times r}$, the natural metric is

$$g_x(\eta_x, \xi_x) = \text{Trace}((\mathbf{G}^T \mathbf{G})^{-1} \eta_{\mathbf{G}}^T \xi_{\mathbf{G}}) + \text{Trace}((\mathbf{H}^T \mathbf{H})^{-1} \eta_{\mathbf{H}}^T \xi_{\mathbf{H}}),$$

where x has the matrix representation $(\mathbf{G}, \mathbf{H}) \in \mathbb{R}_*^{n \times r} \times \mathbb{R}_*^{m \times r}$ and ξ_x, η_x are vectors belonging to the tangent space $\mathbb{R}^{n \times r} \times \mathbb{R}^{m \times r}$, i.e., ξ_x has the matrix representation $(\xi_{\mathbf{G}}, \xi_{\mathbf{H}}) \in \mathbb{R}^{n \times r} \times \mathbb{R}^{m \times r}$.

In contrast, we follow the developments in section 3.3 to propose a family of metrics that takes the problem structure into account by exploiting the structure of the regularized Lagrangian (3.8). Since the set $\mathbb{R}_*^{n \times r} \times \mathbb{R}_*^{m \times r}$ is an open subset of the space $\mathbb{R}^{n \times r} \times \mathbb{R}^{m \times r}$, the Lagrangian only consists of the cost function, i.e.,

$$\begin{aligned} (5.2) \quad & \mathcal{L}(x) = \text{Trace}(\mathbf{H}\mathbf{G}^T \mathbf{A}\mathbf{G}\mathbf{H}^T \mathbf{B})/2 + \text{Trace}(\mathbf{H}\mathbf{G}^T \mathbf{C}) \\ \Rightarrow & \mathcal{L}_x(x) = (\mathbf{A}\mathbf{G}\mathbf{H}^T \mathbf{B}\mathbf{H} + \mathbf{C}\mathbf{H}, \mathbf{B}\mathbf{H}\mathbf{G}^T \mathbf{A}\mathbf{G} + \mathbf{C}^T \mathbf{G}) \\ \Rightarrow & \text{D}^2 \mathcal{L}(x)[\xi_x] = (\mathbf{A}\xi_{\mathbf{G}} \mathbf{H}^T \mathbf{B}\mathbf{H} + 2\mathbf{A}\mathbf{G}\text{Sym}(\mathbf{H}^T \mathbf{B}\xi_{\mathbf{H}}) + \mathbf{C}\xi_{\mathbf{H}}, \\ & \mathbf{B}\xi_{\mathbf{H}} \mathbf{G}^T \mathbf{A}\mathbf{G} + 2\mathbf{B}\mathbf{H}\text{Sym}(\mathbf{G}^T \mathbf{A}\xi_{\mathbf{G}}) + \mathbf{C}^T \xi_{\mathbf{G}}), \end{aligned}$$

where x has the matrix representation $(\mathbf{G}, \mathbf{H}) \in \mathbb{R}_*^{n \times r} \times \mathbb{R}_*^{m \times r}$, ξ_x has the matrix representation $(\xi_{\mathbf{G}}, \xi_{\mathbf{H}}) \in \mathbb{R}^{n \times r} \times \mathbb{R}^{m \times r}$, $\mathcal{L}_x(x)$ is the first-order derivative of $\mathcal{L}(x)$, $\text{D}^2 \mathcal{L}(x)[\xi_x]$ is the second-order derivative of $\mathcal{L}(x)$ applied in the direction ξ_x , and $\text{Sym}(\cdot)$ extracts the symmetric part of a square matrix, i.e., $\text{Sym}(\mathbf{D}) = (\mathbf{D}^T + \mathbf{D})/2$.

It is readily checked that the Lagrangian $\mathcal{L}(x)$ in (5.2) remains unchanged under the transformation $(\mathbf{G}, \mathbf{H}) \mapsto (\mathbf{G}\mathbf{M}^{-1}, \mathbf{H}\mathbf{M}^T)$ for all $\mathbf{M} \in \text{GL}(r)$. Subsequently, we have the following proposition for constructing a family of Riemannian metrics for (5.1) on the fixed-rank quotient manifold.

PROPOSITION 5.1. *There exists a family of Riemannian metrics*

$$\begin{aligned} (5.3) \quad g_x(\xi_x, \eta_x) = & \omega_1 \langle \eta_{\mathbf{G}}, \mathbf{A}\xi_{\mathbf{G}} \mathbf{H}^T \mathbf{B}\mathbf{H} \rangle \\ & + \omega_2 \langle \eta_{\mathbf{G}}, 2\mathbf{A}\mathbf{G}\text{Sym}(\mathbf{H}^T \mathbf{B}\xi_{\mathbf{H}}) + \mathbf{C}\xi_{\mathbf{H}} \rangle \\ & + \omega_3 \langle \eta_{\mathbf{H}}, \mathbf{B}\xi_{\mathbf{H}} \mathbf{G}^T \mathbf{A}\mathbf{G} \rangle \\ & + \omega_4 \langle \eta_{\mathbf{H}}, 2\mathbf{B}\mathbf{H}\text{Sym}(\mathbf{G}^T \mathbf{A}\xi_{\mathbf{G}}) + \mathbf{C}^T \xi_{\mathbf{G}} \rangle \end{aligned}$$

on $\mathbb{R}_*^{n \times r} \times \mathbb{R}_*^{m \times r}$ with $\omega_1, \omega_2, \omega_3, \omega_4 \in [0, 1]$, each of which induces a Riemannian metric on the quotient manifold $\mathbb{R}_*^{n \times r} \times \mathbb{R}_*^{m \times r} / \text{GL}(r)$ in the neighborhood of the local minimum of (5.1). Here $x = (\mathbf{G}, \mathbf{H}) \in \mathcal{M}$, $\mathcal{M} = \mathbb{R}_*^{n \times r} \times \mathbb{R}_*^{m \times r}$, and ξ_x, η_x are vectors in the tangent space $T_x \mathcal{M}$.

Proof. First, we show that the metrics from (5.3) respect the condition in (3.1). In order to show that the metric does not change along the equivalence class $[x] = [(\mathbf{G}, \mathbf{H})] = \{(\mathbf{G}\mathbf{M}^{-1}, \mathbf{H}\mathbf{M}^T) : \mathbf{M} \in \text{GL}(r)\}$ for all $\mathbf{M} \in \text{GL}(r)$, it is equivalent, but simplified following [2, Proposition 3.6.1], to show that the metric for tangent vectors $\xi_x, \eta_x \in T_x \mathcal{M}$ does not change under the transformations $(\mathbf{G}, \mathbf{H}) \mapsto (\mathbf{G}\mathbf{M}^{-1}, \mathbf{H}\mathbf{M}^T)$,

TABLE 4

Optimization-related ingredients for the problem (5.1). The numerical complexity per iteration of the Riemannian steepest-descent algorithm depends on solving for ζ_x for the search direction computation. For example, sparsity in matrices \mathbf{A} and \mathbf{B} considerably reduces the computation cost. The retraction mapping is the Cartesian product of the standard retraction mapping on the manifold $\mathbb{R}_*^{n \times r}$ [2, Example 3.6.4]. A few choices of the regularizing parameters $\omega_1, \omega_2, \omega_3, \omega_4 \in [0, 1]$ for relevant situations are discussed in section 5.1.

	$\min_{\substack{\mathbf{G} \in \mathbb{R}_*^{n \times r} \\ \mathbf{H} \in \mathbb{R}_*^{m \times r}}} \text{Trace}(\mathbf{H}\mathbf{G}^T\mathbf{A}\mathbf{G}\mathbf{H}^T\mathbf{B})/2 + \text{Trace}(\mathbf{H}\mathbf{G}^T\mathbf{C})$
Matrix representation of an element $x \in \mathcal{M}$	$x = (\mathbf{G}, \mathbf{H})$
Computational space \mathcal{M}	$\mathbb{R}_*^{n \times r} \times \mathbb{R}_*^{m \times r}$
Group action	$(\mathbf{G}\mathbf{M}^{-1}, \mathbf{H}\mathbf{M}^T)$ for all $\mathbf{M} \in \text{GL}(r)$
Quotient space	$\mathbb{R}_*^{n \times r} \times \mathbb{R}_*^{m \times r} / \text{GL}(r)$
Tangent vectors in $T_x\mathcal{M}$	$\xi_x = (\xi_{\mathbf{G}}, \xi_{\mathbf{H}}) \in \mathbb{R}^{n \times r} \times \mathbb{R}^{m \times r}$
Metric $g_x(\xi_x, \zeta_x)$ for $\xi_x, \zeta_x \in T_x\mathcal{M}$	$g_x(\xi_x, \eta_x) = \omega_1 \langle \eta_{\mathbf{G}}, \mathbf{A}\xi_{\mathbf{G}}\mathbf{H}^T\mathbf{B}\mathbf{H} \rangle$ $+ \omega_2 \langle \eta_{\mathbf{G}}, 2\mathbf{A}\mathbf{G}\text{Sym}(\mathbf{H}^T\mathbf{B}\xi_{\mathbf{H}}) + \mathbf{C}\xi_{\mathbf{H}} \rangle$ $+ \omega_3 \langle \eta_{\mathbf{H}}, \mathbf{B}\xi_{\mathbf{H}}\mathbf{G}^T\mathbf{A}\mathbf{G} \rangle$ $+ \omega_4 \langle \eta_{\mathbf{H}}, 2\mathbf{B}\mathbf{H}\text{Sym}(\mathbf{G}^T\mathbf{A}\xi_{\mathbf{G}}) + \mathbf{C}^T\xi_{\mathbf{G}} \rangle$ <p>or the metrics proposed in section 5.1</p>
Cost function	$f(x) = \text{Trace}(\mathbf{H}\mathbf{G}^T\mathbf{A}\mathbf{G}\mathbf{H}^T\mathbf{B})/2 + \text{Trace}(\mathbf{H}\mathbf{G}^T\mathbf{C})$
First-order derivative of $f(x)$	$f_x(x) = (\mathbf{S}\mathbf{H}, \mathbf{S}^T\mathbf{G})$, where $\mathbf{S} = \mathbf{A}\mathbf{G}\mathbf{H}^T\mathbf{B} + \mathbf{C}$
Search direction	$\arg \min_{\zeta_x \in T_x\mathcal{M}} f(x) + \langle f_x(x), \zeta_x \rangle + \frac{1}{2}g_x(\zeta_x, \zeta_x)$
Retraction $R_x(\xi_x)$ that maps a search direction ξ_x onto \mathcal{M}	$(\mathbf{G} + \xi_{\mathbf{G}}, \mathbf{H} + \xi_{\mathbf{H}})$

$(\eta_{\mathbf{G}}, \eta_{\mathbf{H}}) \mapsto (\eta_{\mathbf{G}}\mathbf{M}^{-1}, \eta_{\mathbf{H}}\mathbf{M}^T)$, and $(\xi_{\mathbf{G}}, \xi_{\mathbf{H}}) \mapsto (\xi_{\mathbf{G}}\mathbf{M}^{-1}, \xi_{\mathbf{H}}\mathbf{M}^T)$. A few extra computations show that indeed the metrics from (5.3) respect the condition in (3.1).

Second, we show the construction of one particular family of Riemannian metrics. To this end, consider the case in (5.3) where $\omega_1 = 1$ and $\omega_2 = \omega_3 = \omega_4 = 0$. Restricting $\omega \in [0, 1]$ guarantees that in the neighborhood of a minimum $g_x(\zeta_x, \zeta_x) > 0$ for all $\zeta_x \in T_x\mathcal{M}$, satisfying the criterion of positive definiteness on the tangent space.

The proofs for symmetry compatibility and positive definiteness on the tangent space conclude the proof of the proposition. \square

Matrix characterizations of various optimization-related ingredients are summarized in Table 4. The retraction operator is the standard generalization of the retraction operator on the manifold $\mathbb{R}_*^{n \times r}$ defined by Absil, Mahony, and Sepulchre [2, Example 3.6.4].

It should be noted that numerical performance of algorithms depends on com-

putting the Riemannian gradient efficiently with the metric (5.3). This may become a numerically cumbersome task due to a number of coupled terms that are involved in (5.3). However, below we show that the problem structure can be further exploited to decompose the metric (5.3) into a locally dominating part with a simpler metric structure and a *weighted* remainder. The dominant approximation may be preferred in a number of situations.

5.1. Metric tuning and shift policies. It should be emphasized that the cost function in (5.1) is *convex and quadratic* in \mathbf{X} . Consequently, the cost function is also convex and quadratic in the arguments (\mathbf{G}, \mathbf{H}) *individually*. As a consequence, the block diagonal elements of the second-order derivative $\mathcal{L}_{xx}(x)$ of the Lagrangian (5.2) are strictly positive definite. This enables us to construct a family of Riemannian metrics with *shifts* of the form

$$(5.4) \quad g_x(\xi_x, \eta_x) = \omega \langle \eta_{\mathbf{G}}, 2\mathbf{A}\mathbf{G}\text{Sym}(\mathbf{H}^T \mathbf{B}\xi_{\mathbf{H}}) + \mathbf{C}\xi_{\mathbf{H}} \rangle \\ + \omega \langle \eta_{\mathbf{H}}, 2\mathbf{B}\mathbf{H}\text{Sym}(\mathbf{G}^T \mathbf{A}\xi_{\mathbf{G}}) + \mathbf{C}^T \xi_{\mathbf{G}} \rangle \\ + \underbrace{\langle \eta_{\mathbf{G}}, \mathbf{A}\xi_{\mathbf{G}}\mathbf{H}^T \mathbf{B}\mathbf{H} \rangle + \langle \eta_{\mathbf{H}}, \mathbf{B}\xi_{\mathbf{H}}\mathbf{G}^T \mathbf{A}\mathbf{G} \rangle}_{\text{from block diagonal approximation of } \mathcal{L}_{xx}(x)}$$

in the neighborhood of a minimum, where $x = (\mathbf{G}, \mathbf{H}) \in \mathbb{R}_*^{n \times r} \times \mathbb{R}_*^{m \times r}$, ξ_x, η_x are tangent vectors in $\mathbb{R}^{n \times r} \times \mathbb{R}^{m \times r}$, and $\omega \in [0, 1)$. The form in (5.4) is derived from (5.3) by choosing $\omega_1 = \omega_3 = 1$ and $\omega_2 = \omega_4 = \omega$. It should be noted that the case $\omega = 1$ is discarded to prevent singularity of the metric candidate.

Away from the neighborhood, the metric (5.4) with $\omega = 0$ becomes a good metric candidate, as $\mathbf{H}^T \mathbf{B}\mathbf{H}$ and $\mathbf{G}^T \mathbf{A}\mathbf{G}$ are positive definite for all $(\mathbf{G}, \mathbf{H}) \in \mathbb{R}_*^{n \times r} \times \mathbb{R}_*^{m \times r}$. The other benefit of ω being 0 is that the resulting metric has a *simpler matrix characterization*, and hence it may be preferred in numerically demanding instances.

5.2. Symmetric positive definite matrices. A popular subset of fixed-rank matrices is the set of symmetric positive semidefinite matrices [11, 17, 24, 38]. The set $S_+(r, n)$, the set of rank- r symmetric positive semidefinite matrices of size $n \times n$, is equivalent to the set $\mathbb{R}_r^{n \times m}$ with symmetry imposed on the rows and columns, and therefore it admits a number of factorizations similar to those in Figure 5. Consequently, the low-rank parameterization discussed earlier, in the context of the general case, has the counterpart $\mathbf{X} = \mathbf{Y}\mathbf{Y}^T$, where $\mathbf{X} \in S_+(r, n)$ and $\mathbf{Y} \in \mathbb{R}_*^{n \times r}$ (full column rank matrices of size $n \times r$). This parameterization is not unique, as $\mathbf{X} \in S_+(r, n) = \mathbf{Y}\mathbf{Y}^T$ remains unchanged under the transformation $\mathbf{Y} \mapsto \mathbf{Y}\mathbf{O}$ for all $\mathbf{O} \in \mathcal{O}(r)$, where $\mathcal{O}(r)$ is set of orthogonal matrices of size $r \times r$ such that $\mathbf{O}\mathbf{O}^T = \mathbf{O}^T\mathbf{O} = \mathbf{I}$. The resulting search space is, thus, the set of equivalence classes $[\mathbf{Y}] = \{\mathbf{Y}\mathbf{O} : \mathbf{O} \in \mathcal{O}(r)\}$ and is the quotient manifold $\mathbb{R}_*^{n \times r} / \mathcal{O}(r)$ [17].

The following proposition summarizes the discussion on Riemannian metrics for the case of symmetric positive semidefinite matrices.

PROPOSITION 5.2. *Consider the optimization problem*

$$(5.5) \quad \min_{\mathbf{X} \in \mathbb{R}^{n \times n}} \quad \frac{1}{2} \text{Trace}(\mathbf{X}\mathbf{A}\mathbf{X}\mathbf{B}) + \text{Trace}(\mathbf{X}\mathbf{C}) \\ \text{subject to} \quad \mathbf{X} \in S_+(r, n),$$

where $\mathbf{A}, \mathbf{B} \succ 0$ of size $n \times n$ and $\mathbf{C} \in \mathbb{R}^{n \times n}$ is a symmetric matrix. Consider also the factorization $\mathbf{X} = \mathbf{Y}\mathbf{Y}^T$ of rank- r symmetric positive semidefinite matrices to encode the rank constraint, where $\mathbf{Y} \in \mathbb{R}_*^{n \times r}$ (full column rank matrices).

There exists a family of Riemannian metrics

$$(5.6) \quad g_x(\xi_x, \eta_x) = \omega \langle \eta_x, 2\mathbf{A}\mathbf{Y}\text{Sym}(\mathbf{Y}^T \mathbf{B}\xi_x) + 2\mathbf{B}\mathbf{Y}\text{Sym}(\mathbf{Y}^T \mathbf{A}\xi_x) + 2\mathbf{C}\xi_x \rangle + \underbrace{\langle \eta_x, \mathbf{A}\xi_x \mathbf{Y}^T \mathbf{B}\mathbf{Y} + \mathbf{B}\xi_x \mathbf{Y}^T \mathbf{A}\mathbf{Y} \rangle}_{\text{Dominant positive definite approximation of } \mathcal{L}_{xx}(x)}$$

on $\mathbb{R}_*^{n \times r}$ with $\omega \in [0, 1)$, each of which induces a Riemannian metric on $\mathbb{R}_*^{n \times r} / \mathcal{O}(r)$ in the neighborhood of the minimum of the problem (5.5). Here $x = \mathbf{Y} \in \mathbb{R}_*^{n \times r}$, ξ_x, η_x are vectors in the tangent space $\mathbb{R}^{n \times r}$, and $\mathcal{L}_{xx}(x)$ is the second-order derivative of the Lagrangian. Beyond the neighborhood, the metric (5.6) with $\omega = 0$ becomes a good metric candidate, as $\mathbf{Y}^T \mathbf{B}\mathbf{Y}$ and $\mathbf{Y}^T \mathbf{A}\mathbf{Y}$ are positive definite for all $\mathbf{Y} \in \mathbb{R}_*^{n \times r}$.

Proof. The proof follows from the discussion in section 5.1. □

5.3. A numerical illustration. We showcase the Riemannian preconditioning approach for computing low-rank solutions to the generalized Lyapunov equation of the form

$$(5.7) \quad \mathbf{A}\mathbf{X}\mathbf{B} + \mathbf{B}\mathbf{X}\mathbf{A} = \mathbf{C},$$

where $\mathbf{A}, \mathbf{B} \succ 0$, and \mathbf{C} is a low-rank symmetric positive semidefinite matrix. Matrices have appropriate dimensions. \mathbf{A} is referred to as the *system matrix*, and \mathbf{B} is referred to as the *mass matrix*. The solution to (5.7) is expected to be low rank and symmetric positive semidefinite [8, 21, 38].

In order to compute low-rank solutions to (5.7), we minimize the *energy norm* $\text{Trace}(\mathbf{X}\mathbf{A}\mathbf{X}\mathbf{B}) - \text{Trace}(\mathbf{X}\mathbf{C})$ over $S_+(r, n)$ [38]. Proposition 5.2 allows us to characterize a family of metrics in (5.6) for solving the generalized Lyapunov equation. In contrast, an alternative is to consider the Euclidean metric, i.e.,

$$(5.8) \quad g_x(\xi_x, \zeta_x) = \text{Trace}(\zeta_x^T \xi_x),$$

where $x = \mathbf{Y}$ and ξ_x and ζ_x are tangent vectors. This is, for example, the Riemannian metric proposed by Journée et al. [17]. It is invariant to the group action $\mathbf{Y} \mapsto \mathbf{Y}\mathbf{O}$ for all $\mathbf{O} \in \mathcal{O}(r)$. Although the alternative choice (5.8) is appealing for its numerical simplicity, the following test case clearly illustrates the benefits of the Riemannian preconditioning approach.

We consider the standard benchmark problem from [33, Example 2.1] that corresponds to discretization of a one-dimensional heat equation from heat flow in a thin rod. For this example, \mathbf{A} is a tridiagonal matrix of size 500×500 . The main diagonal of \mathbf{A} has all the elements equal to 2. In addition, the first diagonals below and above the main diagonal of \mathbf{A} have all the entries equal to -1 . \mathbf{A} is an ill-conditioned matrix with condition number 10^5 . The mass matrix \mathbf{B} is an identity matrix of size 500×500 . The matrix \mathbf{C} is a rank one matrix of the form ee^T , where e^T is a row vector of length 500 of the form $[0 \ 0 \ \dots \ 0 \ 1]$. We seek to find a rank-5 matrix that best solves the generalized Lyapunov equation (5.7). Both the algorithms are stopped when either the norm of the gradient is below 10^{-8} or when they complete 500 iterations. The plots in Figure 6 show the progress of relative residual $\|\mathbf{A}\mathbf{X}\mathbf{B} + \mathbf{B}\mathbf{X}\mathbf{A} - \mathbf{C}\|_F / \|\mathbf{C}\|_F$ with iterations over 10 random initializations, where $\mathbf{X} = \mathbf{Y}\mathbf{Y}^T$. The Riemannian algorithm with the metric (5.6) and $\omega = 0$ convincingly outperforms the algorithm based on the Euclidean metric (5.8) in Figure 6 for a number of runs, where the initial 300 iterations are shown for clarity.




FIG. 6. The generalized low-rank Lyapunov equation problem (5.7). The test case is the benchmark problem from [33, Example 2.1] with $n = 500$. The proposed Riemannian preconditioning approach with the metric (5.6) and $\omega = 0$ drastically improves the performance over the algorithm based on the Euclidean metric (5.8). Additionally, the choice $\omega = 0$ leads to a simpler metric structure that can be exploited in a large-scale setup. Here we show the convergence of the relative residual $\|\mathbf{AXB} + \mathbf{BXA} - \mathbf{C}\|_F / \|\mathbf{C}\|_F$ (different from the cost function $\text{Trace}(\mathbf{XAXB}) - \text{Trace}(\mathbf{XC})$) that is often used as a measure of recovery.

6. Conclusion. This paper addresses the important issue of selecting a metric in the Riemannian optimization framework on a quotient manifold. We have shown that sequential quadratic programming provides an insight into selecting a family of Riemannian metrics that takes into account second-order information of the problem. Quadratic optimization with orthogonality or rank constraints provides a class of nonconvex problems for which the method is particularly insightful, thanks to local convexity of the cost and constraint when taken separately. In those instances, Riemannian preconditioning connects to a number of existing algorithms and provides a geometric interpretation of a number of “shift” policies in numerical linear algebra.

Acknowledgments. We thank the editor and the two anonymous reviewers for carefully checking the paper and providing a number of helpful remarks.

REFERENCES

- [1] P. A. ABSIL, L. AMODEI, AND G. MEYER, *Two Newton methods on the manifold of fixed-rank matrices endowed with Riemannian quotient geometries*, *Comput. Statist.*, 29 (2014), pp. 569–590, doi:10.1007/s00180-013-0441-6.
- [2] P. A. ABSIL, R. MAHONY, AND R. SEPULCHRE, *Optimization Algorithms on Matrix Manifolds*, Princeton University Press, Princeton, NJ, 2008.
- [3] P.-A. ABSIL, R. MAHONY, R. SEPULCHRE, AND P. VAN DOOREN, *A Grassmann-Rayleigh quotient iteration for computing invariant subspaces*, *SIAM Rev.*, 44 (2002), pp. 57–73, doi:10.1137/S0036144500378648.
- [4] P.-A. ABSIL, R. SEPULCHRE, P. VAN DOOREN, AND R. MAHONY, *Cubically convergent iterations for invariant subspace computation*, *SIAM J. Matrix Anal. Appl.*, 26 (2004), pp. 70–96, doi:10.1137/S0895479803422002.
- [5] P. A. ABSIL, J. TRUMPF, R. MAHONY, AND B. ANDREWS, *All Roads Lead to Newton: Feasible Second-Order Methods for Equality-Constrained Optimization*, Technical report, UCL-INMA-2009.024, Université catholique de Louvain, Louvain-la-Neuve, Belgium, 2009.
- [6] P. A. ABSIL AND P. VAN DOOREN, *Two-sided Grassmann-Rayleigh quotient iteration*, *Numer. Math.*, 114 (2010), pp. 549–571, doi:10.1007/s00211-009-0266-y.

- [7] Y. AMIT, M. FINK, N. SREBRO, AND S. ULLMAN, *Uncovering shared structures in multiclass classification*, in Proceedings of the 24th International Conference on Machine Learning, 2007, pp. 17–24.
- [8] P. BENNER AND J. SAAK, *Numerical solution of large and sparse continuous time algebraic matrix Riccati and Lyapunov equations: A state of the art survey*, GAMM-Mitt., 36 (2013), pp. 32–52, doi:10.1002/gamm.201310003.
- [9] N. BOUMAL AND P. A. ABSIL, *RTRMC: A Riemannian trust-region method for low-rank matrix completion*, in Advances in Neural Information Processing Systems (NIPS) 24, MIT Press, Cambridge, MA, 2011, pp. 406–414.
- [10] N. BOUMAL, B. MISHRA, P. A. ABSIL, AND R. SEPULCHRE, *Manopt: A MATLAB toolbox for optimization on manifolds*, J. Mach. Learn. Res., 15 (2014), pp. 1455–1459.
- [11] S. BURER AND R. MONTEIRO, *A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization*, Math. Program., 95 (2003), pp. 329–357, doi:10.1007/s10107-002-0352-8.
- [12] D. CAI, X. HE, J. HAN, *Efficient kernel discriminant analysis via spectral regression*, in Proceedings of the IEEE International Conference on Data Mining (ICDM), 2007, pp. 427–432.
- [13] 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, doi:10.1137/S0895479895290954.
- [14] L. ELDÉN AND H. PARK, *A Procrustes problem on the Stiefel manifold*, Numer. Math., 82 (1999), pp. 599–619, doi:10.1007/s002110050432.
- [15] D. GABAY, *Minimizing a differentiable function over a differential manifold*, J. Optim. Theory Appl., 37 (1982) p. 177–219, doi:10.1007/BF00934767.
- [16] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, 3rd ed., The Johns Hopkins University Press, Baltimore, MD, 1996.
- [17] M. JOURNÉE, F. BACH, P.-A. ABSIL, AND R. SEPULCHRE, *Low-rank optimization on the cone of positive semidefinite matrices*, SIAM J. Optim., 20 (2010), pp. 2327–2351, doi:10.1137/080731359.
- [18] M. JOURNÉE, Y. NESTEROV, P. RICHTÁRIK, AND R. SEPULCHRE, *Generalized power method for sparse principal component analysis*, J. Mach. Learn. Res., 11 (2010), pp. 517–553.
- [19] R. H. KESHAVAN, A. MONTANARI, AND S. OH, *Matrix completion from a few entries*, IEEE Trans. Inform. Theory, 56 (2010), pp. 2980–2998, doi:10.1109/TIT.2010.2046205.
- [20] B. KULIS, M. SUSTIK, AND I. S. DHILLON, *Low-rank kernel learning with Bregman matrix divergences*, J. Mach. Learn. Res., 10 (2009), pp. 341–376.
- [21] J.-R. LI AND J. WHITE, *Low-rank solution of Lyapunov equations*, SIAM Rev., 46 (2004), pp. 693–713, doi:10.1137/S0036144504443389.
- [22] J. MANTON, *Optimization algorithms exploiting unitary constraints*, IEEE Trans. Signal Process., 50 (2002), pp. 635–650, doi:10.1109/78.984753.
- [23] G. MEYER, *Geometric Optimization Algorithms for Linear Regression on Fixed-Rank Matrices*, Ph.D. thesis, University of Liège, Liège, Belgium, 2011.
- [24] G. MEYER, S. BONNABEL, AND R. SEPULCHRE, *Regression on fixed-rank positive semidefinite matrices: A Riemannian approach*, J. Mach. Learn. Res., 11 (2011), pp. 593–625.
- [25] B. MISHRA, K. ADITHYA APUROOP, AND R. SEPULCHRE, *A Riemannian Geometry for Low-Rank Matrix Completion*, Technical report, arXiv:1211.1550, 2012.
- [26] B. MISHRA, G. MEYER, S. BONNABEL, AND R. SEPULCHRE, *Fixed-rank matrix factorizations and Riemannian low-rank optimization*, Comput. Statist., 29 (2014), pp. 591–621, doi:10.1007/s00180-013-0464-z.
- [27] B. MISHRA, G. MEYER, AND R. SEPULCHRE, *Low-rank optimization for distance matrix completion*, in Proceedings of the 50th IEEE Conference on Decision and Control (CDC-ECC), 2011, pp. 4455–4460.
- [28] B. MISHRA AND R. SEPULCHRE, *R3MC: A Riemannian three-factor algorithm for low-rank matrix completion*, in Proceedings of the 53rd IEEE Conference on Decision and Control (CDC), 2014, pp. 1137–1142.
- [29] B. MISHRA AND B. VANDEREYCKEN, *A Riemannian approach to low-rank algebraic Riccati equations*, in Proceedings of the 21st International Symposium on Mathematical Theory of Networks and Systems (MTNS), 2014, pp. 965–968.
- [30] M. MOAKHER, *Means and averaging in the group of rotations*, SIAM J. Matrix Anal. Appl., 24 (2002), pp. 1–16, doi:10.1137/S0895479801383877.
- [31] T. T. NGO AND Y. SAAD, *Scaled gradients on Grassmann manifolds for matrix completion*, in Advances in Neural Information Processing Systems (NIPS) 25, MIT Press, Cambridge, MA, 2012, pp. 1421–1429.
- [32] J. NOCEDAL AND S. J. WRIGHT, *Numerical Optimization*, 2nd ed., Financ. Eng., Springer, New York, 2006.

- [33] T. PENZL, *A cyclic low-rank Smith method for large sparse Lyapunov equations*, SIAM J. Sci. Comput., 21 (2000), pp. 1401–1418, doi:10.1137/S1064827598347666.
- [34] J. RENNIE AND N. SREBRO, *Fast maximum margin matrix factorization for collaborative prediction*, in Proceedings of the International Conference on Machine Learning (ICML), 2005, pp. 713–719.
- [35] S. T. SMITH, *Optimization techniques on Riemannian manifold*, in Hamiltonian and Gradient Flows, Algorithms and Control, Vol. 3, A. Bloch, ed., American Mathematical Society, Providence, RI, 1994, pp. 113–136.
- [36] F. J. THEIS, T. P. CASON, AND P. A. ABSIL, *Soft dimension reduction for ICA by joint diagonalization on the Stiefel manifold*, in Independent Component Analysis and Signal Separation, Springer, Berlin, Heidelberg, 2009, pp. 354–361, doi:10.1007/978-3-6420-0599-2_45.
- [37] B. VANDEREYCKEN, *Low-rank matrix completion by Riemannian optimization*, SIAM J. Optim., 23 (2013), pp. 1214–1236, doi:10.1137/110845768.
- [38] B. VANDEREYCKEN AND S. VANDEWALLE, *A Riemannian optimization approach for computing low-rank solutions of Lyapunov equations*, SIAM J. Matrix Anal. Appl., 31 (2010), pp. 2553–2579, doi:10.1137/090764566.
- [39] S. J. WRIGHT AND M. J. TENNY, *A feasible trust-region sequential quadratic programming algorithm*, SIAM J. Optim., 14 (2004), pp. 1074–1105, doi:10.1137/S1052623402413227.