EQUALITY CONSTRAINTS, RIEMANNIAN MANIFOLDS AND DIRECT SEARCH METHODS*

David W. Dreisigmeyer†
Los Alamos National Laboratory
Los Alamos, NM 87545

July 27, 2007

Abstract

We present a general procedure for handling equality constraints in optimization problems that is of particular use in direct search methods. First we will provide the necessary background in differential geometry. In particular, we will see what a Riemannian manifold is, what a tangent space is, how to move over a manifold and how to pullback functions from a manifold to the tangent spaces. The central idea of our optimization procedure is to treat the equality constraints as implicitly defining a $\mathcal{C}^2$ Riemannian manifold. Then the function and inequality constraints can be pulled-back to the tangent spaces of this manifold. One then needs to deal with the resulting optimization problem that only involves, at most, inequality constraints. An advantage of this procedure is the implicit reduction in dimensionality of the original problem to that of the manifold. Additionally, under some restrictions, convergence results for the method used to solve the inequality constrained optimization problem can be carried over directly to our procedure.

Key words: Direct search methods, Riemannian manifolds, equality constraints, geodesics

AMS subject classifications: 90C56, 53C21

1 Introduction

Direct search methods are typically designed to work in (a subset of) $\mathbb{R}^k$. So, if we only have inequality constraints present that define a full-dimensional feasible region, some of these methods that have convergence results for inequality constrained problems are potentially useful procedures to employ, e.g., [1, 2, 3, 7, 8, 11, 21]. However, the presence of even one equality constraint defining a feasible set $\mathcal{M}$ of measure zero

*LA-UR-06-7406

†email: dreisigm@lanl.gov

creates difficulties. The LTMADS algorithm [2] provides a typical illustration. The probability that a point on the mesh used in LTMADS also lies on the manifold $\mathcal{M}$ is zero. So every point that LTMADS considers will likely be infeasible. Filter techniques [1, 3, 12] 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 [18] extends the method in [6] to non-smooth problems with equality constraints. ([18] also contains a nice survey of previous methods to deal with equality constraints using direct search methods.) We have in mind something quite different.

The way we will extend direct search methods to equality constrained problems is to employ some techniques from differential geometry. Equality constraints will implicitly define a subset of the ambient space $\mathbb{R}^k$ of our optimization variables. Under some assumptions to be made precise later, this subset can be taken to be the central object of differential geometry: a manifold. Section 2 provides an overview of the required differential geometry. We have tried to make this as self-contained as possible, which has resulted in a very long section. Additionally, the concepts are developed in a non-linear programming 'language' as much as possible. Finally, for ease of exposition, we assume in this section some additional smoothness requirements on our objective function and inequality constraints that will be dropped when we move onto the direct search methods in Section 3. We hope that this will allow the reader to understand our direct search extension as easily as possible.

In Section 3 we outline our general procedure. We borrow some of the computational techniques developed in Section 2 to make sure that every point we consider in our optimization algorithm will always satisfy the equality constraints present. As we will see, this will result in an equivalent problem that only has, at most, inequality constraints, if such constraints were present in the original problem. Let us make the following quite clear from the outset: ***we are only concerned with extending direct search methods to problems with equality constraints.*** As such we assume that these methods are appropriate for the problem at hand. For example, perhaps the objective function is only Lipschitz. Whether the techniques we develop along the way have applications beyond this limited objective will not concern us at all.

An illustrative example is given in Section 4. The reader may find it useful to follow this example as they read the previous sections in order to make the ideas more concrete. A comparison with the generalized reduced gradient (GRG) method is provided in Section 5. Finally, Section 6 is a discussion of our results.

2 Manifolds

The object of our study is the following problem:

$$\min_{\mathbf{x}\in\mathbb{R}^{n+m}} : \quad f(\mathbf{x}) \tag{2.1a}$$

$$\text{subject to} : \quad \mathbf{g}(\mathbf{x}) = \mathbf{0} \tag{2.1b}$$

$$\mathbf{h}(\mathbf{x}) \leq \mathbf{0}. \tag{2.1c}$$

Here $f : \mathbb{R}^{n+m} \to \mathbb{R}$, $\mathbf{g} : \mathbb{R}^{n+m} \to \mathbb{R}^n$ and $\mathbf{h} : \mathbb{R}^{n+m} \to \mathbb{R}^l$. The set of points $\mathcal{M}_{\mathbf{c}} = \{\mathbf{x}|\mathbf{g}(\mathbf{x}) = \mathbf{c}\}$ is called a *level set* of $\mathbf{g}(\mathbf{x})$. For ease of exposition, all of the functions are assumed to be $\mathcal{C}^2$ ***in this section only***. In this section our focus will be on the constraints associated with the nonlinear program in (2.1). Our goal here is to illuminate the geometry associated with the feasible region defined by the equality constraint in (2.1b) and the inequality constraint in (2.1c). So we're not going to concern ourselves with any algorithms for actually solving (2.1). Useful references for this material are [14, 15, 16, 17, 19, 20, 24, 25]. These references can be consulted for derivations of the differential geometry and topology results used in the sequel.

2.1 Implicit Function Theorem and Regular Level Sets

We're going to begin by examining the level set defined by the equation $\mathbf{g}(\mathbf{x}) = \mathbf{0}$. Let's start with something familiar: the implicit function theorem. This will be an essential tool when we begin to look at manifolds. First let us state the theorem, see Figure 1 also:

Theorem 2.1 (The Implicit Function Theorem) *Let $\mathbf{g}(\mathbf{x})$ be a $\mathcal{C}^k$ function, with $k \geq 1$, defined on some open set $U \subset \mathbb{R}^{n+m}$ and taking values in $\mathbb{R}^n$. Assume that $\mathbf{x}_0 = [x_0^1, \ldots, x_0^{n+m}]$ satisfies the equation $\mathbf{g}(\mathbf{x}_0) = \mathbf{0}$. Let $\mathbf{y} = [x^1, \ldots, x^m]$, $\mathbf{y}_0 = [x_0^1, \ldots, x_0^m]$ and $\mathbf{w} = [x^{m+1}, \ldots, x^{n+m}]$. Denote the first-order derivatives of $\mathbf{g}(\mathbf{x})$ with respect to $\mathbf{w}$ by $\nabla_{\mathbf{w}}\mathbf{g}(\mathbf{x})$. Finally assume that $\det(\nabla_{\mathbf{w}}\mathbf{g}(\mathbf{x}_0)) \neq 0$, where $\det(\cdot)$ denotes the determinate. Then there exists open neighborhoods $V \subset \mathbb{R}^{n+m}$ of $\mathbf{x}_0$ and $Y \subset \mathbb{R}^m$ of $\mathbf{y}_0$ and, unique $\mathcal{C}^k$ functions $\psi^1(\mathbf{y}), \ldots, \psi^n(\mathbf{y})$, $\mathbf{y} \in Y$, such that*

$$\mathbf{g}(\mathbf{y}, \psi^1(\mathbf{y}), \ldots, \psi^n(\mathbf{y})) = \mathbf{0}. \tag{2.2}$$

So what is this telling us? For our level set $\mathbf{g}(\mathbf{x}) = \mathbf{0}$ in (2.1b), pick any point $\mathbf{x}_0$ that is on this level set. Now consider the Jacobian $\nabla\mathbf{g}(\mathbf{x}_0)$. If this is full rank, we can choose n columns of the Jacobian that are linearly independent. Then the n variables that correspond to our chosen columns of the Jacobian can be expressed as functions of the remaining m variables. We can always do a change of coordinates so that the n dependent variables correspond to the last n columns of $\nabla\mathbf{g}(\mathbf{x})$.

But we have more than that. Let $\mathcal{M} = \{\mathbf{x}|\mathbf{g}(\mathbf{x}) = \mathbf{0}\}$ and $\mathcal{V} = \mathcal{M} \cap V$, where $V \subset \mathbb{R}^{n+m}$ is the neighborhood of $\mathbf{x}_0$ in Theorem 2.1. Every $\mathbf{x} \in \mathcal{V}$ can be given a unique m-dimensional coordinate by using the $\psi^1(\mathbf{y}), \ldots, \psi^n(\mathbf{y})$ functions in (2.2). Namely, for the point $[\mathbf{y}, \psi^1(\mathbf{y}), \ldots, \psi^n(\mathbf{y})]^T \in \mathcal{V}$, we can assign it the coordinate $\mathbf{y}$.

Letting $\Pi = [I_m\, 0]$, where I_m is the m-dimensional identity matrix, for every $\mathbf{x} \in \mathcal{V}$ we have that $\mathbf{y} = \Pi\mathbf{x}$. So $\Pi : \mathcal{V} \to \mathbb{R}^m$. If we let $\Psi(\mathbf{y}) = [\mathbf{y}, \psi^1(\mathbf{y}), \ldots, \psi^n(\mathbf{y})]^T$, then $\mathbf{y} = (\Pi \circ \Psi)(\mathbf{y})$. That is, Π is the inverse of Ψ restricted to $\mathcal{V}$, denoted $\Pi = \Psi^{-1}|_{\mathcal{V}}$.

Example 2.2 *Let's look at the unit sphere $\mathcal{S}^2$ defined by*

$$g(x, y, z) = x^2 + y^2 + z^2 - 1 = 0. \tag{2.3}$$

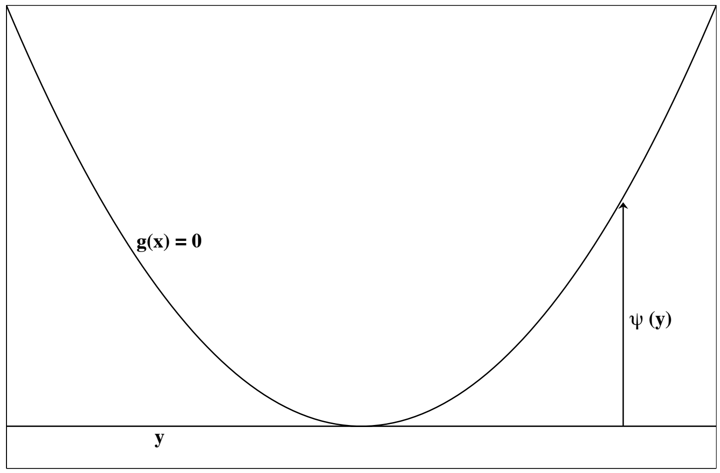

Figure 1: The implicit function theorem.

Let $\mathbf{x}_0 = [0\,0\,1]^T$. *Now,* $\nabla\mathbf{g}(\mathbf{x}_0) = [0\,0\,2]^T$*, so the implicit function theorem tells us that we can express the* z *variable as a function* ψ *of* x *and* y *such that* $g(x, y, \psi(x, y)) = 0$*. The function* $\psi(x, y) = \sqrt{1 - x^2 - y^2}$ *will work and is* $\mathcal{C}^\infty$ *for all* $z > 0$*, i.e., the upper hemisphere of the unit circle not including the equator. We can find the coordinates for the upper hemisphere of* $\mathcal{S}^2$ *by using the projection* $\Pi = [I_2\ \mathbf{0}]$*. Namely, for a point* $\mathbf{x} \in \mathcal{S}^2$ *that lies on the upper hemisphere, we have that* $\Pi\mathbf{x} = [x\ y]^T$*.*

To make the implicit function theorem work for every $\mathbf{x} \in \mathcal{M}$, we would need $\nabla\mathbf{g}(\mathbf{x})$ to have rank $(n - m)$ at every point $\mathbf{x} \in \mathcal{M}$.

Definition 2.3 (Regular Level Sets) *Let* $\nabla\mathbf{g}(\mathbf{x})$ *have full rank at every point on the level set* $\mathbf{g}(\mathbf{x}) = \mathbf{c}$*. Then* $\mathbf{g}(\mathbf{x}) = \mathbf{c}$ *is called a* ***regular level set*** *of* $\mathbf{g}(\mathbf{x})$*.*

Regular level sets will be the object of our attention in the remainder of this paper. Whenever we say level set in the sequel we mean regular level set. They will turn out to be manifolds. But first let us define what manifolds are.

2.2 Manifolds

Let's return to our level set $\mathcal{M}$ defined by (2.1b). Using the implicit function theorem, we found around a point $\mathbf{x}_0 \in \mathcal{M}$ a mapping $\Psi_{\mathbf{x}_0} : \mathbb{R}^m \supset Y \to V_{\mathbf{x}_0} \subset \mathbb{R}^{n+m}$ given by $\Psi_{\mathbf{x}_0}(\mathbf{y}) = [\mathbf{y}, \psi^1(\mathbf{y}), \ldots, \psi^n(\mathbf{y})]$ such that $\Psi_{\mathbf{x}_0}(\mathbf{y}_0) = \mathbf{x}_0$. We also learned that around $\mathbf{x}_0$ we can place local coordinates on $\mathcal{M}$. The neighborhood of $\mathcal{M}$ where this can be done is given by $\mathcal{V}_{\mathbf{x}_0} = \mathcal{M} \cap V_{\mathbf{x}_0}$ and the coordinates are given by $\mathbf{y}$. We call the pair $(\mathcal{V}_{\mathbf{x}_0}, \Psi_{\mathbf{x}_0}^{-1})$ a *chart* or *coordinate system* on $\mathcal{M}$. Note that $\Psi_{\mathbf{x}_0}^{-1}$ maps $\mathcal{V}_{\mathbf{x}_0} \subset \mathcal{M}$ into $\mathbb{R}^m$. So, for a point $\mathbf{x} \in \mathcal{V}_{\mathbf{x}_0}$ we can assign it the local coordinate $\Psi_{\mathbf{x}_0}^{-1}(\mathbf{x}) \in Y \subset \mathbb{R}^m$.

Now, since $\mathcal{M}$ is a regular level set, we can do the same thing for any point $\mathbf{x} \in \mathcal{M}$. Suppose we do this to obtain a whole set of charts $\{(\mathcal{V}_i, \phi_i)\}_{i \in \mathcal{I}}$, where $\mathcal{I}$ is some index set. Remember that $\mathcal{V}_i \subset \mathcal{M}$ and $\phi_i : \mathcal{V}_i \to \mathbb{R}^m$. If $\mathcal{M} = \bigcup_{i \in \mathcal{I}} \mathcal{V}_i$ then $\{(\mathcal{V}_i, \phi_i)\}_{i \in \mathcal{I}}$ is called an *atlas* on $\mathcal{M}$. Now every point on $\mathcal{M}$ must be in some neighborhood $\mathcal{V}_i$ in our atlas. So every point $\mathbf{x} \in \mathcal{M}$ will have at least one local coordinate defined on it. We then call $\mathcal{M}$ a *manifold*. In other words, $\mathcal{M}$ is a set of points in $\mathbb{R}^{n+m}$ that is locally diffeomorphic to $\mathbb{R}^m$.

Let us note that there are other technical requirements for a set of points in $\mathbb{R}^{n+m}$ to be a manifold. We won't concern ourselves with these since exploring them would take us too far afield and will have no impact on our future developments. Instead, we state the following:

Theorem 2.4 (Regular Level Set Manifolds) *Every regular level set of the $\mathcal{C}^2$ function $\mathbf{g}(\mathbf{x})$ is a manifold.*

If we only need a finite number of charts to cover $\mathcal{M}$ so that $\mathcal{M} = \bigcup_{i=0}^{n} \mathcal{V}_i, n < \infty$, we call $\mathcal{M}$ *compact*. Intuitively, a compact manifold is finite: it doesn't have any points 'at infinity'. If every point $\mathbf{x} \in \mathcal{M}$ can be connected to every other point $\mathbf{y} \in \mathcal{M}$ with some path that lies completely in $\mathcal{M}$, we call $\mathcal{M}$ *connected*. That is, $\mathcal{M}$ consists of only one piece.

Example 2.5 *Consider the hyperbola defined by the equation*

$$g_h(x, y) = xy - 1 = 0. \tag{2.4a}$$

The manifold is neither connected nor compact. The parabola defined by

$$g_p(x, y) = x^2 - y = 0 \tag{2.4b}$$

is connected but not compact. The circle

$$g_c(x, y) = x^2 + y^2 - 1 = 0 \tag{2.4c}$$

is both connected and compact.

What about those points on $\mathcal{M}$ that lie in more than one coordinate system: Can we relate the different coordinate systems to each other? What we do here is define an *overlap* or *coordinate change* function. Suppose $\mathbf{x} \in \mathcal{M}$ lies in the neighborhoods $\mathcal{V}_1$ and $\mathcal{V}_2$ on $\mathcal{M}$. Let $\mathcal{W} = \mathcal{V}_1 \cap \mathcal{V}_2 \subset \mathcal{M}$ and, $W_1 = \phi_1(\mathcal{W}) \subset \mathbb{R}^m$ and $W_2 = \phi_2(\mathcal{W}) \subset \mathbb{R}^m$. Consider the function $(\phi_2 \circ \phi_1^{-1}) : W_1 \to W_2$. Let's look at what's going on here. First ϕ_1^{-1} maps $W_1 \subset \mathbb{R}^m$ onto $\mathcal{W} \subset \mathcal{M}$. On the neighborhood $\mathcal{W}$ of

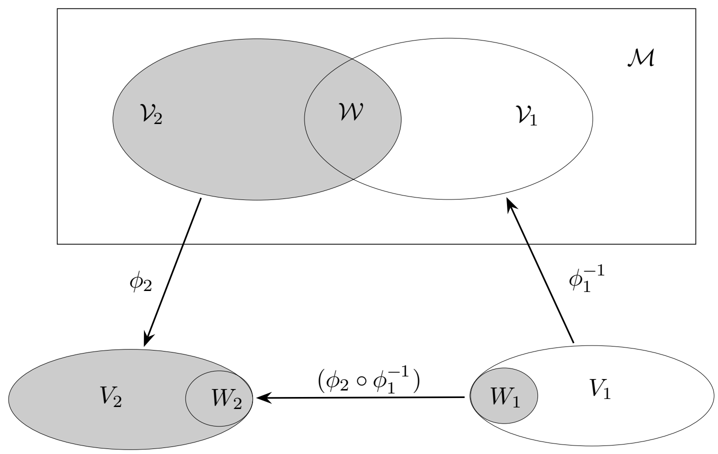

Figure 2: Overlap functions for a manifold.

$\mathcal{M}$, ϕ_2 is well defined because $\mathcal{W} \subset \mathcal{V}_2$. So ϕ_2 is mapping $\mathcal{W}$ back onto $W_2 \subset \mathbb{R}^m$. In this way, we can relate the two different coordinate systems defined on $\mathcal{W}$ by ϕ_1 and ϕ_2. The situation is illustrated in Figure 2

If all of our coordinate change functions are $\mathcal{C}^2$ then $\mathcal{M}$ is called a $\mathcal{C}^2$ manifold. For us, the implicit function theorem guarantees that all of the functions ϕ_i and ϕ_i^{-1} are $\mathcal{C}^2$. So all of our overlap functions are also $\mathcal{C}^2$ and, hence, so is $\mathcal{M}$. In fact, it's easy to see that $\mathcal{M}$ is of the same smoothness class as the $\mathbf{g}(\mathbf{x})$ in (2.1b).

Example 2.6 *Let us return to $\mathcal{S}^2$ given by*

$$g(x, y, z) \;=\; x^2 + y^2 + z^2 - 1 \;=\; 0. \tag{2.5}$$

We already have one chart given by $\mathcal{V}_1 = \{\mathbf{x} \in \mathcal{S}^2 | z > 0\}$ and $\phi_1 = [I_2\ \mathbf{0}]$. A second chart is given by $\mathcal{V}_2 = \{\mathbf{x} \in \mathcal{S}^2 | x > 0\}$ and $\phi_2 = [\mathbf{0}\ I_2]$. Then $\mathcal{W} = \mathcal{V}_1 \cap \mathcal{V}_2 = \{\mathbf{x} \in \mathcal{S}^2 | x, z > 0\}$, $W_1 = \{\mathbf{y} \in \mathbb{R}^2 | \|\mathbf{y}\| < 1 \text{ and } y_1 > 0\}$ and $W_2 = \{\mathbf{z} \in \mathbb{R}^2 | \|\mathbf{z}\| < 1 \text{ and } z_2 > 0\}$. The coordinate change function is given by

$$\begin{aligned} (\phi_2 \circ \phi_1^{-1}) \quad &: \quad W_1 \to W_2 \\ &= \quad \left[y_2 \;\; \sqrt{1 - y_1^2 - y_2^2} \right]^T. \end{aligned} \tag{2.6}$$

2.3 Normal and Tangent Spaces

An important point in the implicit function theorem was that $\nabla\mathbf{g}(\mathbf{x})$ be full rank. But what is $\nabla\mathbf{g}(\mathbf{x})$? The short answer is that $\mathrm{span}([\nabla\mathbf{g}(\mathbf{x})]^T)$ gives us all of the *normal directions* to $\mathcal{M}$ at $\mathbf{x}$, denoted by $\mathcal{N}_\mathbf{x}\mathcal{M}$. This, in turn, implies that $\mathrm{null}(\nabla\mathbf{g}(\mathbf{x}))$ is the *tangent space* to $\mathcal{M}$ at $\mathbf{x}$, which we write as $\mathcal{T}_\mathbf{x}\mathcal{M}$. Let us examine this in a little more detail.

You may have seen the following definition for the tangent cone:

Definition 2.7 (Tangent Vector and Tangent Cone) *The vector $\mathbf{w} \in \mathbb{R}^{n+m}$ is called a **tangent vector** to $\mathcal{M}$ at $\mathbf{x}$ if*

$$\mathbf{w} = \lim_{k\to\infty} \lambda \, \frac{\mathbf{x}_k - \mathbf{x}}{\|\mathbf{x}_k - \mathbf{x}\|}, \tag{2.7}$$

*where $\lambda \geq 0$, $\mathbf{x}_k \in \mathcal{M}$ for every k and, $\mathbf{x}_k \to \mathbf{x}$, $\mathbf{x}_k \neq \mathbf{x}$. The set of all such vectors is called the **tangent cone**.*

This definition coincides with the definition $\mathcal{T}_\mathbf{x}\mathcal{M} = \mathrm{null}(\nabla\mathbf{g}(\mathbf{x}))$.

In differential geometry, one will normally construct $\mathcal{T}_\mathbf{x}\mathcal{M}$ in an alternate way. Let $\mathbf{y}(\tau)$ be a curve on $\mathcal{M}$ defined for $\tau \in [a, b]$, where $a < 0 < b$ and $\mathbf{y}(0) = \mathbf{x}$. Now consider the vector $\dot{\mathbf{y}}(0) \in \mathbb{R}^{n+m}$. This will be a tangent to the curve $\mathbf{y}(\tau)$ at $\tau = 0$. But, since $\mathbf{y}(\tau) \in \mathcal{M}$ for all $\tau \in [a, b]$, $\dot{\mathbf{y}}(0)$ will also be a tangent vector to $\mathcal{M}$. The collection of all such tangent vectors to all of the curves on $\mathcal{M}$ that pass through $\mathbf{x}$ at $\tau = 0$ gives us the tangent space to $\mathcal{M}$ at $\mathbf{x}$.

The normal cone is defined as the orthogonal complement to the tangent cone:

Definition 2.8 (Normal Cone) *The **normal cone** to $\mathcal{M}$ at $\mathbf{x}$ is the collection of all of the vectors $\mathbf{n} \in \mathbb{R}^{n+m}$ such that $\mathbf{n} \cdot \mathbf{w} = 0$ for every $\mathbf{w} \in \mathcal{T}_\mathbf{x}\mathcal{M}$.*

Of course this also corresponds with our original definition of the normal directions, namely $\mathcal{N}_\mathbf{x}\mathcal{M} = \mathrm{span}([\nabla\mathbf{g}(\mathbf{x})]^T)$. Note that we can replace the word *cone* with *subspace* in the above definitions when we are only dealing with the equality constraint $\mathbf{g}(\mathbf{x}) = \mathbf{0}$ that implicitly defines $\mathcal{M}$.

It's easy to see that $\mathbb{R}^{n+m}$ can be decomposed into the direct sum of the tangent and normal spaces to $\mathcal{M}$ at $\mathbf{x}$. That is $\mathbb{R}^{n+m} = \mathcal{T}_\mathbf{x}\mathcal{M} \oplus \mathcal{N}_\mathbf{x}\mathcal{M}$ because, by construction, $\mathcal{T}_\mathbf{x}\mathcal{M} \cap \mathcal{N}_\mathbf{x}\mathcal{M} = \{\mathbf{0}\}$. Here, $\mathbf{0} \in \mathbb{R}^{n+m}$ corresponds to the base point $\mathbf{x} \in \mathcal{M}$ where $\mathcal{T}_\mathbf{x}\mathcal{M}$ and $\mathcal{N}_\mathbf{x}\mathcal{M}$ are 'attached' to $\mathcal{M}$. This is an important point, so let us state it again. The origin $\mathbf{0}$ of $\mathcal{T}_\mathbf{x}\mathcal{M}$ corresponds to $\mathbf{x} \in \mathcal{M}$ and, the origin $\mathbf{0}$ of $\mathcal{N}_\mathbf{x}\mathcal{M}$ also corresponds to the point $\mathbf{x} \in \mathcal{M}$.

An alternate way of thinking about this is as follows. $\mathcal{T}_\mathbf{x}\mathcal{M}$ is an m-dimensional affine subspace of $\mathbb{R}^{n+m}$. So $\mathcal{T}_\mathbf{x}\mathcal{M}$ is a copy of $\mathbb{R}^m$, denoted $\mathcal{T}_\mathbf{x}\mathcal{M} \cong \mathbb{R}^m$. Let T be a full-rank $(n+m)$-by-m matrix who's columns form a basis for $\mathcal{T}_\mathbf{x}\mathcal{M}$. Then the affine map $A_{\mathbf{x}}(\mathbf{w}) = T\mathbf{w} + \mathbf{x}$ maps $\mathbb{R}^m$ into $\mathbb{R}^{n+m}$. It does this in such a way that the image of $\mathbb{R}^m$ corresponds to our intuitive notion of what the tangent space to $\mathcal{M}$ should be. Namely, we have an m-dimensional plane 'attached' to $\mathcal{M}$ at the point $\mathbf{x}$ that is tangent to $\mathcal{M}$ at $\mathbf{x}$. Note in particular that $A_{\mathbf{x}}(\mathbf{0}) = \mathbf{x}$. Here the $\mathbf{w}$ will be our coordinates on $\mathcal{T}_{\mathbf{x}}\mathcal{M}$ and, $T\mathbf{w}$ will be our tangent vectors in the ambient space at the point $\mathbf{x} \in \mathcal{M}$, i.e., $T\mathbf{w} \in \mathcal{T}_{\mathbf{x}}\mathcal{M}$. The situation is illustrated in Figure 3. A similar concept holds for $\mathcal{N}_{\mathbf{x}}\mathcal{M}$.

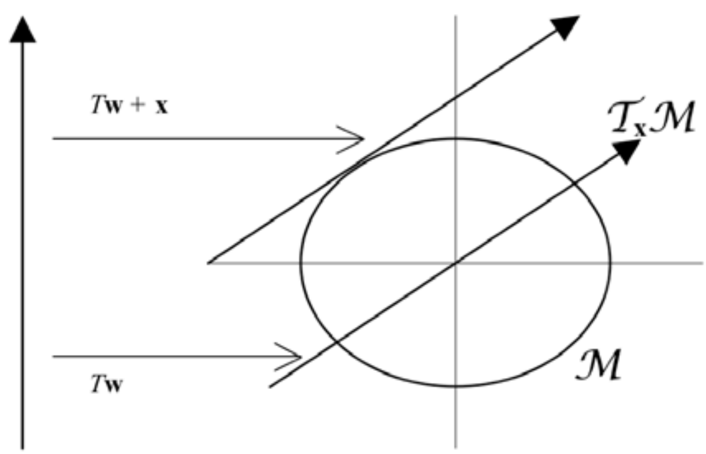

Figure 3: The tangent space $\mathcal{T}_{\mathbf{x}}\mathcal{M} \cong \mathbb{R}$ as a subspace in $\mathbb{R}^2$ and visualized as an affine subspace attached to our manifold $\mathcal{M}$.

2.4 Riemannian Metrics and Manifolds

One thing that is potentially useful to do is to take the inner products of two different tangent vectors $\mathbf{v}, \mathbf{w} \in \mathcal{T}_{\mathbf{x}}\mathcal{M}$ at a given point $\mathbf{x} \in \mathcal{M}$. A *Riemannian metric* is a smooth mapping $m(\cdot, \cdot) : \mathcal{T}_{\mathbf{x}}\mathcal{M} \times \mathcal{T}_{\mathbf{x}}\mathcal{M} \rightarrow \mathbb{R}$ defined over all of $\mathcal{M}$ that is: 1) *symmetric*, i.e., $m(\mathbf{v}, \mathbf{w}) = m(\mathbf{w}, \mathbf{v})$ for every $\mathbf{v}, \mathbf{w} \in \mathcal{T}_{\mathbf{x}}\mathcal{M}$, and; 2) *positive definite*, i.e., $m(\mathbf{w}, \mathbf{w}) > 0$ for all $\mathbf{0} \neq \mathbf{w} \in \mathcal{T}_{\mathbf{x}}\mathcal{M}$. The pair $(\mathcal{M}, m)$ is called a *Riemannian manifold*.

All the manifolds we consider will be subsets of some ambient space $\mathbb{R}^n$. As such, their tangent spaces have a natural metric defined on them that they inherit from $\mathbb{R}^n$, namely, the restriction of the standard Euclidean inner product $\mathbf{w} \cdot \mathbf{v} = \sum_{i=1}^{n} w^i v^i$ to each $\mathcal{T}_{\mathbf{x}}\mathcal{M}$ for every $\mathbf{x} \in \mathcal{M}$. A Riemannian metric derived in this way is called an *induced metric*. In particular, all of our tangent spaces will have a sense of tangent vector lengths and angles between tangent vectors. While all our manifolds will be Riemannian, we will have no explicit use for the metric m in the sequel.

2.5 Manifolds with Boundaries and Corners

Things are complicated slightly when we have inequality constraints present. Let $\Omega = \{\mathbf{x}|\mathbf{x} \text{ satisfies } \mathbf{h}(\mathbf{x}) \leq \mathbf{0}\}$ and $\mathcal{M}_\Omega = \mathcal{M} \cap \Omega$. That is, $\mathcal{M}_\Omega$ is our *feasible set*. Now, if $\mathcal{M}$ is a manifold then $\mathcal{M}_\Omega$ will also be a manifold. However, now we have points on $\mathcal{M}_\Omega$ that are at a boundary. To see this let $\mathbf{x}^* \in \mathcal{M}_\Omega$ satisfy some of our inequality constraints in (2.1c). Around $\mathbf{x}^*$ there will be points on $\mathcal{M}_\Omega$ as well as points that are on $\mathcal{M}$ but not on $\mathcal{M}_\Omega$. These latter points will satisfy the equality constraints in (2.1b) but will violate the inequality constraints in (2.1c). We can fix things up by appropriately defining what we mean by a boundary of a manifold.

Let the m-dimensional *half-space* be given by $\mathbb{H}^m = \{\mathbf{x} \in \mathbb{R}^m | x^1 \geq 0\}$. The boundary of $\mathbb{H}^m$ is given by $\partial\mathbb{H}^m = \{\mathbf{x} \in \mathbb{H}^m | x^1 = 0\}$. To define a boundary for a manifold, we allow our diffeomorphisms ϕ_i in the charts $(\mathcal{V}_i, \phi_i)_{i\in\mathcal{I}}$ to map $\mathcal{V}_i \subset \mathcal{M}_\Omega$ into $\mathbb{H}^m$ instead of $\mathbb{R}^m$. Then the *boundary* of the manifold $\mathcal{M}_\Omega$ is given by the set of points $\partial\mathcal{M}_\Omega = \{\mathbf{x} \in \mathcal{M}_\Omega | \phi_i(\mathbf{x}) \in \partial\mathbb{H}^m \text{ for some } i \in \mathcal{I}\}$. If $\partial\mathcal{M}_\Omega \neq \emptyset$, we call $\mathcal{M}_\Omega$ a *manifold with boundary*. Somewhat more difficult to show is that $\partial\mathcal{M}_\Omega$ is well defined:

Theorem 2.9 *If $\phi_i(\mathbf{x}) \in \partial\mathbb{H}^m$ for $\mathbf{x} \in \mathcal{V}_i$, then $\phi_j(\mathbf{x}) \in \partial\mathbb{H}^m$ if $\mathbf{x} \in \mathcal{V}_j$ with $j \neq i$. That is, a boundary point of $\mathcal{M}_\Omega$ remains a boundary point under a change of coordinates.*

We can also consider more general situations. We can replace $\mathbb{H}^m$ above with $\mathbb{R}^m_+ = \{\mathbf{x} \in \mathbb{R}^m | x^j \geq 0 \text{ for } j = 1, \ldots, m\}$ to define the corners of a manifold. A point $\mathbf{x} \in \mathcal{M}_\Omega$ is a *corner* of $\mathcal{M}_\Omega$ if $\phi_i(\mathbf{x}) = \mathbf{0}$ for some $i \in \mathcal{I}$. If $\mathcal{M}_\Omega$ has a corner, we call it a *cornered manifold*.

Both manifolds with boundary and manifolds with corners are special cases of *manifolds with generalized boundary*. That is, $\mathbb{H}^m$ and $\mathbb{R}^m_+$ are special cases of the more general range space $\mathbb{B}^m_k = \{\mathbf{x} \in \mathbb{R}^m | x^j \geq 0 \text{ for } j = 1, \ldots, k \leq m\}$. The boundary of $\mathbb{B}^m_k$ is given by $\partial\mathbb{B}^m_k = \{\mathbf{x} \in \mathbb{B}^m_k | x^j = 0 \text{ for at least one } 1 \leq j \leq k\}$. Then $\partial\mathcal{M}_\Omega$ is given by those points $\mathbf{x} \in \mathcal{M}_\Omega$ that are mapped by some ϕ_i to $\partial\mathbb{B}^m_k$.

2.6 Linear Independence Constraint Qualification

One common constraint qualification you may have seen is the following:

Definition 2.10 (Linear Independence Constraint Qualification) *Let our constraints be given by $\mathbf{g}(\mathbf{x}) = \mathbf{0}$ and $\mathbf{h}(\mathbf{x}) \leq \mathbf{0}$. At the point $\mathbf{x}^*$, assume that $h_i(\mathbf{x}^*) = 0$ for*

$i \in \mathcal{A}(\mathbf{x}^*)$. The set of indices $\mathcal{A}(\mathbf{x}^*)$ is called the ***active inequality constraint set****. If the set of gradients $\{\nabla \mathbf{g}(\mathbf{x}^*)\} \bigcup \{\nabla h_i^T(\mathbf{x}^*)\}_{i\in\mathcal{A}(\mathbf{x}^*)}$ are linearly independent then the* ***linear independence constraint qualification (LICQ)*** *is said to hold.*

In this section only, for simplicity we will only consider those constraints that satisfy the LICQ everywhere. Appropriately, we'll give the set they define a special name.

Definition 2.11 (Regular Constraint Set) *If the LICQ holds for every feasible point $\mathbf{x} \in \mathcal{M}_\Omega = \mathcal{M} \cap \Omega$, we call $\mathcal{M}_\Omega$ a* ***regular constraint set***

The importance of regular constraint sets is given by the following:

Theorem 2.12 *If $\mathbf{g}(\mathbf{x})$ and $\mathbf{h}(\mathbf{x})$ are $\mathcal{C}^k$, then $\mathcal{M}_\Omega$ is a $\mathcal{C}^k$ manifold with generalized boundary.*

Now let us assume that the LICQ holds at every point on $\mathcal{M}$. Then we can use the slack variables $\mathbf{z} = [z^1, \ldots, z^l]$ to change our optimization problem from one that has both equality and inequality constraints into one that has only equality constraints. If we let $\mathbf{z} \odot \mathbf{z} = [(z^1)^2, \ldots, (z^l)^2]$ and $\mathbf{y} = [\mathbf{x}^T, \mathbf{z}^T]^T$, (2.1) can be restated as

$$\min_{\mathbf{y}\in\mathbb{R}^{n+m+l}} : \quad f(\mathbf{x}) \tag{2.8a}$$

$$\text{subject to} : \quad \mathbf{g}(\mathbf{x}) = \mathbf{0} \tag{2.8b}$$

$$\mathbf{h}(\mathbf{x}) + \mathbf{z} \odot \mathbf{z} = \mathbf{0}. \tag{2.8c}$$

Because the LICQ holds, the set $\widetilde{\mathcal{M}}$ defined by (2.8b) and (2.8c) is a regular level set and, hence, a manifold.

Again let $\Omega = \{\mathbf{x} | \mathbf{x} \text{ satisfies } \mathbf{h}(\mathbf{x}) \leq \mathbf{0}\}$ and $\mathcal{M}_\Omega = \mathcal{M} \cap \Omega$. By converting (2.1) into the equivalent form in (2.8), we are taking multiple copies of $\mathcal{M}_\Omega$ and 'sewing' them together to get $\widetilde{\mathcal{M}}$. To see this, consider some point $\mathbf{x} \in \mathcal{M}_\Omega$. Then $\mathbf{y} \in \widetilde{\mathcal{M}}$ if and only if $z^i = \pm\sqrt{-h_i(\mathbf{x})}$ for $i = 1, \ldots, l$. So we take 2^l copies of $\mathcal{M}_\Omega$ and 'sew' them together at their boundaries to get $\widetilde{\mathcal{M}}$. Note that if $\mathcal{M}_\Omega$ is compact, then $\widetilde{\mathcal{M}}$ is also compact. Also, $\widetilde{\mathcal{M}}$ has no boundary. From the above developments, we see that we can restrict our attention to problems of the form

$$\min_{\mathbf{x}\in\mathbb{R}^{n+m}} : \quad f(\mathbf{x}) \tag{2.9a}$$

$$\text{subject to} : \quad \mathbf{g}(\mathbf{x}) = \mathbf{0}, \tag{2.9b}$$

where $f : \mathbb{R}^{n+m} \to \mathbb{R}$ and $\mathbf{g} : \mathbb{R}^{n+m} \to \mathbb{R}^n$, provided the LICQ holds everywhere.

2.7 Geodesics I

We're done defining the types of manifolds we're going to be dealing with. Of course, just realizing that our constraints define a manifold doesn't get us very much. We want to be able to actually use the concept of a manifold in a practical way. This section marks the beginning of actually working with our manifolds. Here we'll define some special paths on $\mathcal{M}$: geodesics. After developing some additional concepts in later sections, we'll revisit geodesics again where we'll derive an actual formula for calculating them.

In $\mathbb{R}^{n+m}$ the shortest distance between any two points $\mathbf{x}$ and $\mathbf{y}$ is the straight line $\mathbf{z}$ that connects them. We can parameterize the line between $\mathbf{x}$ and $\mathbf{y}$ by letting $\mathbf{z}(0) = \mathbf{x}$ and $\mathbf{z}(1) = \mathbf{y}$, where $\mathbf{z}(\tau)$ is defined for $\tau \in [0, 1]$ and $\dot{\mathbf{z}}(\tau) = \mathbf{y} - \mathbf{x}$. What's special about this line? Let $l(\tau)$, $\tau \in [0, 1]$, be any line that connects $\mathbf{x}$ to $\mathbf{y}$, where $l(0) = \mathbf{x}$ and $l(1) = \mathbf{y}$. Now consider the following integral:

$$L = \int_0^1 \sqrt{\dot{l}(\tau) \cdot \dot{l}(\tau)}\, d\tau. \tag{2.10}$$

The integral tells us the length we travel along $l(\tau)$ as we go from $\mathbf{x}$ to $\mathbf{y}$. Suppose we want to minimize L. Then the first variation of (2.10) must vanish, i.e., $\delta L = 0$, which results in the equation $\ddot{l}(\tau) = 0$ whose solution is $\mathbf{z}(\tau)$.

We can generalize this concept to our manifold $\mathcal{M}$. A *geodesic* between two points $\mathbf{x}$ and $\mathbf{y}$ on $\mathcal{M}$ is a path $\mathbf{z}(\tau)$, $\tau \in [0, 1]$, on $\mathcal{M}$ such that $\mathbf{z}(0) = \mathbf{x}$, $\mathbf{z}(1) = \mathbf{y}$ and, the first variation of the length of $\mathbf{z}(\tau)$ vanishes. We'll be able to derive a formula for finding the geodesics on $\mathcal{M}$ but first we need to develop a little more machinery.

2.8 Lagrangians and Covariant Derivatives

Associated with our optimization problem (2.9) is the Lagrangian

$$\mathcal{L}(\mathbf{x}, \lambda) = f(\mathbf{x}) - \lambda \cdot \mathbf{g}(\mathbf{x}), \tag{2.11}$$

where $\lambda \in \mathbb{R}^n$. Suppose we want to know how $f(\mathbf{x})$ changes around some point $\mathbf{y} \in \mathcal{M}$ as we move away from $\mathbf{y}$ along paths $\mathbf{x}(\tau)$ that are restricted to lie on $\mathcal{M}$ and pass through $\mathbf{y}$ at $\tau = 0$. Now, $\dot{\mathbf{x}}(0) = \mathbf{w} \in \mathcal{T}_{\mathbf{y}}\mathcal{M}$. The *directional derivative* of $f(\mathbf{x})$ in the direction $\mathbf{w}$ at the point $\mathbf{y}$ is given by

$$\begin{aligned} D_{\mathbf{w}} f(\mathbf{y}) &= \nabla f^T(\mathbf{y}) \cdot \mathbf{w} \\ &= \nabla_{\mathbf{x}} \mathcal{L}(\mathbf{y}, \lambda) \cdot \mathbf{w}, \end{aligned} \tag{2.12}$$

where $\nabla f^T(\mathbf{y})$ is the transpose of $\nabla f(\mathbf{y})$ and

$$\lambda = \nabla f^T(\mathbf{y}) \cdot \left[\nabla \mathbf{g}^T(\mathbf{y}) \cdot \left[\nabla \mathbf{g}(\mathbf{y}) \cdot \nabla \mathbf{g}^T(\mathbf{y}) \right]^{-1} \right]. \tag{2.13}$$

We use (2.13) as our λ because it will result in the projection of $\nabla f^T(\mathbf{y})$ onto $\mathcal{T}_{\mathbf{y}}\mathcal{M}$, as we'll see shortly. Notice that

$$\nabla \mathbf{g}^+(\mathbf{y}) = \nabla \mathbf{g}^T(\mathbf{y}) \cdot \left[\nabla \mathbf{g}(\mathbf{y}) \cdot \nabla \mathbf{g}^T(\mathbf{y}) \right]^{-1}, \tag{2.14}$$

is the *pseudo-inverse* of $\nabla \mathbf{g}(\mathbf{y})$ [13] and satisfies

$$\nabla \mathbf{g}(\mathbf{y}) \cdot \nabla \mathbf{g}^{+}(\mathbf{y}) \quad = \quad I_n. \tag{2.15}$$

So (2.13) can be rewritten as

$$\lambda \quad = \quad \nabla f^T(\mathbf{y}) \cdot \nabla \mathbf{g}^{+}(\mathbf{y}). \tag{2.16}$$

In differential geometry the directional derivative $D_{\mathbf{w}} f(\mathbf{y})$, where $\mathbf{w}$ is restricted to lie in $\mathcal{T}_{\mathbf{y}}\mathcal{M}$, is replaced with the *first covariant derivative* of $f(\mathbf{y})$. It is the projection of $\nabla f^T(\mathbf{y})$ onto the tangent plane $\mathcal{T}_{\mathbf{y}}\mathcal{M}$ and will be denoted by $f_{;j}(\mathbf{y})$, $j = 1, \ldots, n+m$. Using (2.13) in (2.12), we have that

$$\begin{aligned} f_{;j}(\mathbf{y}) \quad &= \quad \left[\nabla_{\mathbf{x}} \mathcal{L}(\mathbf{y}, \lambda)\right]_j \\ &= \quad \left[\nabla f^T(\mathbf{y}) \cdot \left(I - \nabla \mathbf{g}^{+}(\mathbf{y}) \cdot \nabla \mathbf{g}(\mathbf{y})\right)\right]_j . \end{aligned} \tag{2.17}$$

Notice that $(I - \nabla \mathbf{g}^{+} \cdot \nabla \mathbf{g})$ is the projection onto $\mathcal{T}_{\mathbf{y}}\mathcal{M}$. For $\mathbf{z} = [z^1, \ldots, z^{n+m}]^T$, where now $\mathbf{z}$ does not have to lie in $\mathcal{T}_{\mathbf{y}}\mathcal{M}$, we'll write the covariant derivative of $f(\mathbf{y})$ in the direction $\mathbf{z}$ as

$$f_{;j} z^j \quad = \quad \sum_{j=1}^{n+m} f_{;j}(\mathbf{y}) z^j. \tag{2.18}$$

In (2.18) we introduced the *Einstein summation convention*. This means that a repeated index is summed over if it occurs once as a superscript and once as a subscript. So, for example, $x^j y_j = \sum_j x^j y_j$. Combinations like $x^j y^j$ and $x_j y_j$ are not allowed and generally indicate an error occurred somewhere in a derivation. If $\mathbf{z} \in \mathcal{T}_{\mathbf{y}}\mathcal{M}$, (2.18) becomes

$$f_{;j} z^j \quad = \quad f_{,j} z^j, \tag{2.19}$$

where $f_{,j}$ is the partial derivative of $f(\mathbf{y})$ in terms of x^j, which is exactly (2.12). The 'comma' and 'semicolon' notation are commonly (though not always) used to indicate partial and covariant derivatives, respectively, in the literature.

Now we'll look at the *second covariant derivative* of $f(\mathbf{y})$, which is the second order changes in $f(\mathbf{y})$ when we are restricted to moving on $\mathcal{M}$. Letting $\mathbf{w} \in \mathcal{T}_{\mathbf{y}}\mathcal{M}$, we have that

$$\begin{aligned} \mathcal{L}(\mathbf{y} + \mathbf{w}, \lambda) \quad \approx \quad & \mathcal{L}(\mathbf{y}, \lambda) + \nabla_{\mathbf{x}} \mathcal{L}(\mathbf{y}, \lambda) \cdot \mathbf{w} \\ & + \frac{1}{2} \mathbf{w} \cdot \left[\nabla_{\mathbf{x}}^2 \mathcal{L}(\mathbf{y}, \lambda)\right] \cdot \mathbf{w}, \end{aligned} \tag{2.20}$$

or, equivalently,

$$f(\mathbf{y} + \mathbf{w}) \quad \approx \quad f(\mathbf{y}) + \nabla_{\mathbf{x}} \mathcal{L}(\mathbf{y}, \lambda) \cdot \mathbf{w} + \frac{1}{2} \mathbf{w} \cdot \left[\nabla_{\mathbf{x}}^2 \mathcal{L}(\mathbf{y}, \lambda)\right] \cdot \mathbf{w}. \tag{2.21}$$

We've already seen that $\nabla_{\mathbf{x}}\mathcal{L}(\mathbf{y},\lambda)$ gives us the first covariant derivative of $f(\mathbf{y})$. One would suspect that $\nabla^2_{\mathbf{x}}\mathcal{L}(\mathbf{y},\lambda)$ will give us the second covariant derivative of $f(\mathbf{y})$

$$\begin{aligned} f_{;j;k}(\mathbf{y}) &= \left[\nabla^2_{\mathbf{x}}\mathcal{L}(\mathbf{y},\lambda)\right]_{jk} \\ &= \left[\nabla^2 f(\mathbf{y}) - \left(\nabla f^T(\mathbf{y})\cdot\nabla\mathbf{g}^+(\mathbf{y})\right)\cdot\nabla^2\mathbf{g}(\mathbf{y})\right]_{jk}, \end{aligned} \tag{2.22}$$

$j,k = 1,\ldots,n+m$. This is in fact correct, which we'll now demonstrate.

We're going to derive the equations for the first and second covariant derivative of $f(\mathbf{y})$ in an intrinsic way. That is, we're going to restrict ourselves to only working on $\mathcal{M}$ by constructing a coordinate chart around $\mathbf{y}\in\mathcal{M}$. Then we can restate everything extrinsically in the ambient space $\mathbb{R}^{n+m}$. Assume that $\mathbf{y}=\mathbf{0}$, which means that $\mathbf{g}(\mathbf{0})=\mathbf{0}$. Additionally, let $\mathcal{T}_{\mathbf{0}}\mathcal{M}$ correspond to the first m coordinates of $\mathbb{R}^{n+m}$. Now we can use the implicit function theorem to find a function $\psi(\mathbf{w})\in\mathbb{R}^n$, $\mathbf{w}\in\mathcal{T}_{\mathbf{0}}\mathcal{M}\cong\mathbb{R}^m$, such that $\mathbf{g}(\mathbf{w},\psi(\mathbf{w}))=\mathbf{0}$. Notice that $\psi(\mathbf{0})=\mathbf{0}$ and, since $\mathcal{T}_{\mathbf{0}}\mathcal{M}$ corresponds to the first m coordinates of $\mathbb{R}^{n+m}$, $\nabla_{\mathbf{w}}\psi(\mathbf{0}) = 0_{n\times m}$, where $0_{n\times m}$ is the n-by-m matrix of all zeros. Now,

$$\begin{aligned} \nabla_{\mathbf{w}}\mathbf{g}(\mathbf{0},\psi(\mathbf{0})) &= \nabla_{\mathbf{x}}\mathbf{g}(\mathbf{0})\cdot\begin{bmatrix} I_m \\ \nabla_{\mathbf{w}}\psi(\mathbf{0})\end{bmatrix} \\ &= \nabla_{\mathbf{x}}\mathbf{g}(\mathbf{0})\cdot\begin{bmatrix} I_m \\ 0_{n\times m}\end{bmatrix} \\ &= 0_{n\times m}, \end{aligned} \tag{2.23}$$

since $\nabla_{\mathbf{x}}\mathbf{g}(\mathbf{0})$ is normal to $\mathcal{M}$ at $\mathbf{0}$ and $\mathcal{T}_{\mathbf{0}}\mathcal{M} = \operatorname{span}([I_m\ 0_{m\times n}]^T)$. This tells us that

$$\nabla_{\mathbf{x}}\mathbf{g}(\mathbf{0}) = \begin{bmatrix} 0_{n\times m} & G(\mathbf{0}) \end{bmatrix}, \tag{2.24}$$

where $G(\mathbf{0})\in\mathbb{R}^{n\times n}$ is nonsingular. We then have

$$\begin{aligned} \nabla^2_{\mathbf{w}}\mathbf{g}(\mathbf{0},\psi(\mathbf{0})) &= \begin{bmatrix} I_m & 0_{m\times n}\end{bmatrix}\cdot\nabla^2_{\mathbf{x}}\mathbf{g}(\mathbf{0})\cdot\begin{bmatrix} I_m \\ 0_{n\times m}\end{bmatrix} \\ &\quad + G(\mathbf{0})\cdot\nabla^2_{\mathbf{w}}\psi(\mathbf{0}) \\ &= 0_{n\times m\times m}, \end{aligned} \tag{2.25}$$

since $\nabla_{\mathbf{w}}\mathbf{g}(\mathbf{w},\psi(\mathbf{w})) = 0_{n\times m}$ for all $\mathbf{w}\in\mathcal{T}_{\mathbf{0}}\mathcal{M}$ sufficiently near $\mathbf{0}$. Solving for $\nabla^2_{\mathbf{w}}\psi(\mathbf{0})$ gives us

$$\nabla^2_{\mathbf{w}}\psi = -G^{-1}\cdot\left(\begin{bmatrix} I_m & 0_{m\times n}\end{bmatrix}\cdot\nabla^2_{\mathbf{x}}\mathbf{g}\cdot\begin{bmatrix} I_m \\ 0_{n\times m}\end{bmatrix}\right). \tag{2.26}$$

Intrinsically, i.e., using the coordinate system we constructed via the implicit function theorem, $f_{;k}$ is given by

$$f_{;k} = \left[\nabla_{\mathbf{w}} f^T(\mathbf{0},\psi(\mathbf{0}))\right]_k, \tag{2.27}$$

for $k = 1, \ldots, m$, where

$$\begin{aligned} \nabla_{\mathbf{w}} f^T(\mathbf{0}, \psi(\mathbf{0})) &= \nabla_{\mathbf{x}} f^T(\mathbf{0}) \cdot \begin{bmatrix} I_m \\ \nabla_{\mathbf{w}} \psi(\mathbf{0}) \end{bmatrix} \\ &= \nabla_{\mathbf{x}} f^T(\mathbf{0}) \cdot \begin{bmatrix} I_m \\ 0_{n \times m} \end{bmatrix}. \end{aligned} \tag{2.28}$$

Then $f_{;k} w^k$ tells us how $f(\mathbf{0})$ changes in the direction $\mathbf{w} = [w^1, \ldots, w^m]^T$ for an arbitrary $\mathbf{w} \in \mathcal{T}_{\mathbf{0}}\mathcal{M}$. Extrinsically, i.e., in the ambient space $\mathbb{R}^{n+m}$, we can restate (2.27) as

$$\begin{aligned} f_{;j} &= \left[\nabla_{\mathbf{x}} f^T(\mathbf{0}) \cdot \left(I - \nabla_{\mathbf{x}} \mathbf{g}^+(\mathbf{0}) \cdot \nabla_{\mathbf{x}} \mathbf{g}(\mathbf{0})\right)\right]_j \\ &= \left[\nabla_{\mathbf{x}} \mathcal{L}(\mathbf{0}, \lambda)\right]_j, \end{aligned} \tag{2.29}$$

where $j = 1, \ldots, n+m$. This is just (2.17).

A similar derivation will give us the second covariant derivative of $f(\mathbf{0})$. Intrinsically we have

$$f_{;k;l} = \left[\nabla^2_{\mathbf{w}} f(\mathbf{0}, \psi(\mathbf{0}))\right]_{kl}, \tag{2.30}$$

for $k, l = 1, \ldots, m$, where

$$\begin{aligned} \nabla^2_{\mathbf{w}} f(\mathbf{0}, \psi(\mathbf{0})) &= \begin{bmatrix} I_m & 0_{m \times n} \end{bmatrix} \cdot \nabla^2_{\mathbf{x}} f(\mathbf{0}) \cdot \begin{bmatrix} I_m \\ 0_{n \times m} \end{bmatrix} \\ &\quad + \nabla_{\mathbf{x}} f^T(\mathbf{0}) \cdot \begin{bmatrix} 0_{m \times m} \\ \nabla^2_{\mathbf{w}} \psi(\mathbf{0}) \end{bmatrix}. \end{aligned} \tag{2.31}$$

Extrinsically, we let

$$\begin{bmatrix} I_m & 0_{m \times n} \end{bmatrix} \cdot \nabla^2_{\mathbf{x}} f(\mathbf{0}) \cdot \begin{bmatrix} I_m \\ 0_{n \times m} \end{bmatrix} \rightarrow \nabla^2_{\mathbf{x}} f(\mathbf{0}) \tag{2.32a}$$

and, using (2.26),

$$\nabla_{\mathbf{x}} f^T(\mathbf{0}) \cdot \begin{bmatrix} 0_{m \times m} \\ \nabla^2_{\mathbf{w}} \psi(\mathbf{0}) \end{bmatrix} \rightarrow -\left(\nabla_{\mathbf{x}} f^T(\mathbf{0}) \cdot \mathbf{g}^+(\mathbf{0})\right) \cdot \nabla^2_{\mathbf{x}} \mathbf{g}(\mathbf{0}). \tag{2.32b}$$

Hence,

$$\begin{aligned} f_{;i;j} &= \left[\nabla^2 f(\mathbf{0}) - \left(\nabla f^T(\mathbf{0}) \cdot \nabla \mathbf{g}^+(\mathbf{0})\right) \cdot \nabla^2 \mathbf{g}(\mathbf{0})\right]_{ij} \\ &= \left[\nabla^2_{\mathbf{x}} \mathcal{L}(\mathbf{0}, \lambda)\right]_{ij}, \end{aligned} \tag{2.33}$$

where $i, j = 1, \ldots, n+m$, which is (2.22).

Equation (2.33) is only correct when used with the vectors $\mathbf{w} \in \mathcal{T}_{\mathbf{0}}\mathcal{M}$. That is, $f_{;i;j} w^i w^j = \mathbf{w} \cdot \nabla^2_{\mathbf{x}} \mathcal{L}(\mathbf{0}, \lambda) \cdot \mathbf{w}$ only if $\mathbf{w} \in \mathcal{T}_{\mathbf{0}}\mathcal{M}$. For a general vector $\mathbf{w} \in \mathbb{R}^{n+m}$ we would first need to project it down onto the tangent space $\mathcal{T}_{\mathbf{0}}\mathcal{M}$. This could be accomplished by letting $f_{;i;j} = [(I - \nabla \mathbf{g}^+ \nabla \mathbf{g}) \cdot \nabla^2_{\mathbf{x}} \mathcal{L} \cdot (I - \nabla \mathbf{g}^+ \nabla \mathbf{g})]_{ij}$. Since we will only be working with tangent vectors, we'll take (2.33) as defining the extrinsic form of $f_{;i;j}$.

2.9 One-forms and the Cotangent Space

The *cotangent space* $\mathcal{T}^*_{\mathbf{x}}\mathcal{M}$ to $\mathcal{M}$ at the point $\mathbf{x}$ is the dual space to $\mathcal{T}_{\mathbf{x}}\mathcal{M}$. An element $\omega \in \mathcal{T}^*_{\mathbf{x}}\mathcal{M}$ is called a *one-form*, which is a linear mapping from $\mathcal{T}_{\mathbf{x}}\mathcal{M}$ to $\mathbb{R}$. So, for $\mathbf{y}, \mathbf{z} \in \mathcal{T}_{\mathbf{x}}\mathcal{M}$ and $a, b \in \mathbb{R}$, we have that

$$\omega(a\mathbf{y} + b\mathbf{z}) \quad = \quad a\omega(\mathbf{y}) + b\omega(\mathbf{z}) \quad \in \quad \mathbb{R}. \tag{2.34}$$

The collection of all such one-forms at $\mathbf{x} \in \mathcal{M}$ gives us the cotangent space $\mathcal{T}^*_{\mathbf{x}}\mathcal{M}$. For us, we will generally take a vector $\mathbf{w} \in \mathcal{T}_{\mathbf{x}}\mathcal{M}$ as being a column vector. Then one-forms can be taken as row vectors.

The components of a vector $\mathbf{w} \in \mathcal{T}_{\mathbf{x}}\mathcal{M}$ will be written with superscripts, so $\mathbf{w} = [w^1, \ldots, w^{n+m}]^T$. For a one-form $\omega \in \mathcal{T}^*_{\mathbf{x}}\mathcal{M}$ the components will be written with subscripts: $\omega = [\omega_1, \ldots, \omega_{n+m}]$. Then the inner product $\omega \cdot \mathbf{w}$ is written as $\omega_j w^j$ using the Einstein summation convention.

We work with one-forms all of the time: $\nabla f(\mathbf{x})$ is a one-form. Let us return to (2.12). We see that taking the inner product of $\nabla f^T(\mathbf{x})$ with the tangent vector $\mathbf{w} \in \mathcal{T}_{\mathbf{x}}\mathcal{M}$ gives us the directional derivative of $f(\mathbf{x})$ at $\mathbf{x} \in \mathcal{M}$ in the direction $\mathbf{w}$. Thus, $\nabla f(\mathbf{x})$ is a linear mapping of $\mathcal{T}_{\mathbf{x}}\mathcal{M}$ into $\mathbb{R}$, so $\nabla f(\mathbf{x})$ is a one-form. Notice that we had to take the transpose of $\nabla f(\mathbf{x})$ in order to make the inner product work. Convention forces us to treat $\nabla f(\mathbf{x})$ as a column vector, though a more sensible convention would be to treat it as a row vector.

There's a nice mental picture of what one-forms are. Around $\mathbf{x} \in \mathbb{R}^{n+m}$ we can model $f(\mathbf{y})$ as a series of planes since $f(\mathbf{y}) \approx f(\mathbf{x}) + \nabla f^T(\mathbf{x}) \cdot \mathbf{w}$, $\mathbf{w} \in \mathbb{R}^{n+m}$. One of the planes satisfies $\nabla f^T(\mathbf{x}) \cdot \mathbf{w} = 0$ and all of the other planes are parallel to this one with $\nabla f(\mathbf{x})$ being their normal direction. So every plane satisfies $\nabla f^T(\mathbf{x}) \cdot \mathbf{w} = c$ for some fixed constant c. We can then imagine $\nabla f(\mathbf{x})$ as being these planes. The situation is illustrated in Figure 4.

2.10 Christoffel Symbols

Let $\omega(\mathbf{x}) = [\omega_1(\mathbf{x}), \ldots, \omega_{n+m}(\mathbf{x})]$ be a one-form field defined on $\mathcal{M}$. That is, $\omega(\mathbf{x}) \in \mathcal{T}^*_{\mathbf{x}}\mathcal{M}$. The equation for the first covariant derivative of $\omega(\mathbf{x})$ one will normally see in a differential geometry book is

$$\omega_{i;j} \quad = \quad \omega_{i,j} - \Gamma^k_{\;ij}\omega_k, \tag{2.35}$$

where the coefficients $\Gamma^k_{\;ij}$ are called the Christoffel symbols. We'll find it very convenient to have a formula for finding the $\Gamma^k_{\;ij}$ given some level set $\mathbf{g}(\mathbf{x}) = \mathbf{0}$. An important one-form field for us is the first covariant derivative of $f(\mathbf{x})$, i.e., $\omega_i = f_{;i}$. Using this in (2.35) gives us

$$\begin{aligned} f_{;i;j} \quad &= \quad f_{;i,j} - \Gamma^k_{\;ij} f_{;k} \\ &= \quad f_{,i,j} - \Gamma^k_{\;ij} f_{,k}, \end{aligned} \tag{2.36}$$

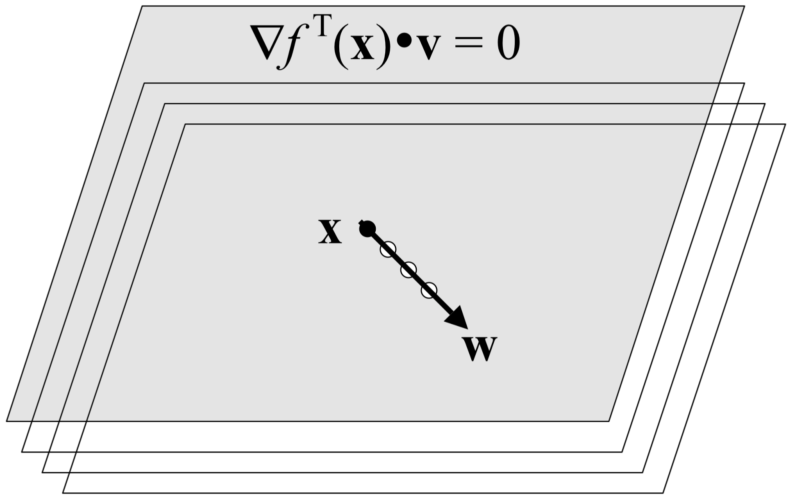

Figure 4: Visualization of a one-form around a point $\mathbf{x} \in \mathcal{M}$.

which is the second covariant derivative of $f(\mathbf{x})$. Comparing (2.36) with (2.33), we see that

$$\begin{aligned} \Gamma^k_{\ ij} &= \left[\nabla \mathbf{g}^+ \cdot \nabla^2 \mathbf{g}\right]^k_{\ ij} \\ &= \frac{\partial \mathbf{g}^T}{\partial x^k} \cdot (\nabla \mathbf{g} \cdot \nabla \mathbf{g}^T)^{-1} \cdot \frac{\partial^2 \mathbf{g}}{\partial x^i \partial x^j}. \end{aligned} \tag{2.37}$$

With the Christoffel symbols in hand, we can now define what the first covariant derivative of the vector field $\mathbf{w}(\mathbf{x}) \in \mathcal{T}_{\mathbf{x}}\mathcal{M}$ is:

$$w^i_{\ ;j} = w^i_{\ ,j} + \Gamma^i_{\ kj} w^k, \tag{2.38}$$

where $\mathbf{w}(\mathbf{x}) = [w^1(\mathbf{x}), \ldots, w^{n+m}(\mathbf{x})]^T$. To see this, let the curve $\mathbf{x}(\tau) \in \mathcal{M}$ satisfy $\mathbf{w}(\mathbf{x}(\tau)) = \dot{\mathbf{x}}(\tau)$, which means we just have to solve an ordinary differential equation (ODE) for the unique solution $\mathbf{x}(\tau)$. Now consider the function

$$\mathcal{L}(\mathbf{x}(\tau), \lambda) = \mathbf{x}(\tau) - \lambda \cdot \mathbf{g}(\mathbf{x}(\tau)), \tag{2.39}$$

where $\lambda = \nabla \mathbf{g}^+(\mathbf{x}(\tau))$, see (2.13). Then

$$\begin{aligned} \dot{\mathbf{x}} &= \left[I - [\nabla \mathbf{g}^+(\mathbf{x})] \cdot \nabla \mathbf{g}(\mathbf{x})\right] \cdot \dot{\mathbf{x}} \\ &= \nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \lambda) \cdot \dot{\mathbf{x}}, \end{aligned} \tag{2.40}$$

because $\nabla \mathbf{g}(\mathbf{x}(\tau))$ is normal to $\mathcal{M}$ while $\dot{\mathbf{x}}(\tau) \in \mathcal{T}_{\mathbf{x}(\tau)}\mathcal{M}$. So now we have

$$\mathbf{w}(\mathbf{x}) = \nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \lambda) \cdot \mathbf{w}(\mathbf{x}), \tag{2.41}$$

Taking the second derivative with respect to τ gives us

$$\begin{aligned}\dot{\mathbf{w}}(\mathbf{x}) &= [\nabla \mathbf{w}(\mathbf{x})] \cdot \dot{\mathbf{x}} \\ &= \dot{\mathbf{x}} \cdot \left[-\nabla \mathbf{g}^{+} \cdot \nabla^2 \mathbf{g}\right] \cdot \mathbf{w} + \left[I - \nabla \mathbf{g}^{+} \cdot \nabla \mathbf{g}\right] \cdot \nabla \mathbf{w} \cdot \dot{\mathbf{x}} \\ &= \mathbf{w} \cdot \nabla^2_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \lambda) \cdot \mathbf{w} + \nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \lambda) \cdot \nabla \mathbf{w} \cdot \mathbf{w}.\end{aligned} \tag{2.42}$$

Rearranging (2.42) gives us

$$\left[I - \nabla \mathbf{g}^{+} \cdot \nabla \mathbf{g}\right] \cdot \nabla \mathbf{w} \cdot \mathbf{w} = \nabla \mathbf{w} \cdot \mathbf{w} + \mathbf{w} \cdot \left(\nabla \mathbf{g}^{+} \cdot \nabla^2 \mathbf{g}\right) \cdot \mathbf{w}. \tag{2.43}$$

Since this holds for an arbitrary $\mathbf{w}(\mathbf{x}(\tau))$, we have that

$$\left[I - \nabla \mathbf{g}^{+} \cdot \nabla \mathbf{g}\right] \cdot \nabla \mathbf{w} = \nabla \mathbf{w} + \left(\nabla \mathbf{g}^{+} \cdot \nabla^2 \mathbf{g}\right) \cdot \mathbf{w}. \tag{2.44}$$

Now, the left-hand side of (2.44) is the projection of $\nabla \mathbf{w}(\mathbf{x}(\tau))$ onto $\mathcal{T}_{\mathbf{x}(\tau)}\mathcal{M}$. So it is the first covariant derivative of $\mathbf{w}(\mathbf{x}(\tau))$. Looking at the individual components of (2.44), we see that it is exactly (2.38).

Now consider some path $\mathbf{x}(\tau)$ defined on $\mathcal{M}$. We can find the first covariant derivative of $\omega(\mathbf{x}(\tau))$ and $\mathbf{w}(\mathbf{x}(\tau))$ along the path $\mathbf{x}(\tau)$ by using the formulas

$$\omega_{i;j}\dot{x}^j = \omega_{i,j}\dot{x}^j - \Gamma^k{}_{ij}\omega_k \dot{x}^j \tag{2.45a}$$

and

$$w^i{}_{;j}\dot{x}^j = w^i{}_{,j}\dot{x}^j + \Gamma^i{}_{kj} w^k \dot{x}^j, \tag{2.45b}$$

respectively. Examples of the above formulas were given in (2.12) and (2.42).

2.11 Parallel Transport

There are special vector fields $\mathbf{w}(\mathbf{x})$ who's first covariant derivative along the path $\mathbf{x}(\tau) \in \mathcal{M}$ vanishes. So, using (2.45b) we see that

$$\begin{aligned} w^i{}_{;j}\dot{x}^j &= 0 \\ &= w^i{}_{,j}\dot{x}^j + \Gamma^i{}_{kj} w^k \dot{x}^j. \end{aligned} \tag{2.46}$$

Such a vector field is said to be *parallel transported* along the curve $\mathbf{x}(\tau)$. This is the natural generalization to a manifold of a concept that is actually very familiar to you. Suppose we want to add some vector $\mathbf{y} \in \mathbb{R}^{n+m}$ to another vector $\mathbf{x} \in \mathbb{R}^{n+m}$. To do this, you place the 'tail' of $\mathbf{y}$ on the 'head' of $\mathbf{x}$ to get $\mathbf{x} + \mathbf{y}$. Originally the tail of $\mathbf{y}$ was at the origin $\mathbf{0}$ of $\mathbb{R}^{n+m}$. When we transported the tail of $\mathbf{y}$ to the head of $\mathbf{x}$, we didn't rotate $\mathbf{y}$. That is, we kept $\mathbf{y}$ constant as we moved it from $\mathbf{0}$ to the point specified by $\mathbf{x}$. So we parallel transported $\mathbf{y}$ along the path specified by $\mathbf{x}$. When $\mathbf{y}$ is parallel transported in this way, we move $\mathbf{y}$ from the tangent space $\mathcal{T}_{\mathbf{0}}\mathbb{R}^{n+m}$ to the new tangent space $\mathcal{T}_{\mathbf{x}}\mathbb{R}^{n+m}$, keeping it constant all along the way. Typically we wouldn't think of it this way because $\mathcal{T}_{\mathbf{x}}\mathbb{R}^{n+m} \cong \mathbb{R}^{n+m}$. Yet, this added complication is what we need

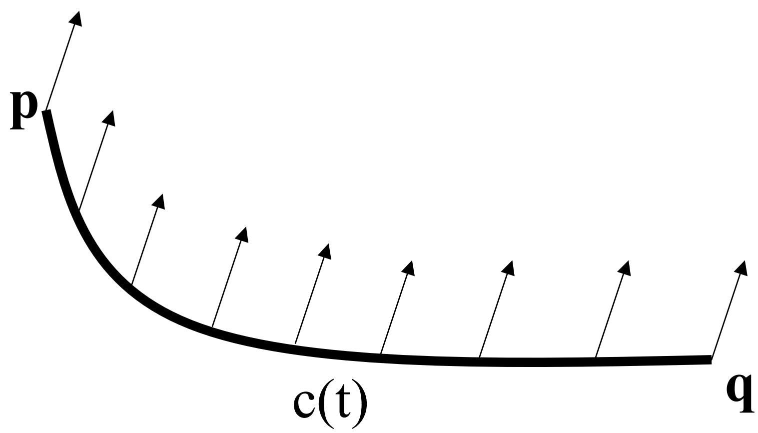

Figure 5: Parallel transporting a vector from $\mathbf{p}$ to $\mathbf{q}$ along the curve $c(t)$ in $\mathbb{R}^2$.

in order to extend what is an easy thing to do in $\mathbb{R}^{n+m}$ to a general manifold $\mathcal{M}$. The concept is illustrated in Figure 5 for a curve in $\mathbb{R}^2$.

Given some curve $\mathbf{x}(\tau) \in \mathcal{M}$ and an initial vector $\mathbf{w}_0 \in \mathcal{T}_{\mathbf{x}(0)}\mathcal{M}$, there is a unique parallel transported vector field $\mathbf{w}(\tau)$ along $\mathbf{x}(\tau)$ such that $\mathbf{w}(0) = \mathbf{w}_0$ and $\mathbf{w}(\tau) \in \mathcal{T}_{\mathbf{x}(\tau)}\mathcal{M}$. Note the requirement that $\mathbf{w}(\tau)$ is always a tangent vector to $\mathcal{M}$. We are not parallel transporting $\mathbf{w}_0$ as a vector in $\mathbb{R}^{n+m}$ as in Figure 5. Looking at (2.46), we see we only need to solve the ODE

$$\dot{w}^i(\mathbf{x}) = w^i_{\ ,j}\dot{x}^j = -\Gamma^i_{\ kj} w^k \dot{x}^j. \tag{2.47}$$

Parallel transport then gives us a way to compare two vectors that live in different tangent spaces. Returning to $\mathbb{R}^{n+m}$ again, suppose we have a vector $\mathbf{y} \in \mathcal{T}_{\mathbf{x}_1}\mathbb{R}^{n+m}$ and another vector $\mathbf{z} \in \mathcal{T}_{\mathbf{x}_2}\mathbb{R}^{n+m}$ that we want to compare. We can't do this directly because they live in different tangent spaces. In order to compare $\mathbf{z}$ to $\mathbf{y}$, we must first parallel transport $\mathbf{z}$ from $\mathbf{x}_2$ to $\mathbf{x}_1$. Now $\mathbf{z}$ is in the tangent space $\mathcal{T}_{\mathbf{x}_1}\mathbb{R}^{n+m}$ and can be directly compared with $\mathbf{y}$ because both vectors are in the same tangent space.

The same concept holds for a general manifold $\mathcal{M}$. Now, if our tangent vectors $\mathbf{y}$ and $\mathbf{z}$ lie in $\mathcal{T}_{\mathbf{x}_1}\mathcal{M}$ and $\mathcal{T}_{\mathbf{x}_2}\mathcal{M}$, respectively, we construct a path $\mathbf{w}(\tau)$ on $\mathcal{M}$ connecting $\mathbf{x}_2$ to $\mathbf{x}_1$. Then $\mathbf{z}$ is parallel transported along $\mathbf{w}(\tau)$ from $\mathbf{x}_2$ to $\mathbf{x}_1$ to get the new vector $\mathbf{z}_*$. Now $\mathbf{z}_*$ and $\mathbf{y}$ can be directly compared because they lie in the same tangent space, namely $\mathcal{T}_{\mathbf{x}_1}\mathcal{M}$.

Generally the $\mathbf{z}_*$ found by parallel transporting $\mathbf{z}$ along the path $\mathbf{w}(\tau)$ will be path dependent. So, if $\widehat{\mathbf{w}}(\tau)$ is a different path connecting $\mathbf{x}_2$ to $\mathbf{x}_1$, the resulting $\widehat{\mathbf{z}}_* \in \mathcal{T}_{\mathbf{x}_1}\mathcal{M}$ will not be the same as $\mathbf{z}_*$. To see this, consider some vector $\mathbf{z}$ at the equator of the earth. We'll parallel transport this vector to the north pole by using two different paths. First, parallel transport $\mathbf{z}$ directly to the north pole to find $\mathbf{z}_*$. For our second path, first parallel transport $\mathbf{z}$ 180° eastward to find $\widetilde{\mathbf{z}}_*$. Now parallel transport $\widetilde{\mathbf{z}}_*$ to the north pole to find $\widehat{\mathbf{z}}_*$. The vectors $\mathbf{z}_*$ and $\widehat{\mathbf{z}}_*$ will point in opposite directions! Such things are the price to be paid when working with general manifolds.

2.12 Geodesics II

Now we'll look at geodesics again. First we'll give an alternate definition of geodesics. Using this new definition, a quasi-linear system of ODEs will be derived for actually finding a geodesic given some initial conditions. Finally, we'll state some properties that geodesics have.

We've already defined geodesics as those paths on $\mathcal{M}$ who's lengths don't change under infinitesimal perturbations. An example of such a path would be the shortest path connecting two points $\mathbf{x} \in \mathcal{M}$ and $\mathbf{y} \in \mathcal{M}$. Now, in $\mathbb{R}^{n+m}$ the shortest path between any two points is the straight line that connects them. Let $\mathbf{z}(\tau)$, $\tau \in [0,1]$, be the straight line of constant velocity $\dot{\mathbf{z}}(\tau) = \mathbf{y} - \mathbf{x}$ that connects $\mathbf{x}$ to $\mathbf{y}$. So $\mathbf{z}(0) = \mathbf{x}$ and $\mathbf{z}(1) = \mathbf{y}$. We can state this as follows: the shortest path $\mathbf{z}(\tau)$ between $\mathbf{x}$ and $\mathbf{y}$ is that which parallel transports it's initial velocity vector $\dot{\mathbf{z}}(0) = \mathbf{y} - \mathbf{x}$ along itself. This is how we'll define our geodesics. The two alternate definitions of a geodesic are equivalent, but it would be somewhat difficult to show that so we won't bother.

Using (2.47), our new definition of a geodesic is a path $\mathbf{x}(\tau)$ on $\mathcal{M}$ that satisfies the equation

$$\ddot{x}^i = -\Gamma^i{}_{kj}\dot{x}^k\dot{x}^j, \tag{2.48}$$

where we used that fact that now $w^k = \dot{x}^k$. Notice that (2.48) is the quasi-linear system of ODEs that we promised. Since we can find the Christoffel symbols using (2.37), if we know $\nabla\mathbf{g}(\mathbf{x})$ and $\nabla^2\mathbf{g}(\mathbf{x})$, we can find the geodesics on our manifold $\mathcal{M}$ implicitly defined by the level set $\mathbf{g}(\mathbf{x}) = \mathbf{0}$. Solving (2.48) allows us to move away from $\mathbf{y} \in \mathcal{M}$ while all the time remaining on $\mathcal{M}$. That is, we can ensure that our constraints in the optimization problem (2.9) are always satisfied provided we start with some initial feasible point $\mathbf{y}$.

Some properties that any geodesic $\mathbf{x}(\tau)$ will satisfy are the following:

$$\|\dot{\mathbf{x}}(0)\| = \|\dot{\mathbf{x}}(\tau)\|, \tag{2.49a}$$
$$\dot{\mathbf{x}}(\tau) \in \mathcal{T}_{\mathbf{x}(\tau)}\mathcal{M}, \tag{2.49b}$$
$$\ddot{\mathbf{x}}(\tau) \in \mathcal{N}_{\mathbf{x}(\tau)}\mathcal{M}, \tag{2.49c}$$

and, the length that $\mathbf{x}(\tau)$ travels on $\mathcal{M}$ from $\tau = 0$ to $\tau = 1$ is given by

$$\int_0^1 \sqrt{\dot{\mathbf{x}}(\tau)\cdot\dot{\mathbf{x}}(\tau)}\,d\tau = \|\dot{\mathbf{x}}(0)\|. \tag{2.49d}$$

Let us revisit covariant derivatives. Let $\mathbf{x}(\tau)$ be a geodesic on $\mathcal{M}$ where $\mathbf{x}(0) = \mathbf{y}$ and $\dot{\mathbf{x}}(0) = \mathbf{w}$ and consider the function $f(\tau) = f(\mathbf{x}(\tau))$. Now, using (2.49b), we have that

$$\begin{aligned} \dot{f}(0) &= [\nabla_{\mathbf{x}} \mathcal{L}(\mathbf{y}, \lambda)]_j \, w^j \\ &= f_{;j} w^j, \end{aligned} \tag{2.50}$$

where λ is given by (2.16). Using (2.49c), we see that

$$\begin{aligned} \ddot{f}(0) &= \left[\nabla^2_{\mathbf{x}} \mathcal{L}(\mathbf{y}, \lambda)\right]_{jk} w^j w^k \\ &= f_{;j;k} w^j w^k. \end{aligned} \tag{2.51}$$

So the first and second covariant derivatives of $f(\mathbf{x})$ tell us how $f(\mathbf{x})$ changes as we move away from $\mathbf{y} \in \mathcal{M}$ on $\mathcal{M}$ via a geodesic $\mathbf{x}(\tau)$.

2.13 Geodesically Complete Manifolds

Here we'll take one final look at the manifolds $\mathcal{M}$ implicitly defined by $\mathbf{g}(\mathbf{x}) = \mathbf{0}$ in our optimization problem (2.9). We already stated that we'll assume $\mathcal{M}$ is compact and without boundary. Is $\mathcal{M}$ connected? Generally the answer is no. $\mathcal{M}$ can consist of multiple pieces that are not connected to each other via a path that lies on $\mathcal{M}$. So let $\mathcal{M} = \bigcup_{i=1}^{N} \mathcal{M}_i$, where the $\mathcal{M}_i$ are the disjoint pieces of $\mathcal{M}$ and, we're assuming there's a finite number of them. Note that the $\mathcal{M}_i$ are themselves manifolds of the same dimension as $\mathcal{M}$ and, since $\mathcal{M}$ is without boundary the $\mathcal{M}_i$ are also without boundary.

Now take an arbitrary point $\mathbf{x} \in \mathcal{M}_K$ for some $K \in \{1, \ldots, N\}$. An important question we'd like to answer is the following: Can we reach any other point $\mathbf{y} \in \mathcal{M}_K$ via a geodesic that starts at the arbitrary point $\mathbf{x} \in \mathcal{M}_K$? The answer is yes. First let us give a definition of the type of manifold we're considering here:

Definition 2.13 (Geodesically Complete Manifolds) *A connected manifold $\mathcal{M}$ is called **geodesically complete** if every geodesic $\mathbf{x}(\tau)$ is defined for all $\tau \in \mathbb{R}$.*

Geodesically complete manifolds have the following important property:

Theorem 2.14 *Any two points on a geodesically complete manifold $\mathcal{M}$ can be joined by a minimal length geodesic.*

We would like to be able to have a useful criterion for deciding if a manifold $\mathcal{M}$ is geodesically complete. The following theorem is particularly relevant for us:

Theorem 2.15 *Let $\mathcal{M}$ be a regular level set that consists of the disjoint connected pieces $\mathcal{M}_i$, $i = 1, \ldots, N$. Then each $\mathcal{M}_i$ is geodesically complete.*

Now you may be asking yourself: Who cares? Well, consider starting our optimization problem in (2.9) with some point $\mathbf{y} \in \mathcal{M}_K$. Since we have a nice equation for the geodesics of $\mathcal{M}_K$ given by (2.48), it makes sense to move over $\mathcal{M}_K$ using geodesics. But imagine the problems that could occur if we couldn't reach some other point $\mathbf{x} \in \mathcal{M}_K$ via a geodesic starting at $\mathbf{y}$. Since we're guaranteed this can't happen, our lives just became significantly easier. Further, the geodesic completeness of the $\mathcal{M}_i$ allows us to extend the idea of basic variables in a rather significant way.

2.14 Basic Variables and Pullbacks

Let $\mathbf{g} : \mathbb{R}^{n+m} \to \mathbb{R}^n$ be an onto function. Suppose we have another function $F : \mathbb{R}^n \to \mathbb{R}$ and consider the composition:

$$(\mathbf{g}^* F)(\mathbf{x}) \quad = \quad (F \circ \mathbf{g})(\mathbf{x}), \tag{2.52}$$

where $\mathbf{x} \in \mathbb{R}^{n+m}$. Then $\mathbf{g}^* F : \mathbb{R}^{n+m} \to \mathbb{R}$ is called the *pullback* of $F(\mathbf{y})$, $\mathbf{y} \in \mathbb{R}^n$, by $\mathbf{g}(\mathbf{x})$, see Figure 6.

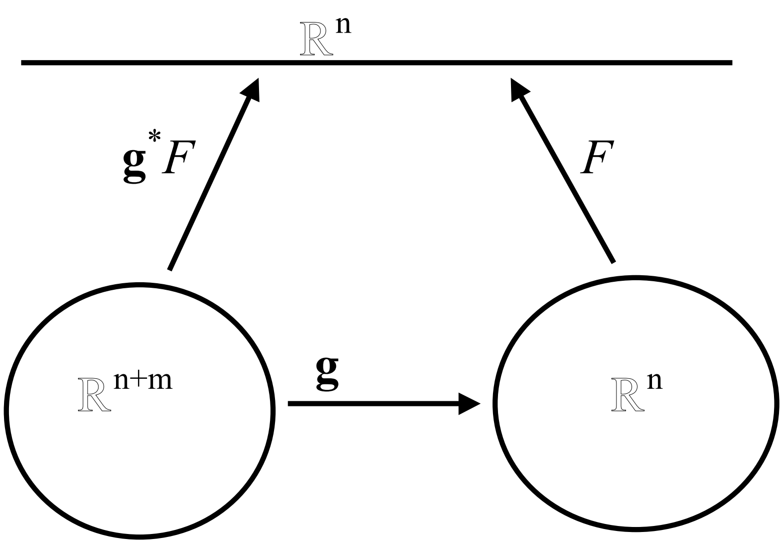

Figure 6: The pullback of F from $\mathbb{R}^n$ to $\mathbb{R}^{n+m}$ by the mapping $\mathbf{g} : \mathbb{R}^{n+m} \to \mathbb{R}^n$.

In the above definition we can replace $\mathbb{R}^n$ with an n-dimensional manifold $\mathcal{M}$. Then $\mathbf{g}^* F$ is the pullback of $F(\mathbf{y})$ from $\mathcal{M}$ to $\mathbb{R}^{n+m}$. Later we will use a pullback procedure like this where $m = 0$. Our manifold $\mathcal{M}$ will be implicitly defined by our equality constraints while our $\mathbb{R}^n$ will be some tangent space $\mathcal{T}_{\mathbf{x}}\mathcal{M}$.

The idea of *basic* and *nonbasic variables* is to express some of our variables (the basic ones) in terms of the remaining variables (the nonbasic ones). An example of this is provided by the implicit function theorem in Theorem 2.1. There we were able to locally express n of our variables in terms of m other variables by using the functions $\psi^i(\mathbf{y}), i = 1, \ldots, n$. This ensured that we would always remain on the level set defining our feasible region.

With a geodesically complete manifold $\mathcal{M}$, we can extend this idea globally over $\mathcal{M}$. By definition, every geodesic $\mathbf{x}(\tau)$ is defined for all $\tau \in \mathbb{R}$. Now, the geodesic $\mathbf{x}(\tau)$ is found by solving the equation (2.48). In order to solve this equation we need two things: 1) An initial starting point $\mathbf{y} \in \mathcal{M}$, and; 2) An initial velocity vector $\mathbf{w} = \dot{\mathbf{x}}(0) \in \mathcal{T}_{\mathbf{y}}\mathcal{M}$. Since our manifold is geodesically complete, this implies that we can also find a new geodesic $\widehat{\mathbf{x}}(\tau)$ where $\widehat{\mathbf{x}}(0) = \mathbf{y}$ and $\dot{\widehat{\mathbf{x}}}(\tau) = \alpha\mathbf{w}$ for any $\alpha \in \mathbb{R}$. Continuing this line of reasoning, we see that all of $\mathcal{T}_{\mathbf{y}}\mathcal{M}$ can be mapped onto $\mathcal{M}$ via geodesics. This mapping is called the *exponential mapping* and is denoted by

$$\mathrm{Exp}_{\mathbf{y}} : \mathcal{T}_{\mathbf{y}}\mathcal{M} \quad \rightarrow \quad \mathcal{M}. \tag{2.53}$$

For every $\mathbf{w} \in \mathcal{T}_{\mathbf{y}}\mathcal{M}$, by definition $\mathrm{Exp}_{\mathbf{y}}(\tau\mathbf{w}) = \mathbf{x}(\tau)$ where $\mathbf{x}(\tau)$ is the geodesic satisfying $\mathbf{x}(0) = \mathbf{y}$ and $\dot{\mathbf{x}}(0) = \mathbf{w}$ and, $\mathrm{Exp}_{\mathbf{y}}(\mathbf{w}) = \mathbf{x}(1)$.

So what we have done is expressed all of our $(n+m)$ variables in terms of the fixed m variables of $\mathcal{T}_{\mathbf{y}}\mathcal{M}$, where the point $\mathbf{y} \in \mathcal{M}$ was arbitrarily chosen. Notice that the $\mathrm{Exp}_{\mathbf{y}}$ mapping is onto but it's not bijective. So, for a point $\mathbf{x} \in \mathcal{M}$ we will generally have multiple coordinates assigned to it.

We can do something very special with the $\mathrm{Exp}_{\mathbf{y}}$ map. Consider a tangent vector $\mathbf{w} \in \mathcal{T}_{\mathbf{y}}\mathcal{M}$. Then $\mathbf{z} = \mathrm{Exp}_{\mathbf{y}}(\mathbf{w}) \in \mathcal{M}$ and, we can evaluate $f(\mathbf{x})$ at $\mathbf{z}$. Now assign the value $f(\mathbf{z})$ to $\mathbf{w}$. What we have just done is pulled-back the function $f(\mathbf{x})$ from $\mathcal{M}$ to $\mathcal{T}_{\mathbf{y}}\mathcal{M}$ by using the composite mapping $(f \circ \mathrm{Exp}_{\mathbf{y}})(\mathbf{w})$ defined on $\mathcal{T}_{\mathbf{y}}\mathcal{M}$. It can be shown that if f is $\mathcal{C}^2$, then $(f \circ \mathrm{Exp}_{\mathbf{y}})$ is also $\mathcal{C}^2$.

Now for the punch line. Since $\mathcal{T}_{\mathbf{y}}\mathcal{M} \cong \mathbb{R}^m$, using the pullback operation we can restate our optimization problem in (2.9) as the equivalent problem:

$$\min_{\mathbf{w} \in \mathbb{R}^m} : \qquad (f \circ \mathrm{Exp}_{\mathbf{y}})(\mathbf{w}). \tag{2.54}$$

on each piece $\mathcal{M}_i$ of our manifold $\mathcal{M}$. So what have we seen? Suppose our original problem included some equality constraints and inequality constraints as in (2.1). By using slack variables in (2.8), we were able to convert (2.1) to an equivalent problem (2.9) that only had equality constraints. Further, if our original problem defined a compact feasible set of dimension m, our new problem also defines a compact feasible set of dimension m. By using the pullback operation, we finally end up with the unconstrained optimization problem in (2.54) with only m variables, versus the original $(n + m)$ variables.

3 Extending direct search methods to equality constrained problems

The problem we will look at now is

$$\min_{\mathbf{x}\in\mathbb{R}^n} \quad f(\mathbf{x}) \tag{3.1a}$$

$$\text{subject to :} \quad \mathbf{g}(\mathbf{x}) = \mathbf{0} \tag{3.1b}$$

$$\mathbf{h}(\mathbf{x}) \leq \mathbf{0}, \tag{3.1c}$$

where $f : \mathbb{R}^n \to \mathbb{R}$ is Lipschitz, $\mathbf{g} : \mathbb{R}^n \to \mathbb{R}^m$ is $\mathcal{C}^2$ and, $\mathbf{h} : \mathbb{R}^n \to \mathbb{R}^l$ is Lipschitz. Let $\Omega = \{\mathbf{x}|\mathbf{h}(\mathbf{x}) \leq \mathbf{0}\}$. We will require that $\nabla\mathbf{g}(\mathbf{x})$ is full rank for every $\mathbf{x} \in \Omega$ that satisfies (3.1b). Our approach to dealing with (3.1) will be to employ some techniques from differential geometry to ensure that (3.1b) is always satisfied. The way we will do this is by treating $\mathbf{g}(\mathbf{x}) = \mathbf{0}$ as implicitly defining a manifold and, require that we always remain on this manifold as the algorithm proceeds. Since this effectively removes the equality constraints from further consideration, we then only need to concern ourselves with solving an inequality constrained optimization problem.

Note that $\nabla\mathbf{g}(\mathbf{x})$ only needs to be full rank on $\mathcal{N} = \mathcal{M} \cap \Omega$ above. The specific reduced dimensional problem we will have to solve is

$$\min_{\mathbf{w}\in\mathbb{R}^{(n-m)}} \quad \widehat{f}(\mathbf{w}) \tag{3.2a}$$

$$\text{subject to :} \quad \widehat{\mathbf{h}}(\mathbf{w}) \leq \mathbf{0}, \tag{3.2b}$$

where the hat denotes the functions $f(\mathbf{x})$ and $\mathbf{h}(\mathbf{x})$ after they are pulled-back to the tangent space $\mathcal{T}_{\mathbf{y}}\mathcal{M}$ and, $\mathbf{w} \in \mathcal{T}_{\mathbf{y}}\mathcal{M}$ is a tangent vector. There are certain advantages to this procedure. First, we have an implicit reduction in the dimensionality of the problem from n to $(n-m)$, the dimension of the manifold. Secondly, with some qualifications, the convergence results associated with the method chosen to solve (3.2) carry over *without modification* to (3.1).

Our procedure would seem to be particularly useful when employed in conjunction with the filter methods in [1, 12], the MADS algorithms in [2, 3], DIRECT [11] or, the frame methods in [7, 8, 21]. These methods can handle inequality constrained problems, i.e., they have convergence results for these types of problems, and, thus, are viable solution techniques for (3.2). By treating the equality constraints as a manifold and remaining on that manifold, we can then extend the above algorithms to (3.1). Additionally, it may be difficult to use any derivative information about $\widehat{f}(\mathbf{w})$ and $\widehat{\mathbf{h}}(\mathbf{w})$ after the pullback operation. So direct search methods may be the only methods that can be effectively employed to solve (3.2). ***The main point of our procedure is to extend direct search methods to problems that include equality constraints.*** If $f(\mathbf{x})$ and $\mathbf{h}(\mathbf{x})$ were smooth, and one had access to their derivative information, the pullback method advocated here would not likely be the best choice for an optimization procedure. But if one has no derivative information about $f(\mathbf{x})$ and/or $\mathbf{h}(\mathbf{x})$ because, e.g., they're only

Lipschitz, there's nothing to be lost, theoretically at least, by employing a pullback method.

All of the pieces are in place to state the procedure that will be used to extend direct search methods to (3.1). Let $\mathcal{M}$ be our $(n-m)$-dimensional Riemannian manifold implicitly defined by (3.1b). We have the following situation:

Choose a piece of $\mathcal{M}$: $\mathcal{M}$ may be composed of disjoint pieces. Pick a point $\mathbf{y}$ on one of these pieces, which we'll call $\widehat{\mathcal{M}}$.

Geodesically complete: $\widehat{\mathcal{M}}$ is geodesically complete, so we can reach any other point $\mathbf{x} \in \widehat{\mathcal{M}}$ via a geodesic starting at $\mathbf{y}$.

Pullback the functions: Because of the geodesic completeness of $\widehat{\mathcal{M}}$, we can pullback the objective and inequality constraints functions from $\widehat{\mathcal{M}}$ to $\mathcal{T}_{\mathbf{y}}\widehat{\mathcal{M}}$. This is done by solving the geodesic equation (2.48) for geodesics starting at $\mathbf{y}$ with an initial velocity $\mathbf{w} \in \mathcal{T}_{\mathbf{y}}\widehat{\mathcal{M}}$.

Minimize the pullback: Since $\mathcal{T}_{\mathbf{y}}\widehat{\mathcal{M}} \cong \mathbb{R}^{(n-m)}$, we're back in a setting that (some) direct search methods can deal with. Namely, minimizing a function subject to some inequality constraints.

Pushforward the solution: Let $\mathbf{w}^* \in \mathcal{T}_{\mathbf{y}}\widehat{\mathcal{M}}$ solve the pulled-back optimization problem above. This defines a geodesic $\mathbf{x}^*(\tau)$, where $\mathbf{x}^*(0) = \mathbf{y}$ and $\dot{\mathbf{x}}^*(0) = \mathbf{w}^*$. Then the solution to our original problem is given by $\mathbf{x}^*(1) \in \widehat{\mathcal{M}}$.

This outline effectively describes our entire method. Now for some details and comments.

To find the initial point $\mathbf{y} \in \mathcal{M}$, one could, for example, perform the following. First, minimize, e.g., $G(\mathbf{x}) = \mathbf{g}(\mathbf{x}) \cdot \mathbf{g}(\mathbf{x})$ to find a point $\mathbf{y}^* \in \mathcal{M}$ where $G(\mathbf{y}^*) = 0$. Now, $\mathbf{y}^*$ may or may not be a feasible point. If it is feasible we can let $\mathbf{y} = \mathbf{y}^*$. If it is an infeasible point, we have a few alternate ways of proceeding. One way is to let $\mathbf{y} = \mathbf{y}^*$ and use a technique like filter MADS [3], that can consider infeasible points, to minimize the pullback. Alternately, one could try to move over $\mathcal{M}$ to find an initial feasible point $\mathbf{y}$ by minimizing, e.g., $H(\mathbf{x}) = \max_i \max(h_i(\mathbf{x}), 0)$, where $i = 1, \ldots, l$ and $\mathbf{x} \in \mathcal{M}$, using the method outlined above. Note that $H(\mathbf{x})$ is Lipschitz because $\mathbf{h}(\mathbf{x})$ is.

Given our $\mathcal{M}$ implicitly defined by (3.1b), the inequality constraints in (3.1c) will confine us to allowed regions of $\mathcal{M}$. We will take (3.1c) as implicitly defining r full dimensional disjoint regions $V_i \subset \mathbb{R}^n$, $i = 1, \ldots, r$, with $V_i \bigcap \mathcal{M} \neq \emptyset$. Let $\mathcal{U}_i = V_i \bigcap \mathcal{M}$, where we will assume that the $\mathcal{U}_i$ have dimension $(n-m)$ everywhere. Then the $\mathcal{U}_i$ are the feasible sets to (3.1). Let us start at an initial point $\mathbf{y} \in \mathcal{U}_j$ for some $j \in [1, \ldots, r]$. Then it is a simple matter to apply, e.g., the pattern search method in [1] to $T_{\mathbf{y}}\mathcal{M}$ and use the pullback procedure described in Section 2.14. That is, we will use the $\mathrm{Exp}_{\mathbf{y}}$ mapping from $\mathcal{T}_{\mathbf{y}}\mathcal{M}$ to $\mathcal{M}$ given by the geodesic equation (2.48) to pullback the objective function values $f(\mathbf{x})$ and the inequality constraints $\mathbf{h}(\mathbf{x})$ from $\mathcal{M}$ to $\mathcal{T}_{\mathbf{y}}\mathcal{M}$. This results in the implicitly defined, dimensionally reduced optimization problem given by (3.2). In particular, the inequality constraints in (3.1c) will become 'black-box' constraints on the tangent vectors $\mathbf{w} \in T_{\mathbf{y}}\mathcal{M}$. It is the feasible tangent vectors that, when used as initial conditions in (2.48), will result in a point $\mathbf{x}(1) \in \mathcal{U}_j$.

Note two things. First, if we start in $\mathcal{U}_j$, we will always remain around $\mathcal{U}_j$, assuming the $\mathcal{U}_i$ are sufficiently separated. Depending on the algorithm used to solve (3.2), it may be possible to jump from $\mathcal{U}_j$ to $\mathcal{U}_l$, $l \in [1, \dots, r]$ and $l \neq j$, if they lie close enough together on the same connected piece of $\mathcal{M}$. This will not typically be the case and, since $\mathcal{M}$ consists of geodesically complete pieces, will not cause any problems even if it does occur. Secondly, given a tangent vector $\mathbf{w}$, we evaluate the corresponding geodesic at $\tau = 1$. Then the length that we travelled from $\mathbf{x}(0) = \mathbf{y}$ to $\mathbf{x}(1)$ will be the same as the Euclidean length of $\mathbf{w}$.

Why should the inequality constraints in (3.1c) be taken as 'black-box' constraints on $T_{\mathbf{y}}\mathcal{M}$? The reason is that it would often be very difficult to restate the constraints that define $\mathcal{U}_j$ directly in terms of constraints on $T_{\mathbf{y}}\mathcal{M}$. Instead, we will solve (2.48) for our geodesic $\mathbf{x}(\tau)$ and find what the constraint value is at our new point $\mathbf{x}(1)$ on the manifold. We will then assign this value, along with any objective function evaluation, to the corresponding tangent vector $\dot{\mathbf{x}}(0)$.

All of the convergence results concerning the method chosen to solve (3.2) carry over to our extension without modification. This is because (2.48) allows us to implicitly restate the problem of minimizing $f(\mathbf{x})$ in (3.1a) over $\mathcal{U}_j \subset \mathcal{M}$ as one of minimizing $\widehat{f}(\mathbf{w})$ over $T_{\mathbf{y}}\mathcal{M} \cong \mathbb{R}^{(n-m)}$, subject to our 'black-box' constraints. The implicitly defined functions $\widehat{f}(\mathbf{w})$ and $\widehat{\mathbf{h}}(\mathbf{w})$ are of the same smoothness class as $f(\mathbf{x})$ and $\mathbf{h}(\mathbf{x})$ since the $\mathrm{Exp}_{\mathbf{y}}$ mapping is smooth [15]. So if $f(\mathbf{x})$ and $\mathbf{h}(\mathbf{x})$ are Lipschitz, $\widehat{f}(\mathbf{w})$ and $\widehat{\mathbf{h}}(\mathbf{w})$ are also Lipschitz.

The general method is given by Procedure 3.1. The key feature for showing convergence results for our method is that the tangent space $\mathcal{T}_{\mathbf{y}}\mathcal{M}$ is fixed in Step 4. By fixing $\mathcal{T}_{\mathbf{y}}\mathcal{M}$ we are fixing our functions $\widehat{f}(\mathbf{w})$ and $\widehat{\mathbf{h}}(\mathbf{w})$ defined on $\mathcal{T}_{\mathbf{y}}\mathcal{M}$. This fixing of $\mathcal{T}_{\mathbf{y}}\mathcal{M}$ creates difficulties if we want to pullback derivative information about $\widehat{f}(\mathbf{w})$ and $\widehat{\mathbf{h}}(\mathbf{w})$. This is examined in Section 5. Since our main concern is extending direct search methods to equality constrained problems, this inability to pullback derivative information is not a particular drawback for us. Additionally, $\mathbf{z}^*$ has the same 'properties' as $\mathbf{w}^*$. By this we mean, e.g., that if $\mathbf{w}^*$ is guaranteed to be a local/global solution to (3.2), then $\mathbf{z}^*$ is guaranteed to be a local/global solution to (3.1).

The requirement that we remain in a single tangent space $\mathcal{T}_{\mathbf{y}}\mathcal{M}$ can be relaxed in a certain way. What we will do is allow ourselves to move around on $\mathcal{M}$ initially. This movement on $\mathcal{M}$ can be done by any procedure the user desires. This freedom is similar to the freedom available in the `SEARCH` step of LTMADS [2]. What we are looking for is a point $\mathbf{y} \in \mathcal{M}$ that is, or is suspected of being, close to the solution $\mathbf{z}^*$ of (3.1). This $\mathbf{y}$ can then be used in Procedure 3.1. The reason for this additional step is to reduce the cost of solving the geodesic equation since we will hopefully only need

Procedure 3.1 ***The General Method for Problem (3.1)***

1. *Let the equality constraints* $\mathbf{g}(\mathbf{x}) = \mathbf{0}$ *implicitly define a manifold* $\mathcal{M}$.
2. *Find an initial feasible point* $\mathbf{y} \in \mathcal{M}$.
3. *Use the geodesic formula (2.48) to pullback the inequality constraints* $\mathbf{h}(\mathbf{x}) \leq \mathbf{0}$ *and, the function* $f(\mathbf{x})$ *from* $\mathcal{M}$ *to* $T_{\mathbf{y}}\mathcal{M}$. *That is, given any tangent vector* $\mathbf{w} \in T_{\mathbf{y}}\mathcal{M}$ *that results in a geodesic* $\mathbf{x}(\tau)$ *such that* $\mathbf{x}(1) = \mathbf{z}$, *associate with* $\mathbf{w}$ *the function values* $\mathbf{h}(\mathbf{z})$ *and* $f(\mathbf{z})$. *Call these pulled-back functions* $\widehat{\mathbf{h}}(\mathbf{w})$ *and* $\widehat{f}(\mathbf{w})$.
4. *Use any desired procedure to solve the resulting reduced dimensional problem given by (3.2).*
5. *Let* $\mathbf{w}^*$ *be the solution to (3.2) found in Step 4. If* $\mathbf{x}^*(\tau)$ *is the geodesic corresponding to* $\mathbf{w}^*$ *via (2.48) and,* $\mathbf{z}^* = \mathbf{x}^*(1)$, *then* $\mathbf{z}^*$ *is the solution to (3.1).*

to look at points on $\mathcal{M}$ that are close to $\mathbf{y}$. It is up to the user to specify how to find such a $\mathbf{y}$ and, when to halt this initial search and enter Procedure 3.1.

One example of a potential way of doing this initial movement over $\mathcal{M}$ is by switching from $\mathcal{T}_{\mathbf{y}}\mathcal{M}$ to $\mathcal{T}_{\mathbf{x}}\mathcal{M}$ if the feasible tangent vector $\mathbf{w} \in \mathcal{T}_{\mathbf{y}}\mathcal{M}$ that corresponds to $\mathbf{x} \in \mathcal{M}$ has norm $\|\mathbf{w}\| > \epsilon$ for some user specified $\epsilon > 0$ and, $f(\mathbf{x}) < f(\mathbf{z})$ for all previously considered $\mathbf{z} \in \mathcal{M}$. Let us use LTMADS [2] in the tangent spaces. Assume that $f(\mathbf{x})$ has a unique global minimum $\mathbf{x}^* \in \mathcal{N} = \mathcal{M} \cap \Omega$ and, we eventually enter into a small enough neighborhood $\mathcal{U}_{\mathbf{x}^*} \subset \mathcal{N}$ of $\mathbf{x}^*$ at the point $\mathbf{y} \in \mathcal{U}_{\mathbf{x}^*}$. Then given some user specified integer $i > 0$, there will be a sequence of improving points $\mathbf{x}_\iota$, $\iota = 1, \ldots, i$, that correspond to feasible tangent vectors $\mathbf{w}_\iota \in \mathcal{T}_{\mathbf{y}}\mathcal{M}$ such that $\|\mathbf{w}_\iota\| \leq \epsilon$ for all $\iota = 1, \ldots, i$. We are guaranteed that this must happen when we enter $\mathcal{U}_{\mathbf{x}^*}$ because $\mathcal{N}$ is compact. When this occurs, we can enter Procedure 3.1 with $\mathbf{y} \in \mathcal{U}_{\mathbf{x}^*}$.

4 An illustrative example

Here we'll look at a relatively simple but illuminating example. Let us consider minimizing some function $f(\mathbf{x})$ over the upper half of $\mathcal{S}^2$, the unit two-dimensional sphere embedded in $\mathbb{R}^3$. So our optimization problem is

$$\min_{\mathbf{x} \in \mathbb{R}^3} \quad f(\mathbf{x}) \tag{4.1a}$$

$$\text{subject to :} \quad \mathbf{x}^T\mathbf{x} = 1 \tag{4.1b}$$

$$x^3 \geq 0, \tag{4.1c}$$

where $\mathbf{x} = [x^1, x^2, x^3]^T$. We will be able to explicitly state the problem that Procedure 3.1 solves implicitly. First we fix a point $\mathbf{y} \in \mathcal{S}^2$, which also fixes the tangent space $\mathcal{T}_\mathbf{y}\mathcal{M}$ that we will work on. Then the mapping

$$\mathrm{Exp}_\mathbf{y} : \mathcal{T}_\mathbf{y}\mathcal{M} \quad \to \quad \mathcal{S}^2, \tag{4.2}$$

is found. This mapping is given by $\mathbf{x}(\mathbf{w}) = \mathrm{Exp}_\mathbf{y}(\mathbf{w}) \in \mathcal{S}^2$, where $\mathbf{w} \in \mathcal{T}_\mathbf{y}\mathcal{M}$. That is, we evaluate the geodesic $\mathbf{x}(\tau) = \mathrm{Exp}_\mathbf{y}(\tau\mathbf{w})$ at $\tau = 1$, where $\mathbf{x}(0) = \mathbf{y}$ and $\dot{\mathbf{x}}(0) = \mathbf{w}$. Then (4.1) can be restated in terms of the $\mathbf{w} \in \mathcal{T}_\mathbf{y}\mathcal{M}$.

First we need to place coordinates on $\mathcal{T}_\mathbf{y}\mathcal{M}$. Fix a point $\mathbf{y} \in \mathcal{S}^2$. The tangent plane $\mathcal{T}_\mathbf{y}\mathcal{S}^2$ is given by $\mathrm{null}(\nabla(\mathbf{x}^T\mathbf{x}))$ evaluated at $\mathbf{y}$. So

$$\mathcal{T}_\mathbf{y}\mathcal{M} \quad = \quad \mathbf{y}^\perp, \tag{4.3}$$

the orthogonal complement to $\mathbf{y}$. This is a 2-plane in $\mathbb{R}^3$. When we say that $\mathcal{T}_\mathbf{y}\mathcal{S}^2$ can be imagined as a copy of $\mathbb{R}^2$ translated to the base point $\mathbf{y} \in \mathcal{M}$ what we mean is that there is an affine transformation $\mathbf{A} : \mathbb{R}^2 \to \mathbb{R}^3$ given by

$$\mathbf{A}(\mathbf{z}) \quad = \quad A\mathbf{z} + \mathbf{y}, \tag{4.4}$$

where the columns of $A \in \mathbb{R}^{3\times 2}$ form a basis for the 2-plane given by (4.3). Then $A\mathbf{z} \in \mathcal{T}_\mathbf{y}\mathcal{S}^2$ and the $\mathbf{z}$ give the coordinates on $\mathcal{T}_\mathbf{y}\mathcal{S}^2$. Note that $\mathbf{A}(\mathbf{0}) = \mathbf{y}$.

Now for the $\mathrm{Exp}_\mathbf{y}$ mapping, i.e., the geodesic equation. The equations for the Christoffel symbols associated with (4.1b) are, see (2.37),

$$\Gamma^k_{ij}(\tau) \quad = \quad x^k(\tau) I_{ij}. \tag{4.5}$$

Then (2.48) is given by

$$\ddot{x}^k(\tau) \quad = \quad -x^k(\tau)[\dot{\mathbf{x}}(\tau) \cdot \dot{\mathbf{x}}(\tau)]. \tag{4.6}$$

Remembering that $\|\dot{\mathbf{x}}(0)\| = \|\dot{\mathbf{x}}(\tau)\|$, we have

$$x^k(\tau) \quad = \quad y^k \cos(\omega\tau) + \frac{w^k}{\omega}\sin(\omega\tau), \tag{4.7}$$

where $\mathbf{x}(0) = \mathbf{y} = [y^1, y^2, y^3]^T$, $\dot{\mathbf{x}}(0) = \mathbf{w} = [w^1, w^2, w^3]^T$, $\omega = \|\mathbf{w}\|$, $\mathbf{y} \in \mathcal{S}^2$ and $\mathbf{w} = A\mathbf{z} \in \mathcal{T}_\mathbf{y}\mathcal{S}^2 \cong \mathbb{R}^2$. Letting $A = [a^k_{\ i}]$, where k is the row index, and $\mathbf{z} = [z^1, z^2]^T$, we can rewrite (4.7) as

$$x^k(\tau; \mathbf{z}) \quad = \quad y^k \cos(\omega\tau) + \sum_{i=1}^{2} \frac{a^k_{\ i} z^i}{\omega}\sin(\omega\tau), \tag{4.8}$$

or, in more formal notation, $\mathbf{x}(\tau; \mathbf{z}) = \mathrm{Exp}_\mathbf{y}(\tau A\mathbf{z})$. Evaluating the mapping in (4.8) at $\tau = 1$ gives us the explicit relationship $\mathbf{x}(1; \mathbf{z}) = \mathbf{x}(\mathbf{z}) = \mathrm{Exp}_\mathbf{y}(A\mathbf{z})$ or

$$x^k(\mathbf{z}) \quad = \quad y^k \cos(\omega) + \sum_{i=1}^{2} \frac{a^k_{\ i} z^i}{\omega}\sin(\omega), \tag{4.9}$$

So we can restate (4.1) in the following equivalent form:

$$\min_{\mathbf{z}\in\mathbb{R}^2} \quad \widehat{f}(\mathbf{z}) \tag{4.10a}$$

$$x^3(\mathbf{z}) \geq 0, \tag{4.10b}$$

where $\widehat{f}(\mathbf{z}) = f(\mathbf{x}(\mathbf{z}))$. That is, we have pulled-back (4.1a) and (4.1c) from $\mathcal{S}^2 \subset \mathbb{R}^3$ to $\mathcal{T}_{\mathbf{y}}\mathcal{M} \cong \mathbb{R}^2$.

Our method in Procedure 3.1 implicitly works with the $\mathbf{z}$ variables, which are related to the tangent vectors $\mathbf{w}$ via the relationship $\mathbf{w} = A\mathbf{z}$. That is, the $\mathbf{z}$ are the coordinates for $\mathcal{T}_{\mathbf{y}}\mathcal{M}$ once we fix a basis given by the columns of A. Also, in this example a closed form solution was given for the $\mathrm{Exp}_{\mathbf{y}}$ mapping in terms of the $\mathbf{z}$. Whether a problem allows this or not has no theoretical impact on our method. Practically though, a closed form solution is useful since one does not then need to solve the geodesic equation (2.48) numerically.

If $f(\mathbf{x})$ in (4.1a) is Lipschitz, then so is $\widehat{f}(\mathbf{z})$ in (4.10a). So now an LTMADS [2] could be done over the $\mathbf{z} \in \mathbb{R}^2$, for example. All of the convergence results for LTMADS carry over to our problem because we are simply minimizing some Lipschitz function subject to some inequality constraints. If $\mathbf{z}$ is a trial point used by LTMADS in $\mathbb{R}^2$, we map this to a point $\mathbf{x}(\mathbf{z}) \in \mathcal{S}^2$ by $\mathrm{Exp}_{\mathbf{y}}(A\mathbf{z})$. Then the values $f(\mathbf{x}(\mathbf{z}))$ and $x^3(\mathbf{z})$ are assigned to $\mathbf{z}$. If $x^3(\mathbf{z})$ violates the constraint in (4.1c), i.e., $\mathbf{z}$ is not a feasible point, we assign the value $\widehat{f}(\mathbf{z}) = \infty$ to $\mathbf{z}$. The fact that the objective function $\widehat{f}(\mathbf{z})$ and the inequality constraints function $\widehat{\mathbf{h}}(\mathbf{z})$ would generally be 'black-boxes' creates no difficulty if one employs a direct search method on the $\mathbf{z}$ variables.

LTMADS is not the only direct search procedure that can be done in $\mathcal{T}_{\mathbf{y}}\mathcal{M}$. If a global solution is desired then one may want to employ, e.g., DIRECT [11] to attempt to find it. If we are willing to allow the inequality constraints to be violated as the algorithm proceeds, a filter version of LTMADS could be employed [3].

5 Relation to the generalized reduced gradient method

Remaining on or near the feasible set is not a novel idea. For example, the generalized reduced gradient method [4] and the filter method [12] both attempt this. What is novel is that (2.48) provides a way for remaining on a manifold $\mathcal{M}$ by solving for the geodesics on $\mathcal{M}$. ***This removes the burden of having to satisfy equality constraints from the optimization method and places it on the particular numerical solver for the ODE system in (2.48).*** It is exactly this shifting of the burden that allows us to trivially extend (at least theoretically) direct search methods to $\mathcal{C}^2$ Riemannian manifolds.

There is nothing particularly special about the $\mathrm{Exp}_{\mathbf{x}}$ mapping. An alternate technique would be to use the implicit function theorem locally to do the pullback procedure on pieces of $\mathcal{M}$. Since $\mathcal{N} = \mathcal{M} \cap \Omega$ is compact, we can perform the implicit function theorem at a finite number of points $\mathbf{x}_i \in \mathcal{M}$, $i = 1, \ldots, \imath$, to construct our neighborhoods $\mathcal{U}_i$ that cover $\mathcal{N}$, that is $\mathcal{N} \subset \bigcup_i \mathcal{U}_i$. Then we can do direct searches on the corresponding tangent spaces $\mathcal{T}_{\mathbf{x}_i}\mathcal{M}$. The algorithm is theoretically more complicated and, is covered separately in [9]. It does allow us to extend the same general idea introduced here to the case when $\mathbf{g}(\mathbf{x})$ in (2.1b) is only Lipschitz.

Using the implicit function theorem technique would make the procedure strongly reminiscent of the generalized reduced gradient (GRG) method [4]. GRG does not use a pullback procedure however. Instead, it minimizes $f(\mathbf{x})$ on a local linear model of $\mathcal{M}$, namely $\mathcal{T}_{\mathbf{x}}\mathcal{M}$, and then attempts to project back onto $\mathcal{M}$ after this minimization. It could do this projection by numerically implementing the implicit function theorem as in [23] (it generally doesn't). The important point is that the projection happens after the original minimization of $f(\mathbf{x})$ on $\mathcal{T}_{\mathbf{x}}\mathcal{M}$.

Let $\mathbf{w} \in \mathcal{T}_{\mathbf{x}}\mathcal{M}$ minimize $f(\mathbf{x})$ on the tangent space. If $\mathbf{w}$ projects to the point $\mathbf{y} \in \mathcal{M}$, GRG (eventually) assumes that $f(\mathbf{x}+\mathbf{w}) \approx f(\mathbf{y})$. Whether this is true or not depends on the behavior of $\mathcal{M}$ and $f(\mathbf{y})$ around $\mathbf{x}$. Regardless, the projection creates difficulties for convergence arguments.

An intrinsic version of GRG could minimize $f(\mathbf{x})$ directly on $\mathcal{M}$ using (2.48), avoiding the need for a final projection [10, intrinsic Newton's method] (see [26] also). Alternately, GRG could use the method in [23] to project each point $\mathbf{w}$ considered in $\mathcal{T}_{\mathbf{y}}\mathcal{M}$ to its corresponding point $\mathbf{x} \in \mathcal{M}$. Then it would evaluate $f(\mathbf{x})$ and $\mathbf{h}(\mathbf{x})$ and assign these values to $\mathbf{w} \in \mathcal{T}_{\mathbf{y}}\mathcal{M}$. This technique is more closely related to the original GRG method.

The method in Section 3 would be the analog of intrinsic GRG using (2.48) when $f(\mathbf{x})$ is only Lipschitz. The crucial difference between Procedure 3.1 and an intrinsic GRG is that in Procedure 3.1 we (eventually) remain in a fixed tangent plane $\mathcal{T}_{\mathbf{y}}\mathcal{M}$ while an intrinsic GRG would continue to move from one tangent plane to another as the algorithm proceeds. We needed to fix $\mathcal{T}_{\mathbf{y}}\mathcal{M}$ in order to derive the convergence results for Procedure 3.1. This has implications for the derivative information available if $f(\mathbf{x})$ were $\mathcal{C}^1$. Of course if this were the case and $\nabla f(\mathbf{x})$ were available, an intrinsic GRG would be a better choice for solving (2.1). Let us examine this in more detail though since it highlights some of the differences between Procedure 3.1 and GRG.

Note that in Step 4 of Procedure 3.1, it may be rather difficult to use any method other than a direct search one. The reason for this is that we will not normally have explicit formulas for $\widehat{f}(\mathbf{w})$ or $\widehat{\mathbf{h}}(\mathbf{w})$. So typically the availability of derivative information will, at best, be limited and/or expensive to calculate. This situation can be mitigated if we can reasonably find or estimate $\nabla_{\mathbf{w}}\mathbf{x}(\tau)|_{\tau=1}$, where $\mathbf{x}(\tau)$ is the geodesic found via (2.48) with initial conditions $\mathbf{x}(0) = \mathbf{y}$ and $\dot{\mathbf{x}}(0) = \mathbf{w}$. That is, $\mathbf{x}(\tau) = \mathrm{Exp}_{\mathbf{y}}(\tau\mathbf{w})$. Then we can find or estimate, e.g.,

$$\begin{aligned} \nabla_{\mathbf{w}}\widehat{f}(\mathbf{w}) &= \nabla_{\mathbf{w}} f(\mathbf{x}(1)) \\ &= [\nabla_{\mathbf{x}} f(\mathbf{x}(1))] \cdot [\nabla_{\mathbf{w}}\mathbf{x}(1)]. \end{aligned} \tag{5.1}$$

For the direct search methods, one could employ the simplex gradient method in [5] for estimating the gradient.

From (2.48) we see that finding $\nabla_{\mathbf{w}} \mathrm{Exp}_{\mathbf{x}}(\mathbf{w})$, $\mathbf{w} \in \mathcal{T}_{\mathbf{y}}\mathcal{M}$, can be an extremely hard proposition. One piece of derivative information that is easily pulled-back is the directional derivative of $f(\mathbf{x}(1))$ in the direction $\dot{\mathbf{x}}(1)$, where $\mathbf{x}(\tau)$ is a geodesic. This tells us how $f(\mathbf{x}(\tau))$ changes along the geodesic $\mathbf{x}(\tau)$ if we continued along that geodesic an infinitesimal time step beyond $\tau = 1$. This corresponds to using an initial velocity $(1+\delta)\mathbf{w} \in \mathcal{T}_{\mathbf{y}}\mathcal{M}$ and letting $\delta \downarrow 0$. So, we can easily assign to $\mathbf{w} \in \mathcal{T}_{\mathbf{y}}\mathcal{M}$ the derivative information

$$\frac{d}{d\delta}\widehat{f}((1+\delta)\mathbf{w})\bigg|_{\delta=0} \quad = \quad [\nabla_{\mathbf{x}} f(\mathbf{x}(1))] \cdot [\dot{\mathbf{x}}(1)]. \tag{5.2}$$

This is just the first covariant derivative of $f(\mathbf{x}(1))$ in the direction $\dot{\mathbf{x}}(1)$. We can do the same thing for the second covariant derivative. Note that we are limited to the single direction $\dot{\mathbf{x}}(1)$ when we evaluate these derivatives at $\mathbf{x}(1)$. Similar comments hold for $\widehat{\mathbf{h}}(\mathbf{w})$.

The situation is quite different for the base point $\mathbf{y}$ of $\mathcal{T}_{\mathbf{y}}\mathcal{M}$. Now we can find the covariant derivatives of $f(\mathbf{x})$ around $\mathbf{y}$ for any direction $\mathbf{w} \in \mathcal{T}_{\mathbf{y}}\mathcal{M}$. These are what GRG uses. They would always be available to GRG because the method models $f(\mathbf{x})$ on $\mathcal{T}_{\mathbf{y}}\mathcal{M}$ where $\mathbf{y}$ is the current iterate.

6 Discussion

We've shown how to deal with $\mathcal{C}^2$ equality constraints in optimization problems by treating them as implicitly defining a Riemannian manifold. This required knowledge of both the Jacobian and Hessian of the equality constraints in order to calculate geodesics on the manifold. By using this framework, the equality constraints can be removed from the original problem. We then have an implicitly defined optimization problem that only has inequality constraints. Then methods like those in [1, 2, 3, 7, 8, 11, 21] can be extended to optimization problems that have both inequality and equality constraints. Additionally, the dimensionality of the original problem is reduced to the dimensionality of the manifold.

The class of problems where this method would most usefully be employed are when the objective function and inequality constraints are only Lipschitz continuous. In fact, both of these functions can be 'black-box'. The procedure is especially effective when the equality constraints define some Lie group [20] since closed form solutions for the geodesic would then typically be available. Lie groups are often met in practice [10, 22], so even if our algorithm were restricted to just that class of problems it would likely be of practical use.

We haven't looked at how to actually solve the quasi-linear system of ODEs for the geodesics given by (2.48). This is obviously an important topic for practically implementing Procedure 3.1 but we feel it is somewhat out of the scope of this article. It is unlikely that any single numerical solution method for (2.48) will always work for every problem. Some numerical procedures are covered in [26].

As in [6, 18], we have avoided any numerical examples of our method. From [18]:

> We agree with the perspective of the authors in [6]:
>
> > We have deliberately not included the results of numerical testing as, in our view, the construction of appropriate software is by no means trivial and we wish to make a thorough job of it. We will report on our numerical experience in due course.
>
> This caution is particularly apt in view of the sort of problems to which pattern search is typically applied.

Since (2.48) must be solved repeatedly in our procedure, an appropriate ODE solver must be chosen that is at once fast and accurate. But this only takes care of the pullback procedure. Now a correct direct search method must be chosen and implemented for the problem at hand. If one understands the example in Section 4, then it should be (relatively) straightforward to implement Procedure 3.1. We plan on testing the method presented on some 'real-world' problems soon and will publish our experiences as they become available.

7 Acknowledgments

The author would like to thank John E. Dennis, JR. (Rice University) for his very helpful discussions. Additional appreciations go to Los Alamos National Laboratory for their postdoctoral program and, the Director's of Central Intelligence Postdoctoral Fellowship program for providing prior research funding, out of which this paper grew. Finally, the author would like to thank Charles Audet and the anonymous reviewers, who's helpful comments made this a much better paper.

References

[1] C. Audet and J. E. Dennis, JR. A pattern search filter method for nonlinear programming without derivatives. *SIAM Journal on Optimization*, 14:980–1010, 2004.

[2] C. Audet and J. E. Dennis, JR. Mesh adaptive direct search algorithms for constrained optimization. *SIAM Journal on Optimization*, 17:188–217, 2006.

[3] C. Audet and J. E. Dennis, JR. Derivative-free nonlinear programming by filters and mesh adaptive direct searches. Unpublished, 2007.

[4] M. S. Bazaraa, H. D. Sherali, and C. M. Shetty. *Nonlinear Programming: Theory and Algorithms*. Wiley, 2nd edition, 1993.

[5] D. M. Bortz and C. T. Kelley. The simplex gradient and noisy optimization problems. In *Computational Methods in Optimal Design and Control*, pages 77–90, 1998.

[6] A. R. Conn, N. I. M. Gould, and P. L. Toint. A globally convergent augmented Lagrangian algorithm for optimization with general constraints and simple bounds. *SIAM Journal on Numerical Analysis*, 28:545–572, 1991.

[7] I. D. Coope and C. J. Price. Frame based methods for unconstrained optimization. *Journal of Optimization Theory and Applications*, 107:261–274, 2000.

[8] J. E. Dennis, JR., C. J. Price, and I. D. Coope. Direct search methods for nonlinearly constrained optimization using filters and frames. *Optimization and Engineering*, 5:123–144, 2004.

[9] D. W. Dreisigmeyer. Direct search methods over Lipschitz manifolds. *Submitted to SIOPT*.

[10] A. Edelman, T. A. Arias, and S. T. Smith. The geometry of algorithms with orthogonality constraints. *SIAM Journal on Matrix Analysis and Applications*, 20:303–353, 1998.

[11] D. E. Finkel and C. T. Kelley. Convergence analysis of the DIRECT algorithm. Technical report CRSC-TR04-28, Center for Research in Scientific Computation, North Carolina State University, 2004.

[12] R. Fletcher and S. Leyffer. Nonlinear programming without a penalty function. *Mathematical Programming, Series A*, 91:239–269, 2002.

[13] G. H. Golub and C. F. van Loan. *Matrix Computations*. Johns Hopkins, 3rd edition, 1996.

[14] M. W. Hirsch. *Differential Topology*. Springer, 1976.

[15] J. M. Lee. *Riemannian manifolds*. Springer, 1997.

[16] J. M. Lee. *Introduction to Topological Manifolds*. Springer, 2000.

[17] J. M. Lee. *Introduction to Smooth Manifolds*. Springer, 2003.

[18] R. M. Lewis and V. Torczon. A globally convergent augmented Lagrangian pattern search algorithm for optimization with general constraints and simple bounds. *SIAM Journal on Optimization*, 12:1075–1089, 2002.

[19] M. Nakahara. *Geometry, Topology and Physics*. Taylor and Francis, 2nd edition, 2003.

[20] P. J. Olver. *Applications of Lie Groups to Differential Equations*. Springer, 2nd edition, 1993.

[21] C. J. Price and I. D. Coope. Frames and grids in unconstrained and linearly constrained optimization: a nonsmooth approach. *SIAM Journal on Optimization*, 14:415–438, 2003.

[22] I. U. Rahman, I. Drori, V. C. Stodden, D. L. Donoho, and P. Schroeder. Multiscale representations for manifold-valued data. *Multiscale Modeling and Simulation*, 4:1201–1232, 2005.

[23] W. C. Rheinboldt. On the computations of multi-dimensional solution manifolds of parametrized equations. *Numerische Mathematik*, 53:165–181, 1988.

[24] B. Schutz. *Geometrical Methods of Mathematical Physics*. Cambridge, 1980.

[25] M. Spivak. *A Comprehensive Introduction to Differential Geometry*, volume 1. Publish or Perish, 3rd edition, 2005.

[26] C. Udriste. *Convex Functions and Optimization Methods on Riemannian Manifolds*. Kluwer, 1994.