

Direct Search Methods on Reductive Homogeneous Spaces

David W. Dreisigmeyer

Received: date / Accepted: date

Abstract Direct search methods are mainly designed for use in problems with no equality constraints. However there are many instances where the feasible set is of measure zero in the ambient space and no mesh point lies within it. There are methods for working with feasible sets that are (Riemannian) manifolds, but not all manifolds are created equal. In particular, reductive homogeneous spaces seem to be the most general space that can be conveniently optimized over. The reason is that a ‘law of motion’ over the feasible region is also given. Examples include \mathbb{R}^n and its linear subspaces, Lie groups and coset manifolds such as Grassmannians and Stiefel manifolds. These are important arenas for optimization, for example, in the areas of image processing and data mining. We demonstrate optimization procedures over general reductive homogeneous spaces utilizing maps from the tangent space to the manifold. A concrete implementation of the probabilistic descent direct search method is shown. This is then extended to a procedure that works solely with the manifold elements, eliminating the need for the use of the tangent space.

Keywords Direct search method · Manifold · Nonlinear optimization

Mathematics Subject Classification (2000) 65K10 · 90C56

Communicated by Alexandru Kristály.

David W. Dreisigmeyer, Corresponding author
United States Census Bureau
Center for Economic Studies
Suitland, MD
Department of Electrical and Computer Engineering
Colorado State University
Fort Collins, CO
david.wayne.dreisigmeyer@census.gov

1 Introduction

Direct search methods are typically designed to work in a full-dimensional subset of \mathbb{R}^n . So, if we only have (at most) inequality constraints present, these methods are potentially useful procedures to employ, e.g., [1–6]. This is especially true for black-box optimization or cases where derivative information is not available. However the presence of even one nonlinear equality constraint defining a feasible set of measure zero creates difficulties. In this case the probability of a mesh point being feasible is zero. It is shown how this difficulty can be avoided in some particularly convenient settings.

The mesh adaptive direct search (MADS) algorithms [2, 7] provide typical illustrations. The probability that a point on the mesh used in MADS also lies in the feasible set is zero. So every point that MADS considers will likely be infeasible. Filter [1, 8] and progressive barrier [9] techniques may be able to alleviate this situation if one is willing to violate the equality constraints somewhat. Also, the augmented Lagrangian pattern search algorithm in [10] extends the method in [11] to non-smooth problems with equality constraints. We have in mind something quite different from these methods.

A recent example of the type of problem we will be addressing is [12]. The method presented there handles linear equalities in the MADS algorithm [2]. The method presented here will deal with subsets of \mathbb{R}^n that can have a much more complicated structure.

Optimization is naturally carried out in \mathbb{R}^n , over Lie groups [13] and within other geometric objects like Grassmannians [14]. A common theme running through these is that of reductive homogeneous (often symmetric) spaces [15]. Globally these have a simple structure and convenient ways to move over them. Exploring how to optimize within them has been a worthwhile pursuit [14, 16].

Following Sharpe [15], Euclidean geometry was generalized through two routes: Riemannian geometry and Klein geometries. Euclidean geometry has: 1) a space \mathbb{R}^n , 2) with a symmetry Lie group, the special Euclidean group, that allows us to reach any point in \mathbb{R}^n from any other point, and 3) a metric. Riemannian geometry generalizes Euclidean geometry by dropping the symmetry Lie group condition allowing arbitrary curvature of our underlying space. Locally the manifold looks like \mathbb{R}^n . However the symmetry group provides a convenient way to navigate within a manifold. A Klein geometry generalizes the symmetry groups that act on the underlying spaces: it specifies a space and a way to move in that space.

Generalizing optimization over \mathbb{R}^n has tended to focus on manifolds with a Riemannian structure [17]. Importantly, Sard's theorem tells us that equality constraints can typically be taken to define a Riemannian manifold. The problem here is that the global structure of this type of manifold can be complicated in which case just moving over them without wandering off of their surface may be difficult.

We will extend direct search methods to constrained problems when the constraints are specified by a Klein geometry. The Klein geometry will implicitly define a subset of the ambient space \mathbb{R}^n of our optimization variables.

This subset is the feasible set, possibly further constrained with inequality constraints. The advantage of using a Klein geometry is that it additionally gives us information on how to move over the feasible set.

The current paper is laid out as follows. Section 2 provides a brief overview of the necessary differential geometry. Lie groups and Lie algebras are introduced in Section 3. Section 4 illustrates our methodology through a series of examples. We return for a look at Lie groups and algebras in Section 5 before presenting the optimization procedure in Section 6. In Section 7 we investigate an optimization routine where one can remain strictly in the Lie group and avoid any use of the Lie algebras. This is perhaps the highlight of the current paper and represents an alternate way of doing optimization. Some data mining applications of the new procedures are looked at in Section 8. A discussion follows in Section 9. Not all possibilities are dealt with: the focus is on what would be considered the easiest spaces that would most likely be used in practice.

2 Manifolds

For us a manifold \mathcal{M} can be thought of as a subset of \mathbb{R}^n that looks locally like \mathbb{R}^m , $n > m$. To formalize this, let \mathcal{I} be an index set and for each $i \in \mathcal{I}$ let $\mathcal{V}_i \subset \mathcal{M}$ be an open subset where $\bigcup_{i \in \mathcal{I}} \mathcal{V}_i = \mathcal{M}$. Additionally, let there be homeomorphisms $\phi_i : \mathcal{V}_i \rightarrow \mathbb{R}^m$. Each pair (\mathcal{V}_i, ϕ_i) is called a *chart* and the set $\{(\mathcal{V}_i, \phi_i)\}_{i \in \mathcal{I}}$ is called an *atlas*. We now have local coordinates etched onto \mathcal{M} .

When two open subsets \mathcal{V}_i and \mathcal{V}_j of \mathcal{M} overlap there will be a homeomorphic transition function $\tau_{ij} : \mathbb{R}^m \rightarrow \mathbb{R}^m$ on the overlap $\mathcal{V}_i \cap \mathcal{V}_j$ defined by $\tau_{ij} = \phi_j \circ \phi_i^{-1}$. The transition functions allow us to stitch together the charts into a coherent whole. This then defines a *topological manifold*.

Extra structure can be placed on the manifold by placing additional requirements on \mathcal{M} . A typical requirement is that the transition functions are differentiable. In this case we have a *differentiable manifold*. If the τ_{ij} are all \mathcal{C}^k , $k \geq 1$, then \mathcal{M} is said to be a \mathcal{C}^k -manifold. If \mathcal{M} is a \mathcal{C}^∞ -manifold then it is called *smooth*. Every \mathcal{C}^k -manifold can be made smooth [18].

At every point $p \in \mathcal{M} \subset \mathbb{R}^n$ a tangent space $\mathcal{T}_p\mathcal{M}$ can be attached. If an inner product $h : \mathcal{T}_p\mathcal{M} \times \mathcal{T}_p\mathcal{M} \rightarrow \mathbb{R}$ is defined for every $p \in \mathcal{M}$ and varies smoothly over \mathcal{M} then we have a *Riemannian manifold*. The function $h(\cdot, \cdot)$ is called a *metric* and with it notions of lengths and angles in $\mathcal{T}_p\mathcal{M}$ are available over \mathcal{M} . We will need surprisingly little direct use of any Riemannian structure.

Manifolds often enter an optimization problem through equality constraints.

Definition 2.1 (Regular level sets) Let $\mathbf{g} : \mathbb{R}^{n+m} \rightarrow \mathbb{R}^m$. For the set $\mathcal{M} = \{\mathbf{x} : \mathbf{g}(\mathbf{x}) = \mathbf{c}\}$ assume $\nabla \mathbf{g}(\mathbf{x})$ is full rank. Then \mathcal{M} is a *regular level set* of $\mathbf{g}(\mathbf{x})$. Additionally, the null space of $\nabla \mathbf{g}(\mathbf{x})$ coincides with the tangent space to \mathcal{M} .

If the function defining the equality constraints is sufficiently differentiable then by Sard's theorem almost every level set is a (Riemannian) manifold.

Theorem 2.1 (Regular level set manifolds and Sard's theorem [19])

Let $\mathbf{g} : \mathbb{R}^{n+m} \rightarrow \mathbb{R}^m$ be a \mathcal{C}^2 function. Then every regular level set of $\mathbf{g}(\mathbf{x})$ is an n -dimensional manifold. Further, if $\mathbf{g}(\mathbf{x})$ is a \mathcal{C}^{n+1} function then almost every level set is a regular level set. A regular level set acquires a metric from its embedding in \mathbb{R}^{n+m} by restricting the standard Euclidean metric to the tangent space $\mathcal{T}_p\mathcal{M}$ given by $\text{null}(\nabla\mathbf{g}(\mathbf{p}))$ for all $p \in \mathcal{M}$.

On a Riemannian manifold \mathcal{M} a *geodesic* is the shortest path between two (nearby) points on \mathcal{M} and generalizes the notion of a straight line. The exact way this is calculated for a Riemannian manifold is not important for us. What is important though is when any arbitrary point $a \in \mathcal{M}$ can be reached from every other point $b \in \mathcal{M}$ by a geodesic. In such a case the manifold \mathcal{M} is said to be *geodesically complete*. This has many nice properties which are covered in more detail in [20]. Not all of the manifolds considered herein will be geodesically complete. When they are not considerably more effort may be required to implement a practical optimization routine.

3 Lie Groups and Lie Algebras

A *Lie group* G is a smooth manifold with the additional structure of being a group. For $g, h \in G$ the function $\phi(g, h) = g^{-1}h$ must additionally be smooth. A *matrix Lie group* is defined as any subgroup of the group of general linear transformations $\text{GL}(n)$ of \mathbb{R}^n , the Lie group of invertible n -by- n matrices. We will always work with matrix Lie groups and simply say Lie group instead of matrix Lie group.

A major advantage of having the group structure is the ability to easily move over the manifold. Every element $g \in G$ defines a diffeomorphism of G by the operation of *left translation* defined by $L_g(h) = g \odot h$ for every $h \in G$. Left translation has moved the element h to the new point $g \odot h$ on the manifold. Often $g \odot h$ is simply matrix multiplication but this is not always the case. For instance symmetric positive definite matrices would have $g \odot h = g^{1/2}hg^{1/2}$. We will drop the \odot with the understanding that the correct multiplication is used when it differs from the usual matrix multiplication.

A *path* in G is given by $L_{h_1} \circ L_{h_2} \circ \dots \circ L_{h_n}(e) = h_1 \cdots h_n = L_{h_1 \cdots h_n}(e)$. The *path-connected identity component* of a Lie group G is the subgroup of those elements of the full group that can be connected to the identity element by a path. We always assume that this is the component that is being used unless stated otherwise.

Since a Lie group is a manifold it will have a tangent space $\mathcal{T}_e\mathcal{M}$ at the identity. $\mathcal{T}_e\mathcal{M}$ is called the *Lie algebra* \mathfrak{g} of the Lie group G . A Lie algebra is a vector space and is a much easier object to work with in a typical optimization procedure. For example, we can lay out a mesh for MADS in a Lie algebra. For a matrix group the elements of the Lie algebra are also given by matrices.

The *Lie bracket* on the Lie algebra is defined as $[\omega_1, \omega_2] = \omega_1\omega_2 - \omega_2\omega_1$ for $\omega_1, \omega_2 \in \mathfrak{g}$.

There is a mapping $\exp : \mathfrak{g} \rightarrow G$ called the *exponential mapping*. For matrix Lie groups this corresponds to the matrix exponential. The following is a key fact for Lie groups:

Theorem 3.1 (Path from the identity [13]) *Every element of a path-connected matrix Lie group is given by a finite product of exponents of elements in the corresponding Lie algebra. That is, if $g \in G$ then $g = \exp[\omega_1] \cdots \exp[\omega_n]$ for some $\omega_i \in \mathfrak{g}$.*

The exponential map $\exp[t\omega]$, $t \in \mathbb{R}$ and $\omega \in \mathfrak{g}$, is a geodesic that sometimes, but not always, corresponds to the Riemannian geodesic. When it doesn't correspond we will use the Lie group exponential. An important special case is when the Lie group G is compact (as a manifold). Then the exponential map is also a Riemannian geodesic and the manifold is geodesically complete.

Given a Lie group G and a subgroup $H \subset G$ we can form the *coset* G/H where every element in G/H is an equivalency class $gH = \{gh : h \in H\}$ for every $g \in G$. We identify $g_1, g_2 \in G$ as the same element if $g_1H = g_2H$ as sets. The space G/H is also a smooth manifold, called the *coset manifold*.

4 Examples

It is important at this point to lay out what is needed to effectively perform an optimization procedure. What is it about \mathbb{R}^n that needs to be retained and what can be thrown away? There are three vital things that need to be done:

1. move in a specified direction;
2. a specified distance, and;
3. evaluate a function at every point.

If we can do these we can optimize. If we can do these efficiently we can optimize efficiently.

For Riemannian geometries all of the three requirements are satisfied. We can move a specified length along a specified direction on the manifold. Having functions defined on a Riemannian manifold also causes no difficulties. The difference from \mathbb{R}^n is that the path we travel along now, our 'direction', is not necessarily a straight line. Instead it's a geodesic, the straightest possible line we could hope for and still remain on the manifold. But this is the difficulty: numerically traveling along the geodesic without straying is hard.

There are some cases of a Riemannian manifold where it is easy to stay on the geodesic. \mathbb{R}^n is the canonical example where the geodesics are simply straight lines and the Euclidean metric gives us our sense of distance. However, it is possible for other Riemannian manifolds to have closed-form solutions for their geodesics and this makes optimizing over them a relatively straightforward procedure.

If the surface we are optimizing over has a simple global structure then we can hope that performing the optimization will be easy. A particular example

of this is when we know the Lie group of symmetries of the manifold. In this case the group elements specify how to move over the manifold. Below we examine this through a few examples of progressive complexity. There will be results and definitions that, while expressed in a specific example, are generic for the spaces and Lie groups we'll consider.

4.1 The Special Euclidean Group $\text{SE}(3)$

Our first case will be \mathbb{R}^3 . Here the orientation preserving symmetry group is given by the special Euclidean group $\text{SE}(3)$ that includes all of the rotations and translations of \mathbb{R}^3 . That is, the space looks the same under any transformation in $\text{SE}(3)$. The typical way we would represent these transformations is by augmenting a vector $\mathbf{x} \in \mathbb{R}^3$ with a final 1 to have $\hat{\mathbf{x}}^T = [\mathbf{x}^T \ 1]$. A matrix representation M of a group element $g \in \text{SE}(3)$ is then given by

$$M = \begin{bmatrix} O & \mathbf{v} \\ \mathbf{0}^T & 1 \end{bmatrix}, \quad (1)$$

where $O \in \text{SO}(3)$ is a 3-by-3 orthogonal matrix with $\det(O) = 1$ and $\mathbf{0}, \mathbf{v} \in \mathbb{R}^3$. Then an explicit representation of the action of $g \in \text{SE}(3)$ on $\mathbf{x} \in \mathbb{R}^3$ is given by $M\hat{\mathbf{x}}$.

With $\text{SE}(3)$ we can reach any point $\mathbf{y} \in \mathbb{R}^3$ from any other point $\mathbf{x} \in \mathbb{R}^3$. However this is not unique since we can always first perform a rotation around \mathbf{x} and then translate from \mathbf{x} to \mathbf{y} . By 'rotation around \mathbf{x} ' we mean translating \mathbf{x} to the origin, performing a rotation and then translating back to \mathbf{x} . There's no reason to consider these (infinite number of) different ways of going from \mathbf{x} to \mathbf{y} as being different. Instead we form equivalency classes where $M_1 \sim M_2$ when

$$M_1 = M_2 \begin{bmatrix} O & \mathbf{0} \\ \mathbf{0}^T & 1 \end{bmatrix}, \quad (2)$$

for $O \in \text{SO}(3)$. Notice that the matrix that multiplies M_2 on the right in (2) is also a representation of $O \in \text{SO}(3)$ since the 'extra' 1 is irrelevant. It is also a member of $\text{SE}(3)$ where there is no translation. We see that $\text{SO}(3) \subset \text{SE}(3)$ and additionally that $\text{SE}(3) = \text{SO}(3) \times \text{T}(3)$, the semi-direct product of rotations and translations in \mathbb{R}^3 . From this we can form the coset manifold $\text{SE}(3)/\text{SO}(3) = \text{T}(3) \simeq \mathbb{R}^3$, getting back a copy of what we started with. The optimization proceeds in the coset manifold.

A *Klein geometry* (G, H) consists of a Lie group G , a subgroup $H \subset G$ and a space (manifold) $\mathcal{M} = G/H$ over which G acts *transitively*: for every $p, q \in \mathcal{M}$ there is a $g \in G$ such that $gp = q$. Alternatively, a Klein geometry can be thought of as a space and a law for moving over this space. For the Klein geometry $(\text{SE}(3), \text{SO}(3))$ the following conditions for the Lie algebras $\mathfrak{se}(3) = \mathfrak{so}(3) \oplus \mathfrak{t}(3)$ hold:

$$[\mathfrak{so}(3), \mathfrak{so}(3)] \subseteq \mathfrak{so}(3), [\mathfrak{so}(3), \mathfrak{t}(3)] \subseteq \mathfrak{t}(3) \text{ and } [\mathfrak{t}(3), \mathfrak{t}(3)] \subseteq \mathfrak{so}(3), \quad (3)$$

where $[\cdot, \cdot]$ is the Lie bracket. The first condition is that $\mathfrak{so}(3)$ is itself a Lie algebra. The second condition tells us that the $\mathfrak{t}(3)$ is invariant under the actions of $\mathfrak{so}(3)$. The final condition states that the Lie bracket of two infinitesimal translations is a rotation. As these conditions are sequentially imposed we move from a *homogeneous space* to a *reductive homogeneous space* to a *symmetric space*.

We are still missing a sense of moving along a straight line in \mathbb{R}^3 by a specified amount. One thing that makes reductive homogeneous spaces convenient to work with is that if the Klein geometry is given by (G, H) , and $\mathfrak{g} = \mathfrak{h} \oplus \mathfrak{m}$ with $[\mathfrak{h}, \mathfrak{h}] \subseteq \mathfrak{h}$ and $[\mathfrak{h}, \mathfrak{m}] \subseteq \mathfrak{m}$, then we can ignore the action of \mathfrak{h} and use only \mathfrak{m} in our optimization. What this means is we can use the same exponential map defined on \mathfrak{g} and simply restrict it to \mathfrak{m} to find the exponential map to G/H . It is the reductive condition that allows this.

In the case of $\mathfrak{t}(3)$ the matrices in the algebra are of the form

$$\tau = \begin{bmatrix} 0_3 & \mathbf{v} \\ \mathbf{0}^T & 0 \end{bmatrix}, \quad (4)$$

and $\exp[\tau] = I_4 + \tau$. This explicitly shows that $T(3) \simeq \mathbb{R}^3$ through the identification $\tau \leftrightarrow \mathbf{v}$. We can travel along a straight line from the point $T_0 \in T(3)$ using $T_0 \exp[t\tau]$, $t \in \mathbb{R}$. As t increases we are traveling further along this geodesic.

A symmetric space is always geodesically complete, which does not necessarily hold for the other homogeneous spaces [21]. Optimizing over a geodesically complete space is easier since this allows one to use a global *pullback procedure* from the manifold to the tangent space as covered in [20]. This means we can pullback the objective function $f : \mathcal{M} \rightarrow \mathbb{R}$, and any inequality constraints, from the manifold \mathcal{M} to the Lie algebra \mathfrak{m} using the function composition $f \circ \exp : \mathfrak{m} \rightarrow \mathbb{R}$. Since the Lie algebra is a normed vector space and the exponential map is smooth, any convergence results about a direct search method in \mathbb{R}^n automatically hold without modification in the Lie algebra. In this case optimizing over the underlying space becomes a standard optimization in Euclidean space apart from a more complicated exponential function.

4.2 The Grassmannian $\mathbb{G}(4, 2)$

Our second example will still be a symmetric space but more abstract. Here we will look at the Grassmannian $\mathbb{G}(4, 2)$ of all two-dimensional planes in \mathbb{R}^4 . We start with the Lie group $\mathfrak{O}(4)$, which is the symmetry group of the 3-sphere $\mathcal{S}^3 \subset \mathbb{R}^4$. A fixed plane in \mathbb{R}^4 can be specified by the subspace normal to it. The subtlety for $\mathbb{G}(4, 2)$ is that we don't want to differentiate between any orthonormal bases that span the plane or the subspace normal to it. Then $\mathbb{G}(4, 2) \simeq \mathfrak{O}(4)/(\mathfrak{O}(2) \times \mathfrak{O}(2))$. The first $\mathfrak{O}(2)$ ignores any changes in the orthonormal basis for the subspace normal to the plane; the second ignores any changes in the orthonormal basis for the plane.

For $\mathbf{G}(4, 2)$, $\mathfrak{o}(4) = \mathfrak{h} \oplus \mathfrak{m}$ where

$$\begin{aligned}\omega &= \eta + \mu \\ &= \begin{bmatrix} A & 0 \\ 0 & B \end{bmatrix} + \begin{bmatrix} 0 & -C \\ C^T & 0 \end{bmatrix},\end{aligned}\tag{5}$$

with A and B skew-symmetric, $\eta \in \mathfrak{h}$ and $\mu \in \mathfrak{m}$. Since

$$[\mathfrak{h}, \mathfrak{h}] \subseteq \mathfrak{h}, \quad [\mathfrak{h}, \mathfrak{m}] \subseteq \mathfrak{m} \quad \text{and} \quad [\mathfrak{m}, \mathfrak{m}] \subseteq \mathfrak{h},\tag{6}$$

$\mathbf{G}(4, 2)$ is a symmetric space and hence geodesically complete. Then the map $\exp[\mu]$, $\mu \in \mathfrak{m}$, allows us to reach any point in $\mathbf{G}(4, 2)$ from any other point. As for \mathbb{R}^3 , we can travel along a geodesic from a point $G_0 \in \mathbf{G}(4, 2)$ using $G_0 \exp[t\mu]$, $t \in \mathbb{R}$.

Here the exponential mapping from the Lie algebra to the manifold is significantly more complex than that for the \mathbb{R}^3 case [14]. Table 1 gives the solution for $\mathbf{G}(n, k)$ along with some other Lie groups and reductive homogeneous spaces. Finding a way to efficiently calculate the exponential map for a reductive homogeneous space is the key to being able to efficiently optimize on it. For a Klein geometry (G, H) it is enough to know the exponential for the Lie group G if the space G/H is a reductive homogeneous space. For us this exponential map is always the matrix exponential.

4.3 The Stiefel Manifold $\mathbf{V}(4, 2)$

Our next example is the Stiefel manifold $\mathbf{V}(4, 2)$ of two-dimensional planes in \mathbb{R}^4 where the orthonormal bases used to span a two-dimensional plane are distinguished. The orthonormal bases that span the normal space to a plane are not important however. So $\mathbf{V}(4, 2) \simeq \mathfrak{O}(4)/\mathfrak{O}(2)$. Now $\mathfrak{o}(4) = \mathfrak{h} \oplus \mathfrak{m}$ is given by

$$\begin{aligned}\omega &= \eta + \mu \\ &= \begin{bmatrix} 0 & 0 \\ 0 & B \end{bmatrix} + \begin{bmatrix} A & -C \\ C^T & 0 \end{bmatrix},\end{aligned}\tag{7}$$

where A and B are skew-symmetric, $\eta \in \mathfrak{h}$ and $\mu \in \mathfrak{m}$.

This is our first example of a reductive homogeneous space that is not a symmetric space:

$$[\mathfrak{h}, \mathfrak{h}] \subseteq \mathfrak{h} \quad \text{and} \quad [\mathfrak{h}, \mathfrak{m}] \subseteq \mathfrak{m} \quad \text{but} \quad [\mathfrak{m}, \mathfrak{m}] \not\subseteq \mathfrak{h}.\tag{8}$$

So we can't say that $\mathbf{V}(4, 2)$ is geodesically complete by appealing to it being a symmetric space. But it is compact and can be made Riemannian so the Hopf-Rinow theorem tells us it is geodesically complete [19]. The same as for \mathbb{R}^3 and $\mathbf{G}(4, 2)$, a geodesic from a point $V_0 \in \mathbf{V}(4, 2)$ is given by $V_0 \exp[t\mu]$, $t \in \mathbb{R}$, which is both a Riemannian geodesic and a matrix exponential map. Table 1 has the closed-form solution for the geodesics.

4.4 A Fixed Subspace

This is a simple example of how equality constraints can be restated as a Klein geometry. (See [12] also for an implementation of this sort of procedure.) Suppose we have a linear constraint $A^T \mathbf{x} = \mathbf{0}$ where $A^T : \mathbb{R}^{n+m} \rightarrow \mathbb{R}^m$. The subspace that $\mathbf{x} \in \mathbb{R}^{n+m}$ must lie in is given by $\text{null}(A^T)$ and A gives the normal directions to the manifold. Let $Q \in \mathbb{R}^{(n+m) \times n}$ be a full rank matrix such that $A^T Q = 0$ so that Q spans the null space. The linear subspace represented by $\text{span}(Q)$ defines the manifold \mathcal{Q} .

Assume now that we additionally have $A^T A = I_m$ and $Q^T Q = I_n$, which can always be done. The symmetry group associated with \mathcal{Q} is the subgroup of $\text{SE}(n+m)$ given by matrices of the form

$$M = \begin{bmatrix} O & Q\mathbf{v} \\ \mathbf{0}^T & 1 \end{bmatrix} \text{ and } O = U \begin{bmatrix} W & 0 \\ 0 & V \end{bmatrix} U^T \quad (9)$$

where $U = [A \ Q] \in \text{SO}(n+m)$, $W \in \text{SO}(m)$, $V \in \text{SO}(n)$ and $\mathbf{v} \in \mathbb{R}^n$. It's easy to check that $M : \mathcal{Q} \rightarrow \mathcal{Q}$ and that all such M form a Lie group G .

Now we have a space (manifold) given by \mathcal{Q} and a law for moving over \mathcal{Q} given by G . This defines a Klein geometry. The last piece we would like is the subgroup $H \subset G$ that will *stabilize* $\mathbf{0} \in \mathcal{Q}$: for all $h \in H$ we must have $h\mathbf{0} = \mathbf{0}$. Then our Klein geometry can be represented as (G, H) . We form equivalence classes the same as in (2) with the O as in (9). Let $H \subset G$ be the subgroup of all matrices of the form

$$N = \begin{bmatrix} O & \mathbf{0} \\ \mathbf{0}^T & 1 \end{bmatrix}. \quad (10)$$

Then H stabilizes $\mathbf{0} \in \mathcal{Q}$ and we have our Klein geometry as (G, H) . Our manifold is now given by $\mathcal{Q} = G/H$. The tangent vectors $q \in \mathcal{T}_{\mathbf{0}}\mathcal{Q}$ have the form

$$q = \begin{bmatrix} 0_{n+m} & Q\mathbf{v} \\ \mathbf{0}^T & 0 \end{bmatrix}, \quad (11)$$

and $\exp[q] = I + q$. Any point $q \in \mathcal{Q}$ can be reached from $\mathbf{0}$ by a path of the form $\exp[q_1] \cdots \exp[q_n]\mathbf{0}$. This is how a path in the subspace would be constructed during an optimization procedure.

The example is illustrative but not practical. No one in their right mind would do this particular optimization problem as above. But the methodology presented is generic for Klein geometries associated with a reductive homogeneous space. As such it does become a practical method as the Klein geometry becomes more complicated.

5 A Second Look at Lie Groups and Lie Algebras

The Klein geometries considered above can all be represented as matrices. Since they're reductive homogeneous spaces we can remain within the space using the matrix exponential of $\exp : \mathfrak{m} \rightarrow G/H$ by restricting the exponential

mapping of \mathfrak{g} to the subspace \mathfrak{m} . For a fixed $\omega \in \mathfrak{m}$ moving in a specific direction is given by $\exp[t\omega]$, where $t \in \mathbb{R}$. It's easier to work with the Lie algebra since it's a vector space. For example, one can place a mesh in \mathfrak{g} and perform MADS there.

Theorem 3.1 does not tell us that the exponential map is surjective but we do have that:

Theorem 5.1 (Exponential map injectivity radius [13]) *The exponential map $\exp : \mathfrak{g} \rightarrow G$ is a diffeomorphism in a neighborhood of $0 \in \mathfrak{g}$ given by $\|\omega\| < \rho$ for some $\rho > 0$ where $\omega \in \mathfrak{g}$. The ρ is called the injectivity radius of the exponential map. Similarly, the matrix logarithm map $\log : G \rightarrow \mathfrak{g}$ is a diffeomorphism in a neighborhood of $e \in G$ given by $\|g - e\| < \epsilon$ for some $\epsilon > 0$ where $g \in G$.*

The ω_i in Theorem 3.1 would all lie within the injectivity radius of the exponential map. Determining the injectivity radius of the exponential map is important since it determines an upper bound on allowable step sizes in any optimization routine. This is not always an easy task. One may be able to appeal to the group being compact or the space being symmetric to show geodesic completeness in which case Theorem 3.1 holds with a single $\exp[\omega]$. Alternatively, one can develop ‘tricks’ on a case by case basis.

To illustrate the latter consider the special linear group $\mathbf{SL}(n)$ where $\exp : \mathfrak{sl}(n) \rightarrow \mathbf{SL}(n)$ is not surjective. A polar decomposition of a matrix $A \in \mathbf{SL}(n)$ shows that $A = QS$ where $Q \in \mathbf{SO}(n)$, $S \in \mathbf{SPD}(n)$ with $\det(S) = 1$, and both Q and S are unique. From Table 1 we see that every $\omega \in \mathfrak{sl}(n)$ can be represented as $\alpha + \sigma$ where $\alpha \in \mathfrak{so}(n)$, $\sigma \in \mathfrak{spd}(n)$ with $\text{Tr}(\sigma) = 0$, and α and σ are also in $\mathfrak{sl}(n)$.

Let \mathfrak{m} be the traceless members of $\mathfrak{spd}(n)$ and $\mathfrak{h} \equiv \mathfrak{so}(n)$. Then $\mathfrak{sl}(n) = \mathfrak{h} \oplus \mathfrak{m}$ and $\mathbf{SL}(n)/\mathbf{SO}(n)$ is a symmetric space. To pick a general member of $\mathbf{SL}(n)$ first select a $\mu \in \mathfrak{m}$ and map this to the base space by $S = \exp[\mu]$. This is a surjective map (since \mathcal{M} is a symmetric space and hence geodesically complete), so we can reach any $S \in \mathbf{SPD}(n)$, $\det(S) = 1$, from the identity element. At S an element $\eta \in \mathfrak{h}$ is mapped into $\mathbf{SO}(n)$ by $Q = \exp[\eta]$. This is also a surjective map. By forming the product QS we see that every member of $\mathbf{SL}(n)$ can be reached by this procedure.

Unfortunately this method can not be employed in general. The reason it works so well here is that when we decompose a matrix $A \in \mathbf{SL}(n)$ into the polar form $A = QS$ both Q and S are also in $\mathbf{SL}(n)$. So we are always working in $\mathbf{SL}(n)$ and do not have to worry about whether the product QS lies in $\mathbf{SL}(n)$ or not.

6 Algorithm and Convergence

The most general optimization problem considered here is stated as

$$\min_{p \in (G, H)} f(p) \text{ subject to } \mathbf{g}(p) \leq \mathbf{0} \quad (12)$$

with (G, H) being a Klein geometry and $\mathfrak{g} = \mathfrak{h} \oplus \mathfrak{m}$. Any inequality constraints are handled purely within the DSM and will not concern us here. They are only allowed if the DSM can handle inequality constraints. In this section we will initially present a procedure that works with a generic DSM. This is then specialized to the probabilistic descent method [22] in Section 7. Section 7 also develops a new method of working directly on manifolds. This new methodology relies on certain properties of the probabilistic descent method.

In the simplest case of a surjective exponential map the procedure is done by laying out the mesh in the Lie subalgebra $\mathfrak{m} \subset \mathfrak{g}$. Since the Lie algebra is a vector space this is a trivial mapping. The exponential map takes the chosen mesh points and maps them into the reductive homogeneous space $\mathcal{M} \subset G$. These points are then fed back into the DSM. (See Figure 1.) Convergence results are the same as whatever are available with the chosen DSM since we have done a global pullback of the objective function (and any inequality constraints) from \mathcal{M} to \mathfrak{m} [20]. In this case, ignoring the inequality constraints, (12) becomes

$$\min_{\omega \in \mathfrak{m}} f \circ \exp(\omega).$$




Fig. 1: Algorithm with a surjective exponential map.

With the surjective exponential map the optimization occurs in the fixed tangent (vector) space \mathfrak{m} . This is why convergence results from a DSM carry over without modification. When the exponential map is not surjective more care must be taken. What we desire is to work in a fixed tangent space after a finite number of steps in the optimization. If this occurs then convergence results are again the same as the DSM employed. This situation is shown in Figure 2. The only difference is that we may move over \mathcal{M} , g_{k-1} and g_k may differ, after picking out new elements from \mathfrak{m} and mapping them to \mathcal{M} . If $g_{k-1} \neq g_k$ then the tangent spaces differ and we essentially start a new optimization with a new mesh.

Recall that ρ is the injectivity radius of the exponential map. For non-surjective exponential maps a maximum step size $0 < s_{max} \leq \rho$ is used. This guarantees that every point on \mathcal{M} in a neighborhood of g_k can be reached from g_k . In order to have convergence results we also use a second radius $0 < s_{fix} < s_{max}$ such that we only move to a new point $\exp[\omega] = g_k \in \mathcal{M}$




Fig. 2: Algorithm without a surjective exponential map.

when the point $\omega \in \mathfrak{m}$ lies in the annulus defined by s_{fix} and s_{max} . (See Figure 3.) Otherwise we set $g_k = g_{k-1}$ and work in the same tangent space.

Suppose g_k becomes fixed after a finite number of steps in the optimization. Since we then work in a single vector space where the exponential map is diffeomorphic all the convergence results carry over from the DSM. We are now doing a pullback procedure around g_k but only locally on \mathcal{M} . If g_k never becomes fixed the optimization does not converge. This is reasonable since for every g_k after a finite number of steps we will move to a new point on \mathcal{M} that is outside a local neighborhood of g_k determined by s_{fix} .

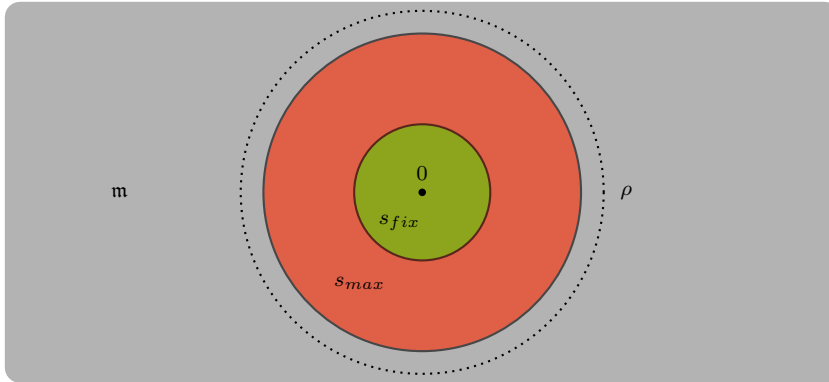


Fig. 3: The area in the vector space \mathfrak{m} where the optimization procedure is performed. This defines the update procedure in Figure 2. If $0 \leq \|\omega\| < s_{fix}$, $\omega \in \mathfrak{m}$, then the tangent space doesn't move. If $s_{fix} \leq \|\omega\| < s_{max} \leq \rho$, ρ the injectivity radius, then the point $g \in \mathcal{M}$ is updated.

A brief note on how we find the tangent space at each point on \mathcal{M} . Every tangent space looks the same as \mathfrak{m} attached at $e \in \mathcal{M}$. If we are at a given point $g \in \mathcal{M}$ we can work with $\mathcal{T}_g\mathcal{M}$ by performing the automorphism $g : \mathcal{M} \rightarrow \mathcal{M}$.

That is, we first pick a point $\omega \in \mathfrak{m}$, map this to a point $\exp[\omega] \in \mathcal{M}$ and then move that point by $L_g(\exp[\omega]) : \mathfrak{m} \rightarrow \mathcal{M}$. This is equivalent to $\exp : \mathcal{T}_g\mathcal{M} \rightarrow \mathcal{M}$ and is computationally convenient.

The final optimization routine is given by Procedure 1. The hardest part of the procedure can be determining what the injectivity radius is if the exponential map is not surjective. See the probabilistic descent DSM [22] in Procedure 2 for a concrete example.

Procedure 1 (Direct searches on reductive homogeneous spaces) *Specify:*

- A direct search method and any required inputs to that method.
- A Klein geometry (G, H) with a reductive homogeneous space $\mathcal{M} = G/H$ and $\mathfrak{g} = \mathfrak{h} \oplus \mathfrak{m}$.
- A mapping of the mesh from the vector space \mathbb{R}^n to the vector space \mathfrak{m} .
- Two radii satisfying $0 < s_{fix} < s_{max} \leq \rho$ if not working with a surjective exponential map. Otherwise set $s_{fix} = s_{max} = \infty$.
- An initial set μ_0^j of mesh points in the Lie algebra \mathfrak{m} where $j = 1, \dots, m$ and m is the number of points required by the DSM.

Let the initial point on the reductive homogeneous space be the identity $g_0 = e$. Finally let $k = 0$.

Iterate Set $k \leftarrow k + 1$.

Move Find $g_k^j = L_{g_{k-1}}(\exp[\mu_{k-1}^j])$ for every j .

DSM Perform the DSM with the g_k^j .

Return If the DSM has converged return the optimal point $\hat{g} \in G/H$ and exit.

Update If the current best μ_k^* has $\|\mu_k^*\| > s_{fix}$ update the current point $g_k \in G/H$. Here $g_k = L_{g_{k-1}}(\exp[\mu_k^*])$. Otherwise set $g_k = g_{k-1}$.

Mesh In \mathfrak{m} lay out the new points μ_k^j returned by the DSM. This is either done in an existing mesh if $g_k = g_{k-1}$ or create a new mesh otherwise.

Goto Goto the **Iterate** step.

7 Working only with the Group

The Lie algebra occupies a central position in the optimization procedure presented in Procedure 1. In this way there is no major difference between what was presented here and the method developed in [20]. What is different is the ease that the manifold can be moved over because the symmetry group of the underlying space is known. As we stated previously, the reductive condition on a Klein geometry's space G/H is what allows this.

But if the group of symmetries is known and the elements provide a way of moving over the space we may ask: Why do we need the Lie algebra? The main reason for the use of a Lie algebra is that as a vector space it is easy to lay out a mesh or pick out a random element. If this convenience can be pushed down onto the group itself then the algebra can be dispensed with. That creates an alternate way of optimizing in reductive homogeneous spaces.

Let's examine what this means in \mathbb{R}^n . Euclidean space is different since it doesn't matter if one looks at the tangent space or the manifold since they

both are globally identical. In subsection 4.1 we saw that a point in the Lie algebra and a point on the Lie group can be represented by, respectively,

$$\tau = \begin{bmatrix} 0_n & \mathbf{v} \\ \mathbf{0}^T & 0 \end{bmatrix} \text{ and } \exp[\tau] = \begin{bmatrix} I_n & \mathbf{v} \\ \mathbf{0}^T & 1 \end{bmatrix}. \quad (13)$$

For moving in the group, if we work with the Lie algebra addition is used (followed by exponentiation) while the Lie group uses multiplication. This gives

$$\tau_1 + \tau_2 = \begin{bmatrix} 0_n & \mathbf{v}_1 + \mathbf{v}_2 \\ \mathbf{0}^T & 0 \end{bmatrix} \text{ and } \exp[\tau_1 + \tau_2] = \begin{bmatrix} I_n & \mathbf{v}_1 + \mathbf{v}_2 \\ \mathbf{0}^T & 1 \end{bmatrix}. \quad (14)$$

(Of course normally we dispense with all of this and just use $\mathbf{v}_1 + \mathbf{v}_2$ not caring if it is being done in the Lie algebra or the Lie group.) The reason the tangent space and the manifold are interchangeable here is that $\exp[\tau_1 + \tau_2] = \exp[\tau_1] \exp[\tau_2]$ which is due to the fact that $[\tau_i, \tau_j] = 0$ for any τ_i and τ_j . A Lie algebra where the Lie bracket is trivial is called *abelian* and has an associated commutative (abelian) connected Lie group. What we wish to do here is to work strictly with the $g = \exp[\tau]$ elements in the Lie group and dispense with the elements τ in the Lie algebra. For abelian Lie groups (algebras) such as \mathbb{R}^n and the n -dimensional tori this is trivial. In fact, \mathbb{R} and the 1-dimensional torus (the circle) effectively exhaust the abelian Lie groups since any such group is the direct product of a finite number of copies of \mathbb{R} and the 1-dimensional torus.

For abelian Lie groups it's the case that $\exp[\omega_1 + \omega_2] = \exp[\omega_1] \exp[\omega_2]$ and so summation in the Lie algebra directly translates to multiplication in the Lie group. It follows that picking some point on a mesh in the Lie algebra, which is a sum, easily translates to the Lie group. This can be extended in a straightforward way whenever sums in the Lie algebra can be easily translated to products in the Lie group regardless if the group is abelian or not. The difficulties arise when this isn't possible.

A probabilistic descent DSM [22] seems to be particularly well suited to the methodology we have in mind. We present an implementation for a reductive homogeneous space in Procedure 2. This relies on the Lie algebra and the exponential map.

In Procedure 2 it's important to notice that the work is being carried out around the $g'_k \in \mathcal{M}$ at each step k even if the current tangent space is being fixed at $g_k \in \mathcal{M}$. We fix the working tangent space in order to trivially carry over convergence results. It is only moved when there's a danger of leaving the area where the exponential map is a diffeomorphism. However, if the procedure is carried out directly on the reductive homogeneous space without reference to \mathfrak{m} we can freely move over \mathcal{M} knowing that at some point the tangent space could be fixed at some point if there is convergence to a solution. The only restriction is that our move is never larger than the injectivity radius of the logarithmic map $\log : \mathcal{M} \rightarrow \mathfrak{m}$. Then we have the same convergence results as if we worked with \mathfrak{m} .

Procedure 2 (Probabilistic descent using the Lie algebra) *Specify:*

- A Klein geometry (G, H) with a reductive homogeneous space $\mathcal{M} = G/H$ where $\mathfrak{g} = \mathfrak{h} \oplus \mathfrak{m}$.
- An objective function $f : \mathcal{M} \rightarrow \mathbb{R}$.
- Two radii satisfying $0 < s_{max} < \rho/2$, ρ the injectivity radius of the exponential map, and $0 < s_{fix} < s_{max}$.
- A forcing function $\varrho : \mathcal{M} \rightarrow \mathbb{R}$ giving sufficient decrease.
- A maximum number of function evaluations.

Let the initial point on the reductive homogeneous space be g_0 . Set $k = 0$, $g'_0 = g_0$, $w_0 = 0 \in \mathfrak{m}$ and $w_{test} = 0 \in \mathfrak{m}$. Pick a $0 < s_0 < s_{max}$.

Choice Pick a random element $\omega_k \in \mathfrak{m}$ where $0 < \|\omega_k\| < s_k$. Set $h^+ = L_{g_k}(\exp[w_k + \omega_k])$ and $h^- = L_{g_k}(\exp[w_k - \omega_k])$.

Evaluate Check:

- IF $f(h^+) < f(g'_k) - \varrho(\|\omega_k\|)$ set $g'_{k+1} = h^+$ and $w_{test} = w_k + \omega_k$
- ELIF $f(h^-) < f(g'_k) - \varrho(\|\omega_k\|)$ set $g'_{k+1} = h^-$ and $w_{test} = w_k - \omega_k$
- ELSE set $g'_{k+1} = g'_k$ and $w_{test} = 0$

Return If the maximum number of function evaluations has been exceeded return g'_{k+1} and exit.

Step size Set $s_{k+1} = \min(s_{max}, 2s_k)$ if the IF or ELIF steps in **Evaluate** held. Otherwise set $s_{k+1} = s_k/2$.

Move If $\|w_{test}\| > s_{fix}$ set $g_{k+1} = g'_{k+1}$ and $w_{k+1} = 0$. Otherwise set $g_{k+1} = g_k$ and $w_{k+1} = w_{test}$.

Iterate Set $k \leftarrow k + 1$.

Goto Goto the **Choice** step.

Assume that we have a black-box matrix generator $\mathbf{B}_{\mathcal{M}}^r$ that can return a random $g \in \mathcal{M}$ satisfying $\|g - e\| < r$, $0 < r < R$, R a constant. Further, $\mathbf{B}_{\mathcal{M}}^r$ must be able to generate all such g that satisfy the norm condition. (Note that R can be the injectivity radius of the logarithmic map but this is not a strict requirement.) The probabilistic descent algorithm that works solely on \mathcal{M} is presented in Procedure 3.

We still need to specify the black boxes $\mathbf{B}_{\mathcal{M}}^r$ in Procedure 3. The obvious thing to use is the exponential map $\exp : \mathfrak{m} \rightarrow \mathcal{M}$ and pick a random $\omega \in \mathfrak{m}$. For that we can use the inequality

$$\|\exp[\omega] - e\| \leq \|\omega\|e^{\|\omega\|} \quad (15)$$

where $\omega \in \mathfrak{m}$ satisfies $\|\omega\| < \rho$. However, the exponential map is being used here only as a convenience and it is not crucial that this is the $\mathbf{B}_{\mathcal{M}}^r$ in Procedure 3. So the algorithm is still independent of \mathfrak{m} . (As a note, we can use the condition $\|\omega\| < \rho$ since \exp and \log are inverses of each other. So working only in the region of \mathfrak{m} where \exp is a diffeomorphism is the same as working in the region on \mathcal{M} where \log is a diffeomorphism.) Any other generator is suitable though. For example, the unitary, orthogonal and symplectic groups are examined in [23–25]. If we can work with orthogonal groups then the Euclidean group can also be worked with. Generating random symmetric positive definite, unipotent, and general and special linear group members are more straightforward prospects. The closed form solutions for the exponential maps of the Grassmannian, Stiefel manifold and unipotent group in

Procedure 3 (Probabilistic descent on reductive homogeneous spaces) *Specify:*

- A Klein geometry (G, H) with a reductive homogeneous space $\mathcal{M} = G/H$.
- A constant $R > 0$.
- A black-box matrix generator $\mathbf{B}_{\mathcal{M}}^r$ that returns a random $g \in \mathcal{M}$ satisfying $\|g - e\| < r < R$. All g satisfying the norm condition must be returnable by $\mathbf{B}_{\mathcal{M}}^r$.
- An objective function $f : \mathcal{M} \rightarrow \mathbb{R}$.
- A forcing function $\varrho : \mathcal{M} \rightarrow \mathbb{R}$ giving sufficient decrease.
- A maximum number of function evaluations.

Let the initial point on the reductive homogeneous space be the identity $g_0 = I$. Set $k = 0$ and pick a $0 < r_0 < \rho$.

Generate Generate a $M = \mathbf{B}_{\mathcal{M}}^{r_k}$ and set $h^+ = L_{g_k}(M)$ and $h^- = L_{g_k}(M^{-1})$.

Evaluate Check:

- IF $f(h^+) < f(g_k) - \varrho(\|M\|)$ set $g_{k+1} = h^+$
- ELIF $f(h^-) < f(g_k) - \varrho(\|M^{-1}\|)$ set $g_{k+1} = h^-$
- ELSE set $g_{k+1} = g_k$

Return If the maximum number of function evaluations has been exceeded return g_{k+1} and exit.

Step size Set $r_{k+1} = \min(R, 2r_k)$ if the IF or ELIF steps in **Evaluate** held. Otherwise set $r_{k+1} = r_k/2$.

Iterate Set $k \leftarrow k + 1$.

Goto Goto the **Generate** step.

Table 1 makes the exponential mapping route more computationally tractable for these manifolds. Of course the norm condition still needs to be taken into account.

One thing we need to consider is the distribution of the random matrices on the manifold \mathcal{M} . Looking at Figure 4 the pullback procedure is to take the function $f : \mathcal{M} \rightarrow \mathbb{R}$ and define a new function on a vector space \mathbb{V} by defining $f \circ \phi : \mathbb{V} \rightarrow \mathbb{R}$ where $\phi : \mathbb{V} \rightarrow \mathcal{M}$. As long as ϕ is sufficiently smooth this creates no difficulties and the optimization can proceed on \mathbb{V} with all of the convergence results of the employed DSM. For a DSM ‘sufficiently smooth’ will typically mean a \mathcal{C}^2 , \mathcal{C}^1 or Lipschitz function that is at least locally surjective. For the probabilistic descent method there are conditions on the distribution




Fig. 4: Pulling back the function $f : \mathcal{M} \rightarrow \mathbb{R}$ from the manifold \mathcal{M} to the vector space \mathbb{V} .

of the random vectors in the vector space \mathbb{V} . These are easily enforced, but while the optimization is carried out in \mathbb{V} the final solution actually lies on \mathcal{M} .

Given the freedom in choosing ϕ it's not obvious how placing restrictions on the vectors chosen in \mathbb{V} carry over to restrictions on the elements of \mathcal{M} . About all that can be stated is that the probability density function of choosing an element must not vanish in any direction from the identity element $e \in \mathcal{M}$.

8 Applications

The procedures developed herein have been applied to two data mining algorithms with positive results. The first data mining application is concerned with dimensionality reduction of manifold-valued data and involves optimizing over Grassmannians [26]. Also, the objective function is not differentiable and a derivative-free optimization method is the only appropriate procedure.

The second application is an extension of the semi-nonnegative matrix factorization and optimizes over hyper-spheres [27]. This additionally involves the use of optional search steps. Equality constraints on vector lengths make the restriction to a manifold implicit. These are naturally enforced by employing the exponential and logarithmic maps from the tangent space to the manifold and vice versa. That is, at the expense of a more complicated objective function the difficulties of working with equality constraints are removed. There is also an inequality constraint on the hyper-area defined by the columns of a matrix and it is convenient to treat this as black-box.

Both of the data mining applications have optimization problems that are naturally stated on manifolds. The exponential and logarithmic maps were used rather than finding points on the manifolds directly. Modifying the solution method for the modified semi-nonnegative matrix factorization in [27] to work solely on the hypersphere \mathcal{S}^{k-1} , $k \in \mathbb{N}$, is relatively straightforward. We show how this can be done presently.

Let $X \in \mathbb{R}^{n \times m}$ be the data matrix where there are m data points in \mathbb{R}^n , and choose a $k \in \mathbb{N}$. Typically $k \ll n$. We're looking for a decomposition $X \approx WH$ where $W \in \mathbb{R}^{n \times k}$ will have normalized columns, $\|\mathbf{w}_i\| = 1$ for $i = 1, \dots, k$, and $H \in \mathbb{R}^{k \times m}$ will have nonnegative entries, denoted by $0 \leq H$. The measure of model tightness is the maximum geodesic distance

$$S(W) = \max_{i < j} \arccos(\mathbf{w}_i \cdot \mathbf{w}_j)$$

over \mathcal{S}^{n-1} between any two columns of W . The optimization problem is

$$\min_{W, 0 \leq H} \|X - WH\|_F^2 \text{ subject to } S(W) \leq \epsilon \text{ and } \|\mathbf{w}_i\| = 1.$$

Take \hat{X} to be the data set X with columns normalized to unit length and $\bar{\mathbf{x}}$ be the Karcher mean of \hat{X} [28]. Find an initial $W_0 = [\mathbf{w}_{0,1} | \dots | \mathbf{w}_{0,k}]$ where $S(W_0) \leq \epsilon$. H_0 is found by solving the nonnegative least squares (NNLS) problem

$$H_0 = \arg \min_{0 \leq H} \|X - W_0 H\|_F^2 \text{ where } \epsilon_0 = \|X - W_0 H_0\|_F^2.$$

Set

- i_{max} as the number of iterations to run the algorithm and initialize $i = 0$,
- $\alpha_{max} = 1$ as the maximum step size,
- $\alpha_0 = \alpha_{max}$ as the initial step size,
- $\theta = 1/2$ as the step size decrease,
- $\gamma = 2$ as the step size increase, and
- $\rho(\alpha) = 10^{-3}\alpha^2$ as the forcing function.

At each step of the algorithm the vectors $\bar{\mathbf{x}}$ and $\mathbf{w}_{i,j}$, $j = 1, \dots, k$, determine a two-dimensional subspace. There is a normalized vector $\mathbf{v}_{i,j}$ in this subspace that satisfies $\bar{\mathbf{x}} \cdot \mathbf{v}_{i,j} = 0$, and an angle $0 < \theta_{i,j} < \pi/2$ such that $\mathbf{w}_{i,j} = \cos(\theta_{i,j})\bar{\mathbf{x}} + \sin(\theta_{i,j})\mathbf{v}_{i,j}$. A **contraction** of $\mathbf{w}_{i,j}$ is a reduction in the value of $\theta_{i,j}$ by some multiplicative factor $0 < \delta_{i,j} < 1$. [Below we'll use $\delta = \delta_{i,j} \equiv 0.99$.] A **dilation** of $\mathbf{w}_{i,j}$ is an increase in the value of $\theta_{i,j}$ by some multiplicative factor $1 < \gamma_{i,j}$. Consider an additional normalized random vector $\mathbf{u}_{i,j}$ where $\bar{\mathbf{x}} \cdot \mathbf{u}_{i,j} = \mathbf{u}_{i,j} \cdot \mathbf{v}_{i,j} = 0$, and an angle $0 < \phi_{i,j} < \pi/2$. A **perturbation** of $\mathbf{w}_{i,j}$ is given by

$$\mathbf{w}_{i,j}^{\pm} = \cos(\phi_{i,j}) \left[\cos(\hat{\theta}_{i,j})\bar{\mathbf{x}} + \sin(\hat{\theta}_{i,j})\mathbf{v}_{i,j} \right] \pm \sin(\phi_{i,j})\mathbf{u}_{i,j}$$

where $\hat{\theta}_{i,j}$ will be in some neighborhood of $\theta_{i,j}$.

The optimization routine is as follows:

Step 1 Set $i \leftarrow i + 1$.

Step 2 (Optional Search Steps) Try the following two searches.

Step 2a (optimal solution with contraction) Set $W' = XH_{i-1}^{\dagger}$, H_{i-1}^{\dagger} the pseudo-inverse of H_{i-1} , and normalize the columns to unit length. Until $S(W') \leq \epsilon$ and $0 < \theta'_j < \pi/2$, iteratively contract the \mathbf{w}'_j by a factor of $\delta = 0.99$. Then find H' by solving the NNLS problem. If $\varepsilon_{i-1} - \varepsilon' > \rho(\alpha_{i-1})$ set $\varepsilon_i \leftarrow \varepsilon'$, $W_i \leftarrow W'$, $H_i \leftarrow H'$, $\alpha_i = \min(\alpha_{max}, \gamma\alpha_{i-1})$ and Goto **Step 5**.

Step 2b (dilation) Form \mathbf{w}'_j by dilating $\mathbf{w}_{i-1,j}$ by a random factor $1 < \gamma'_j < 1 + \alpha_{i-1}$ under the restriction $0 < \theta'_j < \pi/2$. If $S(W') \leq \epsilon$ and $\varepsilon_{i-1} - \varepsilon' > \rho(\alpha_{i-1})$ set $\varepsilon_i \leftarrow \varepsilon'$, $W_i \leftarrow W'$, $H_i \leftarrow H'$, $\alpha_i = \min(\alpha_{max}, \gamma\alpha_{i-1})$ and Goto **Step 5**.

Step 3 (Poll Steps) Pick random angles $0 < \theta'_j < \pi/2$ where $|\theta'_j - \theta_{i-1,j}| < \alpha_{i-1}$, and $0 < \phi'_j < \alpha_{i-1}$.

Step 3a Form \mathbf{w}'_j^+ by perturbing $\mathbf{w}_{i-1,j}$. If $S(W') \leq \epsilon$ and $0 < \arccos(\mathbf{w}'_j^+ \cdot \bar{\mathbf{x}}) < \pi/2$ then find H' and ε' by solving the NNLS problem. If $\varepsilon_{i-1} - \varepsilon' > \rho(\alpha_{i-1})$ set $\varepsilon_i \leftarrow \varepsilon'$, $W_i \leftarrow W'$, $H_i \leftarrow H'$, $\alpha_i = \min(\alpha_{max}, \gamma\alpha_{i-1})$ and Goto **Step 5**.

Step 3b Form \mathbf{w}'_j^- by perturbing $\mathbf{w}_{i-1,j}$. If $S(W') \leq \epsilon$ and $0 < \arccos(\mathbf{w}'_j^- \cdot \bar{\mathbf{x}}) < \pi/2$ then find H' and ε' by solving the NNLS problem. If $\varepsilon_{i-1} - \varepsilon' > \rho(\alpha_{i-1})$ set $\varepsilon_i \leftarrow \varepsilon'$, $W_i \leftarrow W'$, $H_i \leftarrow H'$, $\alpha_i = \min(\alpha_{max}, \gamma\alpha_{i-1})$ and Goto **Step 5**.

Step 4 Set $\varepsilon_i \leftarrow \varepsilon_{i-1}$, $W_i \leftarrow W_{i-1}$, $H_i \leftarrow H_{i-1}$ and $\alpha_i = \theta\alpha_{i-1}$.

Step 5 If $i = i_{max}$ return W_i and H_i , otherwise Goto **Step 1**.

In the above method the tangent space to \mathcal{S}^{k-1} is never used. Rather, we move directly over \mathcal{S}^{k-1} or pick new points on \mathcal{S}^{k-1} . This avoids the need for an exponential (logarithmic) map from (to) the tangent space to (from) the manifold. The search steps also demonstrate the importance of using the probabilistic descent method in [22]. When we let $W = XH^\dagger$ and normalize the columns of W , the results of mapping these points on \mathcal{S}^{k-1} to the tangent space will most likely not lie on any *a priori* existing mesh.

9 Conclusions

If an optimization problem has a feasible set with enough symmetry then it may be possible to find a ‘law of motion’ over the set. This will have a group structure as well as a notion of distance. These two requirements, along with an initial feasible point, are all that are needed for performing an optimization routine. Using the group structure the initial feasible point can be moved over the manifold a specified distance. This movement allows us to (locally) reconstruct the entire feasible set. However a difficulty with direct search methods like MADS is the need for a mesh. The mesh implicitly restricts the group elements and the freedom to reconstruct the feasible set.

When working with nonlinear manifolds, even those with a high degree of symmetry, placing an appropriate mesh on the feasible set can be a daunting challenge. One way to avoid the problem is to have the mesh be in a tangent space to the manifold. This does work but the underlying symmetries of the feasible set are lost. The probabilistic descent method [22] removes the need for a mesh. Because of this one is now free to move over the manifold using the ‘law of motion’ with only a restriction on the distance traveled. Removing the reliance on a mesh allows the symmetries of the feasible set to be fully utilized.

Disclaimer Any opinions and conclusions expressed herein are those of the author and do not necessarily represent the views of the U.S. Census Bureau. The research in this paper does not use any confidential Census Bureau information. This was authored by an employee of the US national government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

Acknowledgements The author would like to thank the editor and reviewers for their helpful comments and suggestions.

References

1. Audet, C., J. E. Dennis, J.: A pattern search filter method for nonlinear programming without derivatives. *SIAM Journal on Optimization* **14**(4), 980–1010 (2004)
2. Audet, C., J. E. Dennis, J.: Mesh adaptive direct search algorithms for constrained optimization. *SIAM Journal on Optimization* **17**(1), 188–217 (2006)
3. Coope, I.D., Price, C.J.: Frame based methods for unconstrained optimization. *Journal of Optimization Theory and Applications* **107**(2), 261–274 (2000)

4. Dennis, J.E., Price, C.J., Coope, I.D.: Direct search methods for nonlinearly constrained optimization using filters and frames. *Optimization and Engineering* **5**(2), 123–144 (2004)
5. Finkel, D.E., Kelley, C.T.: Convergence analysis of the DIRECT algorithm. Tech. rep., Center for Research in Scientific Computation, North Carolina State University (2004). Technical report CRSC-TR04-28
6. Price, C.J., Coope, I.D.: Frames and grids in unconstrained and linearly constrained optimization: A nonsmooth approach. *SIAM Journal on Optimization* **14**(2), 415–438 (2003)
7. Abramson, M.A., Audet, C., J. E. Dennis, J., Digabel, S.L.: Orthomads: A deterministic mads instance with orthogonal directions. *SIAM Journal on Optimization* **20**(2), 948–966 (2009)
8. Fletcher, R., Leyffer, S.: Nonlinear programming without a penalty function. *Mathematical Programming* **91**(2), 239–269 (2002)
9. Audet, C., J. E. Dennis, J.: A progressive barrier for derivative-free nonlinear programming. *SIAM Journal on Optimization* **20**(1), 445–472 (2009)
10. Lewis, R.M., Torczon, V.: A globally convergent augmented Lagrangian pattern search algorithm for optimization with general constraints and simple bounds. *SIAM Journal on Optimization* **12**(4), 1075–1089 (2002)
11. Conn, A.R., Gould, N.I.M., Toint, P.: A globally convergent augmented Lagrangian algorithm for optimization with general constraints and simple bounds. *SIAM Journal on Numerical Analysis* **28**(2), 545–572 (1991)
12. Audet, C., Le Digabel, S., Peyrega, M.: Linear equalities in blackbox optimization. *Computational Optimization and Applications* **61**(1), 1–23 (2015)
13. Stillwell, J.: *Naive Lie Theory*, 1st edn. Undergraduate Texts in Mathematics. Springer-Verlag New York (2008)
14. Edelman, A., Arias, T.A., Smith, S.T.: The geometry of algorithms with orthogonality constraints. *SIAM J. Matrix Anal. Appl.* **20**(2), 303–353 (1999)
15. Sharpe, R.W.: *Differential Geometry: Cartan’s Generalization of Klein’s Erlangen Program*, *Graduate Texts in Mathematics*, vol. 166, 1st edn. Springer-Verlag New York (1997)
16. Smith, S.T.: Optimization techniques on Riemannian manifolds. *Fields Institute Communications* pp. 113–135 (1994)
17. Absil, P.A., Mahony, R., Sepulchre, R.: *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, Princeton, NJ (2008)
18. Hirsch, M.W.: *Differential Topology*. Graduate Texts in Mathematics. Springer New York (1997)
19. Helgason, S.: *Differential geometry, Lie groups, and symmetric spaces*. Graduate Studies in Mathematics. American Mathematical Society (2001)
20. Dreisigmeyer, D.W.: Equality constraints, Riemannian manifolds and direct search methods. http://www.optimization-online.org/DB_HTML/2007/08/1743.html (2006)
21. Sternberg, S.: *Curvature in Mathematics and Physics*. Dover (2012)
22. Gratton, S., Royer, C.W., Vicente, L.N., Zhang, Z.: Direct search based on probabilistic descent. *SIAM Journal on Optimization* **25**(3), 1515–1541 (2015)
23. Diaconis, P., Shahshahani, M.: The subgroup algorithm for generating uniform random variables. *Probability in the Engineering and Informational Sciences* **1**(1), 15–32 (1987)
24. Mezzadri, F.: How to generate random matrices from the classical compact groups. *Notices of the AMS* **54**, 592–604 (2007)
25. Stewart, G.W.: The efficient generation of random orthogonal matrices with an application to condition estimators. *SIAM Journal on Numerical Analysis* **17**(3), 403–409 (1980)
26. Dreisigmeyer, D.W.: A quasi-isometric embedding algorithm. arXiv:1709.01972 (2017)
27. Dreisigmeyer, D.W.: Tight semi-nonnegative matrix factorization. arXiv:1709.04395 (2017)
28. Krakowski, K.A., Manton, J.H.: On the computation of the Karcher mean on spheres and special orthogonal groups. In: *Proc. Workshop Robot. Math. (RoboMat 07)* (2007)
29. Gallier, J., Xu, D.: Computing exponentials of skew symmetric matrices and logarithms of orthogonal matrices. *International Journal of Robotics and Automation* **18**, 10–20 (2000)

-
30. Penec, X., Fillard, P., Ayache, N.: A Riemannian framework for tensor computing. *International Journal of Computer Vision* **66**(1), 41–66 (2006)

Name	\mathcal{M}	$\mathcal{T}_p\mathcal{M}$	$h(\omega_1, \omega_2)$	$\mathbf{Exp}_p(\omega)$	Notes
Grassmannian	$\mathbf{G}(n, k)$	$p^T\omega = 0$	$\text{Tr}(\omega_1^T\omega_2)$	$[pV \cos(\Theta) + U \sin(\Theta)] V^T$	$\omega = U\Theta V^T$, [14]
Stiefel manifold	$\mathbf{V}(n, k)$	$p^T\omega = -\omega^T p$	$\text{Tr}(\omega_1^T [I - \frac{1}{2}pp^T] \omega_2)$	$[p \ Q] \exp \begin{bmatrix} p^T\omega & -R^T \\ R & 0 \end{bmatrix} \begin{bmatrix} I \\ 0 \end{bmatrix}$	$QR = p_\perp p_\perp^T \omega$, [14]
General Linear Group	$\mathbf{GL}^+(n)$	$\mathbf{GL}(n)$	$\text{Tr}(\omega_1^T\omega_2)$	$p \exp(\omega)$	$p \in \mathbf{GL}^+(n)$ means $\det(p) > 0$
Special Linear Group	$\mathbf{SL}(n)$	$\text{Tr}(\omega) = 0$	$\text{Tr}(\omega_1^T\omega_2)$	$p \exp(\omega)$	$\mathbf{Exp}_p(\omega)$ is not surjective
Special Orthogonal Group	$\mathbf{SO}(n)$	$\omega = -\omega^T$	$\text{Tr}(\omega_1^T\omega_2)$	$p \exp(\omega)$	See [29] for $\exp(\omega)$
Special Euclidean Group	$\mathbf{SE}(n)$	$\omega = \begin{bmatrix} A & v \\ 0 & 0 \end{bmatrix}$	$\text{Tr}(\omega_1^T\omega_2)$	$p \exp(\omega)$	$A = -A^T, v \in \mathbb{R}^n$ See [29] for $\exp(\omega)$
Symmetric Positive Definite Group	$\mathbf{SPD}(n)$	$\omega = \omega^T$	$\text{Tr}(p^{-1/2}\omega_1 p^{-1}\omega_2 p^{-1/2})$	$p^{1/2} \exp(p^{-1/2}\omega p^{-1/2}) p^{1/2}$	[30]
Unipotent Group	$\mathbf{UP}(n)$	$\omega = U, U_{ii} = 0$	$\text{Tr}(\omega_1^T\omega_2)$	$\exp(\omega) = \sum_{k=0}^n \omega^k / k!$	U is upper triangular

Table 1: The tangent spaces, metrics and exponential maps for some real Lie groups and matrix manifolds.