Geodesically Convex Optimization and Operator Scaling

Haydn Gwyn, Naveen Durvasula, Frederic Wang

Spring 2021

1 Introduction

In the last several decades, convex functions have become one of the most well-studied concepts in optimization theory. We can efficiently minimize a convex function with basic methods such as gradient descent and line search. Recently, however, optimization over non-convex functions has gained quite a bit of interest, and the nice algorithms of convex optimization tend to offer little more than a local optimum when we pass to the non-convex setting. Fortunately, not all is lost: in many cases, an optimization problem which is apparently non-convex can be realized as *geodesically convex* on a certain Riemannian manifold, which is a topological space with local resemblance to Euclidean space (for example, the unit sphere or the torus) and a clever choice of geometry. Rather than defining convexity in terms of straight lines (which need not be present on a Riemannian manifold), we can instead consider geodesics: the shortest paths between points on the manifold.

Geodesic convexity has emerged as a powerful technique for efficiently solving many non-convex optimization problems. In the present paper, we offer the necessary background to acquaint oneself with geodesic convexity, and describe methods used to solve geodesically convex optimization problems. More specifically, we outline first-order methods for geodesic convexity and their convergence guarantees, and give a more detailed description of a second-order approach as it applies to the operator scaling problem.

2 A primer on Riemannian manifolds

In this section, we give an introduction to the theory of smooth and Riemannian manifolds upon which we will base our treatment of geodesic convexity. Someone with background at the level of Berkeley's Math 140 course can safely skip this section without losing any necessary intuition. We start our discussion assuming that the reader is familiar with metric spaces and topological spaces at the level of Berkeley's Math 104. The unacquainted reader should reference the appendix section A.1 for prerequisite definitions and examples.

Metric spaces possess a useful property pertaining to the separation of points:

Definition 2.1 (Hausdorff space)**.** A topological space $(S,\mathcal{U})$ is **Hausdorff** if for any points $p, q \in S$, there exist open sets $U_p, U_q \in \mathcal{U}$ such that $p \in U_p$, $q \in U_q$ and $U_p \cap U_q = \emptyset$.

Theorem 2.2. *Every metric space is Hausdorff.*

Proof. Let (S, d) be a metric space and let $p, q \in U$ with $p \neq q$. Then $D := d(p, q) > 0$, so define $U_p = B_{D/3}(p)$ and $U_q = B_{D/3}(q)$. Then we have $p \in U_p$, $q \in U_q$, and $U_p, U_q \in \mathcal{U}_{(S,d)}$. The triangle inequality implies that $U_p \cap U_q = \emptyset$, proving that (S, d) is Hausdorff. □

In particular, both $\mathbb{R}$ and $\mathbb{R}^n$, being metric spaces, are Hausdorff. Recall that in Examples A.2 and A.3 and Definition A.4, we constructed a topology on a set as a collection of unions of more "basic" sets. This lends itself to another definition.

Definition 2.3 (Base)**.** Let $(S,\mathcal{U})$ be a topological space. A collection $\mathcal{B} \subseteq \mathcal{U}$ is called a *base* for $\mathcal{U}$ if for any $U \in \mathcal{U}$, there exists a subcollection $\mathcal{A} \subseteq \mathcal{B}$ such that

$$\bigcup_{A \in \mathcal{A}} A = U$$

The collection of open balls provides an obvious base for the topology of a metric space. As a more specific example, the open intervals comprise a base for the topology of $\mathbb{R}$.

A base, in some sense, captures all of the information about a topology, since we can recover the entire topology from a base. We might conclude, then, that the smaller the base for a topology, the simpler the topology. This motivates the following definition.

Definition 2.4. A topological space $(S,\mathcal{U})$ is **second countable** (satisfying the second axiom of countability) if $\mathcal{U}$ has a countable base.

Theorem 2.5. *$\mathbb{R}^n$ is second countable.*

Proof. The countable base for $\mathcal{U}_{\mathbb{R}^n}$ is

$$\mathcal{B} = \{B_r((q_1, \dots, q_n)) \mid r, q_1, \dots, q_n \in \mathbb{Q}\}$$

This collection is in bijection with $\mathbb{Q}^{n+1}$, which is countable. We leave it to the reader to check that this is indeed a base for $\mathcal{U}_{\mathbb{R}^n}$ (recall that $\mathbb{Q}$ is dense in $\mathbb{R}$). $\square$

We have seen that $\mathbb{R}^n$ is both Hausdorff and second countable. Indeed $\mathbb{R}^n$ possesses all sorts of niceties that make it a nice object of study. We now introduce the concept of a topological manifold, which is a topological space that, despite not being homeomorphic to $\mathbb{R}^n$, manages to retain all of the properties locally.

Definition 2.6. A topological space $(M,\mathcal{U})$ is an n**-dimensional topological manifold** if

(i) M is Hausdorff.

(ii) M is second countable.

(iii) M is **locally Euclidean**, meaning that for any $p \in M$, there is an open set $U \ni p$ such that U is homeomorphic to $\mathbb{R}^n$.

A set U as described in (iii) is called a **chart**. We often refer to charts as a pair (U, φ), where $\varphi : U \to \mathbb{R}^n$ is a homeomorphism.

It follows immediately from Theorems 2.2 and 2.5 that $\mathbb{R}^n$ is an n-dimensional topological manifold, as for every point $p \in \mathbb{R}^n$, we have the global chart $\mathbb{R}^n \in \mathcal{U}_{\mathbb{R}^n}$, which is obviously homeomorphic to itself. For a simple example of a non-Euclidean topological manifold, see Example A.15.

The concept of a topological manifold greatly expands our realm of study. With that being said, for what follows, we will need a bit more structure. Recall that a map $\mathbb{R}^m \to \mathbb{R}^n$ is called *smooth* if we can take arbitrary partial derivatives of each of its n component functions.

Let M be a topological manifold, and let (U, φ) and (V, ψ) be charts on M with $U \cap V \neq \emptyset$. Then both φ and ψ restrict to homeomorphisms between $U \cap V$ and some open subset of $\mathbb{R}^n$. The map $\psi \circ \varphi^{-1}$ is called a **transition map**.

Definition 2.7. Let M be a topological manifold and let $\mathcal{C}$ be a collection of charts covering M. Then M is called a **smooth manifold** if the transition maps of $\mathcal{C}$ are all smooth on their domains of definition.

There are additional definitions and technicalities around this definition that we omit.

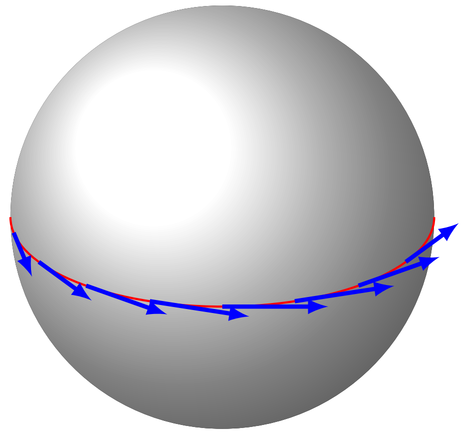

Figure 1: A smooth path on the sphere (red) with several tangent vectors (blue). Note that the tangent vectors are tangent not just to the path, but to the sphere as well.

Example 2.8. The circle $\mathbb{S}^1$ is a smooth manifold with the charts given in Example A.15. In fact, the transition map from U_N to U_S or from U_S to U_n is the smooth function $f(x) = \frac{1}{x}$ (which is its own inverse) defined on $\mathbb{R} \setminus \{0\}$.

In the Euclidean case, the concept of a convex region or a convex function is intimately connected to straight lines. A general manifold, however, may not contain straight lines connecting every point, since not all manifolds are complex (in fact, it's nontrivial that every manifold can be embedded in Euclidean space). For this reason, we need to rethink our notion of straight lines.

Definition 2.9. A function $f : [a, b] \to M$ is called a **smooth path** on M if for every $p \in \operatorname{Im} f$, there is a chart (U, φ) of M containing p such that the composite map $\varphi \circ f$ is smooth. A map $g : M \to \mathbb{R}$ is called a **smooth function** if for every $p \in M$, there is a chart (V, ψ) containing p such that the composite map $g \circ \psi^{-1}$ is smooth.

In our rethinking of the notion of a straight line, the perspective we will pursue derives from the fact that a straight line is the *shortest path* between two points in Euclidean space. In order to extend this idea to general manifolds, we must be able to determine the length of a path. Remember that in normal calculus, the length of a path $\boldsymbol{r}(t)\Big|_{t\in[a,b]}$ is given by

$$\ell(\boldsymbol{r}) = \int_a^b \sqrt{|\boldsymbol{r}'(t)|^2}\, dt$$

In order to extend this to the setting of smooth manifolds, it would seem we need to define two concepts: the derivative of a curve or path on M and the norm of a tangent vector.

Note that the derivative of a path - which is a vector tangent to the path - is naturally tangent to the manifold itself (see Figure 2). For that reason, we introduce the idea of tangent vectors to a smooth manifold.

Definition 2.10. Let $C^\infty(M)$ be the space of smooth functions on M. A function $v : C^\infty(M) \to \mathbb{R}$ is called a **tangent vector to** M **at** p if

(i) v is a linear map: $v(cf + g) = cv(f) + v(g)$ for $f, g \in C^\infty(M)$ and $c \in \mathbb{R}$.

(ii) v satisfies the Leibniz rule at p
$$v(fg) = v(f)g(p) + v(g)f(p)$$

The collection of tangent vectors to M at $p \in M$ is called the **tangent space to** M **at** p and is denoted by T_pM, and the disjoint union of all of the tangent spaces is the **tangent bundle** of M.

It is immediate that T_pM is in fact a vector space. Slightly less obvious is the fact that $\dim T_pM = \dim M$ (see Appendix A for a proof).

The motivation for this definition is that in Euclidean space, tangent vectors represent directions in which we can take directional derivatives. The directional derivative operator D_v in Euclidean space is indeed a linear operator that satisfies the Leibniz rule, offering some retroactive justification of criteria (i) and (ii) in Definition 2.10. This definition equips us to define the derivative of a smooth path on M.

Definition 2.11. Let M be a smooth manifold and $\gamma : [a, b] \to M$ a smooth path on M. Then we define the derivative of γ at $t \in [a, b]$ by

$$\gamma'(t)(f) = (f \circ \gamma)'(t)$$

It's easy enough to check that $\gamma'(t)$ satisfies the criteria of Definition 2.10. This gives us half of what we need to define the length of smooth paths. Now we turn to the other half.

Definition 2.12. Given a smooth manifold M, a **Riemannian metric** on M consists of an inner product (positive definite 2-tensor) $g_p : T_p \times T_p \to \mathbb{R}$ at each point $p \in M$ such that g_p varies smoothly with p (we elaborate on what this means in Appendix A). The pair (M, g) is then a **Riemannian manifold**.

See Appendix A for some examples of Riemannian manifolds.

We have now arrived at the setting in which we can discuss geodesic convexity. We are of course, however, missing the final ingredient: geodesics!

Definition 2.13. Let (M, g) be a Riemannian manifold, and let $\gamma : [a, b] \to M$ be a smooth path. Then the **length of** γ is given by

$$\ell(\gamma) = \int_a^b |g_p(\gamma'(t), \gamma'(t))| \, dt$$

If γ is continuous and only piecewise smooth, then let $a = a_0 < a_1 < \cdots < a_n = b$ be the points where γ is not smooth, and define

$$\ell(\gamma) = \sum_{j=0}^{n-1} \int_{a_j}^{a_{j+1}} |g_p(\gamma'(t), \gamma'(t))| \, dt$$

If γ is such that $\ell(\gamma) \leq \ell(\beta)$ for all paths $\beta : [a, b] \to M$, then γ is a **geodesic** connecting $\gamma(a)$ to $\gamma(b)$.

3 Geodesic convexity

This section gives a brief introduction to the relevant definitions in the geodesic convexity literature. The examples and proofs are adpated from [Vis18]. In standard convex analysis, we think of convex sets as subsets of $\mathbb{R}^n$ where any line segment connecting two points in the set lies within the set. Recall from Section 2 that lines constitute the geodesics of the space $\mathbb{R}^n$. This observation motivates the following more general definition of *geodesic convexity*, which arises when we allow ourselves to introduce a non-Euclidean metric on the space.

Definition 3.1 (Totally geodesically convex set)**.** Let M be a Riemannian manifold with metric g. We say that a set $K \subseteq M$ is totally convex with respect to g if for any two points $p, q \in K$, *any* geodesic γ_{pq} connecting them lies within K.

Observe that if (M, g) is given by Euclidean space, then there is a unique geodesic between any two points p and q, whence the definiton of a totally convex set coincides with our familiar notion of a convex set. However, there are many examples of manifolds M where there is no unique geodesic. Recall that geodesics correspond to paths that *locally* minimize distance, not globally minimize distance. Consider, for example, the 2-sphere $\mathbb{S}^2$. For any two points p and q on the sphere, there are two geodesics that connect the points corresponding to the short and long arcs on the unique great circle that contains p and q (see the diagram below).

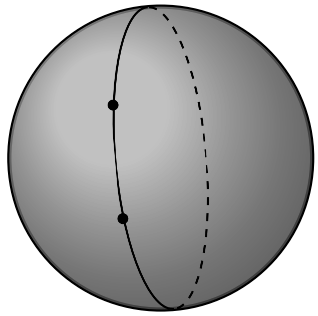

A totally geodesically convex set must contain *both* of these geodesics. We now show an example to demonstrate how expanding the notion of convexity to geodesic convexity (for an appropriate metric) allows us to reason about non-convex sets.

Example 3.2 (A non-convex set that is geodesically convex). Let K_c^n denote the collection of $n \times n$ positive definite matrices that have determinant c. Observe that K_c^n is not convex. As an example, consider the following two matrices in K_2^2

$$P = \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix} \qquad Q = \begin{bmatrix} \sqrt{2} & 0 \\ 0 & \sqrt{2} \end{bmatrix}$$

Observe that $\det(P) = \det(Q) = 2$, and that both P and Q are positive-definite. However, taking the convex combination

$$R = \frac{P+Q}{2} = \begin{bmatrix} 1 + \frac{\sqrt{2}}{2} & \frac{1}{2} \\ \frac{1}{2} & 1 + \frac{\sqrt{2}}{2} \end{bmatrix}$$

we can see that

$$\det(R) = \left(1 + \frac{\sqrt{2}}{2}\right)^2 - \frac{1}{4} \neq 2$$

whence we find that K_2^2 is not convex. However, when we consider the natural geometry of this set, given by the Frobenius inner product $\langle P, Q \rangle$, we can see that the sets K_c^n are indeed totally geodesically convex with respect to this metric. To see this, we use the well-known form for (the unique) geodesic in the space of positive definite matrices

$$\gamma_{PQ}(t) = P^{1/2} \left(P^{-1/2} Q P^{-1/2}\right)^t P^{1/2}$$

It suffices to show that for any t, $\det\left(\gamma_{PQ}(t)\right) = c$. This can be easily computed, as

$$\begin{aligned}
\det\left(\gamma_{PQ}(t)\right) &= \det\left(P^{1/2}\left(P^{-1/2} Q P^{-1/2}\right)^t P^{1/2}\right) \\
&= \det\left(P^{1/2}\right) \det\left(P^{-1/2}\right)^t \det\left(Q\right)^t \det\left(P^{-1/2}\right)^t \det\left(P^{1/2}\right) \\
&= \det\left(P\right)^{1-t} \det\left(Q\right)^t \\
&= c^{1-t} c^t \\
&= c
\end{aligned}$$

as desired.

We now similarly define geodesically convex functions, which generalize the usual definition of a convex function.

Definition 3.3 (Geodesically Convex Function). Let M be a Riemannian manifold with metric g, and let $K \subseteq M$ be a totally convex set with respect to g. We say that a function $f : K \to \mathbb{R}$ is **geodesically convex with respect to** g (or **g-convex**) if for any $p, q \in K$, and for *any* geodesic $\gamma_{pq} : [0,1] \to K$ from p to q,

$$f\left(\gamma_{pq}(t)\right) \leq (1-t)f(p) + tf(q)$$

for all $t \in [0,1]$.

This definition generalizes the notion of having the function lie below the straight line connecting the two function evaluations. As we know from the theory of convex functions, we can equivalently give a first-order characterization of convexity, which states that the linear approximation of the function is an underestimate. Such a characterization exists for geodesically convex functions as well.

Theorem 3.4 (First-order characterization of g-convex functions). *Let M be a Riemannian manifold with metric g, and let $K \subseteq M$ be an open and totally convex set with respect to g. A differentiable function $f : K \to \mathbb{R}$ is geodesically convex with respect to g if and only if for any $p, q \in K$, and for any geodesic $\gamma_{pq} : [0,1] \to K$ connecting p and q,*

$$f(p) + \gamma_{pq}'(t)(f) \leq f(q)$$

where $\gamma_{pq}'(t)(f)$ (as in Definition 2.11) denotes the first derivative of f along the geodesic γ_{pq}.

Proof. First, note that by Definition 3.3, we have that

$$\begin{aligned} f(p) + \frac{f(\gamma_{pq}(t)) - f(p)}{t} \leq f(q) &\implies f(p) + \lim_{t\to 0} \frac{f(\gamma_{pq}(t)) - f(p)}{t} \leq f(q) \\ &\implies f(p) + \gamma_{pq}'(t)(f) \leq f(q) \end{aligned}$$

as desired.

We now show the reverse direction. Let $p, q \in K$ be arbitrary points, and let $\gamma_{pq} : [0,1] \to K$ be a geodesic connecting them. Let $t \in [0,1]$ be a fixed point in the interval, and let $r := \gamma_{pq}(t)$ be its image in the geodesic. Next, consider the following curves $\alpha, \beta : [0,1] \to K$:

$$\alpha(u) := \gamma_{pq}\left(t + u(1-t)\right) \qquad \beta(u) := \gamma_{pq}(t - ut))$$

These curves can be thought of as the "subpaths" of the geodesics that connect r to q and r to p respectively. As such, they must also be geodesics themselves. Their derivatives are given by

$$\alpha'(0) = (1-t)\gamma_{pq}'(t) \qquad \beta'(0) = -t\gamma_{pq}'(t)$$

We then have that

$$f(q) \geq \alpha'(f)(r) + f(r) \implies f(q) \geq f(r) + (1-t)\gamma_{pq}'(f)(r)$$

and similarly,

$$f(q) \geq \beta'(f)(r) + f(r) \implies f(q) \geq f(r) - t\gamma_{pq}'(f)(r)$$

Multiplying the first result by t and the second by $1-t$, and summing the two, we have that

$$tf(q) + (1-t)f(p) \geq f(r) = f(\gamma_{pq}(t))$$

whence f is geodesically convex. ∎

We may similarly make a second-order characterization of geodesic convexity that mirrors the familiar second-derivative condition for convex functions.

Theorem 3.5 (Second-order characterization for g-convex functions)**.** *Let M be a Riemannian manifold with metric g, and let $K \subseteq M$ be an open and totally convex set with respect to g. A twice differentiable function $f : K \to \mathbb{R}$is geodesically convex with respect to g if and only if for any two points $p, q \in K$, and for any geodesic $\gamma_{pq} : [0,1] \to K$ connecting them,*

$$\frac{d^2 f(\gamma_{pq}(t))}{dt^2} \geq 0$$

Proof. Let $p, q \in K$ be arbitrary points. We define $\theta : [0,1] \to \mathbb{R}$ as $\theta(t) := f(\gamma_{pq}(t))$. Geodesic convexity then implies that for all $t \in [0,1]$,

$$\theta(t) \leq (1-t)\theta(0) + t\theta(1)$$

That is, θ is a convex function. The familiar second-order characterization of convex functions leads to

$$0 \geq \frac{d^2\theta(t)}{dt^2} = \frac{d^2 f(\gamma_{pq}(t))}{dt^2}$$

as desired. We now prove the reverse direction by contradiction. Suppose that f is not geodesically convex. Then, there exists $p, q \in K$ and a geodesic γ_{pq} connecting the the two, along with some $t \in [0,1]$ such that

$$f(\gamma_{pq}(t)) > (1-t)f(p) + tf(q) \implies \theta(t)?(1-t)\theta(0) + t\theta(1)$$

whence θ is not convex. It then follows, by the second-order characterization of convex functions, that for some $u \in [0,1]$

$$0 > \frac{d^2\theta(u)}{dt^2} = \frac{d^2 f(\gamma_{pq}(u))}{dt^2}$$

whence the second-order characterization

$$\frac{d^2 f(\gamma_{pq}(t))}{dt^2} \geq 0$$

follows. ☐

We now similarly give an example a function that is non-convex, but is geodesically convex given an apprioriate metric

Example 3.6 (A non-convex, but geodesically convex function)**.** We consider the manifold $\mathbb{R}_{>0}$ given by the positive reals. This manifold can be endowed with non-Euclidean metrics. One such metric arises by considering the Hessian of the log-barrier function $-\log(x)$. The corresponding metric is given by

$$g_p(u, v) = \frac{uv}{p^2}$$

It is easy to see that $g_p(u, v)$ is smooth as a function of p, whence g is a Riemannian metric. While the proof of this is not within the scope of this paper, it can be deduced from the Euler-Lagrange equations for geodesics that geodesics $\gamma_{pq} : [0,1] \to \mathbb{R}^n_{>0}$ joining points p and q under this metric take the form

$$\gamma_{pq}(t) = \exp(\alpha t + \beta)$$

for some constants $\alpha, \beta \in \mathbb{R}$. Next, consider the well known non-convex function $f(x) = \log(x)$. We show that f is indeed geodesically convex. To see this, observe that

$$\begin{aligned} \log(\gamma_{pq}(t)) = \log(\exp(\alpha t + \beta)) = \alpha t + \beta \\ \implies \frac{d}{dt}[\log(\gamma_{pq}(t))] = \alpha \\ \implies \frac{d^2}{dt^2}[\log(\gamma_{pq}(t))] = 0 \end{aligned}$$

whence, by the second-order characterization of geodesically convex functions, the non-convex function $\log(x)$ is indeed geodesically convex!

Function	Algorithm	Stepsize	Convergence Rate	Average Over Iterations?
g-convex L-Lipschitz	Projected Subgradient	$\frac{D}{L\sqrt{ct}}$	$O(\sqrt{\frac{c}{t}})$	Yes
g-convex $\lvert$subgradient$\rvert \leq G$	Projected Stochastic Subgradient	$\frac{D}{G\sqrt{ct}}$	$O(\sqrt{\frac{c}{t}})$	Yes
g-strongly convex L-Lipschitz	Projected Subgradient	$\frac{2}{\mu(s+1)}$	$O(\frac{c}{t})$	Yes
g-strongly convex $\lvert$subgradient$\rvert \leq G$	Projected Stochastic Subgradient	$\frac{2}{\mu(s+1)}$	$O(\frac{c}{t})$	Yes
g-convex L-smooth	Projected Gradient	$\frac{1}{L}$	$O(\frac{c}{c+t})$	No
g-convex, L-smooth $\lvert$variance$\rvert \leq \sigma$	Projected Stochastic Gradient	$\frac{1}{L+\frac{\sigma\sqrt{ct}}{D}}$	$O(\frac{c+\sqrt{ct}}{c+t})$	Yes
g-strongly convex L-smooth	Projected Gradient	$\frac{1}{L}$	$O((1-\min\frac{1}{c},\frac{\mu}{L})^t)$	No

Table 1: **Convergence rates for first-order g-convex optimization algorithms [ZS16]** – here, s denotes the iterate index, t denotes the total number of iterates, D denotes the diameter of the domain, L_f denotes the Lipschitz constant of f, G denotes the upper bound on gradient norms, μ denotes the strong convexity constant of f, L_g denotes the Lipschitz constant of the gradient, σ denotes the square root of the variance of the gradient, and c is a constant that depends on D and the underlying Riemannian metric g.

We now discuss the limitations of geodesic convexity. Indeed, geodesic convexity allows us to apply convex analysis to non-convex functions, but can it be applied to *any* function? In other words, for any function f, can we construct an appropriate Riemmanian metric such that the function f is geodesically convex? We answer this question in the negative. Geodesically convex functions share one key property with standard convex functions – all local minima of the function are also global minima.

Theorem 3.7 (Non-geodesically convex functions). *Let M be a smooth manifold, and let $f : M \to \mathbb{R}$ be a funciton such that there exists some point $p \in M$ and an open neighborhood U_p of p*

$$f(p) = \inf_{q \in U_p} f(q)$$

(i.e. p is a local minima), but

$$f(p) > \inf_{q \in M} f(q)$$

p is not a global minima. Then, there is no metric tensor g on M such that f is geodesically convex with respect to g.

Proof. Suppose for contradiction that a metric g exists such that f is geodesically convex with respect to g. Let $q \in M$ be such that $f(q) < f(p)$, and let $\gamma : [0,1] \to M$ be a geodesic connecting p to q (i.e. $\gamma(0) = p$ and $\gamma(1) = q$). Since f is geodesically convex, we have that

$$f(\gamma(t)) \leq (1-t)f(p) + tf(q) < f(p)$$

for all $t \in [0,1]$. As γ is smooth, for some $t_0 \in (0,1]$, we must have that $\gamma(t) \in U_p$. It then follows, by our premise that p is a local minima, that for all $t \in (0, t_0]$

$$f(\gamma(t)) \geq f(p)$$

whence we have arrived at a contradiction. Thus, no such metric g may exist. ☐

We conclude this section with a brief discussion on how geodesically convex optimization problems are solved. First-order methods for geodesically convex optimization resemble standard gradient descent methods, but steps are taken with respect to the derivative in the direction of the geodesic. There are geodesically convex analogues of projected gradient descent, projected subgradient descent, and projected stochastic subgradient descent. Table 3 shows the convergence rates for different algorithms over different families of functions – these results were derived in [ZS16].

4 Operator Scaling

In this section, we present the Operator Scaling problem, which takes advantage of geodesic convexity. These results are adapted from [All+18]. We define operator scaling as the following:

Definition 4.1 (Operator Scaling). Let $A_1, ..., A_m$ be a set of matrices. Define operator $T : S_+^n \to S_+^n$ such that $T(P) = \sum_{i=1}^m A_i P A_i^\dagger$. Define the **capacity** of T by

$$\text{Capacity}(T) = \min_{P \in S_{++}^n} \frac{\det(T(P))}{\det(P)} = \min_{P \in S_{++}^n, \det P = 1} \det(T(P))$$

and the capacity of a matrix P under T as Capacity$(P) = \frac{\det(T(P))}{\det(P)}$. Our problem is to determine whether the optimal capacity P^* is feasible, because if it is, we can define $Y^* = T(P^*)^{-1}$ such that by scaling our original matrices by $(Y^*)^{1/2}$ and P^* we make T **doubly stochastic**, or that it always outputs a matrix such that every row and column sums to 1. This scaling is given by

$$A_i^* = (Y^*)^{1/2} A_i (P^*)^{1/2}$$

and the new operator becomes

$$T^*(P) = \sum_{i=1}^m A_i^* P \left(A_i^*\right)^\dagger$$

Operator scaling is useful in a variety of contexts, such as optimizing Polynomial Identity Testing, which determines whether two multivariate polynomials are the same, and for efficiently determining if a matrix is invertible.

Without geodesic convex optimization, the best known algorithms to solve this problem were polynomial with respect to n and $1/\epsilon$ given an error bound of $\Theta(\epsilon)$. On the other hand, geodesic convex optimization can solve this problem in polynomial time with respect to n, and logarithmic time with respect to ϵ, a vast improvement. We will go over a brief overview of the algorithm and its intuition in this section.

Definition 4.2 (Positive Definite Characterization of G-Convexity). F is geodesically convex if for every $X \in S_{++}^n$ and Hermitian matrix Δ, $F(X^{1/2} e^{s\Delta} X^{1/2})$ is convex in $s \in \mathbf{R}$.

Definition 4.3 (Geodesically-Second-Order Robust). $F : \mathbf{C}^{nxn} \to \mathbf{R}$ is **g-second-order robust** if for every $X \in S_{++}^n$ and for every Hermitian Δ such that $\|\Delta\|_2 \le 1$, we have

$$\left| \frac{d^3 g}{ds^3} \right| \le 2 \frac{d^2 g}{ds^3}$$

where

$$g(s) = F(X^{1/2} e^{s\Delta} X^{1/2})$$

Definition 4.4 (Doubly Stochastic). If operator T defined by $(A_1, ..., A_m)$ is **doubly stochastic**, then $\sum_{i=1}^m A_i A_i^\dagger = \sum_{i=1}^m A_i^\dagger A_i = I$. In addition, we define the distance for an operator T to being doubly stochastic as

$$d(T) = Tr((\sum_{i=1}^m A_i A_i^\dagger - I)^2) + Tr((\sum_{i=1}^m A_i^\dagger A_i - I)^2)$$

Thus, a doubly stochastic T would have distance 0, and $d(T)$ can be seen as an error margin.

We first present a general second order method to minimize geodesically convex functions. Suppose we are trying to minimize g-convex F; assume we know all gradients and Hessians of F.

$X_0 \leftarrow I$;
for $t = 0, ..., T-1$ **do**
$f^t(\Delta) = F(X_t^{1/2} e^{\Delta} X_t^{1/2})$
$\Delta^t = \arg\min_{\Delta \in \mathbf{C}^{nxn}, \Delta = \Delta^\dagger, \|\Delta\|_2 \leq \frac{1}{2}} Tr(\nabla f^t(0)\Delta) + \frac{1}{2e} Tr(\nabla^2 f^t(0)\Delta \otimes \Delta)$
$X_{t+1} \leftarrow X_t^{1/2} e^{\Delta_t / e^2}$
end for
return X_T

Theorem 4.5. *Suppose F is geodesically convex and geodesically-second-order robust. Then*

$$F(X_T) - F(X*) \leq \epsilon$$

with

$$T = O(R \log \frac{F(I) - F(X*)}{\epsilon})$$

iterations, where $X*$ is the minimizer for F and

$$R = \max_{X | F(X) \leq F(I)} \left\| \log(X^{-1/2} X^* X^{-1/2}) \right\|_2$$

We can extend this algorithm for operator scaling. Note that $\log \text{Capacity}(X)$ is both geodesically convex and geodesically-second-order robust (which can be shown with algebraic matrix manipulations – full proof in [All+18]), and minimizing this function solves Operator Scaling. However, by applying Theorem 4.5 we notice that we cannot bound R with this function; thus, we instead choose to minimize $F(X) = \log \text{Capacity}(X) + \lambda \text{Reg}(X)$, where $\text{Reg}(X)$ is defined as the regularization function $\text{Reg}(X) = Tr(XX^\dagger)Tr((XX^\dagger)^{-1})$ and λ is some adjustable weight. This is also g-convex, g-second-order robust, and allows us to polynomially bound R.

We define $f^t(\Delta) = F(X_t^{1/2} e^{\Delta} X_t^{1/2}) = \log \text{Capacity}(X_t e^{\Delta/2}) + \lambda \text{Reg}(X_t e^{\Delta/2})$. With this g-convex and g-second-order robust definition of f^t, we can also see that our intermediate objective for each iteration t, $\Delta^t = \arg\min_{\Delta \in \mathbf{C}^{nxn}, \Delta = \Delta^\dagger, \|\Delta\|_2 \leq \frac{1}{2}} Tr(\nabla f^t(0)\Delta) + \frac{1}{2e} Tr(\nabla^2 f^t(0)\Delta \otimes \Delta)$, is convex and quadratic in terms of Δ, allowing for efficient optimization.

Theorem 4.6. *Given a choice of error bound ϵ such that $d(T) \leq \epsilon$, and M such that $\|A_i\|_\infty \leq M$ for all $i = 1, ..., m$, if there exists $X_\epsilon^* \in \mathbf{C}^{nxn}$ such that $\log \mathit{Capacity}(X_\epsilon^*) \leq \log \mathit{Capacity}(T) + \epsilon$, we can choose $\lambda = \frac{\epsilon}{n^2(\kappa(X_\epsilon^*))^2}$ and T polylog in n, m, $\kappa(X_\epsilon^*)$, $1/\epsilon$, and M, such that the algorithm converges in time complexity polynomial in n, m, $\log M$, $\log \kappa(X_\epsilon^*)$, and $\log 1/\epsilon$.*

Proofs for Theorems 4.5 and 4.6 can be found in section B of our appendix. Finally, we must show that $\kappa(X_\epsilon^*)$ is polynomially bounded, in case it exponentially grows to infinity as ϵ tends to zero. Fortunately, Theorem 4.7 shows this is not the case:

Theorem 4.7. *Suppose $\|A_i\|_\infty \leq M$ for all $i = 1, ..., m$, and Capacity$(T) > 0$. For all $\epsilon > 0$, there exist $X, Y \in S_+{+}^n$ such that $(B_1, ..., B_m) = (XA_1Y, ..., XA_mY)$ and $T'(P) = \sum_{i=1}^m B_i P B_i^\dagger$ satisfies $\|X\|_2, \|X^{-1}\|_2, \|Y\|_2, \|Y^{-1}\|_2 \leq (mnM)e^{n^{3/2}\log(12mn^4M^2/\epsilon)}$.*

The proof is derived from convergence of continuous gradient flow, and can be found in [All+18]. Thus, $\kappa(X_\epsilon^*)$ is polynomially bounded, and we have seen that by using geodesic convex optimization, we were able to turn Operator Scaling from polynomial in $1/\epsilon$ to polynomial in $\log 1/\epsilon$.

References

[Lee13] John M. Lee. *Introduction to Smooth Manifolds*. New York: Springer Science+Business Media, 2013.

[ZS16] Hongyi Zhang and Suvrit Sra. "First-order methods for geodesically convex optimization". In: *Conference on Learning Theory*. PMLR. 2016, pp. 1617–1638.

[All+18] Z. Allen-Zhu et al. "Operator Scaling via Geodesically Convex Optimization, Invariant Theory and Polynomial Identity Testing". In: *ArXiv e-prints* (Apr. 2018). arXiv: `1804.01076 [cs.DS]`. URL: `https://arxiv.org/pdf/1804.01076.pdf`.

[Vis18] N.K. Vishnoi. "Geodesic Convex Optimization: Differentiation on Manifolds, Geodesics, and Convexity". In: *ArXiv e-prints* (June 2018). arXiv: `1806.06373 [math.OC]`. URL: `https://arxiv.org/pdf/1806.06373.pdf`.

A Supplemental Proofs and Examples to Section 2

In this section, we provide many additional theorems examples to assist with understanding as it relates to topology and the theory of smooth and Riemannian manifolds. It would of course be impossible to provide a complete coverage of the theory of smooth manifolds in a paper of this length; we direct the curious reader to Lee's excellent book [Lee13].

A.1 Topological Spaces

In some proofs, we will choose efficiency over verbosity. For example, we do not define connectedness, but we still make use of it in proofs, because the intuitive idea of what it means to be connected (which is closer to path-connectedness) is more instructive then the way topologists define it.

Definition A.1 (Topological space)**.** A **topological space** $(S, \mathcal{U})$ consists of a set S and a collection $\mathcal{U}$ of **open sets** such that

(i) If $\{U_\alpha\}_{\alpha\in\mathcal{A}}$ is a subset of $\mathcal{U}$ for some arbitrary index set α, then

$$\bigcup_{\alpha\in\mathcal{A}} U_\alpha \in \mathcal{U}$$

(ii) If U_1 and U_2 are open, then so is $U_1 \cap U_2$.

(iii) $\emptyset, S \in \mathcal{U}$.

The collection $\mathcal{U}$ is called the **topology** of the space.

Example A.2. The canonical example of a topological space is $\mathbb{R}$, which has as its topology the collection

$$\mathcal{U}_\mathbb{R} = \left\{ \bigcup_{i=1}^{n} (a_i, b_i) \middle| 0 \le n \le \infty, -\infty \le a_1 < b_1 \le a_2 < b_2 \le \cdots \le a_n < b_n \right\}$$

of disjoint unions of open intervals.

Example A.3. The other example in which we are interested is $\mathbb{R}^n$. As in the previous example, we construct the open subsets of $\mathbb{R}^n$ as unions of more "basic" sets. Define

$$B_r(p) = \{q \in \mathbb{R}^n \mid d(p, q) < r\}$$

where $p \in \mathbb{R}^n$ and $r \in \mathbb{R}$. The topology of $\mathbb{R}^n$ is given by

$$\mathcal{U}_{\mathbb{R}^n} = \left\{ \bigcup_{i=1}^{n} B_{r_i}(p_i) \mid 0 \le n \le \infty, p_i \in \mathbb{R}^n, r_i \in \mathbb{R} \right\}$$

The reader is left to check that the definition of $\mathcal{U}_\mathbb{R}$ given in Example A.2 is equivalent to the definition of $\mathcal{U}_{\mathbb{R}^1}$ given in Example A.3. For further examples of topological spaces, see Appendix A.

The construction given in Example A.3 hints at a more general structure, which we introduce now.

Definition A.4 (Metric space)**.** Let S be a set and let $d : S \times S \to \mathbb{R}$ be a function satisfying

(i) $d(p, q) \ge 0$ and $d(p, q) = d(q, p)$ for all $p, q \in S$.

(ii) $d(p, q) = 0$ if and only if $p = q$.

(iii) For all $p, q, r \in S$, we have the *triangle inequality*

$$d(p, r) \le d(p, q) + d(q, r)$$

Then d is called a **metric** on S, and the **open ball of radius** r **around** p is

$$B_r(p) = \{q \in S \mid d(p,q) < r\}$$

The **topology on** S **induced by** d is

$$\mathcal{U}_{(S,d)} = \left\{ \bigcup_{i=1}^{n} B_{r_i}(p_i) \mid 0 \leq n \leq \infty, p_i \in S, r_i \in \mathbb{R} \right\}$$

The concept of a metric space will help us later when we find that the geodesics determined by a Riemannian metric on a manifold furnish us with a metric on the manifold. Note the similarities between Definition A.4 and Example A.3. Indeed, Example A.3 presents $\mathbb{R}^n$ as a metric space with the Euclidean metric.

Example A.5 (Discrete topology)**.** Let S be any set. Then we can define the **discrete topology** $\mathcal{D}_S$ by

$$\mathcal{D}_S = \mathscr{P}(S)$$

to be the power set of S. In other words, we consider every subset of S to be open. We can check that this satisfies the definition of a topology, as both $\emptyset$ and S itself are subsets of S, making them open. Further, any union of subsets of S is again a subset of S, making it open. Finally, the intersection of two subsets of S is also a subset of S.

As a matter of fact any set S endowed with the discrete topology is in fact a metric space under the **discrete metric**

$$\delta(p,q) = \begin{cases} 0 & p = q \\ 1 & p \neq q \end{cases}$$

Notice that under this metric, we have that

$$B_{1/2}(p) = \{p\}$$

is open. Since the arbitrary union of open sets is open, we find that all subsets of S are open, thereby recovering the discrete topology on S.

Example A.6 (Cofinite topology)**.** Let S be any set. Then we can define the **cofinite topology** $\mathcal{CF}_S$ by

$$\mathcal{CF}_S = \{T \subseteq S \mid |S \setminus T| < \infty \text{ or } T = \emptyset\}$$

That is, $\mathcal{CF}_S$ is the collection of subsets of S whose <u>co</u>mplement is <u>finite</u>. In order for $\mathcal{CF}_S$ to satisfy the definition of a topology, we need to include $\emptyset$ as well. To confirm that $\mathcal{CF}_S$ is indeed a topology, let $\mathcal{C} \subseteq \mathcal{CF}_S$ and note that

$$\bigcup_{T \in \mathcal{C}} T = S \setminus \left(\bigcap_{T \in \mathcal{C}} \underbrace{S \setminus T}_{\text{finite}} \right) = S \setminus \underbrace{\left(\bigcap_{T \in \mathcal{C}} S \setminus T \right)}_{\text{finite}} \in \mathcal{CF}_S$$

(if $\emptyset \in \mathcal{C}$, we can simply remove it without affecting the union) since the intersection of arbitrarily many finite sets is finite. Furthermore, if $T_1, T_2 \in \mathcal{CF}_S$, then

$$T_1 \cap T_2 = S \setminus \left(\underbrace{S \setminus T_1}_{\text{finite}} \cup \underbrace{S \setminus T_2}_{\text{finite}} \right) = S \setminus \underbrace{((S \setminus T_1) \cap (S \setminus T_2))}_{\text{finite}} \in \mathcal{CF}_S$$

(if $\emptyset \in \{T_1, T_2\}$, then the intersection is $\emptyset \in \mathcal{CF}_S$). Since the complement of S is $\emptyset$, which is finite, we can indeed confirm that $\mathcal{CF}_S$ is a topology.

Notice that if S is finite, then every subset of S has finite complement, so $\mathcal{CF}_S = \mathcal{D}_S$.

Theorem A.7. *If S is infinite, then the space $(S, \mathcal{CF}_S)$ is not Hausdorff, and hence $(S, \mathcal{CF}_S)$ is not a metric space.*

Proof. Let $p \neq q$ be two points in S. If $p \in T_p \in \mathcal{CF}_S$ and $q \in T_q \in \mathcal{CF}_S$, then both T_p and T_q are nonempty and hence have finite complements. We found above that the intersection $T_p \cap T_q$ must have finite complement as well. Since S is infinite, this implies that $T_p \cap T_q$ is infinite, and in particular nonempty. As such, there are no sets in $\mathcal{CF}_S$ satisfying the conditions of Definition 2.1, meaning that $(S, \mathcal{CF}_S)$ is not Hausdorff.

We prove in Theorem 2.2 that all metric spaces are Hausdorff. Since $(S, \mathcal{CF}_S)$ is not Hausdorff, it cannot be a metric space. ☐

Definition A.8. Let $(S, \mathcal{U})$ and $(T, \mathcal{V})$ be topological spaces. A function $f : S \to T$ is called **continuous** if for any $V \in \mathcal{V}$, we have $f^{-1}(V) \in \mathcal{U}$, where f^{-1} denotes the preimage (not the inverse of f, which need not be injective).

We offer an alternate conception of continuity (without proof) for metric spaces that makes the connection to $\mathbb{R}^n$ more clear.

Theorem A.9. *If (S, d_S) and (T, d_T) are metric spaces, then a map $f : S \to T$ is continuous if and only if for all $s \in S$ and all $\epsilon > 0$, there exists $\delta > 0$ such that*

$$d_S(s, s') < \delta \implies d_T(f(s), f(s')) < \epsilon$$

Definition A.10. Let $(S, \mathcal{U})$ and $(T, \mathcal{V})$ be topological spaces. A function $f : S \to T$ is called a **homeomorphism** if f is a bijection and both f and f^{-1} are continuous. In this case, the spaces S and T are called **homeomorphic**.

Example A.11. The interval $(0, 1)$ is homeomorphic to $\mathbb{R}$ via the map

$$f(x) = -\log\left(\frac{1-x}{x}\right)$$

whose inverse is the sigmoid function.

To tie up a loose end, we offer a proof of Theorem A.9.

Proof. Let $f : S \to T$ be a continuous map as in Definition A.8, let $s \in S$, and let $\epsilon > 0$. Then $B_\epsilon(f(s))$ is an open set in T, meaning that $f^{-1}(B_\epsilon(f(s)))$ is open in S. It follows that $f^{-1}(B_\epsilon(f(s)))$, which contains s, must contain some open ball $B_\delta(s)$. As such, we have

$$d_S(s, s') < \delta \iff s' \in B_\delta(s) \implies s' \in f^{-1}(B_\epsilon(f(s))) \implies f(s') \in B_\epsilon(f(s)) \iff d_T(f(s), f(s')) < \epsilon$$

Suppose conversely that the implication in Theorem A.9 can be made to hold for any ϵ, and let

$$V = \bigcup_{\alpha \in \mathcal{A}} B_{\epsilon_\alpha}(p_\alpha)$$

For each $p \in f^{-1}(V)$, there is then some α such that $f(p) \in B_{\epsilon_\alpha}(p_\alpha)$. Then choose some ξ such that $B_\xi(f(p)) \subseteq B_{\epsilon_\alpha}(p_\alpha)$, and choose δ_p such that

$$d(p, s') < \delta_p \implies d(f(p), f(s')) < \xi$$

This gives us

$$B_{\delta_p}(p) \subseteq f^{-1}\left(B_\xi(f(p))\right) \subseteq f^{-1}(V)$$

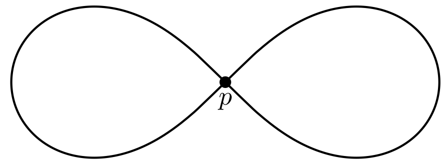

Figure 2: The figure eight metric space with its point of self-intersection, p, labelled.

and hence

$$f^{-1}(V) = \bigcup_{p \in f^{-1}(V)} B_{\delta_p}(p)$$

being a union of open balls, is open. This proves that f is continuous in the sense of Definition A.8, completing the equivalence. □

The following space is often useful for constructing counterexamples to various claims or testing whether or not a claim holds for all metric spaces.

Example A.12 (Figure eight)**.** Consider a figure eight curve (also referred to as a lemniscate) shown in Figure A.12. We can define this space algebraically using polar coordinates

$$L = \left\{(r, \theta) \mid \theta \in (-1.1\pi, 1.1\pi) \wedge r = \sqrt{\cos(2\theta)}\right\}$$

This space is a subset of $\mathbb{R}^2$; we can inherit the Euclidean metric (we would say that L is a *metric subspace* of $\mathbb{R}^2$).

A.2 Topological Manifolds

We have mentioned that the figure eight space is a frequent counterexample in elementary topology. Indeed, we have the following theorem.

Theorem A.13. *The figure eight space is Hausdorff, second countable, and locally Euclidean at every point except for its point of self-intersection, where it is not locally Euclidean*

Proof. Let L denote the figure eight space. Since L is a metric space, it is Hausdorff. To see that it is second countable, notice that the open balls around each point in L are nothing but the intersections of open balls in $\mathbb{R}^2$ with L. Since $\mathbb{R}^2$ is second countable, then, so is L.

Now let $q \in L$ be a point other than the self-intersection point. Then in polar coordinates, we can write

$$q = \left(\sqrt{\cos(2\theta)}, \theta\right)$$

where $\theta \notin \left\{-\frac{3\pi}{4}, -\frac{\pi}{4}, \frac{\pi}{4}, \frac{3\pi}{4}\right\}$. Let I be an open subinterval of $(-1.1\pi, 1.1\pi)$ containing θ but not containing any of these four angles, and observe that

$$G := \{(r, \theta) \in L \mid \theta \in I\}$$

is homeomorphic to I via the map $(r, \theta) \mapsto \theta$. This proves that L is locally Euclidean at least everywhere except the self-intersection point.

To see that L is not locally Euclidean at p (see Figure A.12), note that any neighborhood of p in L is "X-shaped". To put this in mathematical terms, we can say that if N is a sufficiently small neighborhood of p, then N has only one connected component, but $N \setminus \{p\}$ has four. We choose not to introduce additional results (that are covered in any real analysis class) and consider it sufficient to say that if an open subset of $\mathbb{R}$ is connected, then when we remove one point, it cannot have four connected components. This means that no neighborhood of p can be homeomorphic to an open subset of $\mathbb{R}$. □

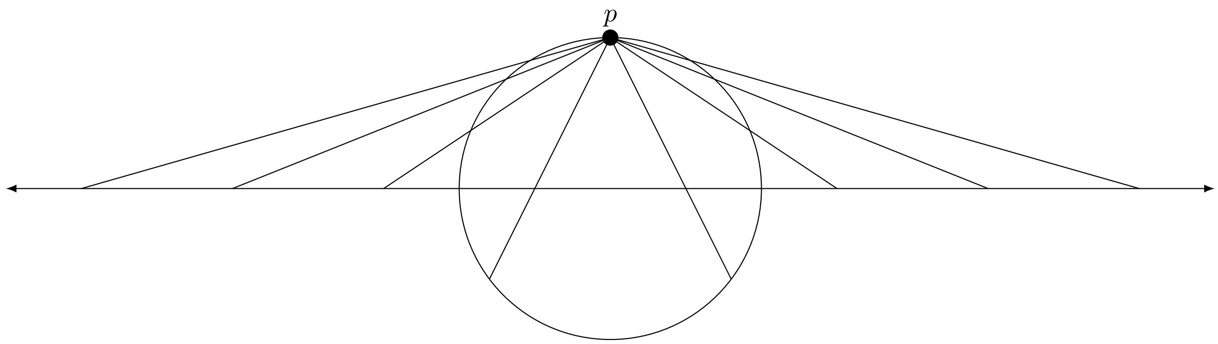

Figure 3: The stereographic projection from bS^1 to $\mathbb{R}$, defined everywhere except at the north pole.

Corollary A.14. *The figure eight space L is not a topological manifold.*

Proof. L is not locally Euclidean at p. □

This shows just how close a space can come to being a topological manifold.

The technique that we mentioned, wherein we consider the number of connected components that result from removing one point from a topological space, is a common method of proof.

Example A.15 (A non-Euclidean topological manifold)**.** The unit circle

$$\mathbb{S}^1 = \{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 = 1\}$$

is *not* homeomorphic to $\mathbb{R}$, but for any $p \in \mathbb{S}^1$, the space $\mathbb{S}^1 \setminus \{p\}$ *is* homeomorphic to $\mathbb{R}$ via the **stereographic projection map** shown in Figure 3. Every line emanating from the north pole p intersects the circle $\mathbb{S}^1$ at one other point q and intersects the line $\mathbb{R}$ at one point r. The stereographic projection takes q to r. This map is undefined at p, hence why $\mathbb{S}^1 \setminus \{p\}$, rather than $\mathbb{S}^1$, is homeomorphic to $\mathbb{R}$. With that in mind, let

$$U_N = \mathbb{S}^1 \setminus \{(0, 1)\} \qquad\qquad U_S = \mathbb{S}^1 \setminus \{(0, -1)\}$$

Both U_N and U_S are homeomorphic to $\mathbb{R}$ via a stereographic projection map, and every point of $\mathbb{S}^1$ is contained in at least one of these two sets. This proves that $\mathbb{S}^1$ is a 1-dimensional topological manifold.

To confirm that $\mathbb{S}^1$ is not homeomorphic to $\mathbb{R}$, note that when we remove a point of $\mathbb{S}^1$, we get a space that is homeomorphic to $\mathbb{R}$ (via the stereographic projection map). I we remove a point of $\mathbb{R}$, then the resulting space is not connected and hence is not homeomorphic to $\mathbb{R}$.

Example A.16 (Torus)**.** The torus is a two-dimensional topological manifold. In Figure 4, we have drawn a torus covered by many rectangles. Each rectangle is homeomorphic to $\mathbb{R}^2$ (or, if you prefer, an open subset thereof), which makes the local Euclidean-ness of the torus quite evident. Second countability and Hausdorff-ness follow from the fact that the torus is a metric subspace of $\mathbb{R}^3$. Some may be aware of the fact that a torus is isomorphic to the quotient of $\mathbb{R}^2$ by a two-dimensional lattice, which can be represented by the gluing diagram in Figure 5 (a reader who has studied some algebraic topology will have seen such a diagram). The gluing diagram furnishes us with a natural doubly periodic continuous surjection from $\mathbb{R}^2$ onto the torus. By selecting a particular period - say $(0, 1) \times (0, 1)$, we in fact obtain a homeomorphism between $(0, 1) \times (0, 1)$ and an open subset of the torus. Choosing another period, such as $(0.5, 1, 5) \times (0.5, 1.5)$, identifies yet another open subset. Each of these open subsets can be taken as a chart, and we can visualize these charts using the gluing diagram, which we have also depicted in Figure 5.

As a final example of a topological manifold, we mention a perhaps slightly more relevant example: the general linear group.

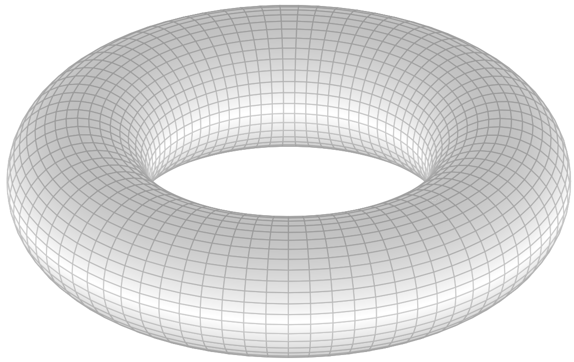

Figure 4: A torus with gridlines. Each rectangle gives an example of a chart homeomorphic to $\mathbb{R}^2$.

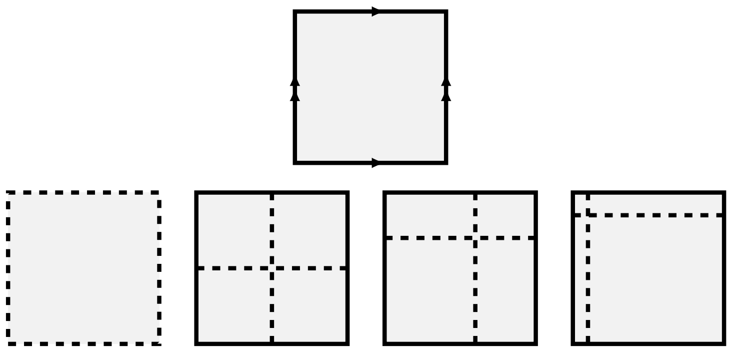

Figure 5: A gluing diagram for the torus together with some natural charts.

Example A.17 (General linear group)**.** Note that the space $\mathbb{R}^{n\times n}$ of $n\times n$ real matrices is naturally isomorphic to $\mathbb{R}^{n^2}$ as both a vector space and a topological space (these are really just two choices of notation meant to fit different contexts). Notice that the determinant function $\det : \mathbb{R}^{n\times n} \to \mathbb{R}$ is a continuous map, which means by Definition A.8 that the preimage of an open set in $\mathbb{R}$ under det is an open subset of $\mathbb{R}^{n\times n}$. This means that the preimage of the open set $(-\infty, 0) \cup (0, \infty)$ - namely, the set of $n \times n$ matrices with nonzero determinant - is an open subset of $\mathbb{R}^{n\times n}$. This set is precisely $\mathrm{GL}(n)$, which we have hence shown to be an n^2-dimensional topological manifold with a global chart.

A.3 Smooth manifolds

As it turns out, all of the topological manifolds that we mentioned in the previous section are smooth manifolds with the charts that we mentioned. In fact, almost every topological manifold (in some sense of the world "almost") can be endowed with a smooth structure. It wasn't until the 1980s that mathematicians produced a topological manifold with no smooth structure. We mentioned in Section 2 that the transition maps of $\mathbb{S}^1$ with stereographic projection charts are both simply the function $f(x) = \frac{1}{x}$. We will carry out the computation of transition maps once more in this section.

Example A.18 ($\mathbb{S}^n$)**.** Let

$$\mathbb{S}^n = \left\{ \left(x^1, \ldots, x^{n+1}\right) \in \mathbb{R}^{n+1} \middle| \sum_{i=1}^{n+1} \left(x^i\right)^2 = 1 \right\}$$

be the n-sphere. Let

$$N = (0, \ldots, 0, 1) \qquad S = (0, \ldots, 0, -1)$$

denote the north and south pole of $\mathbb{S}^n$. We claim that $\mathbb{S}^n$ is a smooth manifold with the charts $\mathbb{S}^n \setminus \{N\}$ and

$\mathbb{S}^n \setminus \{S\}$, where the coordinate maps are the stereographic projections

$$\sigma_N\left(x^1, \ldots, x^{n+1}\right) = \frac{\left(x^1, \ldots, x^n\right)}{1 - x^{n+1}} \qquad \sigma_S\left(x^1, \ldots, x^{n+1}\right) = \frac{\left(x^1, \ldots, x^n\right)}{1 + x^{n+1}}$$

Note that

$$1 - \frac{|u|^2 - 1}{|u|^2 + 1} = \frac{2}{|u|^2 + 1}$$

Define

$$\tau\left(u^1, \ldots, u^n\right) = \frac{\left(2u^1, \ldots, 2u^n, |u|^2 - 1\right)}{|u|^2 + 1}$$

Then this gives us

$$\sigma_N(\tau(u)) = \frac{\left(2u^1, \ldots, 2u^n\right)}{|u|^2 + 1} \cdot \frac{|u|^2 + 1}{2} = (u^1, \ldots, u^n) = u$$

We also have that

$$|\sigma_N(x)| = \frac{\sqrt{(x^1)^2 + \cdots + (x^n)^2}}{1 - x^{n+1}}$$

It follows from plugging this and the other coordinates of $\sigma_N(x)$ into τ that $\tau(\sigma_N(x)) = x$. Hence, $\tau = \sigma_N^{-1}$ and σ_N is a bijection. Both σ_N and σ_N^{-1} can be seen to be continuous, which shows that $\mathbb{S}^n \setminus \{N\}$ can be taken to be a chart of $\mathbb{S}^n$ with coordinate map σ_N. A similar verification shows that σ_S is a homeomorphism, which shows that $\mathbb{S}^n \setminus \{S\}$ is a valid choice of chart.

To verify that $\mathbb{S}^n$ is a smooth manifold with these charts, we compute the transition maps.

Plugging σ_N^{-1} into σ_S, we have

$$\sigma_S \circ \sigma_N^{-1}(u) = \frac{\left(2u^1, \ldots, 2u^n\right)}{|u|^2 + 1} \div \left(\frac{|u|^2 - 1}{|u|^2 + 1} + 1\right) = \frac{\left(2u^1, \ldots, 2u^n\right)}{|u|^2 + 1} \cdot \frac{|u|^2 + 1}{2|u|^2} = \frac{\left(u^1, \ldots, u^n\right)}{|u|^2}$$

Note that

$$\frac{\partial}{\partial u^j} \frac{u^i}{\sum (u^k)^2} = -\frac{2u^i u^j}{\left(\sum (u^k)^2\right)^2} + \frac{\delta_{ij}}{\sum \left(u^k\right)^2}$$

This is smooth for any i and j, whence it follows that the transition map is smooth. Observe that this map is its own inverse, since

$$\left|\sigma_S \circ \sigma_N^{-1}(u)\right| = \frac{|u|}{|u|^2} = \frac{1}{|u|}$$

This proves that the transition map is a diffeomorphism, whence the two charts $(\mathbb{S}^n \setminus \{N\}, \sigma_N)$ and $(\mathbb{S}^n \setminus \{S\}, \sigma_S)$ define a smooth structure on $\mathbb{S}^n$, as desired.

We could perform a similar calculation for any of the topological manifolds mentioned in the previous section. These calculations evidently tend to be very cumbersome, so we satisfy ourselves with just $\mathbb{S}^n$.

A.3.1 Smooth maps

In Definition 2.9, we the ideas of a smooth path and of a smooth function on a manifold. Both of these definitions relied on the smoothness of some composite map between Euclidean spaces, and indeed, the astute reader may note that smooth paths and smooth functions are both examples of a more general concept.

Definition A.19 (Smooth map)**.** Let M and N be smooth manifolds of dimension m and n, respectively. A function $F : M \to N$ is called a **smooth map** if for any $p \in M$, there exists a chart (U, φ) of M containing p and a chart (V, ψ) of N containing $f(U)$ such that the composite function

$$\psi \circ F \circ \varphi^{-1} : \mathbb{R}^m \to \mathbb{R}^n$$

is smooth where it is defined (on $\varphi(U)$). This composite function is called a **coordinate representation** of F.

This definition is quite compatible with the discussion on page 34 of [Lee13].

Theorem A.20. *A smooth function, as defined in Definition 2.9, is a smooth map from M to $\mathbb{R}$. A smooth path, as defined in Definition 2.9, is a smooth map from $[a, b]$ to M.*

Remark. The space $[a, b]$ is in fact not a manifold, but rather a related objected called a *manifold with boundary*. It is not worth elaborating on these spaces because of their extreme similarity to typical smooth manifolds.

Proof. The manifolds $[a, b]$ and $\mathbb{R}$ both have a global chart whose coordinate function is the identity map. In the case of smooth functions, we have $\psi = \mathrm{id}$, and in the case of smooth paths, we have $\varphi = \mathrm{id}$. This reduces Definition A.19 to Definition 2.9. $\square$

Example A.21 (Smooth path)**.** Define

$$M = \mathbb{S}^{2n-1} = \left\{ \left(x^1, \dots, x^{2n}\right) \in \mathbb{R}^{2n} \,\middle|\, \sum_{i=1}^{2n} \left(x^i\right)^2 = 1 \right\} = \left\{ \left(z^1, \dots, z^n\right) \in \mathbb{C}^n \,\middle|\, \sum_{i=1}^{n} \left|z^i\right|^2 = 1 \right\}$$

Let $p_0 = \left(z_0^1, \dots, z_0^n\right)$ be any point in M, and define

$$\gamma(t) = p_0 e^{it} = \left(z_0^1 e^{it}, \dots, z_0^n e^{it}\right)$$

Then for any choice of domain $[a, b]$, $\gamma : [a, b] \to M$ is a smooth path on M. The computation of coordinate representations is roughly as cumbersome as computing transition maps, so we omit the specific calculations to demonstrate this.

A rough visualization of this path is given in Figure 2 (even though $\mathbb{S}^2$ does not have odd dimension as in the above construction).

A.4 Tangent vectors

We remarked in Section 2 that the dimension of the tangent space T_pM as a vector space equals the dimension of M as a manifold. In this section, we offer a proof. First, we need to introduce the concept of the differential of a map, which is analogous to the Jacobian of a map from $\mathbb{R}^m$ to $\mathbb{R}^n$.

Definition A.22. Let M and N be smooth manifolds and let $F : M \to N$ be a smooth map. The **differential of F at p**, denoted $dF_p : T_pM \to T_pN$, is defined by

$$dF_p(v)(f) = v(f \circ F)$$

where f is a smooth function and v is a tangent vector to N at $F(p)$.

Remark. We defined the derivative of a smooth path by

$$\gamma'(t)(f) = (f \circ \gamma)'(t)$$

Notice that the manifold $[a, b]$ has only a single tangent vector at each point up to scalar multiplication, which represents the derivative operator (the directional derivative "in the t direction"). Denoting this vector by v, the above becomes

$$\gamma'(t)(f) = v(f \circ \gamma)(t) = d\gamma_t(v)(f)$$

so the derivative of a smooth path corresponds to viewing the path as a smooth map and taking the differential.

Theorem A.23. *The differential dF_p of a smooth map $F : M \to N$ at a point p is a linear transformation from T_pM to $T_{F(p)}N$.*

Proof. We have

$$\begin{aligned} dF_p(cv+w)(f) &= (cv+w)(f \circ F) \\ &= c \cdot v(f \circ F) + w(f \circ F) \\ &= c \cdot dF_p(v)(f) + dF_p(w)(f) \\ &= (c \cdot dF_p(v) + dF_p(w))(f) \end{aligned}$$

for all smooth functions f, which proves that $dF_p(cv+w) = c \cdot dF_p(v) + dF_p(w)$ for all $c \in \mathbb{R}$ and $v, w \in T_pM$. With the additional observation that $dF_p(0) = 0$, we conclude that the differential is a linear transformation. □

Lemma A.24. *If $F : M \to N$ and $G : N \to M$ are smooth maps satisfying $F \circ G = \mathrm{id}_N$ and $G \circ F = \mathrm{id}_M$ (i.e. F and G are inverse maps), then dF_p and $dG_{F(p)}$ are inverse linear transformations.*

Proof. We have

$$dG_{F(p)}(dF_p(v))(f) = dF_p(v)(f \circ G) = v(f \circ G \circ F) = v(f)$$

and

$$dF_p(dG_{F(p)}(v))(g) = dG_{F(p)}(v)(g \circ F) = v(g \circ F \circ G) = v(g)$$

where f and g are arbitrary smooth functions on M and N, respectively. This gives

$$dF_p \circ dG_{F(p)} = \mathrm{id}_{T_pM} \qquad\qquad dG_{F(p)} \circ dF_p = \mathrm{id}_{T_{F(p)}N}$$

so dF_p and $dG_{F(p)}$ are inverse linear transformations, as desired. □

Theorem A.25. *If M is a smooth manifold of dimension n, then for each point $p \in M$, $\dim T_pM = n$.*

Proof. Let (U, φ) be a chart of M around p. Note that φ is tautologically a smooth map from U to an open subset of $\mathbb{R}^n$ with smooth inverse map. By the previous lemma, the differential of φ is a linear isomorphism of tangent spaces. The tangent spaces T_pU and T_pM coincide (by an elementary argument that is not worth including), and tangent space to $\mathbb{R}^n$ at any point is isomorphic to $\mathbb{R}^n$ as a vector space (another elementary argument: tangent vectors in $\mathbb{R}^n$ correspond to directional derivatives). This proves that the dimension of T_pM equals the dimension of $T_{\varphi(p)}\mathbb{R}^n$, namely n. □

For a more thorough treatment of this idea, see Propositions 3.2 through 3.10 of [Lee13], which build to this result.

A.5 Vector bundles

In Definition 2.12, we noted that the inner product g_p must somehow vary smoothly with p. In this section, we flesh out this remark to formalize the notion of a Riemannian metric. Chapters 10, 12, and 13 of [Lee13] offer a much more robust treatment of the ideas here. We use slightly distinct notation and verbiage from [Lee13] that is more specific to our applications (and makes the significance of, for example, the bundle of covariant 2-tensors on a manifold, more clear).

Definition A.26. Let M and E be smooth manifolds. E is said to be a **smooth rank-k vector bundle on** M if there exists a smooth surjection $\pi : E \to M$ with the following properties:

(i) For each point $p \in M$, $\pi^{-1}(p)$ has the structure of a real vector space of dimension k.

(ii) For each point $p \in M$, there is an open neighborhood U of p and a diffeomorphism (smooth map with smooth inverse) $\Phi : \pi^{-1}(U) \to U \times \mathbb{R}^k$ such that $\pi_U \circ \Phi = \pi$ (where $\pi_U : U \times \mathbb{R}^k \to U$ is the projection map) and $\Phi_{\pi^{-1}(q)}$ is a vector space isomorphism from $\pi^{-1}(q)$ to $\mathbb{R}^k$.

Example A.27 (Tangent bundle)**.** In Definition 2.10, we remarked that the disjoint union of all tangent spaces to a manifold is called the **tangent bundle**. The tangent bundle of a smooth manifold M is denoted by TM. Note that any element can be represented as an ordered pair (p, v), where $p \in M$ and $v \in T_pM$. Then there is a natural projection map $\pi : TM \to M$ which maps (p, v) to p. Indeed, TM can be endowed with a smooth structure in such a way that the map π is smooth (π is trivially surjective).

The concept of a vector bundle underlies what is meant by the phrase "g_p varies smoothly with p." Recall that if V is a real vector space, then V^* - the dual of V - represents the space of linear functionals from V to $\mathbb{R}$. If we take the tensor product $V^* \otimes V^*$, then we get the space of bilinear functionals $V \times V \to \mathbb{R}$.

Definition A.28 (Bundle of 2-tensors)**.** Let M be a smooth manifold. Then the **bundle of covariant 2-tensors on** M is given by

$$T^2T^*M := \bigsqcup_{p \in M} \left(T_p^M\right)^* \otimes \left(T_p^M\right)^*$$

We choose smooth charts and coordinate maps on T^2T^*M to make the projection map onto M smooth. The machinery required to formalize this is out of scope.

With this definition under our belt, we can understand what exactly is meant by the phrase "g_p varies smoothly with p." In particular, the Riemannian metric is a smooth map from M to T^2T^*M (such that g_p is positive definite for each $p \in M$).

Remark. Some authors choose to develop the ideas of geodesics and g-convexity by defining objects called *affine connections* and then restricting attention to a special affine connection known as the *Levi-Civita connection*. We exclude these objects from our treatment, because the Levi-Civita connection is determined by the Riemannian metric on a manifold, meaning that the more general discussion of affine connections is unnecessary. Further still, we can define geodesics with no mention of the Levi-Civita connection, so affine connections as a whole can be avoided altogether.

We offer one further theorem relating to Riemannian manifolds to connect these ideas back to our earlier notion of a metric space.

Theorem A.29. *A Riemannian manifold (M, g) has a natural metric space structure defined as follows: for any $p, q \in M$, let γ be a geodesic from p to q, and define $d(p, q) = \ell(\gamma)$.*

Proof. We check properties (i)-(iii) of Definition A.4. Property (i) follows from the positive definiteness of g_p. For property (ii), it suffices to show, since g_p is positive definite, that if $\gamma(a) \neq \gamma(b)$, then $\gamma'(t)$ is nonzero for some $t \in [a, b]$. Let M have dimension n, and let (U, φ) be a chart of M containing $\gamma(a)$ but not containing $\gamma(b)$. Then γ must take multiple values inside U, so $\varphi \circ \gamma$ takes on multiple values inside $\varphi(U)$. As a result, one of the n coordinates of φ - say, φ^1 - is nonconstant, implying that

$$(\varphi^1 \circ \gamma)'(t) = \gamma'(t)(\varphi^1)$$

is not identically 0, so $\gamma'(t)$ is nonzero for some $t \in [a, b]$. This proves property (ii). For property (iii), let $p, q, r \in M$, let $\alpha : [0, 1] \to M$ be a geodesic from p to q, and let $\beta : [1, 2] \to M$ be a geodesic from q to r. Then

$$(\alpha \cdot \beta)(t) := \begin{cases} \alpha(t) & t \in [0, 1] \\ \beta(t) & t \in [1, 2] \end{cases}$$

is a continuous, piecewise smooth path from p to r, and $\ell(\alpha \cdot \beta) = \ell(\alpha) + \ell(\beta)$. By definition, then, any γ from p to r must have $\ell(\gamma) \leq \ell(\alpha) + \ell(\beta)$, which proves the triangle inequality. $\square$

B Proofs of Theorems 4.5 and 4.6

Proof. To prove Theorem 4.5, it is sufficient to show (inductively) that

$$F(X_{t+1}) \le F(I) \qquad\qquad F(X_{t+1} - F(X^*) \le \left(1 - \frac{1}{2e^2R}\right)^t (F(I) - F(X^*))$$

for all t.

Let $\Delta^* = \log(X_t^{-1/2} X^* X_t^{1/2})$, so $f^t(\Delta^*) = F(X^*)$, and $h(s) = f^t(s\Delta^*)$. Since F is g-convex, so is h, so we know that

$$f^t(0) - f^t(\Delta^*/2R) = h(0) - h(\frac{1}{2R}) \ge \frac{1}{2R}(h(0) - h(1)) = \frac{1}{2R}(f^t(0) - f^t(\Delta^*))$$

.

We know that:

$$\begin{aligned}
&Tr(\nabla f^t(0)\Delta_t) + \frac{1}{2e}Tr(\nabla^2 f^t(0)(\Delta_t \otimes \Delta_t))\\
&\le Tr(\nabla f^t(0)\frac{\Delta^*}{2R}) + \frac{1}{2e}Tr(\nabla^2 f^t(0)(\frac{\Delta^*}{2R} \otimes \frac{\Delta^*}{2R}))\\
&\le -(f^t(0) - f^t(\frac{\Delta^*}{2R})\\
&\le -\frac{1}{2R}(F(X_t) - F(X^*))
\end{aligned}$$

as well as

$$\begin{aligned}
&Tr(\nabla f^t(0)\Delta_t) + \frac{1}{2e}Tr(\nabla^2 f^t(0)(\Delta_t \otimes \Delta_t))\\
&= e^2 Tr(\nabla f^t(0)\Delta_t/e^2) + \frac{e}{2}Tr(\nabla^2 f^t(0)(\frac{\Delta_t}{e^2} \otimes \frac{\Delta_t}{e^2}))\\
&\ge -e^2(f^t(0) - f^t(\frac{\Delta_t}{e^2})\\
&= -e^2(F(X_t) - F(X_{t+1}))
\end{aligned}$$

Thus by rearranging terms, we get $F(X_{t+1}) - F(X^*) \le (1 - \frac{1}{2e^2R})(F(X_t) - F(X^*))$. □

Proof. The proof of Theorem 4.6 is a bit more involved, so we will just go over a brief overview and the full proof can be found in [All+18].

We know

$$\begin{aligned}
\mathrm{Reg}(X_\epsilon^*) &= Tr(X_\epsilon^*(X_\epsilon^*)^\dagger)Tr((X_\epsilon^*(X_\epsilon^*)^\dagger)^{-1})\\
&\le n^2 \frac{\lambda_{max}(X_\epsilon^*(X_\epsilon^*)^\dagger)}{\lambda_{min}(X_\epsilon^*(X_\epsilon^*)^\dagger)}\\
&= n^2(\kappa(X_\epsilon^*))^2
\end{aligned}$$

and hence

$$F(X_\epsilon^*) = \log \mathrm{Capacity}(T) + \epsilon + \lambda \mathrm{Reg}(X_\epsilon^*) \le \log \mathrm{Capacity}(T) + 2\epsilon$$

The following lemma, whose proof we omit, will assist in our proof of Theorem 4.6.

Lemma 1. *If Δ_t and X_{t+1} are calculated exactly for each iteration t, then:*

- $F(X_{t+1}) \le F(X_t)$
- $F(X_{t+1}) - F(X_\epsilon^*) \le \left(1 - \frac{1}{8e^2 \log \kappa_0}\right)(F(X_t) - F(X_\epsilon^*))$

By Lemma 1, by choosing to run the algorithm for

$$T = O(\log \kappa_0 \log(nmM\epsilon^{-1}))$$

iterations, we have

$$F(X_T) - F(X_\epsilon^*) \leq \left(1 - \frac{1}{8e^2 \log(\kappa_0)}\right)^T (F(I) - F(X_\epsilon^*)) \leq \epsilon$$

□