

Part III — Riemannian Geometry

Based on lectures by A. G. Kovalev

Notes taken by Dexter Chua

Lent 2017

These notes are not endorsed by the lecturers, and I have modified them (often significantly) after lectures. They are nowhere near accurate representations of what was actually lectured, and in particular, all errors are almost surely mine.

This course is a possible natural sequel of the course Differential Geometry offered in Michaelmas Term. We shall explore various techniques and results revealing intricate and subtle relations between Riemannian metrics, curvature and topology. I hope to cover much of the following:

A closer look at geodesics and curvature. Brief review from the Differential Geometry course. Geodesic coordinates and Gauss' lemma. Jacobi fields, completeness and the Hopf–Rinow theorem. Variations of energy, Bonnet–Myers diameter theorem and Synge's theorem.

Hodge theory and Riemannian holonomy. The Hodge star and Laplace–Beltrami operator. The Hodge decomposition theorem (with the 'geometry part' of the proof). Bochner–Weitzenböck formulae. Holonomy groups. Interplays with curvature and de Rham cohomology.

Ricci curvature. Fundamental groups and Ricci curvature. The Cheeger–Gromoll splitting theorem.

Pre-requisites

Manifolds, differential forms, vector fields. Basic concepts of Riemannian geometry (curvature, geodesics etc.) and Lie groups. The course Differential Geometry offered in Michaelmas Term is the ideal pre-requisite.

Contents

1	Basics of Riemannian manifolds	3
2	Riemann curvature	7
3	Geodesics	12
3.1	Definitions and basic properties	12
3.2	Jacobi fields	18
3.3	Further properties of geodesics	23
3.4	Completeness and the Hopf–Rinow theorem	27
3.5	Variations of arc length and energy	30
3.6	Applications	34
4	Hodge theory on Riemannian manifolds	44
4.1	Hodge star and operators	44
4.2	Hodge decomposition theorem	47
4.3	Divergence	53
4.4	Introduction to Bochner’s method	56
5	Riemannian holonomy groups	63
6	The Cheeger–Gromoll splitting theorem	68
	Index	71

1 Basics of Riemannian manifolds

Before we do anything, we lay out our conventions. Given a choice of local coordinates $\{x^k\}$, the coefficients X^k for a vector field X are defined by

$$X = \sum_k X^k \frac{\partial}{\partial x^k}.$$

In general, for a tensor field $X \in TM^{\otimes q} \otimes T^*M^{\otimes p}$, we write

$$X = \sum X_{\ell_1 \dots \ell_p}^{k_1 \dots k_q} \frac{\partial}{\partial x^{k_1}} \otimes \dots \otimes \frac{\partial}{\partial x^{k_q}} \otimes dx^{\ell_1} \otimes \dots \otimes dx^{\ell_p},$$

and we often leave out the \otimes signs.

For the sake of sanity, we will often use implicit summation convention, i.e. whenever we write something of the form

$$X_{ijk} Y^{iljk},$$

we mean

$$\sum_{i,j} X_{ijk} Y^{iljk}.$$

We will use upper indices to denote contravariant components, and lower indices for covariant components, as we have done above. Thus, we always sum an upper index with a lower index, as this corresponds to applying a covector to a vector.

We will index the basis elements oppositely, e.g. we write dx^k instead of dx_k for a basis element of T^*M , so that the indices in expressions of the form $A_k dx^k$ seem to match up. Whenever we do not follow this convention, we will write out summations explicitly.

We will also adopt the shorthands

$$\partial_k = \frac{\partial}{\partial x^k}, \quad \nabla_k = \nabla_{\partial_k}.$$

With these conventions out of the way, we begin with a very brief summary of some topics in the Michaelmas Differential Geometry course, starting from the definition of a Riemannian metric.

Definition (Riemannian metric). Let M be a smooth manifold. A *Riemannian metric* g on M is an inner product on the tangent bundle TM varying smoothly with the fibers. Formally, this is a global section of $T^*M \otimes T^*M$ that is fiberwise symmetric and positive definite.

The pair (M, g) is called a *Riemannian manifold*.

On every coordinate neighbourhood with coordinates $x = (x_1, \dots, x_n)$, we can write

$$g = \sum_{i,j=1}^n g_{ij}(x) dx^i dx^j,$$

and we can find the coefficients g_{ij} by

$$g_{ij} = g \left(\frac{\partial}{\partial x^i}, \frac{\partial}{\partial x^j} \right)$$

and are C^∞ functions.

Example. The manifold \mathbb{R}^k has a canonical metric given by the Euclidean metric. In the usual coordinates, g is given by $g_{ij} = \delta_{ij}$.

Does every manifold admit a metric? Recall

Theorem (Whitney embedding theorem). Every smooth manifold M admits an embedding into \mathbb{R}^k for some k . In other words, M is diffeomorphic to a submanifold of \mathbb{R}^k . In fact, we can pick k such that $k \leq 2 \dim M$.

Using such an embedding, we can induce a Riemannian metric on M by restricting the inner product from Euclidean space, since we have inclusions $T_p M \hookrightarrow T_p \mathbb{R}^k \cong \mathbb{R}^k$.

More generally,

Lemma. Let (N, h) be a Riemannian manifold, and $F : M \rightarrow N$ is an immersion, then the pullback $g = F^*h$ defines a metric on M .

The condition of immersion is required for the pullback to be non-degenerate.

In Differential Geometry, if we do not have metrics, then we tend to consider diffeomorphic spaces as being the same. With metrics, the natural notion of isomorphism is

Definition (Isometry). Let (M, g) and (N, h) be Riemannian manifolds. We say $f : M \rightarrow N$ is an *isometry* if it is a diffeomorphism and $f^*h = g$. In other words, for any $p \in M$ and $u, v \in T_p M$, we need

$$h((df)_p u, (df)_p v) = g(u, v).$$

Example. Let G be a Lie group. Then for any x , we have translation maps $L_x, R_x : G \rightarrow G$ given by

$$\begin{aligned} L_x(y) &= xy \\ R_x(y) &= yx \end{aligned}$$

These maps are in fact diffeomorphisms of G .

We already know that G admits a Riemannian metric, but we might want to ask something stronger — does there exist a *left-invariant metric*? In other words, is there a metric such that each L_x is an isometry?

Recall the following definition:

Definition (Left-invariant vector field). Let G be a Lie group, and X a vector field. Then X is *left invariant* if for any $x \in G$, we have $d(L_x)X = X$.

We had a rather general technique for producing left-invariant vector fields. Given a Lie group G , we can define the *Lie algebra* $\mathfrak{g} = T_e G$. Then we can produce left-invariant vector fields by picking some $X_e \in \mathfrak{g}$, and then setting

$$X_a = d(L_a)X_e.$$

The resulting vector field is indeed smooth, as shown in the differential geometry course.

Similarly, to construct a left-invariant metric, we can just pick a metric at the identity and the propagating it around using left-translation. More explicitly, given any inner product on $\langle \cdot, \cdot \rangle$ on $T_e G$, we can define g by

$$g(u, v) = \langle (dL_{x^{-1}})_x u, (dL_{x^{-1}})_x v \rangle$$

for all $x \in G$ and $u, v \in T_x G$. The argument for smoothness is similar to that for vector fields.

Of course, everything works when we replace “left” with “right”. A Riemannian metric is said to be *bi-invariant* if it is both left- and right-invariant. These are harder to find, but it is a fact that every compact Lie group admits a bi-invariant metric. The basic idea of the proof is to start from a left-invariant metric, then integrate the metric along right translations of all group elements. Here compactness is necessary for the result to be finite.

We will later see that we cannot drop the compactness condition. There are non-compact Lie groups that do not admit bi-invariant metrics, such as $\mathrm{SL}(2, \mathbb{R})$.

Recall that in order to differentiate vectors, or even tensors on a manifold, we needed a connection on the tangent bundle. There is a natural choice for the connection when we are given a Riemannian metric.

Definition (Levi-Civita connection). Let (M, g) be a Riemannian manifold. The *Levi-Civita connection* is the unique connection $\nabla : \Omega_M^0(TM) \rightarrow \Omega_M^1(TM)$ on M satisfying

- (i) Compatibility with metric:

$$Zg(X, Y) = g(\nabla_Z X, Y) + g(X, \nabla_Z Y),$$

- (ii) Symmetry/torsion-free:

$$\nabla_X Y - \nabla_Y X = [X, Y].$$

Definition (Christoffel symbols). In local coordinates, the *Christoffel symbols* are defined by

$$\nabla_{\partial_j} \frac{\partial}{\partial x^k} = \Gamma_{jk}^i \frac{\partial}{\partial x^i}.$$

With a bit more imagination on what the symbols mean, we can write the first property as

$$d(g(X, Y)) = g(\nabla X, Y) + g(X, \nabla Y),$$

while the second property can be expressed in coordinate representation by

$$\Gamma_{jk}^i = \Gamma_{kj}^i.$$

The connection was defined on TM , but in fact, the connection allows us to differentiate many more things, and not just tangent vectors.

Firstly, the connection ∇ induces a unique covariant derivative on T^*M , also denoted ∇ , defined uniquely by the relation

$$X\langle \alpha, Y \rangle = \langle \nabla_X \alpha, Y \rangle + \langle \alpha, \nabla_X Y \rangle$$

for any $X, Y \in \mathrm{Vect}(M)$ and $\alpha \in \Omega^1(M)$.

To extend this to a connection ∇ on tensor bundles $\mathcal{T}^{q,p} \equiv (TM)^{\otimes q} \otimes (T^*M)^{\otimes p}$ for any $p, q \geq 0$, we note the following general construction:

In general, suppose we have vector bundles E and F , and $s_1 \in \Gamma(E)$ and $s_2 \in \Gamma(F)$. If we have connections ∇^E and ∇^F on E and F respectively, then we can define

$$\nabla^{E \otimes F}(s_1 \otimes s_2) = (\nabla^E s_1) \otimes s_2 + s_1 \otimes (\nabla^F s_2).$$

Since we already have a connection on TM and T^*M , this allows us to extend the connection to all tensor bundles.

Given this machinery, recall that the Riemannian metric is formally a section $g \in \Gamma(T^*M \otimes T^*M)$. Then the compatibility with the metric can be written in the following even more compact form:

$$\nabla g = 0.$$

2 Riemann curvature

With all those definitions out of the way, we now start by studying the notion of *curvature*. The definition of the curvature tensor might not seem intuitive at first, but motivation was somewhat given in the III Differential Geometry course, and we will not repeat that.

Definition (Curvature). Let (M, g) be a Riemannian manifold with Levi-Civita connection ∇ . The *curvature 2-form* is the section

$$R = -\nabla \circ \nabla \in \Gamma(\wedge^2 T^*M \otimes T^*M \otimes TM) \subseteq \Gamma(T^{1,3}M).$$

This can be thought of as a 2-form with values in $T^*M \otimes TM = \text{End}(TM)$. Given any $X, Y \in \text{Vect}(M)$, we have

$$R(X, Y) \in \Gamma(\text{End } TM).$$

The following formula is a straightforward, and also crucial computation:

Proposition.

$$R(X, Y) = \nabla_{[X, Y]} - [\nabla_X, \nabla_Y].$$

In local coordinates, we can write

$$R = \left(R_{j, k\ell}^i dx^k dx^\ell \right)_{i, j=1, \dots, \dim M} \in \Omega_M^2(\text{End}(TM)).$$

Then we have

$$R(X, Y)_j^i = R_{j, k\ell}^i X^k Y^\ell.$$

The comma between j and $k\ell$ is purely for artistic reasons.

It is often slightly convenient to consider a different form of the Riemann curvature tensor. Instead of having a tensor of type $(1, 3)$, we have one of type $(0, 4)$ by

$$R(X, Y, Z, T) = g(R(X, Y)Z, T)$$

for $X, Y, Z, T \in T_p M$. In local coordinates, we write this as

$$R_{ij, k\ell} = g_{iq} R_{j, k\ell}^q.$$

The first thing we want to prove is that $R_{ij, k\ell}$ enjoys some symmetries we might not expect:

Proposition.

(i)

$$R_{ij, k\ell} = -R_{ij, \ell k} = -R_{ji, k\ell}.$$

(ii) The *first Bianchi identity*:

$$R_{j, k\ell}^i + R_{k, \ell j}^i + R_{\ell, jk}^i = 0.$$

(iii)

$$R_{ij, k\ell} = R_{k\ell, ij}.$$

Note that the first Bianchi identity can also be written for the $(0, 4)$ tensor as

$$R_{ij,kl} + R_{ik,lj} + R_{il,jk} = 0.$$

Proof.

- (i) The first equality is obvious as coefficients of a 2-form. For the second equality, we begin with the compatibility of the connection with the metric:

$$\frac{\partial g_{ij}}{\partial x^k} = g(\nabla_k \partial_i, \partial_j) + g(\partial_i, \nabla_k \partial_j).$$

We take a partial derivative, say with respect to x_ℓ , to obtain

$$\frac{\partial^2 g_{ij}}{\partial x^\ell \partial x^k} = g(\nabla_\ell \nabla_k \partial_i, \partial_j) + g(\nabla_k \partial_i, \nabla_\ell \partial_j) + g(\nabla_\ell \partial_i, \nabla_k \partial_j) + g(\partial_i, \nabla_\ell \nabla_k \partial_j).$$

Then we know

$$0 = \frac{\partial^2 g}{\partial x^\ell \partial x^k} - \frac{\partial^2 g}{\partial x_k \partial x_\ell} = g([\nabla_\ell, \nabla_k] \partial_i, \partial_j) + g(\partial_i, [\nabla_\ell, \nabla_k] \partial_j).$$

But we know

$$R(\partial_k, \partial_\ell) = \nabla_{[\partial_k, \partial_\ell]} - [\nabla_k, \nabla_\ell] = -[\nabla_k, \nabla_\ell].$$

Writing $R_{k\ell} = R(\partial_k, \partial_\ell)$, we have

$$0 = g(R_{k\ell} \partial_i, \partial_j) + g(\partial_i, R_{k\ell} \partial_j) = R_{ji,k\ell} + R_{ij,k\ell}.$$

So we are done.

- (ii) Recall

$$R_{j,k\ell}^i = (R_{k\ell} \partial_j)^i = ([\nabla_\ell, \nabla_k] \partial_j)^i.$$

So we have

$$\begin{aligned} & R_{j,k\ell}^i + R_{k,\ell j}^i + R_{\ell,jk}^i \\ &= [(\nabla_\ell \nabla_k \partial_j - \nabla_k \nabla_\ell \partial_j) + (\nabla_j \nabla_\ell \partial_k - \nabla_\ell \nabla_j \partial_k) + (\nabla_k \nabla_j \partial_\ell - \nabla_j \nabla_k \partial_\ell)]^i. \end{aligned}$$

We claim that

$$\nabla_\ell \nabla_k \partial_j - \nabla_k \nabla_\ell \partial_j = 0.$$

Indeed, by definition, we have

$$(\nabla_k \partial_j)^q = \Gamma_{kj}^q = \Gamma_{jk}^q = (\nabla_j \partial_k)^q.$$

The other terms cancel similarly, and we get 0 as promised.

- (iii) Consider the following octahedron:

The equalities on each vertex is given by (i). By the first Bianchi identity, for each greyed triangle, the sum of the three vertices is zero.

Now looking at the upper half of the octahedron, adding the two greyed triangles shows us the sum of the vertices in the horizontal square is $(-2)R_{ij,kl}$. Looking at the bottom half, we find that the sum of the vertices in the horizontal square is $(-2)R_{kl,ij}$. So we must have

$$R_{ij,kl} = R_{kl,ij}. \quad \square$$

What exactly are the properties of the Levi-Civita connection that make these equality works? The first equality of (i) did not require anything. The second equality of (i) required the compatibility with the metric, and (ii) required the symmetric property. The last one required both properties.

Note that we can express the last property as saying $R_{ij,kl}$ is a symmetric bilinear form on $\wedge^2 T_p^* M$.

Sectional curvature

The full curvature tensor is rather scary. So it is convenient to obtain some simpler quantities from it. Recall that if we had tangent vectors X, Y , then we can form

$$|X \wedge Y| = \sqrt{g(X, X)g(Y, Y) - g(X, Y)^2},$$

which is the area of the parallelogram spanned by X and Y . We now define

$$K(X, Y) = \frac{R(X, Y, X, Y)}{|X \wedge Y|^2}.$$

Note that this is invariant under (non-zero) scaling of X or Y , and is symmetric in X and Y . Finally, it is also invariant under the transformation $(X, Y) \mapsto (X + \lambda Y, Y)$.

But it is an easy linear algebra fact that these transformations generate all isomorphism from a two-dimensional vector space to itself. So $K(X, Y)$ depends only on the 2-plane spanned by X, Y . So we have in fact defined a function on the Grassmannian of 2-planes, $K : \text{Gr}(2, T_p M) \rightarrow \mathbb{R}$. This is called the *sectional curvature* (of g).

It turns out the sectional curvature determines the Riemann curvature tensor completely!

Lemma. Let V be a real vector space of dimension ≥ 2 . Suppose $R', R'' : V^{\otimes 4} \rightarrow \mathbb{R}$ are both linear in each factor, and satisfies the symmetries we found for the Riemann curvature tensor. We define $K', K'' : \text{Gr}(2, V) \rightarrow \mathbb{R}$ as in the sectional curvature. If $K' = K''$, then $R' = R''$.

This is really just linear algebra.

Proof. For any $X, Y, Z \in V$, we know

$$R'(X + Z, Y, X + Z, Y) = R''(X + Z, Y, X + Z, Y).$$

Using linearity of R' and R'' , and cancelling equal terms on both sides, we find

$$R'(Z, Y, X, Y) + R'(X, Y, Z, Y) = R''(Z, Y, X, Y) + R''(X, Y, Z, Y).$$

Now using the symmetry property of R' and R'' , this implies

$$R'(X, Y, Z, Y) = R''(X, Y, Z, Y).$$

Similarly, we replace Y with $Y + T$, and then we get

$$R'(X, Y, Z, T) + R'(X, T, Z, Y) = R''(X, Y, Z, Y) + R''(X, T, Z, Y).$$

We then rearrange and use the symmetries to get

$$R'(X, Y, Z, T) - R''(X, Y, Z, T) = R'(Y, Z, X, T) - R''(Y, Z, X, T).$$

We notice this equation says $R'(X, Y, Z, T) - R''(X, Y, Z, T)$ is invariant under the cyclic permutation $X \rightarrow Y \rightarrow Z \rightarrow X$. So by the first Bianchi identity, we have

$$3(R'(X, Y, Z, T) - R''(X, Y, Z, T)) = 0.$$

So we must have $R' = R''$. □

Corollary. Let (M, g) be a manifold such that for all p , the function $K_p : \text{Gr}(2, T_p M) \rightarrow \mathbb{R}$ is a constant map. Let

$$R_p^0(X, Y, Z, T) = g_p(X, Z)g_p(Y, T) - g_p(X, T)g_p(Y, Z).$$

Then

$$R_p = K_p R_p^0.$$

Here K_p is just a real number, since it is constant. Moreover, K_p is a smooth function of p .

Equivalently, in local coordinates, if the metric at a point is δ_{ij} , then we have

$$R_{ij,ij} = -R_{ij,ji} = K_p,$$

and all other entries all zero.

Of course, the converse also holds.

Proof. We apply the previous lemma as follows: we define $R' = K_p R_p^0$ and $R'' = R_p$. It is a straightforward inspection to see that this R^0 does follow the symmetry properties of R_p , and that they define the same sectional curvature. So $R'' = R'$. We know K_p is smooth in p as both g and R are smooth. □

We can further show that if $\dim M > 2$, then K_p is in fact independent of p under the hypothesis of this function, and the proof requires a second Bianchi identity. This can be found on the first example sheet.

Other curvatures

There are other quantities we can extract out of the curvature, which will later be useful.

Definition (Ricci curvature). The *Ricci curvature* of g at $p \in M$ is

$$\text{Ric}_p(X, Y) = \text{tr}(v \mapsto R_p(X, v)Y).$$

In terms of coordinates, we have

$$\text{Ric}_{ij} = R_{i,jq}^q = g^{pq} R_{pi,jq},$$

where g^{pq} denotes the inverse of g .

This Ric is a symmetric bilinear form on $T_p M$. This can be determined by the quadratic form

$$\text{Ric}(X) = \frac{1}{n-1} \text{Ric}_p(X, X).$$

The coefficient $\frac{1}{n-1}$ is just a convention.

There are still two indices we can contract, and we can define

Definition (Scalar curvature). The *scalar curvature* of g is the trace of Ric respect to g . Explicitly, this is defined by

$$s = g^{ij} \text{Ric}_{ij} = g^{ij} R_{i,jq}^q = R^{qi}_{iq}.$$

Sometimes a convention is to define the scalar curvature as $\frac{s}{n(n-1)}$ instead. In the case of a constant sectional curvature tensor, we have

$$\text{Ric}_p = (n-1)K_p g_p,$$

and

$$s(p) = n(n-1)K_p.$$

Low dimensions

If $n = 2$, i.e. we have surfaces, then the Riemannian metric g is also known as the *first fundamental form*, and it is usually written as

$$g = E du^2 + 2F du dv + G dv^2.$$

Up to the symmetries, the only non-zero component of the curvature tensor is $R_{12,12}$, and using the definition of the scalar curvature, we find

$$R_{12,12} = \frac{1}{2}s(EG - F^2).$$

Thus $s/2$ is also the sectional curvature (there can only be one plane in the tangent space, so the sectional curvature is just a number). One can further check that

$$\frac{s}{2} = K = \frac{LN - M^2}{EG - F^2},$$

the *Gaussian curvature*. Thus, the full curvature tensor is determined by the Gaussian curvature. Also, $R_{12,21}$ is the determinant of the second fundamental form.

If $n = 3$, one can check that $R(g)$ is determined by the Ricci curvature.

3 Geodesics

3.1 Definitions and basic properties

We will eventually want to talk about geodesics. However, the setup we need to write down the definition of geodesics can be done in a much more general way, and we will do that.

The general setting is that we have a vector bundle $\pi : E \rightarrow M$.

Definition (Lift). Let $\pi : E \rightarrow M$ be a vector bundle with typical fiber V . Consider a curve $\gamma : (-\varepsilon, \varepsilon) \rightarrow M$. A *lift* of γ is a map $\gamma^E : (-\varepsilon, \varepsilon) \rightarrow E$ if $\pi \circ \gamma^E = \gamma$, i.e. the following diagram commutes:

$$\begin{array}{ccc} & & E \\ & \nearrow \gamma^E & \downarrow \pi \\ (-\varepsilon, \varepsilon) & \xrightarrow{\gamma} & M \end{array}$$

For $p \in M$, we write $E_p = \pi^{-1}(\{p\}) \cong V$ for the fiber above p . We can think of E_p as the space of some “information” at p . For example, if $E = TM$, then the “information” is a tangent vector at p . In physics, the manifold M might represent our universe, and a point in E_p might be the value of the electromagnetic field at p .

Thus, given a path γ in M , a lift corresponds to providing that piece of “information” at each point along the curve. For example, if $E = TM$, then we can canonically produce a lift of γ , given by taking the derivative of γ at each point.

Locally, suppose we are in some coordinate neighbourhood $U \subseteq M$ such that E is trivial on U . After picking a trivialization, we can write our lift as

$$\gamma^E(t) = (\gamma(t), a(t))$$

for some function $a : (-\varepsilon, \varepsilon) \rightarrow V$.

One thing we would want to do with such lifts is to differentiate them, and see how it changes along the curve. When we have a *section* of E on the whole of M (or even just an open neighbourhood), rather than just a lift along a curve, the connection provides exactly the information needed to do so. It is not immediately obvious that the connection also allows us to differentiate curves along paths, but it does.

Proposition. Let $\gamma : (-\varepsilon, \varepsilon) \rightarrow M$ be a curve. Then there is a uniquely determined operation $\frac{\nabla}{dt}$ from the space of all lifts of γ to itself, satisfying the following conditions:

- (i) For any $c, d \in \mathbb{R}$ and lifts $\tilde{\gamma}^E, \gamma^E$ of γ , we have.

$$\frac{\nabla}{dt}(c\gamma^E + d\tilde{\gamma}^E) = c \frac{\nabla \gamma^E}{dt} + d \frac{\nabla \tilde{\gamma}^E}{dt}$$

- (ii) For any lift γ^E of γ and function $f : (-\varepsilon, \varepsilon) \rightarrow \mathbb{R}$, we have

$$\frac{\nabla}{dt}(f\gamma^E) = \frac{df}{dt} + f \frac{\nabla \gamma^E}{dt}.$$

(iii) If there is a local section s of E and a local vector field V on M such that

$$\gamma^E(t) = s(\gamma(t)), \quad \dot{\gamma}(t) = V(\gamma(t)),$$

then we have

$$\frac{\nabla \gamma^E}{dt} = (\nabla_V s) \circ \gamma.$$

Locally, this is given by

$$\left(\frac{\nabla \gamma^E}{dt} \right)^i = \dot{a}^i + \Gamma_{jk}^i a^j \dot{x}^k.$$

The proof is straightforward — one just checks that the local formula works, and the three properties force the operation to be locally given by that formula.

Definition (Covariant derivative). The uniquely defined operation in the proposition above is called the *covariant derivative*.

In some sense, lifts that have vanishing covariant derivative are “constant” along the map.

Definition (Horizontal lift). Let ∇ be a connection on E with $\Gamma_{jk}^i(x)$ the coefficients in a local trivialization. We say a lift γ^E is *horizontal* if

$$\frac{\nabla \gamma^E}{dt} = 0.$$

Since this is a linear first-order ODE, we know that for a fixed γ , given any initial $a(0) \in E_{\gamma(0)}$, there is a unique way to obtain a horizontal lift.

Definition (Parallel transport). Let $\gamma : [0, 1] \rightarrow M$ be a curve in M . Given any $a_0 \in E_{\gamma(0)}$, the unique horizontal lift of γ with $\gamma^E(0) = (\gamma(0), a_0)$ is called the *parallel transport* of a_0 along $\gamma(0)$. We sometimes also call $\gamma^E(1)$ the parallel transport.

Of course, we want to use this general theory to talk about the case where M is a Riemannian manifold, $E = TM$ and ∇ is the Levi-Civita connection of g . In this case, each curve $\gamma(t)$ has a canonical lift independent of the metric or connection given simply by taking the derivative $\dot{\gamma}(t)$.

Definition (Geodesic). A curve $\gamma(t)$ on a Riemannian manifold (M, g) is called a *geodesic curve* if its canonical lift is horizontal with respect to the Levi-Civita connection. In other words, we need

$$\frac{\nabla \dot{\gamma}}{dt} = 0.$$

In local coordinates, we write this condition as

$$\ddot{x}_i + \Gamma_{jk}^i \dot{x}^j \dot{x}^k = 0.$$

This time, we obtain a second-order ODE. So a geodesic is uniquely specified by the initial conditions $p = x(0)$ and $a = \dot{x}(0)$. We will denote the resulting geodesic as $\gamma_p(t, a)$, where t is the time coordinate as usual.

Since we have a non-linear ODE, existence is no longer guaranteed on all time, but just for some interval $(-\varepsilon, \varepsilon)$. Of course, we still have uniqueness of solutions.

We now want to prove things about geodesics. To do so, we will need to apply some properties of the covariant derivative we just defined. Since we are lazy, we would like to reuse results we already know about the covariant derivative for vector fields. The trick is to notice that locally, we can always extend $\dot{\gamma}$ to a vector field.

Indeed, we work in some coordinate chart around $\gamma(0)$, and we wlog assume

$$\dot{\gamma}(0) = \frac{\partial}{\partial x_1}.$$

By the inverse function theorem, we note that $x_1(t)$ is invertible near 0, and we can write $t = t(x_1)$ for small x_1 . Then in this neighbourhood of 0, we can view x_k as a function of x_1 instead of t . Then we can define the vector field


$$\underline{\dot{\gamma}}(x_1, \dots, x_k) = \dot{\gamma}(x_1, x_2(x_1), \dots, x_k(x_1)).$$

By construction, this agrees with $\dot{\gamma}$ along the curve.

Using this notation, the geodesic equation can be written as

$$\nabla_{\underline{\dot{\gamma}}} \underline{\dot{\gamma}} \Big|_{\gamma(t)} = 0,$$

where the ∇ now refers to the covariant derivative of vector fields, i.e. the connection itself.



Using this, a lot of the desired properties of geodesics immediately follow from well-known properties of the covariant derivative. For example,

Proposition. If γ is a geodesic, then $|\dot{\gamma}(t)|_g$ is constant.

Proof. We use the extension $\underline{\dot{\gamma}}$ around $p = \gamma(0)$, and stop writing the underlines. Then we have

$$\dot{\gamma}(g(\dot{\gamma}, \dot{\gamma})) = g(\nabla_{\dot{\gamma}} \dot{\gamma}, \dot{\gamma}) + g(\dot{\gamma}, \nabla_{\dot{\gamma}} \dot{\gamma}) = 0,$$

which is valid at each $q = \gamma(t)$ on the curve. But at each q , we have

$$\dot{\gamma}(g(\dot{\gamma}, \dot{\gamma})) = \dot{x}^k \frac{\partial}{\partial x_k} g(\dot{\gamma}, \dot{\gamma}) = \frac{d}{dt} |\dot{\gamma}(t)|_g^2$$

by the chain rule. So we are done. \square

At this point, it might be healthy to look at some examples of geodesics.

Example. In \mathbb{R}^n with the Euclidean metric, we have $\Gamma_{jk}^i = 0$. So the geodesic equation is

$$\ddot{x}_k = 0.$$

So the geodesics are just straight lines.

Example. On a sphere S^n with the usual metric induced by the standard embedding $S^n \hookrightarrow \mathbb{R}^{n+1}$. Then the geodesics are great circles.

To see this, we may wlog $p = e_0$ and $a = e_1$, for a standard basis $\{e_i\}$ of \mathbb{R}^{n+1} . We can look at the map

$$\varphi : (x_0, \dots, x_n) \mapsto (x_0, x_1, -x_2, \dots, -x_n),$$

and it is clearly an isometry of the sphere. Therefore it preserves the Riemannian metric, and hence sends geodesics to geodesics. Since it also preserves p and a , we know $\varphi(\gamma) = \gamma$ by uniqueness. So it must be contained in the great circle lying on the plane spanned by e_0 and e_1 .

Lemma. Let $p \in M$, and $a \in T_p M$. As before, let $\gamma_p(t, a)$ be the geodesic with $\gamma(0) = p$ and $\dot{\gamma}(0) = a$. Then

$$\gamma_p(\lambda t, a) = \gamma_p(t, \lambda a),$$

and in particular is a geodesic.

Proof. We apply the chain rule to get

$$\begin{aligned} \frac{d}{dt} \gamma(\lambda t, a) &= \lambda \dot{\gamma}(\lambda t, a) \\ \frac{d^2}{dt^2} \gamma(\lambda t, a) &= \lambda^2 \ddot{\gamma}(\lambda t, a). \end{aligned}$$

So $\gamma(\lambda t, a)$ satisfies the geodesic equations, and have initial velocity λa . Then we are done by uniqueness of ODE solutions. \square

Thus, instead of considering $\gamma_p(t, a)$ for arbitrary t and a , we can just fix $t = 1$, and look at the different values of $\gamma_p(1, a)$. By ODE theorems, we know this depends smoothly on a , and is defined on some open neighbourhood of $0 \in T_p M$.

Definition (Exponential map). Let (M, g) be a Riemannian manifold, and $p \in M$. We define \exp_p by

$$\exp_p(a) = \gamma(1, a) \in M$$

for $a \in T_p M$ whenever this is defined.

We know this function has domain at least some open ball around $0 \in T_p M$, and is smooth. Also, by construction, we have $\exp_p(0) = p$.

In fact, the exponential map gives us a chart around p locally, known as *geodesic local coordinates*. To do so, it suffices to note the following rather trivial proposition.

Proposition. We have

$$(\mathrm{d}\exp_p)_0 = \mathrm{id}_{T_pM},$$

where we identify $T_0(T_pM) \cong T_pM$ in the natural way.

All this is saying is if you go in the direction of $a \in T_pM$, then you go in the direction of a .

Proof.

$$(\mathrm{d}\exp_p)_0(v) = \frac{\mathrm{d}}{\mathrm{d}t} \exp_p(tv) = \frac{\mathrm{d}}{\mathrm{d}t} \gamma(1, tv) = \frac{\mathrm{d}}{\mathrm{d}t} \gamma(t, v) = v. \quad \square$$

Corollary. \exp_p maps an open ball $B(0, \delta) \subseteq T_pM$ to $U \subseteq M$ diffeomorphically for some $\delta > 0$.

Proof. By the inverse mapping theorem. □

This tells us the inverse of the exponential map gives us a chart of M around p . These coordinates are often known as *geodesic local coordinates*.

In these coordinates, the geodesics from p have the very simple form

$$\gamma(t, a) = ta$$

for all $a \in T_pM$ and t sufficiently small that this makes sense.

Corollary. For any point $p \in M$, there exists a local coordinate chart around p such that

- The coordinates of p are $(0, \dots, 0)$.
- In local coordinates, the metric at p is $g_{ij}(p) = \delta_{ij}$.
- We have $\Gamma_{jk}^i(p) = 0$.

Coordinates satisfying these properties are known as *normal coordinates*.

Proof. The geodesic local coordinates satisfies these property, after identifying T_pM isometrically with $(\mathbb{R}^n, \mathrm{eucl})$. For the last property, we note that the geodesic equations are given by

$$\ddot{x}^i + \Gamma_{jk}^i \dot{x}^k \dot{x}^j = 0.$$

But geodesics through the origin are given by straight lines. So we must have $\Gamma_{jk}^i = 0$. □

Such coordinates will be useful later on for explicit calculations, since whenever we want to verify a coordinate-independent equation (which is essentially all equations we care about), we can check it at each point, and then use normal coordinates at that point to simplify calculations.

We again identify $(T_pN, g(p)) \cong (\mathbb{R}^n, \mathrm{eucl})$, and then we have a map

$$(r, \mathbf{v}) \in (0, \delta) \times S^{n-1} \mapsto \exp_p(r\mathbf{v}) \in M^n.$$

This chart is known as *geodesic polar coordinates*. For each fixed r , the image of this map is called a *geodesic sphere* of geodesic radius r , written Σ_r . This is an embedded submanifold of M .

Note that in geodesic local coordinates, the metric at $0 \in T_pN$ is given by the Euclidean metric. However, the metric at other points can be complicated. Fortunately, Gauss' lemma says it is not *too* complicated.

Theorem (Gauss' lemma). The geodesic spheres are perpendicular to their radii. More precisely, $\gamma_p(t, a)$ meets every Σ_r orthogonally, whenever this makes sense. Thus we can write the metric in geodesic polars as

$$g = dr^2 + h(r, \mathbf{v}),$$

where for each r , we have

$$h(r, \mathbf{v}) = g|_{\Sigma_r}.$$

In matrix form, we have

$$g = \begin{pmatrix} 1 & 0 & \cdots & 0 \\ 0 & & & \\ \vdots & & h & \\ 0 & & & \end{pmatrix}$$

The proof is not hard, but it involves a few subtle points.

Proof. We work in geodesic coordinates. It is clear that $g(\partial_r, \partial_r) = 1$.

Consider an arbitrary vector field $X = X(\mathbf{v})$ on S^{n-1} . This induces a vector field on some neighbourhood $B(0, \delta) \subseteq T_p M$ by

$$\tilde{X}(r\mathbf{v}) = X(\mathbf{v}).$$

Pick a direction $\mathbf{v} \in T_p M$, and consider the unit speed geodesic γ in the direction of \mathbf{v} . We define

$$G(r) = g(\tilde{X}(r\mathbf{v}), \dot{\gamma}(r)) = g(\tilde{X}, \dot{\gamma}(r)).$$

We begin by noticing that

$$\nabla_{\partial_r} \tilde{X} - \nabla_{\tilde{X}} \partial_r = [\partial_r, \tilde{X}] = 0.$$

Also, we have

$$\frac{d}{dr} G(r) = g(\nabla_{\dot{\gamma}} \tilde{X}, \dot{\gamma}) + g(\tilde{X}, \nabla_{\dot{\gamma}} \dot{\gamma}).$$

We know the second term vanishes, since γ is a geodesic. Noting that $\dot{\gamma} = \frac{\partial}{\partial r}$, we know the first term is equal to

$$g(\nabla_{\tilde{X}} \partial_r, \partial_r) = \frac{1}{2} \left(g(\nabla_{\tilde{X}} \partial_r, \partial_r) + g(\partial_r, \nabla_{\tilde{X}} \partial_r) \right) = \frac{1}{2} \tilde{X}(g(\partial_r, \partial_r)) = 0,$$

since we know that $g(\partial_r, \partial_r) = 1$ constantly.

Thus, we know $G(r)$ is constant. But $G(0) = 0$ since the metric at 0 is the Euclidean metric. So G vanishes everywhere, and so ∂_r is perpendicular to Σ_g . \square


Corollary. Let $a, w \in T_p M$. Then

$$g((d \exp_p)_a a, (d \exp_p)_a w) = g(a, w)$$

whenever a lives in the domain of the geodesic local neighbourhood.

3.2 Jacobi fields

Fix a Riemannian manifold M . Let's imagine that we have a “manifold” of all smooth curves on M . Then this “manifold” has a “tangent space”. Morally, given a curve γ , a “tangent vector” at γ in the space of curve should correspond to providing a tangent vector (in M) at each point along γ :




Since we are interested in the geodesics only, we consider the “submanifold” of geodesics curves. What are the corresponding “tangent vectors” living in this “submanifold”?

In rather more concrete terms, suppose $f_s(t) = f(t, s)$ is a family of geodesics in M indexed by $s \in (-\varepsilon, \varepsilon)$. What do we know about $\frac{\partial f}{\partial s}\big|_{s=0}$, a vector field along f_0 ?

We begin by considering such families that fix the starting point $f(0, s)$, and then derive some properties of $\frac{\partial f}{\partial s}$ in these special cases. We will then define a *Jacobi field* to be any vector field along a curve that satisfies these properties. We will then prove that these are exactly the variations of geodesics.

Suppose $f(t, s)$ is a family of geodesics such that $f(0, s) = p$ for all s . Then in geodesics local coordinates, it must look like this:



For a fixed p , such a family is uniquely determined by a function

$$a(s) : (-\varepsilon, \varepsilon) \rightarrow T_p M$$

such that

$$f(t, s) = \exp_p(ta(s)).$$

The initial conditions of this variation can be given by $a(0) = a$ and

$$\dot{a}(0) = w \in T_a(T_p M) \cong T_p M.$$

We would like to know the “variation field” of $\gamma(t) = f(t, 0) = \gamma_p(t, a)$ this induces. In other words, we want to find $\frac{\partial f}{\partial s}(t, 0)$. This is not hard. It is just given by

$$(\mathrm{d}\exp_p)_{ta_0}(tw) = \frac{\partial f}{\partial s}(t, 0),$$

As before, to prove something about f , we want to make good use of the properties of ∇ . Locally, we extend the vectors $\frac{\partial f}{\partial s}$ and $\frac{\partial f}{\partial t}$ to vector fields $\frac{\partial}{\partial s}$

and $\frac{\partial}{\partial s}$. Then in this set up, we have

$$\dot{\gamma} = \frac{\partial f}{\partial t} = \frac{\partial}{\partial t}.$$

Note that in $\frac{\partial f}{\partial t}$, we are differentiating f with respect to t , whereas the $\frac{\partial}{\partial t}$ on the far right is just a formal expressions.

By the geodesic equation, we have

$$0 = \frac{\nabla}{dt} \dot{\gamma} = \nabla_{\partial_t} \partial_t.$$

Therefore, using the definition of the curvature tensor R , we obtain

$$\begin{aligned} 0 &= \nabla_{\partial_s} \nabla_{\partial_t} \frac{\partial}{\partial t} = \nabla_{\partial_t} \nabla_{\partial_s} \partial_t - R(\partial_s, \partial_t) \partial_t \\ &= \nabla_{\partial_t} \nabla_{\partial_s} \partial_t + R(\partial_t, \partial_s) \partial_t \end{aligned}$$

We let this act on the function f . So we get

$$0 = \frac{\nabla}{dt} \frac{\nabla}{ds} \frac{\partial f}{\partial t} + R(\partial_t, \partial_s) \frac{\partial f}{\partial t}.$$

We write

$$J(t) = \frac{\partial f}{\partial s}(t, 0),$$

which is a vector field along the geodesic γ . Using the fact that

$$\frac{\nabla}{ds} \frac{\partial f}{\partial t} = \frac{\nabla}{dt} \frac{\partial f}{\partial s},$$

we find that J must satisfy the ordinary differential equation

$$\frac{\nabla^2}{dt^2} J + R(\dot{\gamma}, J) \dot{\gamma} = 0.$$

This is a linear second-order ordinary differential equation.

Definition (Jacobi field). Let $\gamma : [0, L] \rightarrow M$ be a geodesic. A *Jacobi field* is a vector field J along γ that is a solution of the *Jacobi equation* on $[0, L]$

$$\frac{\nabla^2}{dt^2} J + R(\dot{\gamma}, J) \dot{\gamma} = 0. \quad (\dagger)$$

We now embark on a rather technical journey to prove results about Jacobi fields. Observe that $\dot{\gamma}(t)$ and $t\dot{\gamma}(t)$ both satisfy this equation, rather trivially.

Theorem. Let $\gamma : [0, L] \rightarrow N$ be a geodesic in a Riemannian manifold (M, g) . Then

- (i) For any $u, v \in T_{\gamma(0)}M$, there is a unique Jacobi field J along Γ with

$$J(0) = u, \quad \frac{\nabla J}{dt}(0) = v.$$

If

$$J(0) = 0, \quad \frac{\nabla J}{dt}(0) = k\dot{\gamma}(0),$$

then $J(t) = kt\dot{\gamma}(t)$. Moreover, if both $J(0), \frac{\nabla J}{dt}(0)$ are orthogonal to $\dot{\gamma}(0)$, then $J(t)$ is perpendicular to $\dot{\gamma}(t)$ for all $[0, L]$.

In particular, the vector space of all Jacobi fields along γ have dimension $2n$, where $n = \dim M$.

The subspace of those Jacobi fields pointwise perpendicular to $\dot{\gamma}(t)$ has dimensional $2(n - 1)$.

- (ii) $J(t)$ is independent of the parametrization of $\dot{\gamma}(t)$. Explicitly, if $\tilde{\gamma}(t) = \tilde{\gamma}(\lambda t)$, then \tilde{J} with the same initial conditions as J is given by

$$\tilde{J}(\tilde{\gamma}(t)) = J(\gamma(\lambda t)).$$

This is the kind of theorem whose statement is longer than the proof.

Proof.

- (i) Pick an orthonormal basis e_1, \dots, e_n of $T_p M$, where $p = \gamma(0)$. Then parallel transports $\{X_i(t)\}$ via the Levi-Civita connection preserves the inner product.

We take e_1 to be parallel to $\dot{\gamma}(0)$. By definition, we have

$$X_i(0) = e_i, \quad \frac{\nabla X_i}{dt} = 0.$$

Now we can write

$$J = \sum_{i=1}^n y_i X_i.$$

Then taking $g(X_i, \cdot)$ of (†), we find that

$$\ddot{y}_i + \sum_{j=2}^n R(\dot{\gamma}, X_j, \dot{\gamma}, X_i)y_j = 0.$$

Then the claims of the theorem follow from the standard existence and uniqueness of solutions of differential equations.

In particular, for the orthogonality part, we know that $J(0)$ and $\frac{\nabla J}{dt}(0)$ being perpendicular to $\dot{\gamma}$ is equivalent to $y_1(0) = \dot{y}_1(0) = 0$, and then Jacobi's equation gives

$$\ddot{y}_1(t) = 0.$$

- (ii) This follows from uniqueness. □

Our discussion of Jacobi fields so far has been rather theoretical. Now that we have an explicit equation for the Jacobi field, we can actually produce some of them. We will look at the case where we have constant sectional curvature.

Example. Suppose the sectional curvature is constantly $K \in \mathbb{R}$, for $\dim M \geq 3$. We wlog $|\dot{\gamma}| = 1$. We let J along γ be a Jacobi field, normal to $\dot{\gamma}$.

Then for any vector field T along γ , we have

$$\langle R(\dot{\gamma}, J)\dot{\gamma}, T \rangle = K(g(\dot{\gamma}, \dot{\gamma})g(J, T) - g(\dot{\gamma}, J)g(\dot{\gamma}, T)) = Kg(J, T).$$

Since this is true for all T , we know

$$R(\dot{\gamma}, J)\dot{\gamma} = KJ.$$

Then the Jacobi equation becomes

$$\frac{\nabla^2}{dt^2}J + KJ = 0.$$

So we can immediately write down a collection of solutions

$$J(t) = \begin{cases} \frac{\sin(t\sqrt{K})}{\sqrt{K}}X_i(t) & K > 0 \\ tX_i(t) & K = 0 \\ \frac{\sinh(t\sqrt{-K})}{\sqrt{-K}}X_i(t) & K < 0 \end{cases}$$

for $i = 2, \dots, n$, and this has initial conditions

$$J(0) = 0, \quad \frac{\nabla J}{dt}(0) = e_i.$$

Note that these Jacobi fields vanishes at 0.

We can now deliver our promise, proving that Jacobi fields are precisely the variations of geodesics.

Proposition. Let $\gamma : [a, b] \rightarrow M$ be a geodesic, and $f(t, s)$ a variation of $\gamma(t) = f(t, 0)$ such that $f(t, s) = \gamma_s(t)$ is a geodesic for all $|s|$ small. Then

$$J(t) = \frac{\partial f}{\partial s}$$

is a Jacobi field along $\dot{\gamma}$.

Conversely, every Jacobi field along γ can be obtained this way for an appropriate function f .

Proof. The first part is just the exact computation as we had at the beginning of the section, but for the benefit of the reader, we will reproduce the proof again.

$$\begin{aligned} \frac{\nabla^2 J}{dt} &= \nabla_t \nabla_t \frac{\partial f}{\partial s} \\ &= \nabla_t \nabla_s \frac{\partial f}{\partial t} \\ &= \nabla_s \left(\nabla_t \frac{\partial f}{\partial t} \right) - R(\partial_t, \partial_s)\dot{\gamma}_s. \end{aligned}$$

We notice that the first term vanishes, because $\nabla_t \frac{\partial f}{\partial t} = 0$ by definition of geodesic. So we find

$$\frac{\nabla^2 J}{dt} = -R(\dot{\gamma}, J)\dot{\gamma},$$

which is the Jacobi equation.

The converse requires a bit more work. We will write $J'(0)$ for the covariant derivative of J along γ . Given a Jacobi field J along a geodesic $\gamma(t)$ for $t \in [0, L]$, we let $\tilde{\gamma}$ be another geodesic such that

$$\tilde{\gamma}(0) = \gamma(0), \quad \dot{\tilde{\gamma}}(0) = J(0).$$

We take parallel vector fields X_0, X_1 along $\tilde{\gamma}$ such that

$$X_0(0) = \dot{\gamma}(0), \quad X_1(0) = J'(0).$$

We put $X(s) = X_0(s) + sX_1(s)$. We put

$$f(t, s) = \exp_{\tilde{\gamma}(s)}(tX(s)).$$

In local coordinates, for each fixed s , we find

$$f(t, s) = \tilde{\gamma}(s) + tX(s) + O(t^2)$$

as $t \rightarrow 0$. Then we define

$$\gamma_s(t) = f(t, s)$$

whenever this makes sense. This depends smoothly on s , and the previous arguments say we get a Jacobi field

$$\hat{J}(t) = \frac{\partial f}{\partial s}(t, 0)$$

We now want to check that $\hat{J} = J$. Then we are done. To do so, we have to check the initial conditions. We have

$$\hat{J}(0) = \frac{\partial f}{\partial s}(0, 0) = \frac{d\tilde{\gamma}}{ds}(0) = J(0),$$

and also

$$\hat{J}'(0) = \frac{\nabla}{dt} \frac{\partial f}{\partial s}(0, 0) = \frac{\nabla}{ds} \frac{\partial f}{\partial t}(0, 0) = \frac{\nabla X}{ds}(0) = X_1(0) = J'(0).$$


So we have $\hat{J} = J$. □

Corollary. Every Jacobi field J along a geodesic γ with $J(0) = 0$ is given by

$$J(t) = (\mathrm{d} \exp_p)_{t\dot{\gamma}(0)}(tJ'(0))$$

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

This is just a reiteration of the fact that if we pull back to the geodesic local coordinates, then the variation must look like this:



But this corollary is stronger, in the sense that it holds even if we get out of the geodesic local coordinates (i.e. when \exp_p no longer gives a chart).

Proof. Write $\dot{\gamma}(0) = a$, and $J'(0) = w$. By above, we can construct the variation by

$$f(t, s) = \exp_p(t(a + sw)).$$

Then

$$(\mathrm{d} \exp_p)_{t(a+sw)}(tw) = \frac{\partial f}{\partial s}(t, s),$$

which is just an application of the chain rule. Putting $s = 0$ gives the result. \square

It can be shown that in the situation of the corollary, if $a \perp w$, and $|a| = |w| = 1$, then

$$|J(t)| = t - \frac{1}{3!}K(\sigma)t^3 + o(t^3)$$

as $t \rightarrow 0$, where σ is the plane spanned by a and w .

3.3 Further properties of geodesics

We can now use Jacobi fields to prove interesting things. We now revisit the Gauss lemma, and deduce a stronger version.

Lemma (Gauss' lemma). Let $a, w \in T_pM$, and

$$\gamma = \gamma_p(t, a) = \exp_p(ta)$$

a geodesic. Then

$$g_{\gamma(t)}((\mathrm{d} \exp_p)_{ta}a, (\mathrm{d} \exp_p)_{ta}w) = g_{\gamma(0)}(a, w).$$

In particular, γ is orthogonal to $\exp_p\{v \in T_pM : |v| = r\}$. Note that the latter need not be a submanifold.

This is an improvement of the previous version, which required us to live in the geodesic local coordinates.

Proof. We fix any $r > 0$, and consider the Jacobi field J satisfying

$$J(0) = 0, \quad J'(0) = \frac{w}{r}.$$

Then by the corollary, we know the Jacobi field is

$$J(t) = (\mathrm{d} \exp_p)_{ta} \left(\frac{tw}{r} \right).$$

We may write

$$\frac{w}{r} = \lambda a + u,$$

with $a \perp u$. Then since Jacobi fields depend linearly on initial conditions, we write

$$J(t) = \lambda t \dot{\gamma}(t) + J_n(t)$$

for a Jacobi field J_n a normal vector field along γ . So we have

$$g(J(r), \dot{\gamma}(r)) = \lambda r |\dot{\gamma}(r)|^2 = g(w, a).$$

But we also have

$$g(w, a) = g(\lambda r + u, a) = \lambda r|a|^2 = \lambda r|\dot{\gamma}(0)|^2 = \lambda r|\dot{\gamma}(r)|^2.$$

Now we use the fact that

$$J(r) = (\mathrm{d}\exp_p)_{ra}w$$

and

$$\dot{\gamma}(r) = (\mathrm{d}\exp_p)_{ra}a,$$

and we are done. \square

Corollary (Local minimizing of length). Let $a \in T_pM$. We define $\varphi(t) = ta$, and $\psi(t)$ a piecewise C^1 curve in T_pM for $t \in [0, 1]$ such that

$$\psi(0) = 0, \quad \psi(1) = a.$$

Then

$$\mathrm{length}(\exp_p \circ \psi) \geq \mathrm{length}(\exp_p \circ \varphi) = |a|.$$

It is important to interpret this corollary precisely. It only applies to curves with the same end point in T_pM . If we have two curves in T_pM whose end points have the same image in M , then the result need not hold (the torus would be a counterexample).

Proof. We may of course assume that ψ never hits 0 again after $t = 0$. We write

$$\psi(t) = \rho(t)\mathbf{u}(t),$$

where $\rho(t) \geq 0$ and $|\mathbf{u}(t)| = 1$. Then

$$\psi' = \rho'\mathbf{u} + \rho\mathbf{u}'.$$

Then using the extended Gauss lemma, and the general fact that if $\mathbf{u}(t)$ is a unit vector for all t , then $\mathbf{u} \cdot \mathbf{u}' = \frac{1}{2}(\mathbf{u} \cdot \mathbf{u})' = 0$, we have

$$\begin{aligned} \left| \frac{\mathrm{d}}{\mathrm{d}x}(\exp_p \circ \psi)(t) \right|^2 &= |(\mathrm{d}\exp_p)_{\psi(t)}\psi'(t)|^2 \\ &= \rho'(t)^2 + 2g(\rho'(t)\mathbf{u}(t), \rho(t)\mathbf{u}'(t)) + \rho(t)^2|(\mathrm{d}\exp_p)_{\psi(t)}\mathbf{u}'(t)|^2 \\ &= \rho'(t)^2 + \rho(t)^2|(\mathrm{d}\exp_p)_{\psi(t)}\mathbf{u}'(t)|^2, \end{aligned}$$

Thus we have

$$\mathrm{length}(\exp_p \circ \psi) \geq \int_0^1 \rho'(t) \, dt = \rho(1) - \rho(0) = |a|. \quad \square$$

Notation. We write $\Omega(p, q)$ for the set of all piecewise C^1 curves from p to q .

We now wish to define a metric on M , in the sense of metric spaces.

Definition (Distance). Suppose M is connected, which is the same as it being path connected. Let $(p, q) \in M$. We define

$$d(p, q) = \inf_{\xi \in \Omega(p, q)} \mathrm{length}(\xi),$$

where

To see this is indeed a metric, All axioms of a metric space are obvious, apart from the non-negativity part.

Theorem. Let $p \in M$, and let ε be such that $\exp_p|_{B(0,\varepsilon)}$ is a diffeomorphism onto its image, and let U be the image. Then

- For any $q \in U$, there is a unique geodesic $\gamma \in \Omega(p, q)$ with $\ell(\gamma) < \varepsilon$. Moreover, $\ell(\gamma) = d(p, q)$, and is the unique curve that satisfies this property.
- For any point $q \in M$ with $d(p, q) < \varepsilon$, we have $q \in U$.
- If $q \in M$ is any point, $\gamma \in \Omega(p, q)$ has $\ell(\gamma) = d(p, q) < \varepsilon$, then γ is a geodesic.

Proof. Let $q = \exp_p(a)$. Then the path $\gamma(t) = \exp_p(ta)$ is a geodesic from p to q of length $|a| = r < \varepsilon$. This is clearly the only such geodesic, since $\exp_p|_{B(0,\varepsilon)}$ is a diffeomorphism.

Given any other path $\tilde{\gamma} \in \Omega(p, q)$, we want to show $\ell(\tilde{\gamma}) > \ell(\gamma)$. We let

$$\tau = \sup \left\{ t \in [0, 1] : \gamma([0, t]) \subseteq \exp_p(\overline{B(0, r)}) \right\}.$$

Note that if $\tau \neq 1$, then we must have $\gamma(\tau) \in \Sigma_r$, the geodesic sphere of radius r , otherwise we can continue extending. On the other hand, if $\tau = 1$, then we certainly have $\gamma(\tau) \in \Sigma_r$, since $\gamma(\tau) = q$. Then by local minimizing of length, we have

$$\ell(\tilde{\gamma}) \geq \ell(\tilde{\gamma}|_{[0, \tau]}) \geq r.$$

Note that we can always lift $\tilde{\gamma}[0, \tau]$ to a curve from 0 to a in T_pM , since \exp_p is a diffeomorphism in $B(0, \varepsilon)$.

By looking at the proof of the local minimizing of length, and using the same notation, we know that we have equality iff $\tau = 1$ and

$$\rho(t)^2 |(d \exp_p)_{\psi(t)} \psi(t) \mathbf{u}'(t)|^2 = 0$$

for all t . Since $d \exp_p$ is regular, this requires $\mathbf{u}'(t) = 0$ for all t (since $\rho(t) \neq 0$ when $t \neq 0$, or else we can remove the loop to get a shorter curve). This implies $\tilde{\gamma}$ lifts to a straight line in T_pM , i.e. is a geodesic.

Now given any $q \in M$ with $r = d(p, q) < \varepsilon$, we pick $r' \in [r, \varepsilon)$ and a path $\gamma \in \Omega(p, q)$ such that $\ell(\gamma) = r'$. We again let

$$\tau = \sup \left\{ t \in [0, 1] : \gamma([0, t]) \subseteq \exp_p(\overline{B(0, r')}) \right\}.$$

If $\tau \neq 1$, then we must have $\gamma(\tau) \in \Sigma_{r'}$, but lifting to T_pM , this contradicts the local minimizing of length.

The last part is an immediate consequence of the previous two. □

Corollary. The distance d on a Riemannian manifold is a metric, and induces the same topology on M as the C^∞ structure.

Definition (Minimal geodesic). A *minimal geodesic* is a curve $\gamma : [0, 1] \rightarrow M$ such that

$$d(\gamma(0), \gamma(1)) = \ell(\gamma).$$

One would certainly want a minimal geodesic to be an actual geodesic. This is an easy consequence of what we've got so far, using the observation that a sub-curve of a minimizing geodesic is still minimizing.

Corollary. Let $\gamma : [0, 1] \rightarrow M$ be a piecewise C^1 minimal geodesic with constant speed. Then γ is in fact a geodesic, and is in particular C^∞ .

Proof. We wlog γ is unit speed. Let $t \in [0, 1]$, and pick $\varepsilon > 0$ such that $\exp_p|_{B(0, \varepsilon)}$ is a diffeomorphism. Then by the theorem, $\gamma|_{[t, t + \frac{1}{2}\varepsilon]}$ is a geodesic. So γ is C^∞ on $(t, t + \frac{1}{2}\varepsilon)$, and satisfies the geodesic equations there.

Since we can pick ε continuously with respect to t by ODE theorems, any $t \in (0, 1)$ lies in one such neighbourhood. So γ is a geodesic. \square

While it is not true that geodesics are always minimal geodesics, this is locally true:

Corollary. Let $\gamma : [0, 1] \subseteq \mathbb{R} \rightarrow M$ be a C^2 curve with $|\dot{\gamma}|$ constant. Then this is a geodesic iff it is locally a minimal geodesic, i.e. for any $t \in [0, 1)$, there exists $\delta > 0$ such that

$$d(\gamma(t), \gamma(t + \delta)) = \ell(\gamma|_{[t, t + \delta]}).$$

Proof. This is just carefully applying the previous theorem without getting confused.


To prove \Rightarrow , suppose γ is a geodesic, and $t \in [0, 1)$. We wlog γ is unit speed. Then pick U and ε as in the previous theorem, and pick $\delta = \frac{1}{2}\varepsilon$. Then $\gamma|_{[t, t + \delta]}$ is a geodesic with length $< \varepsilon$ between $\gamma(t)$ and $\gamma(t + \delta)$, and hence must have minimal length.

To prove the converse, we note that for each t , the hypothesis tells us $\gamma|_{[t, t + \delta]}$ is a minimizing geodesic, and hence a geodesic, but the previous corollary. By continuity, γ must satisfy the geodesic equation at t . Since t is arbitrary, γ is a geodesic. \square

There is another sense in which geodesics are locally length minimizing. Instead of chopping up a path, we can say it is minimal "locally" in the space $\Omega(p, q)$. To do so, we need to give $\Omega(p, q)$ a topology, and we pick the topology of uniform convergence.

Theorem. Let $\gamma(t) = \exp_p(ta)$ be a geodesic, for $t \in [0, 1]$. Let $q = \gamma(1)$. Assume ta is a regular point for \exp_p for all $t \in [0, 1]$. Then there exists a neighbourhood of γ in $\Omega(p, q)$ such that for all ψ in this neighbourhood, $\ell(\psi) \geq \ell(\gamma)$, with equality iff $\psi = \gamma$ up to reparametrization.

Before we prove the result, we first look at why the two conditions are necessary. To see the necessity of ta being regular, we can consider the sphere and two antipodal points:




Then while the geodesic between them does minimize distance, it does not do so strictly.

We also do not guarantee global minimization of length. For example, we can consider the torus

$$T^n = \mathbb{R}^n / \mathbb{Z}^n.$$

This has a flat metric from \mathbb{R}^n , and the derivative of the exponential map is the “identity” on \mathbb{R}^n at all points. So the geodesics are the straight lines in \mathbb{R}^n . Now consider any two $p, q \in T^n$, then there are infinitely many geodesics joining them, but typically, only one of them would be the shortest.



Proof. The idea of the proof is that if ψ is any curve close to γ , then we can use the regularity condition to lift the curve back up to T_pM , and then apply our previous result.

Write $\varphi(t) = ta \in T_pM$. Then by the regularity assumption, for all $t \in [0, 1]$, we know \exp_p is a diffeomorphism of some neighbourhood $W(t)$ of $\varphi(t) = at \in T_pM$ onto the image. By compactness, we can cover $[0, 1]$ by finitely many such covers, say $W(t_1), \dots, W(t_n)$. We write $W_i = W(t_i)$, and we wlog assume

$$0 = t_0 < t_1 < \dots < t_k = 1.$$

By cutting things up, we may assume

$$\gamma([t_i, t_{i+1}]) \subseteq W_i.$$

We let

$$U = \bigcup \exp_p(W_i).$$

Again by compactness, there is some $\varepsilon < 0$ such that for all $t \in [t_i, t_{i+1}]$, we have $B(\gamma(t), \varepsilon) \subseteq W_i$.

Now consider any curve ψ of distance ε away from γ . Then $\psi([t_i, t_{i+1}]) \subseteq W_i$. So we can lift it up to T_pM , and the end point of the lift is a . So we are done by local minimization of length. \square

Note that the tricky part of doing the proof is to make sure the lift of ψ has the same end point as γ in T_pM , which is why we needed to do it neighbourhood by neighbourhood.

3.4 Completeness and the Hopf–Rinow theorem

There are some natural questions we can ask about geodesics. For example, we might want to know if geodesics can be extended to exist for all time. We might also be interested if distances can always be realized by geodesics. It turns out these questions have the same answer.

Definition (Geodesically complete). We say a manifold (M, g) is *geodesically complete* if each geodesic extends for all time. In other words, for all $p \in M$, \exp_p is defined on *all* of $T_p M$.

Example. The upper half plane

$$H^2 = \{(x, y) : y > 0\}$$

under the induced Euclidean metric is not geodesically complete. However, H^2 and \mathbb{R}^2 are diffeomorphic but \mathbb{R}^2 is geodesically complete.

The first theorem we will prove is the following:

Theorem. Let (M, g) be geodesically complete. Then any two points can be connected by a minimal geodesic.

In fact, we will prove something stronger — let $p \in M$, and suppose \exp_p is defined on all of $T_p M$. Then for all $q \in M$, there is a minimal geodesic between them.

To prove this, we need a lemma

Lemma. Let $p, q \in M$. Let

$$S_\delta = \{x \in M : d(x, p) = \delta\}.$$

Then for all sufficiently small δ , there exists $p_0 \in S_\delta$ such that

$$d(p, p_0) + d(p_0, q) = d(p, q).$$

Proof. For $\delta > 0$ small, we know $S_\delta = \Sigma_\delta$ is a geodesic sphere about p , and is compact. Moreover, $d(\cdot, q)$ is a continuous function. So there exists some $p_0 \in \Sigma_\delta$ that minimizes $d(\cdot, q)$.

Consider an arbitrary $\gamma \in \Omega(p, q)$. For the sake of sanity, we assume $\delta < d(p, q)$. Then there is some t such that $\gamma(t) \in \Sigma_\delta$, and

$$\ell(\gamma) \geq d(p, \gamma(t)) + d(\gamma(t), q) \geq d(p, p_0) + d(p_0, q).$$

So we know

$$d(p, q) \geq d(p, p_0) + d(p_0, q).$$

The triangle inequality gives the opposite direction. So we must have equality. \square

We can now prove the theorem.

Proof of theorem. We know \exp_p is defined on $T_p M$. Let $q \in M$. Let $q \in M$. We want a minimal geodesic in $\Omega(p, q)$. By the first lemma, there is some $\delta > 0$ and p_0 such that

$$d(p, p_0) = \delta, \quad d(p, p_0) + d(p_0, q) = d(p, q).$$

Also, there is some $v \in T_p M$ such that $\exp_p v = p_0$. We let

$$\gamma_p(t) = \exp_p \left(t \frac{v}{|v|} \right).$$

We let

$$I = \{t \in \mathbb{R} : d(q, \gamma_p(t)) + t = d(p, q)\}.$$

Then we know

- (i) $\delta \in I$
- (ii) I is closed by continuity.

Let

$$T = \sup\{I \cap [0, d(p, q)]\}.$$

Since I is closed, this is in fact a maximum. So $T \in I$. We claim that $T = d(p, q)$. If so, then $\gamma_p \in \Omega(p, q)$ is the desired minimal geodesic, and we are done.

Suppose this were not true. Then $T < d(p, q)$. We apply the lemma to $\tilde{p} = \gamma_p(T)$, and q remains as before. Then we can find $\varepsilon > 0$ and some $p_1 \in M$ with the property that

$$\begin{aligned} d(p_1, q) &= d(\gamma_p(T), q) - d(\gamma_p(T), p_1) \\ &= d(\gamma_p(T), q) - \varepsilon \\ &= d(p, q) - T - \varepsilon \end{aligned}$$

Hence we have

$$d(p, p_1) \geq d(p, q) - d(q, p_1) = T + \varepsilon.$$

Let γ_1 be the radial (hence minimal) geodesic from $\gamma_p(T)$ to p_1 . Now we know

$$\ell(\gamma_p|_{[0, T]}) + \ell(\gamma_1) = T + \varepsilon.$$

So γ_1 concatenated with $\gamma_p|_{[0, T]}$ is a length-minimizing geodesic from p to p_1 , and is hence a geodesic. So in fact p_1 lies on γ_p , say $p_1 = \gamma_p(T + s)$ for some s . Then $T + s \in I$, which is a contradiction. So we must have $T = d(p, q)$, and hence

$$d(q, \gamma_p(T)) + T = d(p, q),$$

hence $d(q, \gamma_p(T)) = 0$, i.e. $q = \gamma_p(T)$. □

Corollary (Hopf–Rinow theorem). For a connected Riemannian manifold (M, g) , the following are equivalent:

- (i) (M, g) is geodesically complete.
- (ii) For all $p \in M$, \exp_p is defined on all $T_p M$.
- (iii) For some $p \in M$, \exp_p is defined on all $T_p M$.
- (iv) Every closed and bounded subset of (M, d) is compact.
- (v) (M, d) is complete as a metric space.

Proof. (i) and (ii) are equivalent by definition. (ii) \Rightarrow (iii) is clear, and we proved (iii) \Rightarrow (i).

- (iii) \Rightarrow (iv): Let $K \subseteq M$ be closed and bounded. Then by boundedness, K is contained in $\exp_p(\overline{B(0, R)})$. Let K' be the pre-image of K under \exp_p . Then it is a closed and bounded subset of \mathbb{R}^n , hence compact. Then K is the continuous image of a compact set, hence compact.
- (iv) \Rightarrow (v): This is a general topological fact.
- (v) \Rightarrow (i): Let $\gamma(t) : I \rightarrow \mathbb{R}$ be a geodesic, where $I \subseteq \mathbb{R}$. We wlog $|\dot{\gamma}| \equiv 1$. Suppose $I \neq \mathbb{R}$. We wlog $\sup I = a < \infty$. Then $\lim_{t \rightarrow a} \gamma(t)$ exist by completeness, and hence $\gamma(a)$ exists. Since geodesics are locally defined near a , we can pick a geodesic in the direction of $\lim_{t \rightarrow a} \gamma'(t)$. So we can extend γ further, which is a contradiction. □

3.5 Variations of arc length and energy

This section is mostly a huge computation. As we previously saw, geodesics are locally length-minimizing, and we shall see that another quantity, namely the energy is also a useful thing to consider, as minimizing the energy also forces the parametrization to be constant speed.

To make good use of these properties of geodesics, it is helpful to compute explicitly expressions for how length and energy change along variations. The computations are largely uninteresting, but it will pay off.

Definition (Energy). The *energy function* $E : \Omega(p, q) \rightarrow \mathbb{R}$ is given by

$$E(\gamma) = \frac{1}{2} \int_0^T |\dot{\gamma}|^2 dt,$$

where $\gamma : [0, T] \rightarrow M$.

Recall that $\Omega(p, q)$ is defined as the space of piecewise C^1 curves. Often, we will make the simplifying assumption that all curves are in fact C^1 . It doesn't really matter.

Note that the length of a curve is independent of parametrization. Thus, if we are interested in critical points, then the critical points cannot possibly be isolated, as we can just re-parametrize to get a nearby path with the same length. On the other hand, the energy E *does* depend on parametrization. This does have isolated critical points, which is technically very convenient.

Proposition. Let $\gamma_0 : [0, T] \rightarrow M$ be a path from p to q such that for all $\gamma \in \Omega(p, q)$ with $\gamma : [0, T] \rightarrow M$, we have $E(\gamma) \geq E(\gamma_0)$. Then γ_0 must be a geodesic.

Recall that we already had such a result for length instead of energy. The proof is just the application of Cauchy-Schwartz.

Proof. By the Cauchy-Schwartz inequality, we have

$$\int_0^T |\dot{\gamma}|^2 dt \geq \left(\int_0^T |\dot{\gamma}(t)| dt \right)^2$$

with equality iff $|\dot{\gamma}|$ is constant. In other words,

$$E(\gamma) \geq \frac{\ell(\gamma)^2}{2T}.$$

So we know that if γ_0 minimizes energy, then it must be constant speed. Now given any γ , if we just care about its length, then we may wlog it is constant speed, and then

$$\ell(\gamma) = \sqrt{2E(\gamma)T} \geq \sqrt{2E(\gamma_0)T} = \ell(\gamma_0).$$

So γ_0 minimizes length, and thus γ_0 is a geodesic. \square

We shall consider smooth variations $H(t, s)$ of $\gamma_0(t) = H(t, 0)$. We require that $H : [0, T] \times (-\varepsilon, \varepsilon) \rightarrow M$ is smooth. Since we are mostly just interested

in what happens “near” $s = 0$, it is often convenient to just consider the corresponding vector field along γ :

$$Y(t) = \left. \frac{\partial H}{\partial s} \right|_{s=0} = (dH)_{(t,0)} \frac{\partial}{\partial s},$$

Conversely, given any such vector field Y , we can generate a variation H that gives rise to Y . For example, we can put

$$H(t, s) = \exp_{\gamma_0(t)}(sY(t)),$$

which is valid on some neighbourhood of $[0, T] \times \{0\}$. If $Y(0) = 0 = Y(T)$, then we can choose H fixing end-points of γ_0 .

Theorem (First variation formula).

(i) For any variation H of γ , we have

$$\left. \frac{d}{ds} E(\gamma_s) \right|_{s=0} = g(Y(t), \dot{\gamma}(t)) \Big|_0^T - \int_0^T g \left(Y(t), \frac{\nabla}{dt} \dot{\gamma}(t) \right) dt. \quad (*)$$

(ii) The critical points, i.e. the γ such that

$$\left. \frac{d}{ds} E(\gamma_s) \right|_{s=0}$$

for all (end-point fixing) variation H of γ , are geodesics.

(iii) If $|\dot{\gamma}_s(t)|$ is constant for each fixed $s \in (-\varepsilon, \varepsilon)$, and $|\dot{\gamma}(t)| \equiv 1$, then

$$\left. \frac{d}{ds} E(\gamma_s) \right|_{s=0} = \left. \frac{d}{ds} \ell(\gamma_s) \right|_{s=0}$$

(iv) If γ is a critical point of the length, then it must be a reparametrization of a geodesic.

This is just some calculations.

Proof. We will assume that we can treat $\frac{\partial}{\partial s}$ and $\frac{\partial}{\partial t}$ as vector fields on an embedded submanifold, even though H is not necessarily a local embedding.

The result can be proved without this assumption, but will require more technical work.

(i) We have

$$\begin{aligned} \frac{1}{2} \frac{\partial}{\partial s} g(\dot{\gamma}_s(t), \dot{\gamma}_s(t)) &= g \left(\frac{\nabla}{ds} \dot{\gamma}_s(t), \dot{\gamma}_s(t) \right) \\ &= g \left(\frac{\nabla}{dt} \frac{\partial H}{\partial s}(t, s), \frac{\partial H}{\partial t}(t, s) \right) \\ &= \frac{\partial}{\partial t} g \left(\frac{\partial H}{\partial s}, \frac{\partial H}{\partial t} \right) - g \left(\frac{\partial H}{\partial s}, \frac{\nabla}{dt} \frac{\partial H}{\partial t} \right). \end{aligned}$$

Comparing with what we want to prove, we see that we get what we want by integrating $\int_0^T dt$, and then putting $s = 0$, and then noting that

$$\left. \frac{\partial H}{\partial s} \right|_{s=0} = Y, \quad \left. \frac{\partial H}{\partial t} \right|_{s=0} = \dot{\gamma}.$$

(ii) If γ is a geodesic, then

$$\frac{\nabla}{dt}\dot{\gamma}(t) = 0.$$

So the integral on the right hand side of (*) vanishes. Also, we have $Y(0) = 0 = Y(T)$. So the RHS vanishes.

Conversely, suppose γ is a critical point for E . Then choose H with

$$Y(t) = f(t)\frac{\nabla}{dt}\dot{\gamma}(t)$$

for some $f \in C^\infty[0, T]$ such that $f(0) = f(T) = 0$. Then we know

$$\int_0^T f(t) \left| \frac{\nabla}{dt}\dot{\gamma}(t) \right|^2 dt = 0,$$

and this is true for all f . So we know

$$\frac{\nabla}{dt}\dot{\gamma} = 0.$$

(iii) This is evident from the previous proposition. Indeed, we fix $[0, T]$, then for all H , we have

$$E(\gamma_s) = \frac{\ell(\gamma_s)^2}{2T},$$

and so

$$\frac{d}{ds}E(\gamma_s) \Big|_{s=0} = \frac{1}{T}\ell(\gamma_s) \frac{d}{ds}\ell(\gamma_s) \Big|_{s=0},$$

and when $s = 0$, the curve is parametrized by arc-length, so $\ell(\gamma_s) = T$.

(iv) By reparametrization, we may wlog $|\dot{\gamma}| \equiv 1$. Then γ is a critical point for ℓ , hence for E , hence a geodesic. \square

Often, we are interested in more than just whether the curve is a critical point. We want to know if it maximizes or minimizes energy. Then we need more than the “first derivative”. We need the “second derivative” as well.

Theorem (Second variation formula). Let $\gamma(t) : [0, T] \rightarrow M$ be a geodesic with $|\dot{\gamma}| = 1$. Let $H(t, s)$ be a variation of γ . Let

$$Y(t, s) = \frac{\partial H}{\partial s}(t, s) = (dH)_{(t,s)} \frac{\partial}{\partial s}.$$

Then

(i) We have

$$\frac{d^2}{ds^2}E(\gamma_s) \Big|_{s=0} = g \left(\frac{\nabla Y}{ds}(t, 0), \dot{\gamma} \right) \Big|_0^T + \int_0^T (|Y'|^2 - R(Y, \dot{\gamma}, Y, \dot{\gamma})) dt.$$

(ii) Also

$$\begin{aligned} \frac{d^2}{ds^2}\ell(\gamma_s) \Big|_{s=0} &= g \left(\frac{\nabla Y}{ds}(t, 0), \dot{\gamma}(t) \right) \Big|_0^T \\ &\quad + \int_0^T (|Y'|^2 - R(Y, \dot{\gamma}, Y, \dot{\gamma}) - g(\dot{\gamma}, Y')^2) dt, \end{aligned}$$

where R is the $(4, 0)$ curvature tensor, and

$$Y'(t) = \frac{\nabla Y}{dt}(t, 0).$$

Putting

$$Y_n = Y - g(Y, \dot{\gamma})\dot{\gamma}$$

for the normal component of Y , we can write this as

$$\frac{d^2}{ds^2} \ell(\gamma_s) \Big|_{s=0} = g \left(\frac{\nabla Y_n}{ds}(t, 0), \dot{\gamma}(t) \right) \Big|_0^T + \int_0^T (|Y_n'|^2 - R(Y_n, \dot{\gamma}, Y_n, \dot{\gamma})) dt.$$

Note that if we have fixed end points, then the first terms in the variation formulae vanish.

Proof. We use

$$\frac{d}{ds} E(\gamma_s) = g(Y(t, s), \dot{\gamma}_s(t)) \Big|_{t=0}^{t=T} - \int_0^T g \left(Y(t, s), \frac{\nabla}{dt} \dot{\gamma}_s(t) \right) dt.$$

Taking the derivative with respect to s again gives

$$\begin{aligned} \frac{d^2}{ds^2} E(\gamma_s) &= g \left(\frac{\nabla Y}{ds}, \dot{\gamma} \right) \Big|_{t=0}^T + g \left(Y, \frac{\nabla}{ds} \dot{\gamma}_s \right) \Big|_{t=0}^T \\ &\quad - \int_0^T \left(g \left(\frac{\nabla Y}{ds}, \frac{\nabla}{dt} \dot{\gamma}_s \right) + g \left(Y, \frac{\nabla}{ds} \frac{\nabla}{dt} \dot{\gamma} \right) \right) dt. \end{aligned}$$

We now use that

$$\begin{aligned} \frac{\nabla}{ds} \frac{\nabla}{dt} \dot{\gamma}_s(t) &= \frac{\nabla}{dt} \frac{\nabla}{ds} \dot{\gamma}_s(t) + R \left(\frac{\partial H}{\partial s}, \frac{\partial H}{\partial t} \right) \dot{\gamma}_s \\ &= \left(\frac{\nabla}{dt} \right)^2 Y(t, s) + R \left(\frac{\partial H}{\partial s}, \frac{\partial H}{\partial t} \right) \dot{\gamma}_s. \end{aligned}$$

We now set $s = 0$, and then the above gives

$$\begin{aligned} \frac{d^2}{ds^2} E(\gamma_s) \Big|_{s=0} &= g \left(\frac{\nabla Y}{ds}, \dot{\gamma} \right) \Big|_0^T + g \left(Y, \frac{\nabla}{ds} \dot{\gamma} \right) \Big|_0^T \\ &\quad - \int_0^T \left[g \left(Y, \left(\frac{\nabla}{dt} \right)^2 Y \right) + R(\dot{\gamma}, Y, \dot{\gamma}, Y) \right] dt. \end{aligned}$$

Finally, applying integration by parts, we can write

$$- \int_0^T g \left(Y, \left(\frac{\nabla}{dt} \right)^2 Y \right) dt = - g \left(Y, \frac{\nabla}{dt} Y \right) \Big|_0^T + \int_0^T \left| \frac{\nabla Y}{dt} \right|^2 dt.$$

Finally, noting that

$$\frac{\nabla}{ds} \dot{\gamma}(s) = \frac{\nabla}{dt} Y(t, s),$$

we find that

$$\frac{d^2}{ds^2} E(\gamma_s) \Big|_{s=0} = g \left(\frac{\nabla Y}{ds}, \dot{\gamma} \right) \Big|_0 + \int_0^T (|Y'|^2 - R(Y, \dot{\gamma}, Y, \dot{\gamma})) dt.$$

It remains to prove the second variation of length. We first differentiate

$$\frac{d}{ds} \ell(\gamma_s) = \int_0^T \frac{1}{2\sqrt{g(\dot{\gamma}_s, \dot{\gamma}_s)}} \frac{\partial}{\partial s} g(\dot{\gamma}_s, \dot{\gamma}_s) dt.$$

Then the second derivative gives

$$\frac{d^2}{ds^2} \ell(\gamma_s) \Big|_{s=0} = \int_0^T \left[\frac{1}{2} \frac{\partial^2}{\partial s^2} g(\dot{\gamma}_s, \dot{\gamma}_s) \Big|_{s=0} - \frac{1}{4} \left(\frac{\partial}{\partial s} g(\dot{\gamma}_s, \dot{\gamma}_s) \right)^2 \Big|_{s=0} \right] dt,$$

where we used the fact that $g(\dot{\gamma}, \dot{\gamma}) = 1$.

We notice that the first term can be identified with the derivative of the energy function. So we have

$$\frac{d^2}{ds^2} \ell(\gamma_s) \Big|_{s=0} = \frac{d^2}{ds^2} E(\gamma_s) \Big|_{s=0} - \int_0^T \left(g \left(\dot{\gamma}_s, \frac{\nabla}{ds} \dot{\gamma}_s \right) \Big|_{s=0} \right)^2 dt.$$

So the second part follows from the first. \square

3.6 Applications

This finally puts us in a position to prove something more interesting.

Synge's theorem

We are first going to prove the following remarkable result relating curvature and topology:

Theorem (Synge's theorem). Every compact orientable Riemannian manifold (M, g) such that $\dim M$ is even and has $K(g) > 0$ for all planes at $p \in M$ is simply connected.

We can see that these conditions are indeed necessary. For example, we can consider $\mathbb{RP}^2 = S^2 / \pm 1$ with the induced metric from S^2 . Then this is compact with positive sectional curvature, but it is not orientable. Indeed it is not simply connected.

Similarly, if we take \mathbb{RP}^3 , then this has odd dimension, and the theorem breaks.

Finally, we do need strict inequality, e.g. the flat torus is not simply connected.

We first prove a technical lemma.

Lemma. Let M be a compact manifold, and $[\alpha]$ a non-trivial homotopy class of closed curves in M . Then there is a closed minimal geodesic in $[\alpha]$.

Proof. Since M is compact, we can pick some $\varepsilon > 0$ such that for all $p \in M$, the map $\exp_p|_{B(0,p)}$ is a diffeomorphism.

Let $\ell = \inf_{\gamma \in [\alpha]} \ell(\gamma)$. We know that $\ell > 0$, otherwise, there exists a γ with $\ell(\gamma) < \varepsilon$. So γ is contained in some geodesic coordinate neighbourhood, but then α is contractible. So ℓ must be positive.

Then we can find a sequence $\gamma_n \in [\alpha]$ with $\gamma_n : [0, 1] \rightarrow M$, $|\dot{\gamma}|$ constant, such that

$$\lim_{n \rightarrow \infty} \ell(\gamma_n) = \ell.$$

Choose

$$0 = t_0 < t_1 < \cdots < t_k = 1$$

such that

$$t_{i+1} - t_i < \frac{\varepsilon}{2\ell}.$$

So it follows that

$$d(\gamma_n(t_i), \gamma_n(t_{i+1})) < \varepsilon$$

for all n sufficiently large and all i . Then again, we can replace $\gamma_n|_{[t_i, t_{i+1}]}$ by a radial geodesic without affecting the limit $\lim \ell(\gamma_n)$.

Then we exploit the compactness of M (and the unit sphere) again, and pass to a subsequence of $\{\gamma_n\}$ so that $\gamma_n(t_i), \dot{\gamma}_n(t_i)$ are all convergent for every fixed i as $n \rightarrow \infty$. Then the curves converges to some

$$\gamma_n \rightarrow \hat{\gamma} \in [\alpha],$$

given by joining the limits $\lim_{n \rightarrow \infty} \gamma_n(t_i)$. Then we know that the length converges as well, and so we know $\hat{\gamma}$ is minimal among curves in $[\alpha]$. So $\hat{\gamma}$ is locally minimal, hence a geodesic. So we can take $\gamma = \hat{\gamma}$, and we are done. \square

Proof of Synge's theorem. Suppose M satisfies the hypothesis, but $\pi_1(M) \neq \{1\}$. So there is a path α with $[\alpha] \neq 1$, i.e. it cannot be contracted to a point. By the lemma, we pick a representative γ of $[\alpha]$ that is a closed, minimal geodesic.

We now prove the theorem. We may wlog assume $|\dot{\gamma}| = 1$, and t ranges in $[0, T]$. Consider a vector field $X(t)$ for $0 \leq t \leq T$ along $\gamma(t)$ such that

$$\frac{\nabla X}{dt} = 0, \quad g(X(0), \dot{\gamma}(0)) = 0.$$

Note that since g is a geodesic, we know

$$g(X(t), \dot{\gamma}(t)) = 0,$$

for all $t \in [0, T]$ as parallel transport preserves the inner product. So $X(T) \perp \dot{\gamma}(T) = \dot{\gamma}(0)$ since we have a closed curve.

We consider the map P that sends $X(0) \mapsto X(T)$. This is a linear isometry of $(\dot{\gamma}(0))^\perp$ with itself that preserves orientation. So we can think of P as a map

$$P \in \text{SO}(2n-1),$$

where $\dim M = 2n$. It is an easy linear algebra exercise to show that every element of $\text{SO}(2n-1)$ must have an eigenvector of eigenvalue 1. So we can find $v \in T_p M$ such that $v \perp \dot{\gamma}(0)$ and $P(v) = v$. We take $X(0) = v$. Then we have $X(T) = v$.

Consider now a variation $H(t, s)$ inducing this $X(t)$. We may assume $|\dot{\gamma}_s|$ is constant. Then

$$\frac{d}{ds} \ell(\gamma_s)|_{s=0} = 0$$

as γ is minimal. Moreover, since it is a minimum, the second derivative must be positive, or at least non-negative. Is this actually the case?

We look at the second variation formula of length. Using the fact that the loop is closed, the formula reduces to

$$\left. \frac{d^2}{ds^2} \ell(\gamma_s) \right|_{s=0} = - \int_0^T R(X, \dot{\gamma}, X, \dot{\gamma}) dt.$$

But we assumed the sectional curvature is positive. So the second variation is negative! This is a contradiction. \square

Conjugate points

Recall that when a geodesic starts moving, for a short period of time, it is length-minimizing. However, in general, if we keep on moving for a long time, then we cease to be minimizing. It is useful to characterize when this happens.

As before, for a vector field J along a curve $\gamma(t)$, we will write

$$J' = \frac{\nabla J}{dt}.$$

Definition (Conjugate points). Let $\gamma(t)$ be a geodesic. Then

$$p = \gamma(\alpha), \quad q = \gamma(\beta)$$

are *conjugate points* if there exists some non-trivial J such that $J(\alpha) = 0 = J(\beta)$.

It is easy to see that this does not depend on parametrization of the curve, because Jacobi fields do not.

Proposition.

- (i) If $\gamma(t) = \exp_p(ta)$, and $q = \exp_p(\beta a)$ is conjugate to p , then q is a singular value of \exp .
- (ii) Let J be as in the definition. Then J must be pointwise normal to $\dot{\gamma}$.

Proof.

- (i) We wlog $[\alpha, \beta] = [0, 1]$. So $J(0) = 0 = J(1)$. We $a = \dot{\gamma}(0)$ and $w = J'(0)$. Note that a, w are both non-zero, as Jacobi fields are determined by initial conditions. Then $q = \exp_p(a)$.

We have shown earlier that if $J(0) = 0$, then

$$J(t) = (d \exp_p)_{ta}(tw)$$

for all $0 \leq t \leq 1$. So it follows $(d \exp_p)_a(w) = J(1) = 0$. So $(d \exp_p)_a$ has non-trivial kernel, and hence isn't surjective.

- (ii) We claim that any Jacobi field J along a geodesic γ satisfies

$$g(J(t), \dot{\gamma}(t)) = g(J'(0), \dot{\gamma}(0))t + g(J(0), \dot{\gamma}(0)).$$

To prove this, we note that by the definition of geodesic and Jacobi fields, we have

$$\frac{d}{dt} g(J', \dot{\gamma}) = g(J'', \dot{\gamma}(0)) = -g(R(\dot{\gamma}, J), \dot{\gamma}, \dot{\gamma}) = 0$$

by symmetries of R . So we have

$$\frac{d}{dt}g(J, \dot{\gamma}) = g(J'(t), \dot{\gamma}(t)) = g(J'(0), \dot{\gamma}(0)).$$

Now integrating gives the desired result.

This result tells us $g(J(t), \dot{\gamma}(t))$ is a linear function of t . But we have

$$g(J(0), \dot{\gamma}(0)) = g(J(1), \dot{\gamma}(1)) = 0.$$

So we know $g(J(t), \dot{\gamma}(t))$ is constantly zero. \square

From the proof, we see that for any Jacobi field with $J(0) = 0$, we have

$$g(J'(0), \dot{\gamma}(0)) = 0 \iff g(J(t), \dot{\gamma}(t)) = \text{constant}.$$

This implies that the dimension of the normal Jacobi fields along γ satisfying $J(0) = 0$ is $\dim M - 1$.

Example. Consider $M = S^2 \subseteq \mathbb{R}^3$ with the round metric, i.e. the “obvious” metric induced from \mathbb{R}^3 . We claim that $N = (0, 0, 1)$ and $S = (0, 1, 0)$ are conjugate points.


To construct a Jacobi field, instead of trying to mess with the Jacobi equation, we construct a variation by geodesics. We let

$$f(t, s) = \begin{pmatrix} \cos s \sin t \\ \sin s \sin t \\ \cos t \end{pmatrix}.$$


We see that when $s = 0$, this is the great-circle in the (x, z) -plane. Then we have a Jacobi field

$$J(t) = \left. \frac{\partial f}{\partial s} \right|_{s=0} = \begin{pmatrix} 0 \\ \sin t \\ 0 \end{pmatrix}.$$

This is then a Jacobi field that vanishes at N and S .



When we are *at* the conjugate point, then there are many adjacent curves whose length is equal to ours. If we extend our geodesic *beyond* the conjugate point, then it is no longer even locally minimal:



We can push the geodesic slightly over and the length will be shorter. On the other hand, we proved that up to the conjugate point, the geodesic is always locally minimal.

It turns out this phenomenon is generic:

Theorem. Let $\gamma : [0, 1] \rightarrow M$ be a geodesic with $\gamma(0) = p$, $\gamma(1) = q$ such that p is conjugate to some $\gamma(t_0)$ for some $t_0 \in (0, 1)$. Then there is a piecewise smooth variation of $f(t, s)$ with $f(t, 0) = \gamma(t)$ such that

$$f(0, s) = p, \quad f(1, s) = q$$

and $\ell(f(\cdot, s)) < \ell(\gamma)$ whenever $s \neq 0$ is small.

The proof is a generalization of the example we had above. We know that up to the conjugate point, we have a Jacobi field that allows us to vary the geodesic without increasing the length. We can then give it a slight “kick” and then the length will decrease.

Proof. By the hypothesis, there is a $J(t)$ defined on $t \in [0, 1]$ and $t_0 \in (0, 1)$ such that

$$J(t) \perp \dot{\gamma}(t)$$

for all t , and $J(0) = J(t_0) = 0$ and $J \not\equiv 0$. Then $J'(t_0) \neq 0$.

We define a parallel vector field Z_1 along γ by $Z_1(t_0) = -J'(t_0)$. We pick $\theta \in C^\infty[0, 1]$ such that $\theta(0) = \theta(1) = 0$ and $\theta(t_0) = 1$.

Finally, we define

$$Z = \theta Z_1,$$

and for $\alpha \in \mathbb{R}$, we define

$$Y_\alpha(t) = \begin{cases} J(t) + \alpha Z(t) & 0 \leq t \leq t_0 \\ \alpha Z(t) & t_0 \leq t \leq 1 \end{cases}.$$

We notice that this is not smooth at t_0 , but is just continuous. We will postpone the choice of α to a later time.

We know $Y_\alpha(t)$ arises from a piecewise C^∞ variation of γ , say $H_\alpha(t, s)$. The technical claim is that the second variation of length corresponding to $Y_\alpha(t)$ is negative for some α .

We denote by $I(X, Y)_T$ the symmetric bilinear form that gives rise to the second variation of length with fixed end points. If we make the additional assumption that X, Y are normal along γ , then the formula simplifies, and reduces to

$$I(X, Y)_T = \int_0^T (g(X', Y') - R(X, \dot{\gamma}, Y, \dot{\gamma})) dt.$$

Then for $H_\alpha(t, s)$, we have

$$\begin{aligned} \left. \frac{d^2}{ds^2} \ell(\gamma_s) \right|_{s=0} &= I_1 + I_2 + I_3 \\ I_1 &= I(J, J)_{t_0} \\ I_2 &= 2\alpha I(J, Z)_{t_0} \\ I_3 &= \alpha^2 I(Z, Z)_1. \end{aligned}$$

We look at each term separately.

We first claim that $I_1 = 0$. We note that

$$\frac{d}{dt}g(J, J') = g(J', J') + g(J, J''),$$

and $g(J, J'')$ added to the curvature vanishes by the Jacobi equation. Then by integrating by parts and applying the boundary condition, we see that I_1 vanishes.

Also, by integrating by parts, we find

$$I_2 = 2\alpha g(Z, J')\Big|_0^{t_0}.$$

Whence

$$\frac{d^2}{ds^2}\ell(\gamma_s)\Big|_{s=0} = -2\alpha|J'(t_0)|^2 + \alpha^2 I(Z, Z)_1.$$

Now if $\alpha > 0$ is very very small, then the linear term dominates, and this is negative. Since the first variation vanishes (γ is a geodesic), we know this is a local maximum of length. \square

Note that we made a compromise in the theorem by doing a piecewise C^∞ variation instead of a smooth one, but of course, we can fix this by making a smooth approximation.

Bonnet–Myers diameter theorem

We are going to see yet another application of our previous hard work, which may also be seen as an interplay between curvature topology. In case it isn't clear, all our manifolds are connected.

Definition (Diameter). The *diameter* of a Riemannian manifold (M, g) is

$$\text{diam}(M, g) = \sup_{p, q \in M} d(p, q).$$

Of course, this definition is valid for any metric space.

Example. Consider the sphere

$$S^{n-1}(r) = \{x \in \mathbb{R}^n : |x| = r\},$$

with the induced “round” metric. Then

$$\text{diam}(S^{n-1}(r)) = \pi r.$$

It is an exercise to check that

$$K \equiv \frac{1}{r^2}.$$

We will also need the following notation:

Notation. Let h, \hat{h} be two symmetric bilinear forms on a real vector space. We say $h \geq \hat{h}$ if $h - \hat{h}$ is non-negative definite.

If $h, \hat{h} \in \Gamma(S^2 T^* M)$ are fields of symmetric bilinear forms, we write $h \geq \hat{h}$ if $h_p \geq \hat{h}_p$ for all $p \in M$.

The following will also be useful:

Definition (Riemannian covering map). Let (M, g) and (\tilde{M}, \tilde{g}) be two Riemannian manifolds, and $f : \tilde{M} \rightarrow M$ be a smooth covering map. We say f is a *Riemannian covering map* if it is a local isometry. Alternatively, $f^*g = \tilde{g}$. We say \tilde{M} is a Riemannian cover of M .

Recall that if f is in fact a universal cover, i.e. \tilde{M} is simply connected, then we can (non-canonically) identify $\pi_1(M)$ with $f^{-1}(p)$ for any point $p \in M$.

Definition (Bonnet–Myers diameter theorem). Let (M, g) be a complete n -dimensional manifold with

$$\text{Ric}(g) \geq \frac{n-1}{r^2}g,$$

where $r > 0$ is some positive number. Then

$$\text{diam}(M, g) \leq \text{diam } S^n(r) = \pi r.$$

In particular, M is compact and $\pi_1(M)$ is finite.

Proof. Consider any $L < \text{diam}(M, g)$. Then by definition (and Hopf–Rinow), we can find $p, q \in M$ such that $d(p, q) = L$, and a minimal geodesic $\gamma \in \Omega(p, q)$ with $\ell(\gamma) = d(p, q)$. We parametrize $\gamma : [0, L] \rightarrow M$ so that $|\dot{\gamma}| = 1$.

Now consider any vector field Y along γ such that $Y(p) = 0 = Y(q)$. Since Γ is a minimal geodesic, it is a critical point for ℓ , and the second variation $I(Y, Y)_{[0, L]}$ is non-negative (recall that the second variation has fixed end points).

We extend $\dot{\gamma}(0)$ to an orthonormal basis of $T_p M$, say $\dot{\gamma}(0) = e_1, e_2, \dots, e_n$. We further let X_i be the unique vector field such that

$$X_i' = 0, \quad X_i(0) = e_i.$$

In particular, $X_1(t) = \dot{\gamma}(t)$.

For $i = 2, \dots, n$, we put

$$Y_i(t) = \sin\left(\frac{\pi t}{L}\right) X_i(t).$$

Then after integrating by parts, we find that we have

$$I(Y_i, Y_i)_{[0, L]} = - \int_0^L g(Y_i'' + R(\dot{\gamma}, Y_i)Y_i, \dot{\gamma}) dt$$

Using the fact that X_i is parallel, this can be written as

$$= \int_0^L \sin^2 \frac{\pi t}{L} \left(\frac{\pi^2}{L^2} - R(\dot{\gamma}, X_i, \dot{\gamma}, X_i) \right) dt,$$

and since this is length minimizing, we know this is ≥ 0 .

We note that we have $R(\dot{\gamma}, X_1, \dot{\gamma}, X_1) = 0$. So we have

$$\sum_{i=2}^n R(\dot{\gamma}, X_i, \dot{\gamma}, X_i) = \text{Ric}(\dot{\gamma}, \dot{\gamma}).$$

So we know

$$\sum_{i=2}^n I(Y_i, Y_i) = \int_0^L \sin^2 \frac{\pi t}{L} \left((n-1) \frac{\pi^2}{L} - \text{Ric}(\dot{\gamma}, \dot{\gamma}) \right) dt \geq 0.$$

We also know that

$$\text{Ric}(\dot{\gamma}, \dot{\gamma}) \geq \frac{n-1}{r^2}$$

by hypothesis. So this implies that

$$\frac{\pi^2}{L^2} \geq \frac{1}{r^2}.$$

This tells us that

$$L \leq \pi r.$$

Since L is any number less than $\text{diam}(M, g)$, it follows that

$$\text{diam}(M, g) \leq \pi r.$$

Since M is known to be complete, by Hopf-Rinow theorem, any closed bounded subset is compact. But M itself is closed and bounded! So M is compact.

To understand the fundamental group, we simply have to consider a universal Riemannian cover $f: \tilde{M} \rightarrow M$. We know such a topological universal covering space must exist by general existence theorems. We can then pull back the differential structure and metric along f , since f is a local homeomorphism. So this gives a universal Riemannian cover of M . But the hypothesis of the theorem is local, so it is also satisfied for \tilde{M} . So it is also compact. Since $f^{-1}(p)$ is a closed discrete subset of a compact space, it is finite, and we are done. \square

It is an easy exercise to show that the hypothesis on the Ricci curvature cannot be weakened to just saying that the Ricci curvature is positive definite.

Hadamard–Cartan theorem

To prove the next result, we need to talk a bit more about coverings.

Proposition. Let (M, g) and (N, h) be Riemannian manifolds, and suppose M is complete. Suppose there is a smooth surjection $f: M \rightarrow N$ that is a local diffeomorphism. Moreover, suppose that for any $p \in M$ and $v \in T_p M$, we have $|\text{d}f_p(v)|_h \geq |v|$. Then f is a covering map.

Proof. By general topology, it suffices to prove that for any smooth curve $\gamma: [0, 1] \rightarrow N$, and any $q \in M$ such that $f(q) = \gamma(0)$, there exists a lift of γ starting from q .

$$\begin{array}{ccc} & & M \\ & \nearrow \tilde{\gamma} & \downarrow f \\ [0, 1] & \xrightarrow{\gamma} & N \end{array}$$

From the hypothesis, we know that $\tilde{\gamma}$ exists on $[0, \varepsilon_0]$ for some “small” $\varepsilon_0 > 0$. We let

$$I = \{0 \leq \varepsilon \leq: \tilde{\gamma} \text{ exists on } [0, \varepsilon]\}.$$

We immediately see this is non-empty, since it contains ε_0 . Moreover, it is not difficult to see that I is open in $[0, 1]$, because f is a local diffeomorphism. So it suffices to show that I is closed.

We let $\{t_n\}_{n=1}^\infty \subseteq I$ be such that $t_{n+1} > t_n$ for all n , and

$$\lim_{n \rightarrow \infty} t_n = \varepsilon_1.$$

Using Hopf-Rinow, either $\{\tilde{\gamma}(t_n)\}$ is contained in some compact K , or it is unbounded. We claim that unboundedness is impossible. We have

$$\begin{aligned} \ell(\gamma) &\geq \ell(\gamma|_{[0, t_n]}) = \int_0^{t_n} |\dot{\gamma}| \, dt \\ &= \int_0^{t_n} |df_{\tilde{\gamma}(t)} \dot{\tilde{\gamma}}(t)| \, dt \\ &\geq \int_0^{t_n} |\dot{\tilde{\gamma}}| \, dt \\ &= \ell(\tilde{\gamma}|_{[0, t_n]}) \\ &\geq d(\tilde{\gamma}(0), \tilde{\gamma}(t_n)). \end{aligned}$$

So we know this is bounded. So by compactness, we can find some x such that $\tilde{\gamma}(t_{n_\ell}) \rightarrow x$ as $\ell \rightarrow \infty$. There exists an open $x \in V \subseteq M$ such that $f|_V$ is a diffeomorphism.

Since there are extensions of $\tilde{\gamma}$ to each t_n , eventually we get an extension to within V , and then we can just lift directly, and extend it to ε_1 . So $\varepsilon_1 \in I$. So we are done. \square

Corollary. Let $f : M \rightarrow N$ be a local isometry onto N , and M be complete. Then f is a covering map.

Note that since N is (assumed to be) connected, we know f is necessarily surjective. To see this, note that the completeness of M implies completeness of $f(M)$, hence $f(M)$ is closed in N , and since it is a local isometry, we know f is in particular open. So the image is open and closed, hence $f(M) = N$.

For a change, our next result will assume a *negative* curvature, instead of a positive one!

Theorem (Hadamard–Cartan theorem). Let (M^n, g) be a complete Riemannian manifold such that the sectional curvature is always non-positive. Then for every point $p \in M$, the map $\exp_p : T_p M \rightarrow M$ is a covering map. In particular, if $\pi_1(M) = 0$, then M is diffeomorphic to \mathbb{R}^n .

We will need one more technical lemma.

Lemma. Let $\gamma(t)$ be a geodesic on (M, g) such that $K \leq 0$ along γ . Then γ has no conjugate points.

Proof. We write $\gamma(0) = p$. Let $I(t)$ be a Jacobi field along γ , and $J(0) = 0$. We claim that if J is not identically zero, then J does not vanish everywhere else.

We consider the function

$$f(t) = g(J(t), J(t)) = |J(t)|^2.$$

Then $f(0) = f'(0) = 0$. Consider

$$\frac{1}{2}f''(t) = g(J''(t), J(t)) + g(J'(t), J'(t)) = g(J', J') - R(\dot{\gamma}, J, \dot{\gamma}, J) \geq 0.$$

So f is a convex function, and so we are done. \square

We can now prove the theorem.

Proof of theorem. By the lemma, we know there are no conjugate points. So we know \exp_p is regular everywhere, hence a local diffeomorphism by inverse function theorem. We can use this fact to pull back a metric from M to T_pM such that \exp_p is a local isometry. Since this is a local isometry, we know geodesics are preserved. So geodesics originating from the origin in T_pM are straight lines, and the speed of the geodesics under the two metrics are the same. So we know T_pM is complete under this metric. Also, by Hopf–Rinow, \exp_p is surjective. So we are done. \square

4 Hodge theory on Riemannian manifolds

4.1 Hodge star and operators

Throughout this chapter, we will assume our manifolds are oriented, and write n of the dimension. We will write $\varepsilon \in \Omega^n(M)$ for a non-vanishing form defining the orientation.

Given a coordinate patch $U \subseteq M$, we can use Gram-Schmidt to obtain a positively-oriented orthonormal frame e_1, \dots, e_n . This allows us to dualize and obtain a basis $\omega_1, \dots, \omega_n \in \Omega^1(M)$, defined by

$$\omega_i(e_j) = \delta_{ij}.$$

Since these are linearly independent, we can multiply all of them together to obtain a non-zero n -form

$$\omega_1 \wedge \dots \wedge \omega_n = a\varepsilon,$$

for some $a \in C^\infty(U)$, $a > 0$. We can do this for any coordinate patches, and the resulting n -form agrees on intersections. Indeed, given any other choice $\omega'_1, \dots, \omega'_n$, they must be related to the original $\omega_1, \dots, \omega_n$ by an element $\Phi \in \text{SO}(n)$. Then by linear algebra, we have

$$\omega'_1 \wedge \dots \wedge \omega'_n = \det(\Phi) \omega_1 \wedge \dots \wedge \omega_n = \omega_1 \wedge \dots \wedge \omega_n.$$

So we can patch all these forms together to get a global n -form $\omega_g \in \Omega^n(M)$ that gives the same orientation. This is a canonical such n -form, depending only on g and the orientation chosen. This is called the (Riemannian) *volume form* of (M, g) .

Recall that the ω_i are orthonormal with respect to the natural dual inner product on T^*M . In general, g induces an inner product on $\bigwedge^p T^*M$ for all $p = 0, 1, \dots, n$, which is still denoted g . One way to describe this is to give an orthonormal basis on each fiber, given by

$$\{\omega_{i_1} \wedge \dots \wedge \omega_{i_p} : 1 \leq i_1 < \dots < i_p \leq n\}.$$

From this point of view, the volume form becomes a form of “unit length”.

We now come to the central definition of Hodge theory.

Definition (Hodge star). The *Hodge star operator* on (M^n, g) is the linear map

$$\star : \bigwedge^p(T_x^*M) \rightarrow \bigwedge^{n-p}(T_x^*M)$$

satisfying the property that for all $\alpha, \beta \in \bigwedge^p(T_x^*M)$, we have

$$\alpha \wedge \star\beta = \langle \alpha, \beta \rangle_g \omega_g.$$

Since g is non-degenerate, it follows that this is well-defined.

How do we actually compute this? Since we have vector spaces, it is natural to consider what happens in a basis.

Proposition. Suppose $\omega_1, \dots, \omega_n$ is an orthonormal basis of T_x^*M . Then we claim that

$$\star(\omega_1 \wedge \dots \wedge \omega_p) = \omega_{p+1} \wedge \dots \wedge \omega_n.$$

We can check this by checking all basis vectors of $\bigwedge^p M$, and the result drops out immediately. Since we can always relabel the numbers, this already tells us how to compute the Hodge star of all other basis elements.

We can apply the Hodge star twice, which gives us a linear endomorphism $\star\star : \bigwedge^p T_x^* M \rightarrow \bigwedge^p T_x^* M$. From the above, it follows that

Proposition. The double Hodge star $\star\star : \bigwedge^p(T_x^* M) \rightarrow \bigwedge^p(T_x^* M)$ is equal to $(-1)^{p(n-p)}$.

In particular,

$$\star 1 = \omega_g, \quad \star \omega_g = 1.$$

Using the Hodge star, we can define a differential operator:

Definition (Co-differential (δ)). We define $\delta : \Omega^p(M) \rightarrow \Omega^{p-1}(M)$ for $0 \leq p \leq \dim M$ by

$$\delta = \begin{cases} (-1)^{n(p+1)+1} \star d\star & p \neq 0 \\ 0 & p = 0 \end{cases}.$$

This is (sometimes) called the *co-differential*.

The funny powers of (-1) are chosen so that our future results work well. We further define

Definition (Laplace–Beltrami operator Δ). The *Laplace–Beltrami operator* is

$$\Delta = d\delta + \delta d : \Omega^p(M) \rightarrow \Omega^p(M).$$

This is also known as the (Hodge) *Laplacian*.

We quickly note that

Proposition.

$$\star\Delta = \Delta\star.$$

Consider the special case of $(M, g) = (\mathbb{R}^n, \text{eucl})$, and $p = 0$. Then a straightforward calculation shows that

$$\Delta f = -\frac{\partial^2 f}{\partial x_1^2} - \cdots - \frac{\partial^2 f}{\partial x_n^2}$$

for each $f \in C^\infty(\mathbb{R}^n) = \Omega^0(\mathbb{R}^n)$. This is just the usual Laplacian, except there is a negative sign. This is there for a good reason, but we shall not go into that.

More generally, metric $g = g_{ij} dx^i dx^j$ on \mathbb{R}^n (or alternatively a coordinate patch on any Riemannian manifold), we have

$$\omega_g = \sqrt{|g|} dx^1 \wedge \cdots \wedge dx^n,$$

where $|g|$ is the determinant of g . Then we have

$$\Delta_g f = -\frac{1}{\sqrt{|g|}} \partial_j (\sqrt{|g|} g^{ij} \partial_i f) = -g^{ij} \partial_i \partial_j f + \text{lower order terms}.$$

How can we think about this co-differential δ ? One way to understand it is that it is the “adjoint” to d .

Proposition. δ is the formal adjoint of d . Explicitly, for any compactly supported $\alpha \in \Omega^{p-1}$ and $\beta \in \Omega^p$, then

$$\int_M \langle d\alpha, \beta \rangle_g \omega_g = \int_M \langle \alpha, \delta\beta \rangle_g \omega_g.$$

We just say it is a formal adjoint, rather than a genuine adjoint, because there is no obvious Banach space structure on $\Omega^p(M)$, and we don't want to go into that. However, we can still define

Definition (L^2 inner product). For $\xi, \eta \in \Omega^p(M)$, we define the L^2 inner product by

$$\langle \langle \xi, \eta \rangle \rangle_g = \int_M \langle \xi, \eta \rangle_g \omega_g,$$

where $\xi, \eta \in \Omega^p(M)$.

Note that this may not be well-defined if the space is not compact.

Under this notation, we can write the proposition as

$$\langle \langle d\alpha, \beta \rangle \rangle_g = \langle \langle \alpha, \delta\beta \rangle \rangle_g.$$

Thus, we also say δ is the L^2 adjoint.

To prove this, we need to recall Stokes' theorem. Since we don't care about manifolds with boundary in this course, we just have

$$\int_M d\omega = 0$$

for all forms ω .

Proof. We have

$$\begin{aligned} 0 &= \int_M d(\alpha \wedge \star\beta) \\ &= \int_M d\alpha \wedge \star\beta + \int_M (-1)^{p-1} \alpha \wedge d\star\beta \\ &= \int_M \langle d\alpha, \beta \rangle_g \omega_g + (-1)^{p-1} (-1)^{(n-p+1)(p-1)} \int_M \alpha \wedge \star\star d\star\beta \\ &= \int_M \langle d\alpha, \beta \rangle_g \omega_g + (-1)^{(n-p)(p-1)} \int_M \alpha \wedge \star\star d\star\beta \\ &= \int_M \langle d\alpha, \beta \rangle_g \omega_g - \int_M \alpha \wedge \star\delta\beta \\ &= \int_M \langle d\alpha, \beta \rangle_g \omega_g - \int_M \langle \alpha, \delta\beta \rangle_g \omega_g. \quad \square \end{aligned}$$

This result explains the funny signs we gave δ .

Corollary. Δ is formally self-adjoint.

Similar to what we did in, say, IB Methods, we can define

Definition (Harmonic forms). A *harmonic form* is a p -form ω such that $\Delta\omega = 0$. We write

$$\mathcal{H}^p = \{\alpha \in \Omega^p(M) : \Delta\alpha = 0\}.$$

We have a further corollary of the proposition.

Corollary. Let M be compact. Then

$$\Delta\alpha = 0 \Leftrightarrow d\alpha = 0 \text{ and } \delta\alpha = 0.$$

We say α is closed and *co-closed*.

Proof. \Leftarrow is clear. For \Rightarrow , suppose $\Delta\alpha = 0$. Then we have

$$0 = \langle\langle\alpha, \Delta\alpha\rangle\rangle = \langle\langle\alpha, d\delta\alpha + \delta d\alpha\rangle\rangle = \|\delta\alpha\|_g^2 + \|d\alpha\|_g^2.$$

Since the L^2 norm is non-degenerate, it follows that $\delta\alpha = d\alpha = 0$. \square

In particular, in degree 0, co-closed is automatic. Then for all $f \in C^\infty(M)$, we have

$$\Delta f = 0 \Leftrightarrow df = 0.$$

In other words, harmonic functions on a compact manifold must be constant. This is a good way to demonstrate that the compactness hypothesis is required, as there are many non-trivial harmonic functions on \mathbb{R}^n , e.g. x .

Some of these things simplify if we know about the parity of our manifold. If $\dim M = n = 2m$, then $\star\star = (-1)^p$, and

$$\delta = -\star d\star$$

whenever $p \neq 0$. In particular, this applies to complex manifolds, say $\mathbb{C}^n \cong \mathbb{R}^{2n}$, with the Hermitian metric. This is to be continued in sheet 3.

4.2 Hodge decomposition theorem

We now work towards proving the Hodge decomposition theorem. This is a very important and far-reaching result.

Theorem (Hodge decomposition theorem). Let (M, g) be a compact oriented Riemannian manifold. Then

- For all $p = 0, \dots, \dim M$, we have $\dim \mathcal{H}^p < \infty$.
- We have

$$\Omega^p(M) = \mathcal{H}^p \oplus \Delta\Omega^p(M).$$

Moreover, the direct sum is orthogonal with respect to the L^2 inner product. We also formally set $\Omega^{-1}(M) = 0$.

As before, the compactness of M is essential, and cannot be dropped.

Corollary. We have orthogonal decompositions

$$\begin{aligned} \Omega^p(M) &= \mathcal{H}^p \oplus d\delta\Omega^p(M) \oplus \delta d\Omega^p(M) \\ &= \mathcal{H}^p \oplus d\Omega^{p-1}(M) \oplus \delta\Omega^{p+1}(M). \end{aligned}$$

Proof. Now note that for an α, β , we have

$$\langle \langle d\delta\alpha, \delta d\beta \rangle \rangle_g = \langle \langle dd\delta\alpha, d\beta \rangle \rangle_g = 0.$$

So

$$d\delta\Omega^p(M) \oplus \delta d\Omega^p(M)$$

is an orthogonal direct sum that clearly contains $\Delta\Omega^p(M)$. But each component is also orthogonal to harmonic forms, because harmonic forms are closed and co-closed. So the first decomposition follows.

To obtain the final decomposition, we simply note that

$$d\Omega^{p-1}(M) = d(\mathcal{H}^{p-1} \oplus \Delta\Omega^{p-1}(M)) = d(\delta d\Omega^{p-1}(M)) \subseteq d\delta\Omega^p(M).$$

On the other hand, we certainly have the other inclusion. So the two terms are equal. The other term follows similarly. \square

This theorem has a rather remarkable corollary.

Corollary. Let (M, g) be a compact oriented Riemannian manifold. Then for all $\alpha \in H_{\text{dR}}^p(M)$, there is a unique $\alpha \in \mathcal{H}^p$ such that $[\alpha] = a$. In other words, the obvious map

$$\mathcal{H}^p \rightarrow H_{\text{dR}}^p(M)$$

is an isomorphism.

This is remarkable. On the left hand side, we have \mathcal{H}^p , which is a completely analytic thing, defined by the Laplacian. On the other hand, the right hand side involves the de Rham cohomology, which is just a topological, and in fact homotopy invariant.

Proof. To see uniqueness, suppose $\alpha_1, \alpha_2 \in \mathcal{H}^p$ are such that $[\alpha_1] = [\alpha_2] \in H_{\text{dR}}^p(M)$. Then

$$\alpha_1 - \alpha_2 = d\beta$$

for some β . But the left hand side and right hand side live in different parts of the Hodge decomposition. So they must be individually zero. Alternatively, we can compute

$$\|d\beta\|_g^2 = \langle \langle d\beta, \alpha_1 - \alpha_2 \rangle \rangle_g = \langle \langle \beta, \delta\alpha_1 - \delta\alpha_2 \rangle \rangle_g = 0$$

since harmonic forms are co-closed.

To prove existence, let $\alpha \in \Omega^p(M)$ be such that $d\alpha = 0$. We write

$$\alpha = \alpha_1 + d\alpha_2 + \delta\alpha_3 \in \mathcal{H}^p \oplus d\Omega^{p-1}(M) \oplus \delta\Omega^{p+1}(M).$$

Applying d gives us

$$0 = d\alpha_1 + d^2\alpha_2 + d\delta\alpha_3.$$

We know $d\alpha_1 = 0$ since α_1 is harmonic, and $d^2 = 0$. So we must have $d\delta\alpha_3 = 0$. So

$$\langle \langle \delta\alpha_3, \delta\alpha_3 \rangle \rangle_g = \langle \langle \alpha_3, d\delta\alpha_3 \rangle \rangle_g = 0.$$

So $\delta\alpha_3 = 0$. So $[\alpha] = [\alpha_1]$ and α has a representative in \mathcal{H}^p . \square

We can also heuristically justify why this is true. Suppose we are given some de Rham cohomology class $a \in H_{\text{dR}}^p(M)$. We consider

$$B_a = \{\xi \in \Omega^p(M) : d\xi = 0, [\xi] = a\}.$$

This is an infinite dimensional affine space.

We now ask ourselves — which $\alpha \in B_a$ minimizes the L^2 norm? We consider the function $F : B_a \rightarrow \mathbb{R}$ given by $F(\alpha) = \|\alpha\|^2$. Any minimizing α is an extremum. So for any $\beta \in \Omega^{p-1}(M)$, we have

$$\left. \frac{d}{dt} \right|_{t=0} F(\alpha + t d\beta) = 0.$$

In other words, we have

$$0 = \left. \frac{d}{dt} \right|_{t=0} (\|\alpha\|^2 + 2t \langle \alpha, d\beta \rangle_g + t^2 \|d\beta\|^2) = 2 \langle \alpha, d\beta \rangle_g.$$

This is the same as saying

$$\langle \delta\alpha, \beta \rangle_g = 0.$$

So this implies $\delta\alpha = 0$. But $d\alpha = 0$ by assumption. So we find that $\alpha \in \mathcal{H}^p$. So the result is at least believable.

The proof of the Hodge decomposition theorem involves some analysis, which we are not bothered to do. Instead, we will just quote the appropriate results. For convenience, we will use $\langle \cdot, \cdot \rangle$ for the L^2 inner product, and then $\|\cdot\|$ is the L^2 norm.

The first theorem we quote is the following:

Theorem (Compactness theorem). If a sequence $\alpha_n \in \Omega^n(M)$ satisfies $\|\alpha_n\| < C$ and $\|\Delta\alpha_n\| < C$ for all n , then α_n contains a Cauchy subsequence.

This is almost like saying $\Omega^n(M)$ is compact, but it isn't, since it is not complete. So the best thing we can say is that the subsequence is Cauchy.

Corollary. \mathcal{H}^p is finite-dimensional.

Proof. Suppose not. Then by Gram–Schmidt, we can find an infinite orthonormal sequence e_n such that $\|e_n\| = 1$ and $\|\Delta e_n\| = 0$, and this certainly does not have a Cauchy subsequence. \square

A large part of the proof is trying to solve the PDE

$$\Delta\omega = \alpha,$$

which we will need in order to carry out the decomposition. In analysis, one useful idea is the notion of weak solutions. We notice that if ω is a solution, then for any $\varphi \in \Omega^p(M)$, we have

$$\langle \omega, \Delta\varphi \rangle = \langle \Delta\omega, \varphi \rangle = \langle \alpha, \varphi \rangle,$$

using that Δ is self-adjoint. In other words, the linear form $\ell = \langle \omega, \cdot \rangle : \Omega^p(M) \rightarrow \mathbb{R}$ satisfies

$$\ell(\Delta\varphi) = \langle \alpha, \varphi \rangle.$$

Conversely, if $\langle \omega, \cdot \rangle$ satisfies this equation, then ω must be a solution, since for any β , we have

$$\langle \Delta\omega, \beta \rangle = \langle \omega, \Delta\beta \rangle = \langle \alpha, \beta \rangle.$$

Definition (Weak solution). A weak solution to the equation $\Delta\omega = \alpha$ is a linear functional $\ell : \Omega^p(M) \rightarrow \mathbb{R}$ such that

- (i) $\ell(\Delta\varphi) = \langle \alpha, \varphi \rangle$ for all $\varphi \in \Omega^p(M)$.
- (ii) ℓ is *bounded*, i.e. there is some C such that $|\ell(\beta)| < C\|\beta\|$ for all β .

Now given a weak solution, we want to obtain a genuine solution. If $\Omega^p(M)$ were a Hilbert space, then we are immediately done by the Riesz representation theorem, but it isn't. Thus, we need a theorem that gives us what we want.

Theorem (Regularity theorem). Every weak solution of $\Delta\omega = \alpha$ is of the form

$$\ell(\beta) = \langle \omega, \beta \rangle$$

for $\omega \in \Omega^p(M)$.

Thus, we have reduced the problem to finding weak solutions. There is one final piece of analysis we need to quote. The definition of a weak solution only cares about what ℓ does to $\Delta\Omega^p(M)$. And it is easy to define what ℓ should do on $\Delta\Omega^p(M)$ — we simply define

$$\ell(\Delta\eta) = \langle \eta, \alpha \rangle.$$

Of course, for this to work, it must be well-defined, but this is not necessarily the case in general. We also have to check it is bounded. But suppose this worked. Then the remaining job is to extend this to a bounded functional on all of $\Omega^p(M)$ in *whatever* way we like. This relies on the following (relatively easy) theorem from analysis:

Theorem (Hahn–Banach theorem). Let L be a normed vector space, and L_0 be a subspace. We let $f : L_0 \rightarrow \mathbb{R}$ be a bounded linear functional. Then f extends to a bounded linear functional $L \rightarrow \mathbb{R}$ with the same bound.

We can now begin the proof.

Proof of Hodge decomposition theorem. Since \mathcal{H}^p is finite-dimensional, by basic linear algebra, we can decompose

$$\Omega^p(M) = \mathcal{H}^p \oplus (\mathcal{H}^p)^\perp.$$

Crucially, we know $(\mathcal{H}^p)^\perp$ is a *closed* subspace. What we want to show is that

$$(\mathcal{H}^p)^\perp = \Delta\Omega^p(M).$$

One inclusion is easy. Suppose $\alpha \in \mathcal{H}^p$ and $\beta \in \Omega^p(M)$. Then we have

$$\langle \alpha, \Delta\beta \rangle = \langle \Delta\alpha, \beta \rangle = 0.$$

So we know that

$$\Delta\Omega^p(M) \subseteq (\mathcal{H}^p)^\perp.$$

The other direction is the hard part. Suppose $\alpha \in (\mathcal{H}^p)^\perp$. We may assume α is non-zero. Since our PDE is a linear one, we may wlog $\|\alpha\| = 1$.

By the regularity theorem, it suffices to prove that $\Delta\omega = \alpha$ has a *weak* solution. We define $\ell : \Delta\Omega^p(M) \rightarrow \mathbb{R}$ as follows: for each $\eta \in \Omega^p(M)$, we put

$$\ell(\Delta\eta) = \langle \eta, \alpha \rangle.$$

We check this is well-defined. Suppose $\Delta\eta = \Delta\xi$. Then $\eta - \xi \in \mathcal{H}^p$, and we have

$$\langle \eta, \alpha \rangle - \langle \xi, \alpha \rangle = \langle \eta - \xi, \alpha \rangle = 0$$

since $\alpha \in (\mathcal{H}^p)^\perp$.

We next want to show the boundedness property. We now claim that there exists a positive $C > 0$ such that

$$\ell(\Delta\eta) \leq C\|\Delta\eta\|$$

for all $\eta \in \Omega^p(M)$. To see this, we first note that by Cauchy–Schwartz, we have

$$|\langle \alpha, \eta \rangle| \leq \|\alpha\| \cdot \|\eta\| = \|\eta\|.$$

So it suffices to show that there is a $C > 0$ such that

$$\|\eta\| \leq C\|\Delta\eta\|$$

for every $\eta \in \Omega^p(M)$.

Suppose not. Then we can find a sequence $\eta_k \in (\mathcal{H}^p)^\perp$ such that $\|\eta_k\| = 1$ and $\|\Delta\eta_k\| \rightarrow 0$.

But then $\|\Delta\eta_k\|$ is certainly bounded. So by the compactness theorem, we may wlog η_k is Cauchy. Then for any $\psi \in \Omega^p(M)$, the sequence $\langle \psi, \eta_k \rangle$ is Cauchy, by Cauchy–Schwartz, hence convergent.

We define $a : \Omega^p(M) \rightarrow \mathbb{R}$ by

$$a(\psi) = \lim_{k \rightarrow \infty} \langle \psi, \eta_k \rangle.$$

Then we have

$$a(\Delta\psi) = \lim_{k \rightarrow \infty} \langle \eta_k, \Delta\psi \rangle = \lim_{k \rightarrow \infty} \langle \Delta\eta_k, \psi \rangle = 0.$$

So we know that a is a weak solution of $\Delta\xi = 0$. By the regularity theorem again, we have

$$a(\psi) = \langle \xi, \psi \rangle$$

for some $\xi \in \Omega^p(M)$. Then $\xi \in \mathcal{H}^p$.

We claim that $\eta_k \rightarrow \xi$. Let $\varepsilon > 0$, and pick N such that $n, m > N$ implies $\|\eta_n - \eta_m\| < \varepsilon$. Then

$$\|\eta_n - \xi\|^2 = \langle \eta_n - \xi, \eta_n - \xi \rangle \leq |\langle \eta_m - \xi, \eta_n - \xi \rangle| + \varepsilon\|\eta_n - \xi\|.$$

Taking the limit as $m \rightarrow \infty$, the first term vanishes, and this tells us $\|\eta_n - \xi\| \leq \varepsilon$. So $\eta_n \rightarrow \xi$.

But this is bad. Since $\eta_k \in (\mathcal{H}^p)^\perp$, and $(\mathcal{H}^p)^\perp$ is closed, we know $\xi \in (\mathcal{H}^p)^\perp$. But also by assumption, we have $\xi \in \mathcal{H}^p$. So $\xi = 0$. But we *also* know $\|\xi\| = \lim \|\eta_k\| = 1$, which is a contradiction. So ℓ is bounded.

We then extend ℓ to any bounded linear map on $\Omega^p(M)$. Then we are done. \square

That was a correct proof, but we just pulled a bunch of theorems out of nowhere, and non-analysts might not be sufficiently convinced. We now look at an explicit example, namely the torus, and sketch a direct proof of Hodge decomposition. In this case, what we needed for the proof reduces to the fact Fourier series and Green's functions work, which is IB Methods.

Consider $M = T^n = \mathbb{R}^n / (2\pi\mathbb{Z})^n$, the n -torus with flat metric. This has local coordinates (x_1, \dots, x_n) , induced from the Euclidean space. This is convenient because $\bigwedge^p T^*M$ is trivialized by $\{dx^{i_1} \wedge \dots \wedge dx^{i_p}\}$. Moreover, the Laplacian is just given by

$$\Delta(\alpha dx^{i_1} \wedge \dots \wedge dx^{i_p}) = - \sum_{i=1}^n \frac{\partial^2 \alpha}{\partial x_i^2} dx^{i_1} \wedge \dots \wedge dx^{i_p}.$$

So to do Hodge decomposition, it suffices to consider the case $p = 0$, and we are just looking at functions $C^\infty(T^n)$, namely the 2π -periodic functions on \mathbb{R} .

Here we will quote the fact that Fourier series work.

Fact. Let $\varphi \in C^\infty(T^n)$. Then it can be (uniquely) represented by a convergent Fourier series

$$\varphi(x) = \sum_{k \in \mathbb{Z}^n} \varphi_k e^{ik \cdot x},$$

where k and x are vectors, and $k \cdot x$ is the standard inner product, and this is uniformly convergent in all derivatives. In fact, φ_k can be given by

$$\varphi_k = \frac{1}{(2\pi)^n} \int_{T^n} \varphi(x) e^{-ik \cdot x} dx.$$

Consider the inner product

$$\langle \varphi, \psi \rangle = (2\pi)^n \sum \bar{\varphi}_k \psi_k.$$

on ℓ^2 , and define the subspace

$$H_\infty = \{(\varphi_k) \in \ell^2 : \varphi_k = o(|k|^m) \text{ for all } m \in \mathbb{Z}\}.$$

Then the map

$$\begin{aligned} \mathcal{F} : C^\infty(T^n) &\rightarrow \ell^2 \\ \varphi &\mapsto (\varphi_k). \end{aligned}$$

is an isometric bijection onto H_∞ .

So we have reduced our problem of working with functions on a torus to working with these infinite series. This makes our calculations rather more explicit.

The key property is that the Laplacian is given by

$$\mathcal{F}(\Delta\varphi) = (-|k|^2 \varphi_k).$$

In some sense, \mathcal{F} “diagonalizes” the Laplacian. It is now clear that

$$\begin{aligned} \mathcal{H}^0 &= \{\varphi \in C^\infty(T^n) : \varphi_k = 0 \text{ for all } k \neq 0\} \\ (\mathcal{H}^0)^\perp &= \{\varphi \in C^\infty(T^n) : \varphi_0 = 0\}. \end{aligned}$$

Moreover, since we can divide by $|k|^2$ whenever k is non-zero, it follows that $(\mathcal{H}^0)^\perp = \Delta C^\infty(T^n)$.

4.3 Divergence

In ordinary multi-variable calculus, we had the notion of the divergence. This makes sense in general as well. Given any $X \in \text{Vect}(M)$, we have

$$\nabla X \in \Gamma(TM \otimes T^*M) = \Gamma(\text{End } TM).$$

Now we can take the trace of this, because the trace doesn't depend on the choice of the basis.

Definition (Divergence). The *divergence* of a vector field $X \in \text{Vect}(M)$ is

$$\text{div } X = \text{tr}(\nabla X).$$

It is not hard to see that this extends the familiar definition of the divergence. Indeed, by definition of trace, for any local frame field $\{e_i\}$, we have

$$\text{div } X = \sum_{i=1}^n g(\nabla_{e_i} X, e_i).$$

It is straightforward to prove from definition that

Proposition.

$$\text{div}(fX) = \text{tr}(\nabla(fX)) = f \text{div } X + \langle df, X \rangle.$$

The key result about divergence is the following:

Theorem. Let $\theta \in \Omega^1(M)$, and let $X_\theta \in \text{Vect}(M)$ be such that $\langle \theta, V \rangle = g(X_\theta, V)$ for all $V \in TM$. Then

$$\delta \theta = -\text{div } X_\theta.$$

So the divergence isn't actually a new operator. However, we have some rather more explicit formulas for the divergence, and this helps us understand δ better.

To prove this, we need a series of lemmas.

Lemma. In local coordinates, for any p -form ψ , we have

$$d\psi = \sum_{k=1}^n dx^k \wedge \nabla_k \psi.$$

Proof. We fix a point $x \in M$, and we may wlog we work in normal coordinates at x . By linearity and change of coordinates, we wlog

$$\psi = f dx^1 \wedge \cdots \wedge dx^p.$$

Now the left hand side is just

$$d\psi = \sum_{k=p+1}^n \frac{\partial f}{\partial x^k} dx^k \wedge dx^1 \wedge \cdots \wedge dx^p.$$

But this is also what the RHS is, because $\nabla_k = \partial_k$ at p . □

To prove this, we first need a lemma, which is also useful on its own right.

Definition (Interior product). Let $X \in \text{Vect}(M)$. We define the *interior product* $i(X) : \Omega^p(M) \rightarrow \Omega^{p-1}(M)$ by

$$(i(X)\psi)(Y_1, \dots, Y_{p-1}) = \psi(X, Y_1, \dots, Y_{p-1}).$$

This is sometimes written as $i(X)\psi = X \lrcorner \psi$.

Lemma. We have

$$(\text{div} X) \omega_g = d(i(X) \omega_g),$$

for all $X \in \text{Vect}(M)$.

Proof. Now by unwrapping the definition of $i(X)$, we see that

$$\nabla_Y(i(X)\psi) = i(\nabla_Y X)\psi + i(X)\nabla_Y\psi.$$

From example sheet 3, we know that $\nabla \omega_g = 0$. So it follows that

$$\nabla_Y(i(X) \omega_g) = i(\nabla_Y X) \omega_g.$$

Therefore we obtain

$$\begin{aligned} & d(i(X)\omega_g) \\ &= \sum_{k=1}^n dx^k \wedge \nabla_k(i(X)\omega_g) \\ &= \sum_{k=1}^n dx^k \wedge i(\nabla_k X)\omega_g \\ &= \sum_{k=1}^n dx^k \wedge i(\nabla_k X)(\sqrt{|g|}dx^1 \wedge \dots \wedge dx^n) \\ &= dx^k(\nabla_k X) \omega_g \\ &= (\text{div} X) \omega_g. \end{aligned}$$

Note that this requires us to think carefully how wedge products work ($i(X)(\alpha \wedge \beta)$ is not just $\alpha(X)\beta$, or else $\alpha \wedge \beta$ would not be anti-symmetric). \square

Corollary (Divergence theorem). For any vector field X , we have

$$\int_M \text{div}(X) \omega_g = \int_M d(i(X) \omega_g) = 0.$$

We can now prove the theorem.

Theorem. Let $\theta \in \Omega^1(M)$, and let $X_\theta \in \text{Vect}(M)$ be such that $\langle \theta, V \rangle = g(X_\theta, V)$ for all $V \in TM$. Then

$$\delta\theta = -\text{div} X_\theta.$$

Proof. By the formal adjoint property of δ , we know that for any $f \in C^\infty(M)$, we have

$$\int_M g(df, \theta) \omega_g = \int_M f \delta \theta \omega_g.$$

So we want to show that

$$\int_M g(df, \theta) \omega_g = - \int_M f \operatorname{div} X_\theta \omega_g.$$

But by the product rule, we have

$$\int_M \operatorname{div}(f X_\theta) \omega_g = \int_M g(df, \theta) \omega_g + \int_M f \operatorname{div} X_\theta \omega_g.$$

So the result follows by the divergence theorem. \square

We can now use this to produce some really explicit formulae for what δ is, which will be very useful next section.

Corollary. If θ is a 1-form, and $\{e_k\}$ is a local orthonormal frame field, then

$$\delta \theta = - \sum_{k=1}^n i(e_k) \nabla_{e_i} \theta = - \sum_{k=1}^n \langle \nabla_{e_k} \theta, e_k \rangle.$$

Proof. We note that

$$\begin{aligned} e_i \langle \theta, e_i \rangle &= \langle \nabla_{e_i} \theta, e_i \rangle + \langle \theta, \nabla_{e_i} e_i \rangle \\ e_i g(X_\theta, e_i) &= g(\nabla_{e_i} X_\theta, e_i) + g(X_\theta, \nabla_{e_i} e_i). \end{aligned}$$

By definition of X_θ , this implies that

$$\langle \nabla_{e_i} \theta, e_i \rangle = g(\nabla_{e_i} X_\theta, e_i).$$

So we obtain

$$\delta \theta = - \operatorname{div} X_\theta = - \sum_{i=1}^n g(\nabla_{e_i} X_\theta, e_i) = - \sum_{k=1}^n \langle \nabla_{e_k} \theta, e_k \rangle, \quad \square$$

We will assume a version for 2-forms (the general result is again on the third example sheet):

Proposition. If $\beta \in \Omega^2(M)$, then

$$(\delta \beta)(Y) = - \sum_{k=1}^n (\nabla_{e_k} \beta)(e_k, Y).$$

In other words,

$$\delta \beta = - \sum_{k=1}^n i(e_k) (\nabla_{e_k} \beta).$$

4.4 Introduction to Bochner's method

How can we apply the Hodge decomposition theorem? The Hodge decomposition theorem tells us the de Rham cohomology group is the kernel of the Laplace–Beltrami operator Δ . So if we want to show, say, $H_{\text{dR}}^1(M) = 0$, then we want to show that $\Delta\alpha \neq 0$ for all non-zero $\alpha \in \Omega^1(M)$. The strategy is to show that

$$\langle\langle \alpha, \Delta\alpha \rangle\rangle \neq 0$$

for all $\alpha \neq 0$. Then we have shown that $H_{\text{dR}}^1(M) = 0$. In fact, we will show that this inner product is positive. To do so, the general strategy is to introduce an operator T with adjoint T^* , and then write

$$\Delta = T^*T + C$$

for some operator C . We will choose T cleverly such that C is very simple.

Now if we can find a manifold such that C is always positive, then since

$$\langle\langle T^*T\alpha, \sigma \rangle\rangle = \langle\langle T\alpha, T\alpha \rangle\rangle \geq 0,$$

it follows that Δ is always positive, and so $H_{\text{dR}}^1(M) = 0$.

Our choice of T will be the covariant derivative ∇ itself. We can formulate this more generally. Suppose we have the following data:

- A Riemannian manifold M .
- A vector bundle $E \rightarrow M$.
- An inner product h on E .
- A connection $\nabla = \nabla^E : \Omega^0(E) \rightarrow \Omega^1(E)$ on E .

We are eventually going to take $E = T^*M$, but we can still proceed in the general setting for a while.

The formal adjoint $(\nabla^E)^* : \Omega^1(E) \rightarrow \Omega^0(E)$ is defined by the relation

$$\int_M \langle \nabla\alpha, \beta \rangle_{E,g} \omega_g = \int_M \langle \alpha, \nabla^*\beta \rangle_E \omega_g$$

for all $\alpha \in \Omega^0(E)$ and $\beta \in \Omega^1(E)$. Since h is non-degenerate, this defines ∇^* uniquely.

Definition (Covariant Laplacian). The *covariant Laplacian* is

$$\nabla^*\nabla : \Gamma(E) \rightarrow \Gamma(E)$$

We are now going to focus on the case $E = T^*M$. It is helpful to have the following explicit formula for ∇^* , which we shall leave as an exercise:

As mentioned, the objective is to understand $\Delta - \nabla^*\nabla$. The theorem is that this difference is given by the Ricci curvature.

This can't be quite right, because the Ricci curvature is a bilinear form on TM^2 , but $\Delta - \nabla^*\nabla$ is a linear endomorphism $\Omega^1(M) \rightarrow \Omega^1(M)$. Thus, we need to define an alternative version of the Ricci curvature by “raising indices”. In coordinates, we consider $g^{jk} \text{Ric}_{ik}$ instead.

We can also define this Ric_{ik} without resorting to coordinates. Recall that given an $\alpha \in \Omega^1(M)$, we defined $X_\alpha \in \text{Vect}(M)$ to be the unique field such that

$$\alpha(z) = g(X_\alpha, Z)$$

for all $Z \in \text{Vect}(M)$. Then given $\alpha \in \Omega^1(M)$, we define $\text{Ric}(\alpha) \in \Omega^1(M)$ by

$$\text{Ric}(\alpha)(X) = \text{Ric}(X, X_\alpha).$$

With this notation, the theorem is

Theorem (Bochner–Weitzenböck formula). On an oriented Riemannian manifold, we have

$$\Delta = \nabla^* \nabla + \text{Ric}.$$

Before we move on to the proof of this formula, we first give an application.

Corollary. Let (M, g) be a compact connected oriented manifold. Then

- If $\text{Ric}(g) > 0$ at each point, then $H_{\text{dR}}^1(M) = 0$.
- If $\text{Ric}(g) \geq 0$ at each point, then $b^1(M) = \dim H_{\text{dR}}^1(M) \leq n$.
- If $\text{Ric}(g) \geq 0$ at each point, and $b^1(M) = n$, then g is flat.

Proof. By Bochner–Weitzenböck, we have

$$\begin{aligned} \langle \Delta \alpha, \alpha \rangle &= \langle \nabla^* \nabla \alpha, \alpha \rangle + \int_M \text{Ric}(\alpha, \alpha) \omega_g \\ &= \|\nabla \alpha\|_2^2 + \int_M \text{Ric}(\alpha, \alpha) \omega_g. \end{aligned}$$

- Suppose $\text{Ric} > 0$. If $\alpha \neq 0$, then the RHS is strictly positive. So the left-hand side is non-zero. So $\Delta \alpha \neq 0$. So $\mathcal{H}_M^1 \cong H_{\text{dR}}^1(M) = 0$.
- Suppose α is such that $\Delta \alpha = 0$. Then the above formula forces $\nabla \alpha = 0$. So if we know $\alpha(x)$ for some fixed $x \in M$, then we know the value of α everywhere by parallel transport. Thus α is determined by the initial condition $\alpha(x)$. Thus there are $\leq n = \dim T_x^* M$ linearly independent such α .
- If $b^1(M) = n$, then we can pick a basis $\alpha_1, \dots, \alpha_n$ of \mathcal{H}_M^1 . Then as above, these are parallel 1-forms. Then we can pick a dual basis $X_1, \dots, X_n \in \text{Vect}(M)$. We claim they are also parallel, i.e. $\nabla X_i = 0$. To prove this, we note that

$$\langle \alpha_j, \nabla X_i \rangle + \langle \nabla \alpha_j, X_i \rangle = \nabla \langle \alpha_j, X_i \rangle.$$

But $\langle \alpha_j, X_i \rangle$ is constantly 0 or 1 depending on i and j , So the RHS vanishes. Similarly, the second term on the left vanishes. Since the α_j span, we know we must have $\nabla X_i = 0$.

Now we have

$$R(X_i, X_j)X_k = (\nabla_{[X_i, X_j]} - [\nabla X_i, \nabla X_j])X_k = 0,$$

Since this is valid for all i, j, k , we know R vanishes at each point. So we are done. \square

Bochner–Weitzenböck can be exploited in a number of similar situations.

In the third part of the theorem, we haven't actually proved the optimal statement. We can deduce more than the flatness of the metric, but requires some slightly advanced topology. We will provide a sketch proof of the theorem, making certain assertions about topology.

Proposition. In the case of (iii), M is in fact isometric to a flat torus.

Proof sketch. We fix $p \in M$ and consider the map $M \rightarrow \mathbb{R}^n$ given by

$$x \mapsto \left(\int_p^x \alpha_i \right)_{i=1, \dots, n} \in \mathbb{R}^n,$$

where the α_i are as in the previous proof. The integral is taken along any path from p to x , and this is not well-defined. But by Stokes' theorem, and the fact that $d\alpha_i = 0$, this only depends on the homotopy class of the path.

In fact, \int_p^x depends only on $\gamma \in H_1(M)$, which is finitely generated. Thus, $\int_p^x \alpha_i$ is a well-defined map to $S^1 = \mathbb{R}/\lambda_i\mathbb{Z}$ for some $\lambda_i \neq 0$. Therefore we obtain a map $M \rightarrow (S^1)^n = T^n$. Moreover, a bit of inspection shows this is a local diffeomorphism. But since the spaces involved are compact, it follows by some topology arguments that it must be a covering map. But again by compactness, this is a finite covering map. So M must be a torus. So we are done. \square

We only proved this for 1-forms, but this is in fact valid for forms of any degree. To do so, we consider $E = \bigwedge^p T^*M$, and then we have a map

$$\nabla : \Omega_M^0(E) \rightarrow \Omega_M^1(E),$$

and this has a formal adjoint

$$\nabla^* : \Omega_M^1(E) \rightarrow \Omega_M^0(E).$$

Now if $\alpha \in \Omega^p(M)$, then it can be shown that

$$\Delta\alpha = \nabla^*\nabla\alpha + \mathfrak{R}(\alpha),$$

where \mathfrak{R} is a linear map $\Omega^p(M) \rightarrow \Omega^p(M)$ depending on the curvature. Then by the same proof, it follows that if $\mathfrak{R} > 0$ at all points, then $\mathcal{H}^k(M) = 0$ for all $k = 1, \dots, n-1$.

If $\mathfrak{R} \geq 0$ only, which in particular is the case if the space is flat, then we have

$$b^k(M) \leq \binom{n}{k} = \dim \bigwedge^k T^*M,$$

and moreover $\Delta\alpha = 0$ iff $\nabla\alpha = 0$.

Proof of Bochner–Weitzenböck

We now move on to actually prove Bochner–Weitzenböck. We first produce an explicit formula for ∇^* , and hence $\nabla^*\nabla$.

Proposition. Let e_1, \dots, e_n be an orthonormal frame field, and $\beta \in \Omega^1(T^*M)$. Then we have

$$\nabla^*\beta = -\sum_{i=1}^n i(e_i)\nabla_{e_i}\beta.$$

Proof. Let $\alpha \in \Omega^0(T^*M)$. Then by definition, we have

$$\langle \nabla \alpha, \beta \rangle = \sum_{i=1}^n \langle \nabla_{e_i} \alpha, \beta(e_i) \rangle.$$

Consider the 1-form given by

$$\theta(Y) = \langle \alpha, \beta(Y) \rangle.$$

Then we have

$$\begin{aligned} \operatorname{div} X_\theta &= \sum_{i=1}^n \langle \nabla_{e_i} X_\theta, e_i \rangle \\ &= \sum_{i=1}^n \nabla_{e_i} \langle X_\theta, e_i \rangle - \langle X_\theta, \nabla_{e_i} e_i \rangle \\ &= \sum_{i=1}^n \nabla_{e_i} \langle \alpha, \beta(e_i) \rangle - \langle \alpha, \beta(\nabla_{e_i} e_i) \rangle \\ &= \sum_{i=1}^n \langle \nabla_{e_i} \alpha, \beta(e_i) \rangle + \langle \alpha, \nabla_{e_i} (\beta(e_i)) \rangle - \langle \alpha, \beta(\nabla_{e_i} e_i) \rangle \\ &= \sum_{i=1}^n \langle \nabla_{e_i} \alpha, \beta(e_i) \rangle + \langle \alpha, (\nabla_{e_i} \beta)(e_i) \rangle. \end{aligned}$$

So by the divergence theorem, we have

$$\int_M \langle \nabla \alpha, \beta \rangle \omega_g = \int_M \sum_{i=1}^n \langle \alpha, (\nabla_{e_i} \beta)(e_i) \rangle \omega_g.$$

So the result follows. \square

Corollary. For a local orthonormal frame field e_1, \dots, e_n , we have

$$\nabla^* \nabla \alpha = - \sum_{i=1}^n \nabla_{e_i} \nabla_{e_i} \alpha.$$

We next want to figure out more explicit expressions for $d\delta$ and δd . To make our lives much easier, we will pick a normal frame field:

Definition (Normal frame field). A local orthonormal frame $\{e_k\}$ field is *normal* at p if further

$$\nabla e_k|_p = 0$$

for all k .

It is a fact that normal frame fields exist. From now on, we will fix a point $p \in M$, and assume that $\{e_k\}$ is a normal orthonormal frame field at p . Thus, the formulae we derive are only valid at p , but this is fine, because p was arbitrary.

The first term $d\delta$ is relatively straightforward.

Lemma. Let $\alpha \in \Omega^1(M)$, $X \in \operatorname{Vect}(M)$. Then

$$\langle d\delta \alpha, X \rangle = - \sum_{i=1}^n \langle \nabla_X \nabla_{e_i} \alpha, e_i \rangle.$$

Proof.

$$\begin{aligned}
\langle d\delta\alpha, X \rangle &= X(\delta\alpha) \\
&= -\sum_{i=1}^n X\langle \nabla_{e_i}\alpha, e_i \rangle \\
&= -\sum_{i=1}^n \langle \nabla_X \nabla_{e_i}\alpha, e_i \rangle. \quad \square
\end{aligned}$$

This takes care of one half of Δ for the other half, we need a bit more work. Recall that we previously found a formula for δ . We now re-express the formula in terms of this local orthonormal frame field.

Lemma. For any 2-form β , we have

$$(\delta\beta)(X) = \sum_{k=1}^n -e_k(\beta(e_k, X)) + \beta(e_k, \nabla_{e_k}X).$$

Proof.

$$\begin{aligned}
(\delta\beta)(X) &= -\sum_{k=1}^n (\nabla_{e_k}\beta)(e_k, X) \\
&= \sum_{k=1}^n -e_k(\beta(e_k, X)) + \beta(\nabla_{e_k}e_k, X) + \beta(e_k, \nabla_{e_k}X) \\
&= \sum_{k=1}^n -e_k(\beta(e_k, X)) + \beta(e_k, \nabla_{e_k}X). \quad \square
\end{aligned}$$

Since we want to understand $\delta d\alpha$ for α a 1-form, we want to find a decent formula for $d\alpha$.

Lemma. For any 1-form α and vector fields X, Y , we have

$$d\alpha(X, Y) = \langle \nabla_X\alpha, Y \rangle - \langle \nabla_Y\alpha, X \rangle.$$

Proof. Since the connection is torsion-free, we have

$$[X, Y] = \nabla_X Y - \nabla_Y X.$$

So we obtain

$$\begin{aligned}
d\alpha(X, Y) &= X\langle \alpha, Y \rangle - Y\langle \alpha, X \rangle - \langle \alpha, [X, Y] \rangle \\
&= \langle \nabla_X\alpha, Y \rangle - \langle \nabla_Y\alpha, X \rangle. \quad \square
\end{aligned}$$

Finally, we can put these together to get

Lemma. For any 1-form α and vector field X , we have

$$\langle \delta d\alpha, X \rangle = -\sum_{k=1}^n \langle \nabla_{e_k} \nabla_{e_k} \alpha, X \rangle + \sum_{k=1}^n \langle \nabla_{e_k} \nabla_X \alpha, e_k \rangle - \sum_{k=1}^n \langle \nabla_{\nabla_{e_k} X} \alpha, e_k \rangle.$$

Proof.

$$\begin{aligned}
\langle \delta d\alpha, X \rangle &= \sum_{k=1}^n \left[-e_k(d\alpha(e_k, X)) + d\alpha(e_k, \nabla_{e_k} X) \right] \\
&= \sum_{k=1}^n \left[-e_k(\langle \nabla_{e_k} \alpha, X \rangle - \langle \nabla_X \alpha, e_k \rangle) \right. \\
&\quad \left. + \langle \nabla_{e_k} \alpha, \nabla_{e_k} X \rangle - \langle \nabla_{\nabla_{e_k} X} \alpha, e_k \rangle \right] \\
&= \sum_{k=1}^n \left[-\langle \nabla_{e_k} \nabla_{e_k} \alpha, X \rangle - \langle \nabla_{e_k} \alpha, \nabla_{e_k} X \rangle + \langle \nabla_{e_k} \nabla_X \alpha, e_k \rangle \right. \\
&\quad \left. + \langle \nabla_{e_k} \alpha, \nabla_{e_k} X \rangle - \langle \nabla_{\nabla_{e_k} X} \alpha, e_k \rangle \right] \\
&= -\sum_{k=1}^n \langle \nabla_{e_k} \nabla_{e_k} \alpha, X \rangle + \sum_{k=1}^n \langle \nabla_{e_k} \nabla_X \alpha, e_k \rangle - \sum_{k=1}^n \langle \nabla_{\nabla_{e_k} X} \alpha, e_k \rangle. \quad \square
\end{aligned}$$

What does this get us? The first term on the right is exactly the $\nabla^* \nabla$ term we wanted. If we add $d\delta\alpha$ to this, then we get

$$\sum_{k=1}^n \langle ([\nabla_{e_k}, \nabla_X] - \nabla_{\nabla_{e_k} X}) \alpha, e_k \rangle.$$

We notice that

$$[e_k, X] = \nabla_{e_k} X - \nabla_X e_k = \nabla_{e_k} X.$$

So we can alternatively write the above as

$$\sum_{k=1}^n \langle ([\nabla_{e_k}, \nabla_X] - \nabla_{[e_k, X]}) \alpha, e_k \rangle.$$

The differential operator on the left looks just like the Ricci curvature. Recall that

$$R(X, Y) = \nabla_{[X, Y]} - [\nabla_X, \nabla_Y].$$

Lemma (Ricci identity). Let M be any Riemannian manifold, and $X, Y, Z \in \text{Vect}(M)$ and $\alpha \in \Omega^1(M)$. Then

$$\langle ([\nabla_X, \nabla_Y] - \nabla_{[X, Y]}) \alpha, Z \rangle = \langle \alpha, R(X, Y)Z \rangle.$$

Proof. We note that

$$\langle \nabla_{[X, Y]} \alpha, Z \rangle + \langle \alpha, \nabla_{[X, Y]} Z \rangle = [X, Y] \langle \alpha, Z \rangle = \langle [\nabla_X, \nabla_Y] \alpha, Z \rangle + \langle \alpha, [\nabla_X, \nabla_Y] Z \rangle.$$

The second equality follows from writing $[X, Y] = XY - YX$. We then rearrange and use that $R(X, Y) = \nabla_{[X, Y]} - [\nabla_X, \nabla_Y]$. \square

Corollary. For any 1-form α and vector field X , we have

$$\langle \Delta \alpha, X \rangle = \langle \nabla^* \nabla \alpha, X \rangle + \text{Ric}(\alpha)(X).$$

This is the theorem we wanted.

Proof. We have found that

$$\langle \Delta\alpha, X \rangle = \langle \nabla^* \nabla \alpha, X \rangle + \sum_{i=1}^n \langle \alpha, R(e_i, X)e_i \rangle.$$

We have

$$\sum_{i=1}^n \langle \alpha, R(e_i, X)e_i \rangle = \sum_{i=1}^n g(X_\alpha, R(e_i, X)e_i) = \text{Ric}(X_\alpha, X) = \text{Ric}(\alpha)(X).$$

So we are done. □

5 Riemannian holonomy groups

Again let M be a Riemannian manifold, which is always assumed to be connected. Let $x \in M$, and consider a path $\gamma \in \Omega(x, y)$, $\gamma : [0, 1] \rightarrow M$. At the rather beginning of the course, we saw that γ gives us a parallel transport from $T_x M$ to $T_y M$. Explicitly, given any $X_0 \in T_x M$, there exists a unique vector field X along γ with

$$\frac{\nabla X}{dt} = 0, \quad X(0) = X_0.$$

Definition (Holonomy transformation). The *holonomy transformation* $P(\gamma)$ sends $X_0 \in T_x M$ to $X(1) \in T_y M$.

We know that this map is invertible, and preserves the inner product. In particular, if $x = y$, then $P(\gamma) \in \text{O}(T_x M) \cong \text{O}(n)$.

Definition (Holonomy group). The *holonomy group* of M at $x \in M$ is

$$\text{Hol}_x(M) = \{P(\gamma) : \gamma \in \Omega(x, x)\} \subseteq \text{O}(T_x M).$$

The group operation is given by composition of linear maps, which corresponds to composition of paths.

We note that this group doesn't really depend on the point x . Given any other $y \in M$, we can pick a path $\beta \in \Omega(x, y)$. Writing P_β instead of $P(\beta)$, we have a map

$$\begin{aligned} \text{Hol}_x(M) &\longrightarrow \text{Hol}_y(M) \\ P_\gamma &\longmapsto P_\beta \circ P_\gamma \circ P_{\beta^{-1}} \in \text{Hol}_y(M) \end{aligned}$$

So we see that $\text{Hol}_x(M)$ and $\text{Hol}_y(M)$ are isomorphic. In fact, after picking an isomorphism $\text{O}(T_x M) \cong \text{O}(T_y M) \cong \text{O}(n)$, these subgroups are conjugate as subgroups of $\text{O}(n)$. We denote this class by $\text{Hol}(M)$.

Note that depending of what we want to emphasize, we write $\text{Hol}(M, g)$, or even $\text{Hol}(g)$ instead.

Really, $\text{Hol}(M)$ is a *representation* (up to conjugacy) induced by the standard representation of $\text{O}(n)$ on \mathbb{R}^n .

Proposition. If M is simply connected, then $\text{Hol}_x(M)$ is path connected.

Proof. $\text{Hol}_x(M)$ is the image of $\Omega(x, x)$ in $\text{O}(n)$ under the map P , and this map is continuous from the standard theory of ODE's. Simply connected means $\Omega(x, x)$ is path connected. So $\text{Hol}_x(M)$ is path connected. \square

It is convenient to consider the *restricted* holonomy group.

Definition (Restricted holonomy group). We define

$$\text{Hol}_x^0(M) = \{P(\gamma) : \gamma \in \Omega(x, x) \text{ nullhomotopic}\}.$$

As before, this group is, up to conjugacy, independent of the choice of the point in the manifold. We write this group as $\text{Hol}^0(M)$.

Of course, $\text{Hol}^0(M) \subseteq \text{Hol}(M)$, and if $\pi_1(M) = 0$, then they are equal.

Corollary. $\text{Hol}^0(M) \subseteq \text{SO}(n)$.

Proof. $\text{Hol}^0(M)$ is connected, and thus lies in the connected component of the identity in $O(n)$. \square

Note that there is a natural action of $\text{Hol}_x(M)$ and $\text{Hol}_x^0(M)$ on T_x^*M , $\bigwedge^p T_x^*M$ for all p , and more generally tensor products of T_xM .

Fact.

- $\text{Hol}^0(M)$ is the connected component of $\text{Hol}(M)$ containing the identity element.
- $\text{Hol}^0(M)$ is a *Lie subgroup* of $SO(n)$, i.e. it is a subgroup and an immersed submanifold. Thus, the Lie algebra of $\text{Hol}^0(M)$ is a Lie subalgebra of $\mathfrak{so}(n)$, which is just the skew-symmetric $n \times n$ matrices.

This is a consequence of Yamabe theorem, which says that a path-connected subgroup of a Lie group is a Lie subgroup.

We will not prove these.

Proposition (Fundamental principle of Riemannian holonomy). Let (M, g) be a Riemannian manifold, and fix $p, q \in \mathbb{Z}_+$ and $x \in M$. Then the following are equivalent:

- (i) There exists a (p, q) -tensor field α on M such that $\nabla\alpha = 0$.
- (ii) There exists an element $\alpha_0 \in (T_xM)^{\otimes p} \otimes (T_x^*M)^{\otimes q}$ such that α_0 is invariant under the action of $\text{Hol}_x(M)$.

Proof. To simplify notation, we consider only the case $p = 0$. The general case works exactly the same way, with worse notation. For $\alpha \in (T_x^*M)^q$, we have

$$(\nabla_X\alpha)(X_1, \dots, X_q) = X(\alpha(X_1, \dots, X_q)) - \sum_{i=1}^q \alpha(X_1, \dots, \nabla_X X_i, \dots, X_q).$$

Now consider a loop $\gamma : [0, 1] \rightarrow M$ be a loop at x . We choose vector fields X_i along γ for $i = 1, \dots, q$ such that

$$\frac{\nabla X_i}{dt} = 0.$$

We write

$$X_i(\gamma(0)) = X_i^0.$$

Now if $\nabla\alpha = 0$, then this tells us

$$\frac{\nabla\alpha}{dt}(X_1, \dots, X_q) = 0.$$

By our choice of X_i , we know that $\alpha(X_1, \dots, X_q)$ is constant along γ . So we know

$$\alpha(X_1^0, \dots, X_q^0) = \alpha(P_\gamma(X_1^0), \dots, P_\gamma(X_q^0)).$$

So α is invariant under $\text{Hol}_x(M)$. Then we can take $\alpha_0 = \alpha_x$.

Conversely, if we have such an α_0 , then we can use parallel transport to transfer it to everywhere in the manifold. Given any $y \in M$, we define α_y by

$$\alpha_y(X_1, \dots, X_q) = \alpha_0(P_\gamma(X_1), \dots, P_\gamma(X_q)),$$

where γ is any path from y to x . This does not depend on the choice of γ precisely because α_0 is invariant under $\text{Hol}_x(M)$.

It remains to check that α is C^∞ with $\nabla\alpha = 0$, which is an easy exercise. \square

Example. Let M be oriented. Then we have a volume form ω_g . Since $\nabla\omega_g = 0$, we can take $\alpha = \omega_g$. Here $p = 0$ and $q = n$. Also, its stabilizer is $H = \text{SO}(n)$. So we know $\text{Hol}(M) \subseteq \text{SO}(n)$ if (and only if) M is oriented.

The “only if” part is not difficult, because we can use parallel transport to transfer an orientation at a particular point to everywhere.

Example. Let $x \in M$, and suppose

$$\text{Hol}_x(M) \subseteq \text{U}(n) = \{g \in \text{SO}(2n) : gJ_0g^{-1} = J_0\},$$


where

$$J_0 = \begin{pmatrix} 0 & I \\ -I & 0 \end{pmatrix}.$$

By looking at $\alpha_0 = J_0$, we obtain $\alpha = J \in \Gamma(\text{End } TM)$ with $\nabla J = 0$ and $J^2 = -1$. This is a well-known standard object in complex geometry, and such a J is an instance of an *almost complex structure* on M .

Example. Recall (from the theorem on applications of Bochner–Weitzenböck) that a Riemannian manifold (M, g) is flat (i.e. $R(g) \equiv 1$) iff around each point $x \in M$, there is a parallel basis of parallel vector fields. So we find that (M, g) is flat iff $\text{Hol}^0(M, g) = \{id\}$.

It is essential that we use $\text{Hol}^0(M, g)$ rather than the full $\text{Hol}(M, g)$. For example, we can take the Klein bottle



with the flat metric. Then parallel transport along the closed loop γ has

$$P_\gamma = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

In fact, we can check that $\text{Hol}(K) = \mathbb{Z}_2$. Note that here K is non-orientable.

Since we know that $\text{Hol}(M)$ is a Lie group, we can talk about its Lie algebra.

Definition (Holonomy algebra). The *holonomy algebra* $\mathfrak{hol}(M)$ is the Lie algebra of $\text{Hol}(M)$.

Thus $\mathfrak{hol}(M) \leq \mathfrak{so}(n)$ up to conjugation.

Now consider some open coordinate neighbourhood $U \subseteq M$ with coordinates x_1, \dots, x_n . As before, we write

$$\partial_i = \frac{\partial}{\partial x_i}, \quad \nabla_i = \nabla_{\partial_i}.$$


The curvature may also be written in terms of coordinates $R = R_{j,k\ell}^i$, and we also have

$$R(\partial_k, \partial_\ell) = -[\nabla_k, \nabla_\ell].$$

Thus, $\mathfrak{hol}(M)$ contains

$$\left. \frac{d}{dt} \right|_{t=0} P(\gamma_t),$$

where γ_t is the square



By a direct computation, we find

$$P(\gamma_t) = I + \lambda t R(\partial_k, \partial_\ell) + o(t).$$

Here $\lambda \in \mathbb{R}$ is some non-zero absolute constant that doesn't depend on anything (except convention).

Differentiating this with respect to t and taking the limit $t \rightarrow 0$, we deduce that at for $p \in U$, we have

$$R_p = (R_{j,k\ell}^i)_p \in \wedge^2 T_p^* M \otimes \mathfrak{hol}_p(M),$$

where we think $\mathfrak{hol}_p(M) \subseteq \text{End } T_p M$. Recall we also had the $R_{ij,k\ell}$ version, and because of the nice symmetric properties of R , we know

$$(R_{ij,k\ell})_p \in S^2 \mathfrak{hol}_p(M) \subseteq \wedge^2 T_p^* M \otimes \wedge^2 T_p^* M.$$

Strictly speaking, we should write

$$(R_{j\ell}^ik)_p \in S^2 \mathfrak{hol}_p(M),$$

but we can use the metric to go between $T_p M$ and $T_p^* M$.

So far, what we have been doing is rather tautological. But it turns out this allows us to decompose the de Rham cohomology groups.

In general, consider an arbitrary Lie subgroup $G \subseteq \text{GL}_n(\mathbb{R})$. There is a standard representation ρ of $\text{GL}_n(\mathbb{R})$ on \mathbb{R}^n , which restricts to a representation (ρ, \mathbb{R}^n) of G . This induces a representation $(\rho^k, \wedge^k(\mathbb{R}^n))$ of G .

This representation is in general not irreducible. We decompose it into irreducible components (ρ_i^k, W_i^k) , so that

$$\wedge^k(\mathbb{R}^n) = \bigoplus_i W_i^k.$$

We can do this for bundles instead of just vector spaces. Consider a manifold M with a G -structure, i.e. there is an atlas of coordinate neighbourhoods where the transition maps satisfy

$$\begin{pmatrix} \partial x_\alpha \\ \partial x'_\beta \end{pmatrix}_p \in G$$

for all p . Then we can use this to split our bundle of k -forms into well-defined vector sub-bundles with typical fibers W_i^k :

$$\bigwedge^k T^*M = \bigoplus_i \Lambda_i^k.$$

We can furthermore say that every G -equivariant linear map $\varphi : W_i^k \rightarrow W_j^\ell$ induces a morphism of vector bundles $\phi : \Lambda_i^k \rightarrow \Lambda_j^\ell$.

Now suppose further that $\text{Hol}(M) \leq G \leq O(n)$. Then we can show that parallel transport preserves this decomposition into sub-bundles. So ∇ restricts to a well-defined connection on each Λ_i^k .

Thus, if $\xi \in \Gamma(\Lambda_i^k)$, then $\nabla\xi \in \Gamma(T^*M \otimes \Lambda_i^k)$, and then we have $\nabla^*\nabla\xi \in \Gamma(\Lambda_i^k)$. But we know the covariant Laplacian is related to Laplace–Beltrami via the curvature. We only do this for 1-forms for convenience of notation. Then if $\xi \in \Omega^1(M)$, then we have

$$\Delta\xi = \nabla^*\nabla\xi + \text{Ric}(\xi).$$

We can check that Ric also preserves these subbundles. Then it follows that $\Delta : \Gamma(\Lambda_j^1) \rightarrow \Gamma(\Lambda_j^1)$ is well-defined.

Thus, we deduce

Theorem. Let (M, g) be a connected and oriented Riemannian manifold, and consider the decomposition of the bundle of k -forms into irreducible representations of the holonomy group,

$$\bigwedge^k T^*M = \bigoplus_i \Lambda_i^k.$$

In other words, each fiber $(\Lambda_i^k)_x \subseteq \bigwedge^k T_x^*M$ is an irreducible representation of $\text{Hol}_x(g)$. Then

(i) For all $\alpha \in \Omega_i^k(M) \equiv \Gamma(\Lambda_i^k)$, we have $\Delta\alpha \in \Omega_i^k(M)$.

(ii) If M is compact, then we have a decomposition

$$H_{\text{dR}}^k(M) = \bigoplus H_{i,\text{dR}}^k(M),$$

where

$$H_{i,\text{dR}}^k(M) = \{[\alpha] : \alpha \in \Omega_i^k(M), \Delta\alpha = 0\}.$$

The dimensions of these groups are known as the *refined Betti numbers*.

We have only proved this for $k = 1$, but the same proof technique can be used to do it for arbitrary k .

Our treatment is rather abstract so far. But for example, if we are dealing with complex manifolds, then we know that $\text{Hol}(M) \leq U(n)$. So this allows us to have a canonical refinement of the de Rham cohomology, and this is known as the Lefschetz decomposition.

6 The Cheeger–Gromoll splitting theorem

We will talk about the Cheeger–Gromoll splitting theorem. This is a hard theorem, so we will not prove it. However, we will state it, and discuss a bit about it. To state the theorem, we need some preparation.

Definition (Ray). Let (M, g) be a Riemannian manifold. A *ray* is a map $r(t) : [0, \infty) \rightarrow M$ if r is a geodesic, and minimizes the distance between any two points on the curve.

Definition (Line). A *line* is a map $\ell(t) : \mathbb{R} \rightarrow M$ such that $\ell(t)$ is a geodesic, and minimizes the distance between any two points on the curve.

We have seen from the first example sheet that if M is a complete unbounded manifold, then M has a ray from each point, i.e. for all $x \in M$, there exists a ray r such that $r(0) = x$.

Definition (Connected at infinity). A complete manifold is said to be *connected at infinity* if for all compact set $K \subseteq M$, there exists a compact $C \supseteq K$ such that for every two points $p, q \in M \setminus C$, there exists a path $\gamma \in \Omega(p, q)$ such that $\gamma(t) \in M \setminus K$ for all t .

We say M is *disconnected at infinity* if it is not connected at infinity.

Note that if M is disconnected at infinity, then it must be unbounded, and in particular non-compact.

Lemma. If M is disconnected at infinity, then M contains a line.

Proof. Note that M is unbounded. Since M is disconnected at infinity, we can find a compact subset $K \subseteq M$ and sequences $p_m, q_m \rightarrow \infty$ as $m \rightarrow \infty$ (to make this precise, we can pick some fixed point x , and then require $d(x, p_m), d(x, q_m) \rightarrow \infty$) such that every $\gamma_m \in \Omega(p_m, q_m)$ passes through K .

In particular, we pick γ_m to be a minimal geodesic from p_m to q_m parametrized by arc-length. Then γ_m passes through K . By reparametrization, we may assume $\gamma_m(0) \in K$.

Since K is compact, we can pass to a subsequence, and wlog $\gamma_m(0) \rightarrow x \in K$ and $\dot{\gamma}_m(0) \rightarrow a \in T_x M$ (suitably interpreted).

Then we claim the geodesic $\gamma_{x,a}(t)$ is the desired line. To see this, since solutions to ODE's depend smoothly on initial conditions, we can write the line as

$$\ell(t) = \lim_{m \rightarrow \infty} \gamma_m(t).$$

Then we know

$$d(\ell(s), \ell(t)) = \lim_{m \rightarrow \infty} d(\gamma_m(s), \gamma_m(t)) = |s - t|.$$

So we are done. □

Let's look at some examples.

Example. The elliptic paraboloid

$$\{z = x^2 + y^2\} \subseteq \mathbb{R}^3$$

with the induced metric does not contain a line. To prove this, we can show that any geodesic that is not a meridian must intersect itself.

Example. Any complete metric g on $S^{n-1} \times \mathbb{R}$ contains a line since it is disconnected at ∞ .

Theorem (Cheeger–Gromoll line-splitting theorem (1971)). If (M, g) is a complete connected Riemannian manifold containing a line, and has $\text{Ric}(g) \geq 0$ at each point, then M is isometric to a Riemannian product $(N \times \mathbb{R}, g_0 + dt^2)$ for some metric g_0 on N .

We will not prove this, but just see what the theorem can be good for.

First of all, we can slightly improve the statement. After applying the theorem, we can check again if N contains a line or not. We can keep on proceeding and splitting lines off. Then we get

Corollary. Let (M, g) be a complete connected Riemannian manifold with $\text{Ric}(g) \geq 0$. Then it is isometric to $X \times \mathbb{R}^q$ for some $q \in \mathbb{N}$ and Riemannian manifold X , where X is complete and does not contain any lines.

Note that if X is zero-dimensional, since we assume all our manifolds are connected, then this implies M is flat. If $\dim X = 1$, then $X \cong S^1$ (it can't be a line, because a line contains a line). So again M is flat.

Now suppose that in fact $\text{Ric}(g) = 0$. Then it is not difficult to see from the definition of the Ricci curvature that $\text{Ric}(X) = 0$ as well. If we know $\dim X \leq 3$, then M has to be flat, since in dimensions ≤ 3 , the Ricci curvature determines the full curvature tensor.

We can say a bit more if we assume more about the manifold. Recall (from example sheets) that a manifold is *homogeneous* if the group of isometries acts transitively. In other words, for any $p, q \in M$, there exists an isometry $\phi : M \rightarrow M$ such that $\phi(p) = q$. This in particular implies the metric is complete.

It is not difficult to see that if M is homogeneous, then so is X . In this case, X must be compact. Suppose not. Then X is unbounded. We will obtain a line on X .

By assumption, for all $n = 1, 2, \dots$, we can find p_n, q_n with $d(p_n, q_n) \geq 2n$. Since X is complete, we can find a minimal geodesic γ_n connecting these two points, parametrized by arc length. By homogeneity, we may assume that the midpoint $\gamma_n(0)$ is at a fixed point x_0 . By passing to a subsequence, we wlog $\dot{\gamma}_n(0)$ converges to some $a \in T_{x_0}(X)$. Then we use a as an initial condition for our geodesic, and this will be a line.

A similar argument gives

Lemma. Let (M, g) be a compact Riemannian manifold, and suppose its universal Riemannian cover (\tilde{M}, \tilde{g}) is non-compact. Then (\tilde{M}, \tilde{g}) contains a line.

Proof. We first find a compact $K \subseteq \tilde{M}$ such that $\pi(K) = M$. Since \tilde{M} must be complete, it is unbounded. Choose p_n, q_n, γ_n like before. Then we can apply deck transformations so that the midpoint lies inside K , and then use compactness of K to find a subsequence so that the midpoint converges. \square

We do more applications.

Corollary. Let (M, g) be a compact, connected manifold with $\text{Ric}(g) \geq 0$. Then

- The universal Riemannian cover is isometric to the Riemannian product $X \times \mathbb{R}^N$, with X compact, $\pi_1(X) = 1$ and $\text{Ric}(g_X) \geq 0$.

- If there is some $p \in M$ such that $\text{Ric}(g)_p > 0$, then $\pi_1(M)$ is finite.
- Denote by $I(\tilde{M})$ the group of isometries $\tilde{M} \rightarrow \tilde{M}$. Then $I(\tilde{M}) = I(X) \times E(\mathbb{R}^q)$, where $E(\mathbb{R}^q)$ is the group of rigid Euclidean motions,

$$\mathbf{y} \mapsto A\mathbf{y} + \mathbf{b}$$

where $\mathbf{b} \in \mathbb{R}^n$ and $A \in O(q)$.

- If \tilde{M} is homogeneous, then so is X .

Proof.

- This is direct from Cheeger–Gromoll and the previous lemma.
- If there is a point with strictly positive Ricci curvature, then the same is true for the universal cover. So we cannot have any non-trivial splitting. So by the previous part, \tilde{M} must be compact. By standard topology, $|\pi_1(M)| = |\pi^{-1}(\{p\})|$.
- We use the fact that $E(\mathbb{R}^q) = I(\mathbb{R}^q)$. Pick a $g \in I(\tilde{M})$. Then we know g takes lines to lines. Now use that all lines in $\tilde{M} \times \mathbb{R}^q$ are of the form $p \times \mathbb{R}$ with $p \in X$ and $\mathbb{R} \subseteq \mathbb{R}^q$ an affine line. Then

$$g(p \times \mathbb{R}) = p' \times \mathbb{R},$$

for some p' and possibly for some other copy of \mathbb{R} . By taking unions, we deduce that $g(p \times \mathbb{R}^q) = p' \times \mathbb{R}^q$. We write $h(p) = p'$. Then $h \in I(X)$.

Now for any $X \times \mathbf{a}$ with $\mathbf{a} \in \mathbb{R}^q$, we have $X \times \mathbf{a} \perp p \times \mathbb{R}^q$ for all $p \in X$. So we must have

$$g(X \times \mathbf{a}) = X \times \mathbf{b}$$

for some $\mathbf{b} \in \mathbb{R}^q$. We write $e(\mathbf{a}) = \mathbf{b}$. Then

$$g(p, a) = (h(p), e(a)).$$

Since the metric of X and \mathbb{R}^q are decoupled, it follows that h and e must separately be isometries. \square

We can look at more examples.

Proposition. Consider $S^n \times \mathbb{R}$ for $n = 2$ or 3 . Then this does not admit any Ricci-flat metric.

Proof. Note that $S^n \times \mathbb{R}$ is disconnected at ∞ . So any metric contains a line. Then by Cheeger–Gromoll, \mathbb{R} splits as a Riemannian factor. So we obtain $\text{Ric} = 0$ on the S^n factor. Since we are in $n = 2, 3$, we know S^n is flat, as the Ricci curvature determines the full curvature. So S^n is the quotient of \mathbb{R}^n by a discrete group, and in particular $\pi_1(S^n) \neq 1$. This is a contradiction. \square

Let G be a Lie group with a bi-invariant metric g . Suppose the center $Z(G)$ is finite. Then the center of \mathfrak{g} is trivial (since it is the Lie algebra of $G/Z(G)$, which has trivial center). From sheet 2, we find that $\text{Ric}(g) > 0$ implies $\pi_1(G)$ is finite. The converse is also true, but is harder. This is done on Q11 of sheet 3 — if $\pi_1(G)$ is finite, then $Z(G)$ is finite.

Index

- L^2 inner product, 46
- Δ , 45
- $\text{Hol}(N)$, 63
- $\text{Hol}^0(M)$, 63
- $\text{Hol}_x^0(M)$, 63
- $\text{Hol}_x(M)$, 63
- $\Omega(p, q)$, 24
- Σ_r , 16
- $\langle\langle \xi, \eta \rangle\rangle_g$, 46
- δ , 45
- \mathcal{H}^p , 46
- $\mathcal{T}^{q,p}M$, 5
- ∇_k , 3
- ∂_k , 3
- $i(X)$, 54
- almost complex structure, 65
- bi-invariant metric, 5
- Bianchi identity
 - first, 7
- Bochner–Weitzenböck formula, 57
- Bonnet–Myers theorem, 40
- bounded, 50
- Cheeger–Gromoll line-splitting theorem, 69
- Christoffel symbols, 5
- co-closed, 47
- co-differential, 45
- compatible connection, 5
- complete
 - geodesically, 28
- conjugate points, 36
- connected at infinity, 68
- connection
 - compatible with metric, 5
 - symmetric, 5
 - torsion-free, 5
- covariant derivative, 13
- covariant Laplacian, 56
- cover
 - Riemannian, 40
- covering map
 - Riemannian, 40
- curvature, 7
 - Ricci, 11
 - curvature 2-form, 7
- diameter, 39
- distance, 24
- divergence, 53
- divergence theorem, 54
- energy, 30
- exponential map, 15
- first Bianchi identity, 7
- first fundamental form, 11
- first variation formula
 - geodesic, 31
- Gauss’ lemma, 23
- Gaussian curvature, 11
- geodesic, 13
 - first variation formula, 31
 - second variation formula, 32
- geodesic local coordinates, 16
- geodesic polar coordinates, 16
- geodesic sphere, 16
- geodesically complete, 28
- Hadamard–Cartan theorem, 42
- harmonic form, 46
- Hodge decomposition theorem, 47
- Hodge Laplacian, 45
- Hodge star, 44
- holonomy algebra, 65
- holonomy group, 63
 - restricted, 63
- holonomy transformation, 63
- homogeneous, 69
- Hopf–Rinow theorem, 29
- horizontal lift, 13
- interior product, 54
- isometry, 4
- Jacobi equation, 19
- Jacobi field, 19
- Laplace–Beltrami operator, 45
- Laplacian, 45
- left-invariant metric?, 4
- left-invariant vector field, 4

- Levi-Civita connection, 5
- Lie algebra, 4
- lift, 12
 - horizontal, 13
- line, 68
- minimal geodesic, 25
- normal coordinates, 16
- normal frame field, 59
- parallel transport, 13
- ray, 68
- refined Betti numbers, 67
- restricted holonomy group, 63
- Ricci curvature, 11
- Ricci identity, 61
- Riemannian cover, 40
- Riemannian covering, 40
- Riemannian covering map, 40
- Riemannian manifold, 3
- Riemannian metric, 3
- scalar curvature, 11
- second variation formula
 - geodesic, 32
- sectional curvature, 9
- symmetric connection, 5
- Synge's theorem, 34
- torsion-free connection, 5
- vector field
 - left invariant, 4
- volume form, 44
- weak solution, 50
- Whitney embedding theorem, 4