Lectures on Differential Geometry

Paul Bryan

June 13, 2018

Contents

1 Week 01

1.1 Week 01, Lecture 01: Introduction

1.1.1 Recommended Reading

1. Text Books

   - Do Carmo: Riemannian Geometry (a classic text that is certainly relevant today but sometimes considered a little terse. Does include material on differentiable manifolds.)
   - Lee: Riemannian Manifolds: An Introduction to Curvature (very readable, possibly a little elementary in places)
   - Chavel: Riemannian Geometry: A Modern Introduction (more advanced, extensive discussion of many aspects of Riemannian Geometry)
   - Petersen: Riemannian Geometry (more advanced, slightly non-standard approach definitely worth a look at some point)
   - Gallot, Hulin, Lafontaine: Riemannian Geometry (more advanced, but very nice development of the formalism of Riemannian Geometry)

2. Lecture Notes

   - Lectures on Differential Geometry by Ben Andrews (I learned from these notes)

3. Differentiable Manifolds Some sources for differential manifolds. There are many resources available, and some of the resources listed above treat this topic before moving on to Riemannian Geometry. The following should be sufficient background reading.

   - Lee: Introduction to Smooth Manifolds
   - Hitchin: Differentiable Manifolds

1.1.2 Course Summary

This course is about *Riemannian geometry*, that is the extension of geometry to spaces where differential/integral calculus is possible, namely to manifolds. We will study how to define the notions of length, angle and area on a smooth manifold, which leads to the definition of a *Riemannian Manifold.* Important concepts are

- Riemannian metrics
- Connections (differentiation of vector fields)
- Length, angle
- Geodesics (shortest paths)
- Area
- Curvature

At the end of the course, we will touch on some *global* aspects of Riemannian geometry, and discuss a little about the interaction between curvature, geometry and topology. Such interaction was studied heavily in the mid to late 20'th century and is currently still an active area of research. A famous example is the Hamilton-Perelman resolution of the Poincaré conjecture, one of the Clay Foundation's seven Millennium Prizes, was resolved only this century.

1.1.3 What is Riemannian Geometry?

What follows is an imprecise overview of the basic ideas behind Riemannian Geometry. No proofs, or references are given. That will happen throughout the course! For now, we just want to a basic feel for the topics to be studied in this course.

1. Curves and Surfaces in Euclidean Space (The Genesis of Riemannian Geometry)

   (a) Curves
   $\gamma : [a, b] \to \mathbb{R}^2$ (or $\mathbb{R}^3$) that are differentiable and *regular* ($gamma' \neq 0$). Why regular? We really want to study the *image*

$C = \gamma([a,b])$; γ is only a parmetrisation of this set that imbues it with a structure allowing for differential and integral calculus on C. Without regularity, the calculus on $[a,b]$ does not transfer to C via γ: if $f : [a,b] \to \mathbb{R}$ is differentiable, then we can define $\bar{f} : C \to \mathbb{R}$ by $\bar{f}(\gamma(t)) = f(t)$, or (slightly informally, as $\bar{f}(x) = f(\gamma^{-1}(x))$ for $x \in C$). But notice that (formally) by the chain-rule,

$$\partial_t \bar{f}(\gamma(t)) = \bar{f}'(\gamma(t)) \cdot \gamma'(t)$$

will be zero at *irregular points* t_0 ($\gamma'(t_0) = 0$) even if $f' \neq 0$ at such a point. Thus the function $\bar{f} \circ \gamma(t)$ may have zero derivative at t_0 even if this is not a critical point of the function f, in which case the first derivative test for extremal points would fail! The regularity assumption rules out this "bad" behaviour.

(b) Surfaces

A set $S \subset \mathbb{R}^3$ covered by *local parametrisations* $\{\phi_\alpha : U_\alpha \subset \mathbb{R}^2 \to S \subset \mathbb{R}^3\}$ where each U_α is open and each ϕ_α is C^1 as a map $\mathbb{R}^2 \to \mathbb{R}^3$ and is a homeomorphism with it's image (the image is equipped with the subspace topology). In general, we cannot expect a single, global parametrisation to cover S (e.g. the sphere requires at least two local parametrisations to cover it).

Calculus on S is achieved by doing calculus on U_α and transferring it to S by means of U_α. Again we require *regularity*: $D\phi_\alpha$ is a rank 2 matrix everywhere on U_α so that calculus on S can be identified with calculus on U_α.

A major advance was the realisation in the early 20th century that the parametrisations ϕ_α should also be *compatible*. For convenience, let us write $V_\alpha = \phi_\alpha[U_\alpha] = \{\phi_\alpha(u) : u \in U_\alpha\}$. Then compatibility means that for each α, β,

$$\phi_\alpha^{-1} \circ \phi_\beta : \phi_\beta^{-1}[V_\alpha] \subset \mathbb{R}^2 \to \phi_\alpha^{-1}[V_\beta] \subset \mathbb{R}^2$$

is a *diffeomorphism* (differentiable with differentiable inverse) between open sets in $\mathbb{R}^2$. This requirement ensures that our calculus on S is independent of the choice of local parametrisation. Such luminaries as Gauss, Riemann and even Einstein did have this compatability condition at their disposal and considerable effort was required in order to show that calculations were independent of the choice of parametrisation. It was Herman Weyl (1912) who gave the first rigourous definition of a smooth manifold, and which wasn't widely accepted until the work of Hassler Whitney in the 1930's.

2. Differentiable manifolds Forget S is a subset of $\mathbb{R}^3$: Simply require the compatibility condition to obtain the definition of a 2-dimensional differentiable manifold! An n-dimensional manifold is defined analogously. This is the *intrinsic* definition of a manifold.

   What about geometry? The length of a curve $\gamma : [a, b] \to S$ is defined as

   $$L[\gamma] = \int_a^b |\gamma'(t)|\, dt$$

   where $\gamma' \in \mathbb{R}^3$ is the tangent vector to the curve thought of as a curve in $\mathbb{R}^3$. But how do we make this intrinsic?

3. Riemannian Manifolds

   Consider a local parametrisation $\phi_\alpha : U_\alpha \to S$. Remember, $U_\alpha \subset \mathbb{R}^2$ is open and so let's introduce *coordinates*, (x, y) on $\mathbb{R}^2$. Fix a point $(x_0, y_0) \in U_\alpha$ and consider the coordinate curves through (x_0, y_0),

   $$\gamma_x(t) = (x_0 + t, y_0) \text{ and } \gamma_y(t) = (x_0, y_0 + t)$$

   which remain in U_α provided t is small enough since U_α is open. Now we observe that the length of a curve depends on integrating the *size* $|\gamma'|$ of the tangent vector and so we define

   $$g_{xx}(x_0, y_0) = \langle \gamma_x'(x_0, y_0), \gamma_x'(x_0, y_0)\rangle, \quad g_{yy}(x_0, y_0) = \langle \gamma_y'(x_0, y_0), \gamma_y'(x_0, y_0)\rangle.$$

   By varying the point (x_0, y_0) in U_α, we obtain two functions on U_α and the length of the x-coordinate curve is

   $$\int \sqrt{g_{xx}(\gamma_x(t))} dt.$$

   Similarly for the y-coordinate curve. Now to obtain the length of an arbitrary curve, we note that the tangent vector γ' may be written uniquely as a linear combination of γ_x' and γ_y',

   $$\gamma' = \gamma^x \gamma_x' + \gamma^y \gamma_y'$$

for $\gamma^x, \gamma^y \in \mathbb{R}$. The length of γ' is obtained from

$$|\gamma'|^2 = \langle \gamma', \gamma' \rangle = (\gamma_x)^2 g_{xx} + (\gamma_y)^2 g_{yy} + 2\gamma_x \gamma_y \langle \gamma'_x, \gamma'_y \rangle$$

Or defining two new functions,

$$g_{xy} = g_{yx} = \langle \gamma'_x, \gamma'_y \rangle$$

we can write

$$|\gamma'|^2 = (\gamma_x)^2 g_{xx} + (\gamma_y)^2 g_{yy} + 2\gamma_x \gamma_y g_{xy}.$$

The symmetric matrix,

$$g = \begin{pmatrix} g_{xx} & g_{xy} \\ g_{yx} & g_{yy} \end{pmatrix}$$

is classically referred to as the *First Fundamental Form*. These days, it is simply referred to as the *metric*, the **central object in the study of Riemannian geometry**. This matrix valued function, defined on U_α has no reference to the fact that S is a subset of $\mathbb{R}^3$ and again we can forget that we got it from working in $\mathbb{R}^3$, and use g to define the notion of a Riemannian Manifold: *a Riemannian manifold is a manifold M with a positive definite, symmetric, matrix valued function $g(x)$.*

This idea, that geometry can be defined in terms of the metric is due to Gauss.

4. Intrinsic Geometry

   Euclid's 5th postulate: Distinct parallel, straight lines never intersect. For many years, mathematicians tried to understand whether Euclid's 5th postulate could be proven from the prior four postulates, or whether it was somehow independent. Major breakthroughs came about through the work of Lobachevsky, Bolyai, and Gauss in the 19th century: **There are geometries satisfying the first four axioms, but not the fifth**. These were the classical, constant negative curvature surfaces (hyperbolic plane): *There exist two lines parallel to a given line through a given point not on the line.* For constant positive curvature (the sphere) we have: *Every line through a point not on a given line meets that given line.* The sphere does not satisfy all the other four axioms of Euclid so this particular geometry probably wasn't so disconcerting. However, by passing to projective space we obtain a geometry satisfying the first four of Euclid's axioms but not the fifth: *There do not exist parallel lines! All lines intersect.*

Riemann submitted for Habilitation on the 10th of June, 1854, in order to become a lecturer. The usual procedure was for the candidate to submit three potential topics on which to lecture, the informal understanding being that the first topic would be accepted by the evaluating committee. Riemann's *third* and thus lowest preference topic was on the intrinsic geometry of surfaces. Gauss was serving on the committee and (somewhat unexpectedly) insisted that Riemann spoke on this third topic as he had worked in the problem and wanted to hear what Riemann had to say about it. Fortunately, Riemann had quite a lot to say! His lecture, "Über die Hypothesen welche der Geometrie zu Grunde liegen" is one of the most famous in mathematics history. We'll find out in this course what he had to say, though considerable developments have been made since his time, in particular the precise definitions smooth manifolds, smooth metrics, connections (Levi-Civita, Christoffel) and curvature, and the tensor calculus developed by Ricci and Shouton.

5. General Relativity (Intrinsic Geometry of Space-Time)

   Speaking of tensor calculus which was developing around the turn of the 20'th century, Einstein was developing his General Theory of Relativity at about this time too. Historically it has taken some time for the mathematical developments to percolate throughout physics, but part of Einstein's genius was to realise that the cutting edge of mathematics was precisely the necessary formalism to develop a *geometric theory of gravity.*

   Roughly speaking, mass curves space-time and curvature of space-time determines the paths of motion (massive and non-massive!) objects will follow in the absence of other forces. We've since observed these effects for instance in the gravitational lensing of light around massive objects such as stars. Light follows a *curved path* induced by the mass of the star.

   These curved paths are known as *geodesics* and are the shortest path between two points (a slight technicality in General Relativity requires we talk about *null-geodesics*, though we won't worry about the distinction here). A beam of light should take the shortest possible path and this is not always what we consider a straight line. Consider the sphere for example: the shortest path between two points on the sphere, where the path is constrained to stay on the surface of the sphere, is along a great circle. These are the geodesics of the sphere. If a body is constrained by some forces to remain on the surface of the sphere and is acted on by no other forces, e.g. a cart rolling freely, it will trace out a path along a geodesic.

Einstein's equation,

$$G + \Lambda g = T$$

roughly speaking, says that the deviation of curvature from *isotropic* (the same in all directions) equals the stress on space-time itself plus a cosmological constant.

Riemannian Geometry (more precisely, Lorentzian Geometry) is the formalism on which the General Theory of Relativity rests! Although we will not discuss Lorentzian geometry in this course, much of the formalism developed here carries over to the Lorentzian setting.

1.2 Week 01, Lecture 02: Overview of Curvature and Review of Manifolds

1.2.1 What Is Curvature?

1. Extrinsic Curvature

   - A curve $\gamma : [a, b] \to \mathbb{R}^2$ has curvature $\kappa = \gamma''$ (assuming $|\gamma'| = 1$). This curvature measures how far away the curve bends from a straight line (in fact from it's tangent line).
   - A surface $S \subset \mathbb{R}^3$ has curvature in 2-dimensions "worth" of directions. Take any path $\gamma : [a, b] \to S \subset \mathbb{R}^3$ (especially a path obtained by intersecting S with a plane) and compute it's curvature in $\mathbb{R}^3$. This is the curvature *of the surface* along the path γ. Thus a surface has two dimensional curvature since at any point, the tangent plane is spanned by the tangent vectors to two curves. Consider for example the curvature of the equator on the sphere, or the curvature of a line of latitude.
   - Examples:
     - A veritable zoo of curves: straight lines, circles, cardiod, ...
     - Plane, sphere, cylinder, torus (doughnut), minimal surfaces,...

2. Intrinsic Curvature

   The discussion above relates to *extrinsic curvature*, i.e. the curvature induced by the manner in which the curve or surface "sits" in Euclidean space. Think of a curve as a piece of string. It has no intrinsic curvature and one way to see this is to note that the geometry along a piece of string is not affected by how you bend the string: the distance (along the string) between two points on the string is no different, no matter how you place it space (provided you don't stretch it). All the curvature comes from how the string is place in space, one can always straighten out the string without breaking it.

   A more illustrative pair of examples, is that of a sphere and a cylinder. Take a piece of paper and draw a straight line on it. Now roll the paper up into a cylinder and measure the length of the curve in space formed by the line on the paper. You'll find it's exactly the same length as the original straight line! A flat sheet of paper obviously has zero curvature, and it turns out although the cylinder has curvature, our discussion about length implies that all this curvature is extrinsic, none of it comes from the *intrinsic geometry of the cylinder* which is exactly that of the flat plane. On the other hand, the sphere has intrinsic curvature, you can't flatten it out without tearing and distorting it.

This is an example of the famous Gauss Theorema Egrigium (Remarkable Theorem), one of the problems that made Gauss so interested in Riemann's work. It states that the curvature (which we have not yet defined and is more precisely referred to as the *Gauss Curvature*) of a surface depends only on the metric (i.e. the matrix g above) and not on the way in which the surface is situated in space. In other words, if two surfaces have the same *geometry* (length, angles, area) then they have the same intrinsic curvature! If we are somehow able to put our two surfaces in space in such a way as to preserve length and angles, then they will have the same intrinsic curvature, though they may have different extrinsic curvature, such as in the example of a cylinder and a plane.

3. Gauss Curvature

   Let B_r denote a ball of radius $r > 0$ on a surface S measured with respect to a metric g (recall the metric allows us to define the notion of length). The *Gauss curvature* K, may be defined by the asymptotic expansion

   $$L[\partial B_r] = 2\pi r(1 - \frac{\mathrm{K}}{6}r^2 + \mathcal{O}(r^4)).$$

   The geometry determines the curvature!

   - Gauss curvature of a plane: $L[\partial B_r] = 2\pi r$ and hence $\mathrm{K} = 0$.
   - Gauss curvature of the unit sphere:

     $$L[\partial B_r] = 2\pi \sin(r) = 2\pi r(1 - \frac{1}{6}r^2 + \mathcal{O}(r^4)),$$

     hence $K = 1$.

4. Global Questions.

   There are many interesting connections between curvature and topology. Here is a brief sample:

- Gauss-Bonnet Formula: $\int_\Omega \mathrm{K} = 2\pi\chi$ where Ω is a compact 2-dimensional manifold (without boundary) and χ denotes it's *Euler Characteristic* (a topological invariant).
- Uniformisation Theorem: A compact 2-dimensional manifold (without boundary) has a metric of constant curvature, and this constant depends only on the genus of the surface (how many holes) or equivalently on the Euler characteristic. Can you work out this constant in terms of χ?
- Cartan-Hadamard Manifolds: A *complete* (to be defined later), simply connected manifold (no holes) with non-positive curvature is diffeomorphic to $\mathbb{R}^n$. It will in general, have a metric different from the usual Euclidean metric.
- Bonnet-Myers: A manifold with curvature bounded below by a positive constant is compact.
- More generally, one can compare manifolds with curvature bounded below (sometimes above) by a constant with so-called, constant curvature models. One obtains geometric inequalities such as volume and/or diameter comparisons, and may also obtain topological information (such as the dimension of the cohomology groups, known as Gromov's Betti numbers estimate). This is the fascinating field on *comparison geometry* which is beyond the scope of this course. An open problem (Milnor's conjecture) says that if a manifold has non-negative curvature, then it is equal (isometric) to the interior of a compact manifold with boundary.

1.2.2 Intrinsic Definition

Definition 1.1. A topological manifold, M is a second countable, Hausdorff topological space, that is *locally Euclidean*. That is, M is covered by open sets U_α and homeomorphisms $\phi_\alpha : U_\alpha \to V_\alpha \subset_{\text{open}} \mathbb{R}^n$. The maps ϕ_α are called *local charts* for M.

Remark 1.2. If M is connected, then the dimension n is constant. More generally, the dimension is constant on the connected components. Conventions vary here, but typically the dimension is *required to be constant* in the definition and this is the convention we will adopt. We will mostly assume that our manifolds are in fact connected, and I'll state when this is not the case.

Definition 1.3. A smooth manifold is a topological manifold equipped with a smooth structure. Namely M has an maximal smooth atlas,

$$\mathcal{A} = \{\phi_\alpha U_\alpha \to V_\alpha\}.$$

A smooth atlas $\mathcal{A}$ is a collection of local charts, in which all the transition maps

$$\tau_{\alpha,\beta} = \phi_\beta \circ \phi_\alpha^{-1} : \phi_\alpha^{-1}[V_\beta] \to \phi_\beta^{-1}[V_\alpha]$$

are smooth (infinitely differentiable) maps between open sets of $\mathbb{R}^n$. The requirement that $\mathcal{A}$ be maximal says that it contains every compatible chart, where a compatible chart $\phi : U \to V$ is any chart such that every transition maps $\phi \circ \phi_\alpha^{-1}$, $\phi_\alpha \circ \phi^{-1}$ are smooth.

Remark 1.4. A surprising result of Milnor (1956) is that a manifold may have more than one smooth structure. With Kervaire (1963) He showed that $\mathbb{S}^7$ possesses 28 incompatible smooth structures, where $\mathcal{A}_1$ is incompatible with $\mathcal{A}_2$ if there are local charts $\phi_i \in \mathcal{A}_i$, with τ_{ij} not differentiable. The spheres equipped with smooth structures not given by the standard structure (inherited from Euclidean space) are referred to as Exotic Spheres.

Definition 1.5. Local coordinates. Given a local chart ϕ_α write $\phi_\alpha = (x^1, \cdots, x^n)$ for the components. The functions $x^i : U_\alpha \to \mathbb{R}$ are called local coordinates.

We often think of $(x_1, \cdots, x_n)$ as the coordinates for a point $x \in M$. Changing to a different local chart gives different coordinates for the point x.

1.2.3 Examples

- Euclidean space! The identity $\mathbb{R}^n \to \mathbb{R}^n$ determines a smooth atlas.
- Regular curves C. In turns out there either there is a single chart for C (C is homeomorphic to $\mathbb{R}$) or C is homeomorphic with $\mathbb{S}^1$.
- Regular surfaces S. Charts are given by the inverse of the local parametrisations.
  - This includes spheres, planes, tori, paraboloids, ...
- Open sets of Euclidean space.

- Since $M_{n,m}(\mathbb{R})$, the set of $n \times m$ manifolds can be identified with $\mathbb{R}^{nm}$, this set is a smooth manifold. The open set $GL_n(\mathbb{R}) = \{M \in M_{n,n} : \det M \neq 0\}$ is thus a smooth manifold.
- The complement of the Cantor set, i.e. all real numbers with a base 3 expansion including the digit 1. Not connected!

- Projective space: $\mathbb{RP}^n = (\mathbb{R}^{n+1} - \{0\})/\sim$ where $v \sim u$ if $v = \lambda u$ for some non-zero $\lambda \neq 0$. This is the set of lines in $\{\mathbb{R}^{n+1}\}$.
- Grassmanians: More generally, the set $G_{kn}(\mathbb{R})$ of real k-planes in $\mathbb{R}^n$ are smooth manifolds.
- Products: $M \times N$ is a smooth manifold with charts $\phi_\alpha \times \psi_i : U_\alpha \times U_i \to V_\alpha \times V_i \subseteq \mathbb{R}^m \times \mathbb{R}^n = \mathbb{R}^{m+n}$ where ϕ_α is a chart for M and ψ_i is a chart for N.

1.3 Week 01, Lecture 03: Riemannian Surfaces, Tangent Vectors and Smooth Maps

1.3.1 Geometry of Surfaces

Recall that a regular, smooth surface S, is a smooth manifold via the smooth local parametrisations, $\phi_\alpha : U_\alpha \to S$, and that on each U_α we have a smooth matrix valued function g, the metric. We can use the metric to measure the length of, and angle between tangent vectors.

A tangent vector X to S at a point x_0 is the derivative $X = \gamma'(0)$ of a curve $\gamma : (-\epsilon, \epsilon) \to S \subset \mathbb{R}^3$ with $\gamma(0) = x_0$. But how do we interpret the derivative? Certainly as a curve in $\mathbb{R}^3$, $\gamma'(0)$ makes sense, but we want to work intrinsically on S without reference to the ambient $\mathbb{R}^3$.

The way forward is to note that given any curve $\gamma : (-\epsilon, \epsilon) \to S$, the tangent vector $\gamma'(0)$ lies in the tangent plane, $T_{x_0}S$ at x_0. Pick a local parametrisation $\phi : U \to V$ with $x_0 \in V$. For convenience, we may assume $\phi(0,0) = x_0$. Then $T_{x_0}S$ is the plane through the point x_0 and spanned by the vectors,

$$\{\phi_u = \tfrac{\partial \phi}{\partial u} = d\phi \cdot e_u, \phi_v = \tfrac{\partial \phi}{\partial v} = d\phi \cdot e_v\}$$

where (u, v) are coordinates on $U \subset \mathbb{R}^2$ and $e_u = (1, 0)$, $e_v = (0, 1)$ are the standard basis vectors on $\mathbb{R}^2$. Recall that the definition of a regular surface requires that $d\phi : \mathbb{R}^2 \to \mathbb{R}^3$ has rank 2 and hence ϕ_u and ϕ_v are linearly independent, i.e. $T_{x_0}S$ really is a plane.

In particular, the map

$$d\phi : \mathbb{R}^2 \to T_{p_0}S$$

is a linear isomorphism. Therefore, we may write

$$X = \gamma'(0) = c_1\phi_u + c_2\phi_v$$

for unique constants $c_1, c_2 \in \mathbb{R}$. But now observe that the curve $\gamma_\phi(t) = \phi^{-1} \circ \gamma(t) \subseteq U$ has tangent vector

$$\gamma'_\phi(0) = c_1 e_u + c_2 e_v$$

and that by definition of g,

$$\langle \gamma', \gamma' \rangle_{\mathbb{R}^3} = c_1^2 g_{uu}(0,0) + 2c_1c_2 g_{uv}(0,0) + c_2^2 g_{vv}(0,0).$$

Thus our expression for the length of the tangent vector $X = \gamma'(0) = c_1 e_u + c_2 e_v$ may be written

$$|X|^2 = g(X, X) = \begin{pmatrix} c_1 & c_2 \end{pmatrix} \begin{pmatrix} g_{uu} & g_{uv} \\ g_{vu} & g_{vv} \end{pmatrix} \begin{pmatrix} c_1 \\ c_2 \end{pmatrix}.$$

The angle θ between two tangent vectors $X = c_1 e_u + c_2 e_v$ and $Y = b_1 e_u + b_2 e_v$ is given by

$$\cos(\theta) = g(X, Y) = \begin{pmatrix} c_1 & c_2 \end{pmatrix} \begin{pmatrix} g_{uu} & g_{uv} \\ g_{vu} & g_{vv} \end{pmatrix} \begin{pmatrix} b_1 \\ b_2 \end{pmatrix}.$$

Notice that since $g_{uv} = g_{vu}$, the metric g is symmetric: $g(X, Y) = g(Y, X)$.

1.3.2 Smooth maps

Definition 1.6. Let M^m and N^n be smooth manifolds of dimension m and n respectively. A map $f : M \to N$ is smooth if for every chart ϕ on M and every smooth chart ψ on N, the map

$$\psi \circ f \circ \phi^{-1} : U \subset \mathbb{R}^m \to V \subset \mathbb{R}^n$$

is smooth.

Let us write

$$\psi \circ f \circ \phi^{-1}(x^1, \cdots, x^m) = (f_1(x^1, \cdots, x^m), \cdots, f_n(x^1, \cdots, x^m)).$$

To say that f is smooth is equivalent to the requirement that each f_j has smooth partial derivatives of all orders. The composition of smooth maps is easily seen to be smooth from the definition. Write it out!

Definition 1.7. Let $f : M \to N$ be a smooth map. Then we say f is a diffeomorphism if f has a smooth inverse $f^{-1} : N \to M$. That is $f \circ f^{-1} = \mathrm{Id}_N$ and $f^{-1} \circ f = \mathrm{Id}_M$. Two manifolds are diffeomorphic if there exists a diffeomorphism between them. Typically we identify such manifolds.

Remark 1.8. Two atlases $\mathcal{A}_1$ and $\mathcal{A}_2$ are compatible if and only if the identity map $(M, \mathcal{A}_1) \to (M, \mathcal{A}_2)$ is a diffeomorphism. Why? What does this mean for exotic spheres?

There are some variants of diffeomorphism: immersions and submersions characterised by maximality of the rank of the differential and embeddings, realising a manifold as a sub-manifold of another manifold.

Definition 1.9. A smooth map $f : M \to N$ is an *immersion*, if given any charts as above, the map $\psi \circ f \circ \phi^{-1}$ has injective differential. It is a *submersion* if the differential is surjective. The map is an embedding if it is an immersion, and is a homeomorphism with it's image equipped with the subspace topology.

Recall that the differential $d\bar{f}_{x_0}$ at x_0 of the function $\bar{f} = \psi \circ f \circ \phi^{-1}$ between (open sets of) $\mathbb{R}^m$ and $\mathbb{R}^n$ is the unique linear map $d\bar{f}$ such that

$$f(x) = f(x_0) + d\bar{f}_{x_0} \cdot (x - x_0) + \mathrm{o}(|x - x_0|).$$

It can be realised explicitly in two common ways:

1. $d\bar{f}_{x_0} \cdot v = \partial_t|_{t=0} (\bar{f}(\gamma(t))$ where $\gamma(t)$ is any curve with $\gamma(0) = x_0$ and $\gamma'(0) = v$, e.g. $\gamma(t) = x_0 + tv$.

2. $$d\bar{f}_{x_0} = \begin{pmatrix} \frac{\partial f_1}{\partial x^1} & \cdots & \frac{\partial f_1}{\partial x^m} \\ \vdots & \ddots & \vdots \\ \frac{\partial f_n}{\partial x^1} & \cdots & \frac{\partial f_n}{\partial x^m} \end{pmatrix}$$
   where, as above, $\bar{f}(x^1, \cdots, x^m) = (f_1(x^1, \cdots, x^m), \cdots, f_n(x^1, \cdots, x^m))$.

Notice that the definition of immersion may be restated as $m \leq n$ and the rank of $d(\psi \circ f \circ \phi^{-1})$ is m, and the definition of sumbersion may be stated as $m \geq n$ and the rank of $d(\psi \circ f \circ \phi^{-1})$ is n. We therefore make the definition

$$\mathrm{rank}(f) = \mathrm{rank}(d(\psi \circ f \circ \phi^{-1}))$$

for charts ψ, ϕ. The rank of a smooth map is well defined, independent of the choice of charts. Suppose for instance that $U_\alpha \cap U_\beta \neq \emptyset$ for charts U_α, U_β. Then on the intersection,

$$\psi \circ f \circ \phi_\beta^{-1} = \psi \circ f \circ \phi_\alpha^{-1} \circ \phi_\alpha \circ \phi_\beta^{-1} = \psi \circ f \circ \phi_\alpha^{-1} \circ \tau_{\beta,\alpha}$$

The rank of the right hand side is obviously equal to the rank of the left hand side. But now, since $\tau_{\beta,\alpha}$ is an diffeomorphism, the rank of the right hand side is the same as the rank of $\psi \circ f \circ \phi_\alpha^{-1}$ as claimed. Similarly, one obtains the rank is independent of the choice of ψ.

Remark 1.10. Thanks to the Implicit Function Theorem, immersions locally look like inclusions,

$$\begin{aligned}\iota :& \mathbb{R}^k \to \mathbb{R}^{n+k} \\ &(x^1, \cdots, x^k) \mapsto (x^1, \cdots, x^k, 0, \cdots, 0)\end{aligned}$$

and submersions locally look like projections

$$\begin{aligned}\pi :& \mathbb{R}^{n+k} \to \mathbb{R}^k \\ &(x^1, \cdots, x^k, x^{k+1}, \cdots, x^{n+k}) \mapsto (x^1, \cdots, x^k).\end{aligned}$$

An immersion is locally an embedding, meaning that M can be covered by open sets U_α (not necessarily charts) on which $f|_{U_\alpha}$ is an embedding.

1.3.3 Tangent Vectors

A tangent vector may be defined in one of three ways, as tangent vectors in charts, as equivalence relations of curves equal to first order and as derivations. The first is essentially the definition as given above for regular surfaces.

Fix a point $x \in M$ in each of the definitions below.

Definition 1.11. Let ϕ_α, ϕ_β be charts for M with $\phi_\alpha(x^\alpha) = x = \phi_\beta(x^\beta)$ with $x^\alpha \in U_\alpha$ and $x^\beta \in U_\beta$. Define an equivalence relation on pairs $(\phi_\alpha, u \in \mathbb{R}^n)$ by

$$(\phi_\alpha, u) \sim (\phi_\beta, v) \text{ if } v = d\tau_{\alpha,\beta} \cdot u.$$

A tangent vector at $x \in M$ is such an equivalence class.

Definition 1.12. Define an equivalence class on curves γ where

$$\gamma \sim \nu \text{ if } \gamma(0) = \nu(0) = x, \text{ and in a chart }, (\phi \circ \gamma)'(0) = (\phi \circ \nu)'(0).$$

Such an equivalence class is a tangent vector at $x \in M$.

Note that by applying the transition map between charts, if $(\phi \circ \gamma)'(0) = (\phi \circ \nu)'(0)$ in one chart, then the analogous equality holds in any chart. Write it out to see explicitly!

Definition 1.13. Let $C^\infty(M) = C^\infty(M, \mathbb{R})$ denote the ring of smooth, real valued functions on M. A derivation at x, d_x is a map $C^\infty(M) \to \mathbb{R}$ such that

1. $d(c_1 f + c_2 g) = c_1 d(f) + c_2 d(g)$ for $f, g \in C^\infty(M, \mathbb{R})$ and $c_1, c_2 \in \mathbb{R}$ (linearity over $\mathbb{R}$,
2. $d(fg) = d(f)g(x) + f(x)d(g)$ (Leibniz rule).

A derivation at x is a tangent vector.

To see the equivalence of the first two definitions, take a pair ϕ_α, u and define consider the curve $\gamma(t) = \phi_\alpha(x + tu)$. The equivalence class of the first is mapped to the equivalence class of the second and this establishes an isomorphism between the first and second definitions (check this!). The equivalence of the second and third definitions is established by mapping a curve γ to the derivation $d(f) = \frac{\partial}{\partial_t}\Big|_{t=0} (f \circ \gamma(t))$. This establishes an isomorphism between the second and third definitions (check this also!).

Essentially, the first definition says we can define tangent vectors as tangent vectors in charts (which we already understand), and we identify tangent vectors in different charts via the transition maps. The second definition says tangent vectors should be the tangents to curves. The third definition says that tangent vectors uniquely determine directional derivatives: the derivations.

Choose local coordinates $\phi : U \subset M \to V \subset \mathbb{R}^n$ around x (i.e. $x \in U$), and write $(x^1, \cdots, x^n)$ for the coordinates on $\mathbb{R}^n$. Let us again assume that $\phi(0, \cdots, 0) = x$. The first definition says that a tangent vector X can be written

$$v = \sum_i v^i e_i$$

for coefficients $X^i \in \mathbb{R}$ and $e_i = (0, \cdots, 0, 1, 0, \cdots 0)$ the i'th standard basis element for $\mathbb{R}^n$ with the 1 in the i'th place.

The second definition says that $v = \gamma'(0)$ where

$$\gamma(t) = tv = (tv^1, \cdots, tv^n).$$

The third definition says that v acts as derivation on smooth functions. If $f : M \to \mathbb{R}$ is a smooth function, define $f_\phi = f \circ \phi^{-1} : V \to \mathbb{R}$ to be the representation of f in a chart. Then as a derivation,

$$v(f) = \sum_i v^i \frac{\partial f}{\partial x^i}(0, \cdots, 0).$$

For this reason, it is conventional to write $\frac{\partial}{\partial x^i}$ for e_i and we will often abbreviate this as ∂_i or ∂_{x^i}.

2 Week 02

2.1 Week 02, Lecture 01: Push Forward and The Tangent Bundle

2.1.1 Push Forward

Definition 2.1. Let $F : M \to N$ be a smooth map. The *push forward* (or differential), $F_*(v)$ of a vector $v \in T_xM$ by F, is the tangent vector in $T_{F(x)}N$ defined to be the derivation,

$$(F_*(v))(f) = v(f \circ F).$$

Alternatively, realising v as an equivalence class of curves, $v = [\gamma]$,

$$F_*(v) = \left.\frac{d}{dt}\right|_{t=0} (F \circ \gamma(t)).$$

In local coordinates, this is

$$F_*(v) = \begin{pmatrix} \frac{\partial F^1}{\partial x^1} & \cdots & \frac{\partial F^1}{\partial x^n} \\ \vdots & \ddots & \vdots \\ \frac{\partial F^m}{\partial x^1} & \cdots & \frac{\partial F^m}{\partial x^n} \end{pmatrix} \begin{pmatrix} v^1 \\ \vdots \\ v^n \end{pmatrix}$$

where in local coordinates $F = (F_1(x_1, \cdots, x_n), \cdots, F_m(x_1, \cdots, x_n))$.

Remark 2.2. Immersions, submersions and embeddings may now be defined in terms of the push-forward. The map $F : M \to N$ is an immersion if $F_* : T_xM \to T_{F(x)}N$ is injective at every point. It is a submersion if this map is surjective at every point an is an embedding if it's an immersion and a homeomorphic with it's image equipped with the subspace topology.

Example 2.3. Let $\phi : U \to V$ be a local chart for M. Then the map $\phi^{-1} : V \subset \mathbb{R}^n \to M$ is a diffeomorphism with it's image. What do we need to show to prove this statement? We need to show that ϕ^{-1} is differentiable and that it has a differentiable inverse. We already have an inverse, ϕ itself so we need to show that both these maps are differentiable.

To show ϕ^{-1} is differentiable we need to show that given any chart, ψ, the map $\psi \circ \phi^{-1}$ is differentiable. But this is immediate since $\psi \circ \phi^{-1}$ is precisely the transition map between the charts.

To show ϕ is differentiable, we need to show that given any chart ψ, the map $\phi \circ \psi^{-1}$ is differentiable and again this is immediate since it's the reverse transition map.

Now, tangent vectors with respect to the chart ϕ, i.e. tangent vectors on V are just $v_\phi = v_\phi^i \partial_i$ for real numbers v_ϕ^i. Since ϕ^{-1} is a diffeomorphism, $d\phi^{-1}$ must be an isomorphism. This latter statement follows from the chain rule

$$\mathrm{Id}_V = d(\phi \circ \phi^{-1}) = d\phi \cdot d\phi^{-1}$$

so that

$$(d\phi^{-1})^{-1} = d\phi.$$

Thus any tangent vector $v \in T_x M$ may be written

$$v = d\phi^{-1} \cdot v_\phi.$$

In other words, we can do all our calculations on V and then transfer the result to M by $d\phi^{-1}$. Thus we often identify M with V via the chart ϕ, and abuse notation by not distinguishing between the two. In particular, on V we have the canonical basis vector fields $\partial_i = e_i = (0, \cdots, 0, 1, 0, \cdots, 0)$ which determine vector fields on U (same definition as for vector fields on M but these are defined only on $U \subseteq M$): $d\phi^{-1}\partial_i = [\phi^{-1}(x + te_i)]$ **and these are commonly also denoted** ∂_i where $[\cdot]$ denotes an equivalence class of curves.

Example 2.4. It is rumoured that Archimedes requested a sphere and cylinder to be placed on his tombstone, to honour his discovery of the properties of the projection from a cylinder onto a sphere. Consider the bounded, open ended cylinder

$$C = \{x^2 + y^2 = 1, -1 < z < 1\}$$

and the sphere

$$\mathbb{S}^2 = \{x^2 + y^2 + z^2 = 1\}.$$

The projection that so interested Archimedes is the map,

$$(x, y, z) \in C \mapsto (\sqrt{1 - z^2}x, \sqrt{1 - z^2}y, z)$$

taking points in horizontal slices $\{z = \text{ constant}\}$ of the cylinder to the nearest point in the same slice on the sphere. This is an injective map with image $\mathbb{S}^2 - \{(0, 0, \pm 1)\}$, the entire sphere minus the north and south poles.

How do we calculate the differential of this map? Let's consider two possibilities. One is to compute the differential as map $\mathbb{R}^3 \to \mathbb{R}^3$ restricted to tangent vectors of the cylinder, and the other is to express it in local coordinates.

For the first method, the differential is represented by the matrix

$$\begin{pmatrix} \sqrt{1-z^2} & 0 & -\frac{xz}{\sqrt{1-z^2}} \\ 0 & \sqrt{1-z^2} & -\frac{yz}{\sqrt{1-z^2}} \\ 0 & 0 & 1 \end{pmatrix}.$$

At a point $(x, y, z) \in C$ a basis for the tangent space to C is

$$\{(-y, x, 0), (0, 0, 1)\}.$$

The differential acting on these vectors is

$$\begin{pmatrix} \sqrt{1-z^2} & 0 & -\frac{xz}{\sqrt{1-z^2}} \\ 0 & \sqrt{1-z^2} & -\frac{yz}{\sqrt{1-z^2}} \\ 0 & 0 & 1 \end{pmatrix}, \begin{pmatrix} -y \\ x \\ 0 \end{pmatrix} = \begin{pmatrix} -y\sqrt{1-z^2} \\ x\sqrt{1-z^2} \\ 0 \end{pmatrix} = \{\sqrt{1-z^2} \begin{pmatrix} -y \\ x \\ 0 \end{pmatrix}$$

and

$$\begin{pmatrix} \sqrt{1-z^2} & 0 & -\frac{xz}{\sqrt{1-z^2}} \\ 0 & \sqrt{1-z^2} & -\frac{yz}{\sqrt{1-z^2}} \\ 0 & 0 & 1 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} = \begin{pmatrix} -\frac{xz}{\sqrt{1-z^2}} \\ -\frac{yz}{\sqrt{1-z^2}} \\ 1 \end{pmatrix}.$$

The first vector is simply scaled according to height, z. The second vector is rotated from the vertical towards the origin until it is tangent to the sphere. Draw the picture! Now that you have the picture, it should be clear what's going on qualitatively. What's going on quantitatively comes from the calculation. Understanding how these two relate leads to deeper insights!

The second method is to use coordinates and obtain a coordinate expression for the differential. Taking coordinates

$$(z, \theta) \in (-1, 1) \times (0, 2\pi) \mapsto (\cos\theta, \sin\theta, z)$$

for C, the projection may be written

$$\varphi : (z, \theta) \mapsto (\sqrt{1-z^2}\cos\theta, \sqrt{1-z^2}\sin\theta, z).$$

Taking coordinates

$$(w, \phi) \in (-1, 1) \times (0, 2\pi) \mapsto (\sqrt{1-w^2}\cos\phi, \sqrt{1-w^2}\sin\phi, w)$$

for $\mathbb{S}^2$, the projection becomes

$$(w, \phi) = (z, \theta)$$

and the differential is just the 2×2 identity in these coordinates!

To understand what's going on here, observe that the tangent vector ∂_z is mapped to the vector ∂_z via the parametrisation of the cylinder above while the tangent vector ∂_θ is mapped to $-\sin\theta\partial_x + \cos\theta\partial_y$.

For the sphere, the tangent vector ∂_w is mapped to $-\frac{\cos\phi w}{\sqrt{1-w^2}}\partial_x, -\frac{\sin\phi w}{\sqrt{1-w^2}}\partial_y + \partial_z$ while ∂_ϕ is mapped to $-\sqrt{1-w^2}\sin\phi\partial_x + \sqrt{1-w^2}\cos\phi\partial_y$.

Now for the cylinder, make the substitutions $x = \cos\theta$, $y = \sin\theta$ and $z = z$ to obtain expressions for the tangent vectors given in the first method. For the sphere, make the substitutions, $x = \sqrt{1-w^2}\cos\phi$, $y = \sqrt{1-w^2}\sin\phi$ and $z = w$ to obtain the expressions for the tangent vectors on the sphere and trace everything back to see you have exactly the same thing. Therefore, this particular map is particularly easy to work with in the right coordinates, where it's just the identity.

Understanding how the two approaches relate most certainly leads to deeper insights. We'll see below an alternative way to view this situation, and how the metrics on the cylinder and sphere relate to these coordinates. This latter way of viewing the situation will lead to yet deeper insights, and highlights some important aspects of approach in this course.

2.1.2 The Tangent Bundle

Definition 2.5. A smooth vector field, X is smooth choice $X(x)$ of tangent vector at each point $x \in M$. By smooth, we mean that in any local coordinate representation,

$$X(x) = \sum_i X^i(x)\partial_i$$

where the X^i are smooth real valued functions. Another name for a vector field is *smooth section of the tangent bundle.*

Definition 2.6. The tangent bundle TM of M is the set of all tangent vectors at any point,

$$TM = \sqcup_{x \in M} T_x M$$

where $\sqcup$ denotes the disjoint union.

The tangent bundle is an example of a vector bundle (more on this next lecture), and can be given the structure of a smooth manifold as follows. Let $\phi : U \to V$ be a chart on M and define a map $\Phi : TU = \sqcup_{x\in U} T_xM \to V \times \mathbb{R}^n \subseteq_{\text{open}} \mathbb{R}^{2n}$ by,

$$\Phi : v = [\gamma] \mapsto (\phi(\gamma(0)), (\phi \circ \gamma)'(0))$$

where we may choose any representation of the equivalence class $[\gamma]$ since by definition of this class $x = \gamma(0) = \sigma(0)$ and $(\phi\circ\gamma)'(0) = (\phi\circ\sigma)'(0)$ for $\gamma \sim \sigma$.

Now I claim that Φ is a bijection.

Let $(x, v) \in V \times \mathbb{R}^n$ and define the curve

$$\gamma(t) = \phi^{-1}(x + tv).$$

Then $\gamma(0) = \phi^{-1}(x)$ and $\phi\circ\gamma(t) = x+tv$ clearly satisfies $(\phi\circ\gamma)'(0) = x$ and hence $\Phi([\gamma]) = (x, v)$ and Φ is surjective.

To show Φ is injective, suppose that $\Phi([\gamma]) = \Phi([\sigma])$. That is,

$$(\phi(\gamma(0)), (\phi \circ \gamma)'(0)) = (\phi(\sigma(0)), (\phi \circ \sigma)'(0)).$$

Then by definition, $[\gamma] = [\sigma]$.

Now we may equip TM with a topology by specifying the sets

$$\Phi^{-1}(W)$$

as open for all charts ϕ and all open sets $W \subset V \times \mathbb{R}^n$. This generates a topology for which each map Φ is a homeomorphism, and hence TM has the structure of a topological manifold. The differentiable structure is obtained from the atlas

$$\mathcal{A} = \{\Phi : TU \to V \times \mathbb{R}^n | \phi : U \to V \text{ is a chart}\}.$$

The transition maps are

$$\Phi \circ \Psi^{-1} = (\phi \circ \psi^{-1}, d\phi \cdot d\psi^{-1})$$

which is smooth since $\phi \circ \psi^{-1}$ is smooth. The inverse is $\Psi \circ \Phi^{-1}$ which is smooth for the same reason. Hence the transition maps are diffeomorphisms, and our atlas determines a differentiable structure on TM, and therefore is a differentiable manifold.

Moreover, if $v \in TM$ is in the tangent space T_xM, then the map $\pi(v) = x$ from $TM \to M$ is a submersion. This follows immediately from the fact that in local charts it is just the projection

$$V \times \mathbb{R}^n \to V.$$

2.2 Week 02, Lecture 02: Riemannian Metrics

2.2.1 Riemannian Metrics

Definition 2.7. A *Riemannian metric*, g is a smooth choice g_x of an inner-product on T_xM at each point $x \in M$.

Recall that an inner product on a vector space is a positive definite, symmetric bilinear form, i.e.

- $g_x(c_1X_1 + c_2X_2, Y) = c_1g_x(X_1, Y) + c_2g_x(X_2, Y)$
- $g_x(X, Y) = g(Y, X)$
- $g_x(X, X) \geq 0$
- $g_x(X, X) > 0$ if $X \neq 0$

To say that g is smooth is to say that for any smooth vector fields X and Y, the function

$$x \in M \mapsto g_x(X(x), Y(x))$$

is a smooth, real valued function on M. Equivalently, in any local coordinate chart, $\phi : U \to V$ if we define $g_{ij}(y) = g_{\phi^{-1}(y)}((\phi^{-1})_*\partial_i, (\phi^{-1})_*\partial_j)$ for $y \in V$, then g can be represented as

$$g(X, X) = \begin{pmatrix} X^1 & \vdots & X^n \end{pmatrix} \begin{pmatrix} g_{11} & \cdots & g_{1n} \\ \vdots & \ddots & \vdots \\ g_{n1} & \cdots & g_{nn} \end{pmatrix} \begin{pmatrix} X^1 \\ \vdots \\ X^n \end{pmatrix}$$

with each g_{ij} a smooth real valued function on $V \subset \mathbb{R}^n$ and where $X = \sum X^i\partial_i$.

Remark 2.8. Recall that we often abuse notation as write ∂_i for $(\phi^{-1})_*\partial_i$. In this convention we then simply write $g_{ij} = g(\partial_i, \partial_j)$ and this may be considered either as a smooth function on U or a smooth function on V as circumstances dictate. This is a very common abuse, but usually does not cause great difficulty for long if one keeps in mind that smooth functions on U are in one-to-one correspondence with smooth functions on V via ϕ.

Proposition 2.9. *Let g be a Riemannian metric. Then the following are equivalent,*

1. *$x \mapsto g_x(X(x), Y(x))$ is a smooth function for every $x \in M$ and every pair of smooth vector fields X, Y on M,*

2. *for any chart $\phi : U \to V$ the n^2 functions $y \in V \mapsto g_{ij}(y)$ are smooth,*

3. *the function $y \in V \mapsto (g_{ij}(y)) \in \mathbb{R}^{n^2}$ is smooth where (a_{ij}) denotes a the matrix with entries a_{ij}.*

Proof.

- $2 \Leftrightarrow 3$: Immediate since a function $\mathbb{R}^n \to \mathbb{R}^{n^2}$ is smooth if and only if the component functions are smooth.

- $1 \Rightarrow 2$: This is almost immediate, since applying 1 to the vector fields ∂_i, ∂_j we find the function $g_{ij} = g(\partial_i, \partial_j)$ is smooth. A slight technicality here is that the vector fields $\{\partial_i\}$ are not defined on all of M, but only on the open set U. To get around this technicality we use a bump function. The question of smoothness is local, namely it is enough to show for each $x \in U$, there is an open neighbourhood $W \subset U$ of x on which g_{ij} is smooth. Therefore, fix x and choose any open neighbourhood $W \subsetneq U$ with $x \in W$. Let $\phi : M \to \mathbb{R}$ be a smooth bump function, identically equal to 1 on W and identically equal to 0 outside of U.

  Define vector fields $X_i = \phi\partial_i$ on U and the zero vector outside of U. By the definition of ϕ, these vector fields are smooth and hence $g(X_i, X_j)$ is a smooth function by the assumption 1. Now just observe that since $\phi \equiv 1$ on W, we have $g_{ij} = g(X_i, X_j)$ on W and therefore g_{ij} is smooth on W as required.

- $(2 \Rightarrow 1)$: Again the question is local, namely it is enough to show that the function $g_x(X(x), Y(x))$ is smooth function on each U_α where U_α is an open cover of M. Choose U_α to be an open cover by charts. Choose any chart and write $X = \sum X^i\partial_i$, $Y = \sum Y^i\partial_i$ with each X^i, Y^i a smooth function for $i = 1, \cdots, n$. Then

  $$g_x(X(x), Y(x)) = g_x(\sum_i X^i(x)\partial_i, \sum_j Y^j(x)\partial_j) = \sum_{i,j} X^i(x)Y^j(x)g_{ij}(x)$$

  by bilinearity. The right hand side is smooth since the functions X^i, Y^j and g_{ij} are smooth. Since this is true in any chart, we're done.

□

Example 2.10. On $\mathbb{R}^n$, with the single, global chart, $g_{ij} = \delta_{ij}$ is constant, hence smooth.

Example 2.11. On a regular surface S, a curve $\gamma : I \to S$ may also be considered as a curve $\gamma : I \to \mathbb{R}^3$ and hence a tangent vector to S may also be considered as a tangent vector to $\mathbb{R}^3$. Define a Riemannian metric,

$$g(X,Y) = \langle \gamma'(0), \sigma'(0) \rangle_{\mathbb{R}^3}$$

where we represent $X = [\gamma]$, $Y = [\sigma]$ as equivalence classes of curves, and on the right hand side we think of these as curves in $\mathbb{R}^3$. It's not hard to check that this definition is independent of the curve in the equivalence class. In a local parametrisation $\phi : U \subset \mathbb{R}^2 \to S \subset \mathbb{R}^3$,

$$g_{ij} = \langle \frac{\partial \phi}{\partial x^i}, \frac{\partial \phi}{\partial x^j} \rangle_{\mathbb{R}^3}$$

is a smooth function on U since ϕ is smooth.

Example 2.12. On the sphere, locally parametrised by

$$u(\phi, \theta) = (\sin\phi\cos\theta, \sin\phi\sin\theta, \cos\phi)$$

we have

$$\frac{\partial u}{\partial \theta} = (-\sin\phi\sin\theta, \sin\phi\cos\theta, 0)$$

and

$$\frac{\partial u}{\partial \phi} = (\cos\phi\cos\theta, \cos\phi\sin\theta, -\sin\phi)$$

. Therefore,

$$g_{\theta\theta} = \sin^2\phi \quad g_{\phi\phi} = 1 \quad g_{\theta\phi} = g_{\phi\theta} = 0.$$

In matrix form,

$$g = \begin{pmatrix} \sin^2\phi & 0 \\ 0 & 1 \end{pmatrix}.$$

2.3 Week 02, Lecture 03: Vector Bundles

2.3.1 Vector Bundles

The definition in the previous lecture a Riemannian metric is perfectly correct, but relies too heavily on local coordinates. One of the aims of this course is for you to think more globally. The kind of construction above is typical of early 20'th century differential geometry. The principal draw back is that it constrains our thought processes too much: with it we'll constantly deal with technicalities and won't be able to see the forest for the trees. So let us once and for all develop the formalism necessary to allow us to see the splendours of the forest of differential geometry. That is, we need the notion of a *vector bundle.*

Definition 2.13. A smooth *Vector Bundle* $\pi : E \to M$ of rank k, is a triple (π, E, M) with E and M smooth manifolds and π a smooth, surjective map where locally E looks like a product $U \times \mathbb{R}^k$ with $U \subseteq_{\text{open}} M$. More precisely, there exists an open cover $\{U_\alpha\}$ of M and *local trivialisations*

$$\phi_\alpha : \pi^{-1}[U_\alpha] \subset E \to U_\alpha \times \mathbb{R}^k$$

satisfying

1. ϕ_α is a diffeomorphism,

2. $p \circ \phi_\alpha = \pi$ where $p : U_\alpha \times \mathbb{R}^k \to U_\alpha$ is the projection onto the first factor, and

3. The *transition maps*,

   $$\tau_{\alpha\beta} = \phi_\beta \circ \phi_\alpha^{-1} : U_\alpha \cap U_\beta \times \mathbb{R}^k \to U_\alpha \cap U_\beta \times \mathbb{R}^k$$

   are of the form
   $$\tau_{\alpha\beta}(x, v) = (x, A_{\alpha\beta}(x) \cdot v)$$
   where $A_{\alpha\beta} : U \to GL_n(\mathbb{R})$ is a smooth map, so that $A_{\alpha\beta}(x)$ is an invertible, linear transformation for each x.

Some terminology: E is called the *total space*, M is called the *base space* and π is called the projection.

Remark 2.14. For each $x \in M$, the set $E_x = \pi^{-1}[\{x\}]$ is called the *fibre* of E over x. Condition 2 implies that in any local trivialisation,

$$\phi_\alpha : E_x \to \{x\} \times \mathbb{R}^k$$

is a bijection and hence may be equipped with a natural vector space structure.

This structure smoothly depends on x. Via ϕ_α, we identify $\pi^{-1}[U_\alpha] \subset E$ with $U_\alpha \times \mathbb{R}^k$ and the fibres E_x with $\{x\} \times \mathbb{R}^k$. From this identification, we may think of elements of $\pi^{-1}[U_\alpha]$ as vectors $v \in \mathbb{R}^k$ based at points $x \in U_\alpha$. The fibres vary smoothly with x in the sense that if we choose any fixed vector $v \in \mathbb{R}^k$, then the map $x \in U_\alpha \mapsto \phi_\alpha^{-1}(x, v) \in E$ is smooth.

What if we use different local trivialisations to identify elements of E with the pair (x, v)? That is, at a point x, we may identify $v \in E_x$ with a vector in $\mathbb{R}^k$ using different local trivialisations. The transition maps allow us to identify these different local representations of v. Namely, if $\phi_\alpha(v) = (x, v_\alpha)$ and $\phi_\beta(v) = (x, v_\beta)$ with $v_\alpha, v_\beta \in \mathbb{R}^{\rceil}$, then these two vectors relate by the transition maps,

$$v_\beta = A_{\alpha\beta}(x) \cdot v_\alpha.$$

Moreover, since $A_{\alpha\beta}$ is a smooth function, this relation varies smoothly with the point x in a similar way to how fibres vary smoothly as described above. Thus we think of local trivialisations as a choice of local basis for E and the transition maps are just the change of basis.

Lets look at some examples, to see if we can unpack the definition a little and discover what it means exactly.

Definition 2.15. A *Trivial Bundle* is a bundle of the form $\pi : M \times \mathbb{R}^k \to M$ with π the projection onto the first factor. Here $E = M \times \mathbb{R}^k$ is certainly a smooth manifold (the product of two smooth manifolds is again a smooth manifold) and π is a smooth surjection. For our open cover, we may choose the single open set M and $\phi : E \to M \times \mathbb{R}^k$ the identity. Lets check this satisfies the requirements:

1. The identity is certainly a diffeomorphism!
2. we already took $\pi = p$ and our local trivialisation is just the identity, so $\pi = \mathrm{Id} \circ p$ is true.
3. We only have one local trivialisation and the transition map is just the identity.

Example 2.16. Consider the cylinder $C = \mathbb{S}^1 \times \mathbb{R}$. This is a trivial bundle over $\mathbb{S}^1$ with $\pi : (\theta, t) \in C \mapsto \theta \in \mathbb{S}^1$. In particular, for any $x_0 \in \mathbb{R}$, the map $\mathbb{S}^1 \to C$,

$$s_{x_0} : \theta \in \mapsto (\theta, x_0)$$

is smooth and satisfies $\pi \circ s_{x_0} = \mathrm{Id}_{\mathbb{S}^1}$.

Example 2.17. Consider the Möbius strip, $\mathcal{M} = [0, 2\pi] \times \mathbb{R}$ modulo the equivalence relation $(0, x) \sim (2\pi, -x)$. Define $\pi : (\theta, x) \in \mathcal{M} \mapsto \theta \in \mathbb{S}^1$ where we think of $\mathbb{S}^1 = [0, 2\pi]$ modulo $0 \sim 2\pi$. As an exercise, check that this is in fact a vector bundle over $\mathbb{S}^1$. Note that any continuous function $s : \mathbb{S}^1 \to \mathcal{M}$ with the property that $\pi \circ s = \mathrm{Id}_{\mathbb{S}^1}$ must be zero somewhere: $s(\theta_0) = (\theta_0, 0)$ for some $\theta_0 \in \mathbb{S}^1$. This follows from the intermediate value theorem and the equivalence relation defining the Möbius strip. The Möbius strip is therefore, not a trivial bundle. Why? *Hint*: compare with the cylinder.

Definition 2.18. A *smooth section* of E is a smooth map $s : M \to E$ such that $\pi \circ s = \mathrm{Id}_M$. The set of all smooth sections is often denoted $\Gamma(M, E)$, or just $\Gamma(E)$ when M is clear from context. A *smooth, local section* is a smooth map $s : U \subseteq_{\text{open}} M \to E$ such that $\pi \circ s = \mathrm{Id}_M$. The set of local sections over an open set U is denoted $\Gamma(U, E)$, or sometime $\Gamma_U(E)$.

Recall that E is itself a smooth manifold and thus the notion of smooth map $M \to E$ is well defined. The condition, $\pi \circ s = \mathrm{Id}_M$ implies that $s(x) \in E_x$, or in other words, smooth sections are a smoothly varying choice of vector $s(x) \in E_x$ for each $x \in M$ (or $x \in U$ for local sections).

Example 2.19. The tangent bundle, TM to M is a smooth vector bundle of rank n over M. An element $v \in TM$ is an element of one of the fibres T_xM and the projection is $\pi(v) = x$. The local trivialisations are given by the charts $\{U_\alpha\}$ for M with transition maps $A_{\alpha\beta} = d(\tau_{\alpha\beta})$. A smooth section is simply a vector field X, and the set of all vector fields may be denoted $\Gamma(M, TM)$ as above, but also commonly as $\mathfrak{X}(M)$. This bundle is the most important bundle in differential geometry.

Example 2.20. A tangent vector on $\mathbb{R}^n$ is the tangent vector to a curve $\gamma'(0)$ based at a point x. Translating this to the origin, we may think of $\gamma'(0)$ as an element of $\mathbb{R}^n$. In this way we see that $T\mathbb{R}^n \simeq \mathbb{R}^n \times \mathbb{R}^n$ is a trivial vector bundle.

Example 2.21. The amusingly named, *Hairy Ball Theorem* states that the tangent bundle on $\mathbb{S}^2$ is not trivial. Even more striking is the fact that there is no non-vanishing, continuous vector field on $\mathbb{S}^2$. This implies for instance, that if the force of the wind on the surface of the earth at a fixed time is modelled as a continuous vector field, then there is always a point on the earth with no wind!

On the other hand, the tangent bundle to $\mathbb{S}^1$ may be identified with the cylinder, and as such, is a trivial vector bundle. The 3-sphere, $\mathbb{S}^3$ and the 7-sphere, $\mathbb{S}^7$ also have trivial tangent bundles. This follows, for those that are interested, for example by identifying the 3-sphere as the *unit Quaternions* giving it a Lie Group structure, and identifying the 7-sphere as the *unit Octonions* which has a non-associative multiplication structure similar to a Lie Group (the only diffrence being the lack of associativity). I a similar vein, $\mathbb{S}^1$ may be identified with the unit complex numbers and a Lie group structure induced thereon.

All other spheres, have non-trivial tangent bundle. This is deep result in topology which may be proven using the Poincaré-Hopf index theorem.

A result in functional analysis says that the only *normed division algebras* over $\mathbb{R}$ are the real numbers, the complex numbers (both commutative and associative), the quaternions (assosiative, but not commutative) and the octonions (neither associated nor commutative), which result being closely related to the triviality of the tangent bundle on spheres. It may be used to prove the triviality of the aforementioned spheres, but not the non-triviality of the remaining spheres (at least as far as I know).

Definition 2.22. A *morphism* of smooth vector bundles $\pi_1 : E_1 \to M$, $\pi_2 : E_2 \to M$ is smooth map $F : E_1 \to E_2$ such that

1. $\pi_2 \circ F = \pi_1$,
2. $F_x : (E_1)_x \to (E_2)_x$ is a linear map, where $F_x = F|_{(E_1)_x}$ is the restriction of F to the fibre.

An *isomorphism* of smooth vector bundles is a morphism such that there exists a two-sided inverse morphism.

Note that the first condition says that F maps fibres to fibres. The second condition says that this map is linear. Equivalently in local trivialisations, $\phi_i : \pi_i^{-1}[U_i] \to U_i \times \mathbb{R}^{k_i}$, $i = 1, 2$ the map

$$\pi_2 \circ F \circ \pi_1^{-1} : (F^{-1}[U_2] \cap U_1) \times \mathbb{R}^{k_1} \to (F^{-1}[U_2] \cap U_1) \times \mathbb{R}^{k_2}$$

has the form

$$\pi_2 \circ F \circ \pi_2(x, v) = (x, A(x)v)$$

for $A : (F^{-1}[U_2] \cap U_1) \to \mathrm{Hom}(\mathbb{R}^{k_1}, \mathbb{R}^{k_2})$ a smooth assignment of linear map $A(x) : \mathbb{R}^{k_1} \to \mathbb{R}^{k_2}$ for each x.

In particular, when we say a vector bundle is trivial, we typically mean it is isomorphic with a trivial bundle. For example, the tangent bundle to the circle $\mathbb{S}^1$ is isomorphic to the cylinder (as a bundle over $\mathbb{S}^1$) which is trivial. Similarly, the tangent bundle to $\mathbb{R}^n$, as the disjoint union of the tangent spaces is isomorphic to a trivial bundle.

3 Week 03

3.1 Week 03, Lecture 01: Multi-linear Algebra

Before returning to geometry, we need to develop some multi-linear algebra. This will enable us to define various tensor bundles that play an important role in Riemannian geometry. It will allow us to give a precise, clean definition of a smooth metric as a section of a bundle, define various differential operators acting on sections of bundles and define the curvature tensor, also as a section of a vector bundle. Throughout we will work with *real* vector spaces, and simply call them vector spaces.

Below I will give many details of multi-linear algebra (more so than often encountered in a first course of differential geometry). Many of the details will not be covered in the lectures, but are contained here as a reference. It is important to obtain a good grounding in multi-linear algebra in order to undertake the tensor calculus calculations, through which modern differential geometry is expressed. These calculations allow us to "lift ourselves out of the mud" of local coordinate calculations, and think more globally; to understand *geometry in the large*.

3.1.1 Hom and The Dual Space

Definition 3.1. Let V, W be vector spaces. The set of linear transformations $V \to W$ is denotes $\mathrm{Hom}(V, W)$ (or $\mathrm{Hom}_{\mathbb{R}}(V, W)$ to emphasise that these are real, linear transformations). In particular, when $W = \mathbb{R}$, we write $V^* = \mathrm{Hom}(V, \mathbb{R})$ for the set of real valued, linear functions on V. The space V^* is called the *dual space*.

The set $\mathrm{Hom}(V, W)$ has the natural vector space structure, given by pointwise addition and scalar multiplication in W. Let $c_1, c_2 \in \mathbb{R}$, $\phi_1, \phi_2 \in \mathrm{Hom}(V, W)$ and define $c_1\phi_1 + c_2\phi_2 \in \mathrm{Hom}(V, W)$ by

$$(c_1\phi_1 + c_2\phi_2)(v) = c_1\phi_2(v) + c_2\phi_2(v)$$

for $v \in V$ and where the right hand side operations (addition, scalar multiplication) are taken in W. You can check that the axioms for a vector space are satisfied by this definition.

Lemma 3.2. *If V, W are finite dimensional vector spaces of dimensions m and n respectively, then Hom(V, W) is finite dimensional with dimension $m \cdot n$. In particular dim $V^* =$ dim V.*

Proof. Let $\{e_i\}_{i=1,\cdots,m}$ be a basis for V and $\{f_j\}_{j=1\cdots,n}$ be a basis for W, and define elements $\theta^i_j \in \mathrm{Hom}(V, W)$, $1 \leq i \leq m$, $1 \leq j \leq n$ by

$$\theta^i_j(e_k) = (\delta^i_k) f_j$$

and extended by linearity to all of V: writing $v = v^1 e_1 + \cdots + v^m e_m \in V$, we have

$$\begin{aligned}
\theta^i_j(v^1 e_1 + \cdots + v^m e_m) &= v^1 \theta^i_j(e_1) + \cdots v^m \theta^i_j(e_m) \\
&= (v^1 \delta^i_1 + \cdots + v^m \delta^i_m) f_j \\
&= v^i f_j.
\end{aligned}$$

You can check directly that this is a basis. □

Remark 3.3. For any $\alpha \in \mathrm{Hom}(V, W)$, writing $v = v^1 e_1 + \cdots + v^m e_m \in V$ and for each i writing $\alpha(e_i) = \alpha^1_i f_1 + \cdots + \alpha^n_i f_n \in W$, then we may write

$$\begin{aligned}
\alpha(v) = \alpha(v^1 e_1 + \cdots + v^m e_m) &= v^1 \alpha(e_1) + \cdots + v^m \alpha(e_m) \\
&= v^1(\alpha^1_1 f_1 + \cdots \alpha^n_1 f_n) + \cdots + v^m(\alpha^1_m f_1 + \cdots \alpha^n_m f_n) \\
&= \left(\sum_{i=1}^m v^i \alpha^1_i\right) f_1 + \cdots + \left(\sum_{i=1}^m v^i \alpha^n_i\right) f_n \\
&= \alpha^1(v) f_1 + \cdots + \alpha^n(v) f_n.
\end{aligned}$$

In matrix form,

$$\alpha(v) = \begin{pmatrix} \alpha^1(v) \\ \vdots \\ \alpha^n(v) \end{pmatrix} = \begin{pmatrix} \alpha^1_1 & \cdots & \alpha^1_m \\ \vdots & \ddots & \vdots \\ \alpha^n_1 & \cdots & \alpha^n_m \end{pmatrix} \begin{pmatrix} v^1 \\ \vdots \\ v^m \end{pmatrix}.$$

In particular, the basis elements θ^i_j are matrices with 1 in the j'th row of the i'th column and 0 everywhere else.

In terms of the basis θ^i_j,

$$\alpha = \sum_{1 \leq i \leq m, 1 \leq j \leq n} \alpha^j_i \theta^i_j$$

with the real numbers α^j_i the entries of the matrix above.

Remark 3.4. For the dual space, $n = 1$ and so we only have one component $j = 1$. A basis of V^* is given by

$$\theta_1^i(e_k) = \delta_k^i f_1.$$

Typically we drop the superfluous 1 and simply write $\theta^i(e_k) = \delta_k^i$. The basis $\{\theta^i\}$ is referred to as the *dual basis* to $\{e_i\}$. With respect to this basis, an arbitrary $\alpha \in V^*$ has the unique expression

$$\alpha = \alpha_1\theta^1 + \cdots + \alpha_m\theta^m$$

with $\alpha_i \in \mathbb{R}$, $1 \leq i \leq m$.

As a matrix, if we write vectors $v \in V$ as row vectors, then we write α as the row vector, $\alpha = (\alpha^1, \cdots, \alpha^m)$, giving

$$\alpha(v) = \begin{pmatrix} \alpha^1 & \cdots & \alpha^m \end{pmatrix} \begin{pmatrix} v_1 \\ \vdots \\ v_m \end{pmatrix}$$

In particular, the dual basis is given by $\theta^i = (0, \cdots, 0, 1, 0, \cdots, 0)$ where 1 is in the i'th place.

Also not that given any $v = v^1e_1 + \cdots + v^me_m \in V$, the coefficients v^i are given by

$$v^i = \theta^i(v).$$

Remark 3.5. An element $\alpha = \alpha^1 f_1 + \cdots + \alpha^m f_m \in \mathrm{Hom}(V, W)$, may thought of as an m-tuple of elements, $\alpha^j \in V^*$. That is, the *components* of α are themselves, real valued, linear functions on V (check this!).

3.1.2 Remarks on Indices and the (Einstein) Summation Convention

You may have noticed that some of the indices above are superscripts (*upper indices*) and some are subscripts (*lower indices*). There is a pattern to the choice of which goes where. We write lower indices for basis elements of V and upper indices for the coefficients $\{v^i\} \subseteq \mathbb{R}$ in the expression $v = v^1e_1 + \cdots + v^me_m$. For dual elements, this is reversed, $\alpha = \alpha_1\theta^1 + \cdots + \alpha_m\theta^m$. Try to see how this forces is to make the choices above for upper and lower indices of elements $\alpha \in \mathrm{Hom}(V, W)$.

Note, this is just a convention, and we could easily do it the other way round (basis elements of V with upper indices, and basis elements of V^* as lower indices). However, this particular convention is extremely firmly established that one flaunts it at great peril to one's reputation!

A very useful convention (for which it's important to strictly follow the upper and lower index convention) is the *(Einstein) summation convention.* The rule is, if an index is repeated, once as an upper index, and once as a lower index, it is implicitly summed over. This produces rather compact notation. Without the convention, for example we would write,

$$v = v^1 e_1 + \cdots + v^m e_m = \sum_{i=1}^{m} v^i e_i.$$

In the summation convention, we simply drop the $\sum$ symbol and write

$$v = v^i e_i.$$

The i appears exactly twice, once as upper and once as lower and this implicitly means exactly the same thing as the prior equation. To see how this can be really useful, let $\alpha \in \mathrm{Hom}(V, W)$. Without the convention we would write

$$\alpha(v) = \sum_{j=1}^{n} \sum_{i=1}^{m} v^i \alpha_i^j f_j$$

and now we may write this much more briefly as

$$\alpha(v) = v^i \alpha_i^j f_j.$$

This notation is particularly useful when dealing with higher order tensors which may involve a sum over three or four (or more!) indices, which becomes quite messy to write out in full (i.e. without the summation convention). The curvature tensor, a central object of study in Riemannian geometry requires four indices for example.

Remark 3.6. Remember, the rule is repeated indices, appearing *once above, and once below* are summed over. Repeated indices both above or both below are not to be summed over in this convention. Later we will see a way to convert between upper and lower indices, that is fundamental in Riemannian geometry (*metric contraction*). Some authors use this conversion to some over repeated indices, even if they are both upper or both lower, with the implicit understanding that one should first convert one lower index to an upper index (or vice versa). We will avoid this practice, as even though it leads to clear, compact notation, it can be confusing to the beginner (and to the adept!).

3.1.3 Tensor Products

Definition 3.7. Let S be a set. The *free vector space over S* is the vector space $F(S)$ of all formal, finite, linear combinations of elements of S. A general element $x \in F(S)$ may thus be written

$$x = \sum_{a=1}^{N} \lambda^a s_a$$

for any finite subset $\{s_a \in S\}_{a=1}^N$ with the coefficients, $\lambda^a \in \mathbb{R}$. In other words, $F(S)$ is a vector space with basis S.

Example 3.8. If S itself is a finite set $S = \{s_1, \cdots, s_N\}$, then $F(S) \simeq \mathbb{R}^N$. If S is an infinite set, then $F(S)$ is an infinite dimensional vector space (the dimension is equal to the cardinality of the set S).

Remark 3.9. The free vector space $F(S)$ can be realised concretely as the set of all *co-finitely zero* functions $S \to \mathbb{R}$,

$$F(S) = \{x : S \to \mathbb{R} : |\{s : x(s) \neq 0\}| < \infty\}.$$

That is, co-finitely zero means all but finitely many values of x are 0. If we write $\{s_1, \cdots, s_N\} = \{s : x(s) \neq 0\}$, then the coefficients λ^a of x in the definition above are then just $\lambda^a = x(s_a)$. The vector space structure on $F(s)$ is then realised as the pointwise vector space structure:

$$(c_1 x_1 + c_2 x_2)(s) = c_1 x_1(s) + c_2 x_2(s).$$

Note that the function $c_1 x_1 + c_2 x_2$ is zero for all but finitely many s since this is true of both x_1 and x_2, hence it is indeed an element of $F(s)$.

Another way to think of elements $x \in F(S)$ is as families of real numbers

$$x = \{\lambda^s \in \mathbb{R} : s \in S, |\{\lambda^s \neq 0\}| < \infty\},$$

index by S where all but finitely many of these real numbers are zero.

Since all but finitely many of the coefficients are zero, we may also write,

$$x = \sum_{x \in S} \lambda^s s.$$

There is a natural, injective map (not a linear map since S is just a set),

$$\iota : s \in S \mapsto x \in F(S)$$

where x has coefficients

$$\lambda^t = \begin{cases} 1, & t = s \\ 0, & \text{otherwise} \end{cases}$$

Thus we may regard an element $s \in S$ as an element of $F(S)$ via this map. As a function $S \to \mathbb{R}$, $x(t) = \delta_{ts}$ is zero for all $t \in S$ except for s at which $x(s) = 1$.

Given vector spaces V, W, denote by $F(V, W)$ the free vector space over $S = V \times W$. This is the set of all formal, finite, linear combinations, $\sum_a \lambda^a (v_a, w_a)$ of pairs $(v_a, w_a) \in V \times W$.

Definition 3.10. Let V, W be a vector spaces. The *tensor product* of V and W, written $V \bigotimes W$ is the vector space given by the quotient

$$F(V, W)/R(V, W)$$

where $R(V, W)$ (*the subspace of relations*) is the subspace generated by elements of the form

1. $(\lambda v, w) - \lambda(v, w)$,
2. $(v, \lambda w) - \lambda(v, w)$,
3. $(v_1 + v_2, w) - (v_1, w) - (v_2, w)$,
4. $(v, w_1 + w_2) - (v, w_1) - (v, w_2)$

with $\lambda \in \mathbb{R}$, $v, v_1, v_2 \in V$ and $w, w_1, w_2 \in W$.

Recall that to say $R(V, W)$ is generated by the elements above is to say that $R(V, W)$ is the smallest subspace containing all elements of $F(V, W)$ of the form 1-4. Equivalently, it is the set of all finite linear combinations of elements of the form 1-4. Check this!

Letting $p : F(V, W) \to V \bigotimes W$ denote the quotient map, and $\iota : V \times W \to F(V, W)$ the natural map defined above, we write

$$v \otimes w = p \circ \iota(v, w).$$

That is $v \otimes w$ is the equivalence class represented by (v, w) thought of as element of $F(V, W)$. The assignment $(v, w) \mapsto v \otimes w$ defines a bilinear map $\pi : V \times W \to V \bigotimes W$.

The relations defining $R(V, W)$ may now be expressed as,

1. $(\lambda v) \otimes w = \lambda(v \otimes w)$,
2. $v \otimes (\lambda w) = \lambda(v \otimes w)$,
3. $(v_1 + v_2) \otimes w = v_1 \otimes w + v_2 \otimes w$,
4. $v \otimes (w_1 + w_2) = v \otimes w_1 + v \otimes w_2$.

That is the tensor product $\otimes$ distributes over vector addition and scalar multiplication.

Remark 3.11. The map $\pi : V \times W \to V \bigotimes W$ is neither injective not surjective. To see it is not injective, just note that for any $v \in V$, $v \neq 0$ we have $(v, 0) \sim 0 \cdot (v, 0) \sim (0 \cdot v, 0) = (0, 0)$ so that $\pi(v, 0) = \pi(0, 0)$. To why π is not surjective, let $v_1, v_2 \in V$ be linearly independent and $w_1, w_2 \in W$ also be linearly independent. Then $v_1 \otimes w_1 + v_2 \otimes w_2$ is not equivalent to $v \otimes w$ for any $(v, w) \in V \otimes W$. As an exercise, try to prove this claim.

Nevertheless, elements of the form $v \otimes w$ play an important role. Such elements are the so-called *indecomposable* elements. In particular, if $\{e_i\}_{1 \leq i \leq m}$ is a basis for V and $\{f_j\}_{1 \leq j \leq n}$ is a basis for W, the indecomposable elements $e_i \otimes f_j$ form a basis for $V \bigotimes W$.

Lemma 3.12. *Let V be a m-dimensional vector space, and W an n-dimensional vector space. Then dim $V \bigotimes W = mn$.*

Proof. Let $\{e_i\}_{1 \leq i \leq m}$ be a basis for V and $\{f_j\}_{1 \leq j \leq n}$ be a basis for W. Then to see that a basis for $V \otimes W$ is $\{e_i \otimes f_j\}$, choose an arbitrary element,

$$\sum_{a=1}^{N} \lambda^a v_a \otimes w_a \in V \bigotimes W$$

and write each $v_a = v_a^i e_i$ and $w_a = w_a^j f_j$ in terms of the bases. The distribution law for the tensor product implies that

$$\sum_{a=1}^{N} \lambda^a v_a \otimes w_a = \sum_{i,j} \left(\sum_a \lambda^a v_a^i w_a^j \right) e_i \otimes f_j.$$

□

Therefore, an arbitrary element $T \in V \otimes W$ may be written

$$T = \sum_{i=1}^{m} \sum_{j=1}^{n} T^{ij} e_i \otimes f_j,$$

or more compactly using the summation convention,

$$T = T^{ij} e_i \otimes f_j.$$

The most important property of the tensor product is the following *universal property.* The universal property characterises the tensor product uniquely and in practice is the way we work with the tensor product rather than directly by the definition.

Theorem 3.13. *Let V, W, Z be vector spaces. Then the pair $(V \otimes W, \pi : V \times W \to V \otimes W)$ is the unique pair (up to isomorphism) such that given any bi-linear map,*

$$\phi : V \times W \to Z$$

there is a unique, linear map

$$\psi : V \bigotimes W \to Z$$

such that $\psi \circ \pi = \phi$.

The proof here is a standard proof in algebra and the ideas are quite general, useful in many different situations. **If you are not algebraically inclined, feel free to skip the proof.** The important thing, as far as this course goes, is that you are aware of the universal property. For the algebraically inclined, here's the sketch of the proof (some of the details such as verifying that certain maps are well defined are left as an exercise).

Proof. The proof consists in showing that

1. $V \bigotimes W$ satisfies the universal property,

2. If any other pair of a vector space U and a map $p : V \times W \to U$ satisfies the universal property, then there is an isomorphism $T : U \to V \bigotimes W$ such that $p = T \circ \pi$.

Let us begin with the first item.

1. Given ϕ, define ψ on indecomposable elements $v \otimes w$ by

$$\psi(v \otimes w) = \phi(v, w).$$

   Check that this is well defined, independent of the indecomposable element representing the equivalence of $v \otimes w$. For example, given any $\lambda \neq 0$, $v \otimes w \sim (\lambda v) \otimes (\lambda^{-1} w)$.

   Next, a general element of $x \in V \bigotimes W$ may be written (non-uniquely!) as a finite sum of indecomposable elements,

$$\sum_i v_i \otimes w_i.$$

   We extend ψ to all of $V \bigotimes W$ by linearity:

$$\psi(x) = \sum_i \psi(v_i \otimes w_i).$$

   Again, we need to check this is well defined, independent of the equivalence class, i.e. x can be written as a finite linear combination of indecomposable elements in more than one way.

   That $\psi \circ \pi = \phi$ follows immediately, since the image of π is the precisely the set of indecomposable elements, hence

$$\psi \circ \pi(v, w) = \psi(v \otimes w) = \phi(v, w)$$

   with the first equality the definition of π and the second equality the definition of ψ.

   Lastly, to see that ψ is the unique map such that $\psi \circ \pi = \phi$, just note that if there is any other such map η, then on indecomposable elements

$$\eta(v \otimes w) = \phi(v, w) = \psi(v \otimes w)$$

   so that $\eta = \psi$ on indecomposable elements. By linearity of both η and ψ, they must agree on all elements.

2. To show uniqueness of the tensor product, suppose that there exists a vector space U and a map $p : V \times W \to U$ satisfying the universal property. Then apply this universal property to the map $\phi : V \times W \to Z$ with $Z = V \bigotimes W$ and $\phi = \pi$. This gives a unique, linear map $\psi : U \to V \bigotimes W$. Now reverse the roles of U and $V \bigotimes W$ to obtain a unique linear map $V \bigotimes W \to U$. Using uniqueness, show that these maps are inverse to each other.

□

The non-algebraically inclined should resume reading here again.

The essence of the theorem says that to define a linear map $\phi : V \bigotimes W \to Z$, we need to define it on indecomposable elements $v \otimes w$. Provided this map satisfies $\phi((\lambda v) \otimes w) = \phi(v \otimes (\lambda w)) = \lambda\phi(v \otimes w)$, then it is well defined by the universal property.

The universal property may be used to prove the following very useful facts.

Lemma 3.14. *There is a natural isomorphism,*

$$V \bigotimes W \to W \bigotimes V$$

which takes an indecomposable element $v \otimes w$ to $w \otimes v$. This is called reordering the factors.

Lemma 3.15. *The tensor product is associate in the sense that given vector spaces V, W, Z, there is a natural isomorphism*

$$(V \bigotimes W) \bigotimes Z \to V \bigotimes (W \bigotimes Z),$$

taking $(v \otimes w) \otimes z$ to $v \otimes (w \otimes z)$.

Hence we may unambiguously (up to isomorphisms) write $V \bigotimes W \bigotimes Z$. We can of course, repeat this as many times as desired: $V_1 \bigotimes \cdots \bigotimes V_n$ is well defined up to isomorphism regardless of the order in which we take the tensor products.

The following two lemmas will be used later when working with tensor fields.

Lemma 3.16. *There is a natural isomorphism,*

$$V^* \otimes W \to Hom\ (V, W)$$

taking $\theta \otimes w$ to $u \in V \mapsto \theta(u)w$.

In particular $End\ (V) = Hom\ (V, V) \simeq V^* \bigotimes V$.

Lemma 3.17. *There is a natural isomorphism,*

$$V^* \otimes W^* \to B(V, W)$$

where $B(V, W)$ denotes the set of bilinear maps $V \times W \to \mathbb{R}$, taking $\alpha \otimes \beta$ to $(v, w) \in V \times W \mapsto \alpha(v)\beta(w)$.

In particular $B\ (V) =\ B(V, V) \simeq V^* \bigotimes V^*$.

These first lemma provides an alternative method of proving finite dimensionality of Hom from finite dimensionality of the tensor product and vice-versa. Analogous results apply to for bilinear forms by the second lemma.

In terms of bases $\{\theta^i\}$ for V^* and $\{f_j\}$ for W, a basis for $V^* \bigotimes W$ is $\{\theta^i \otimes f_j\}$. Then the isomorphism $V^* \otimes W \to \text{Hom}\ (V, W)$ satisfies

$$\theta^i \otimes f_j \mapsto \theta^i_j$$

where $\{\theta^i_j\}$ is the basis of Hom (V, W) we defined earlier.

If $\{\theta^i\}$ is a basis for V^* and $\{\phi^j\}$ is a basis for W^*, what is the image of the basis $\{\theta^i \otimes \phi^j\}$ under the isomorphism $V^* \otimes W^* \to \mathrm{B}(V, W)$?

The first of these two lemmas allows us to define a very important operation: contraction.

Definition 3.18. The contraction is the linear map

$$V^* \bigotimes V \to \mathbb{R}$$

defined on indecomposable elements by

$$\theta \otimes v \mapsto \theta(v).$$

Via the isomorphism, $V^* \otimes V \to \text{End}\ (V)$, a linear map $T : V \to V$ may be considered as an element $T \in V^* \otimes V$. This isomorphism allows us to define the trace of a linear operator as the contraction.

Lemma 3.19. *Let $T : V \to V$ be a linear map, let $\{e_i\}$ be a basis for V and let T^i_j be the matrix representing T with respect to this basis. Then as an element of $V^* \otimes V$,*

$$T = T^i_j \theta^i \otimes e_j$$

and the contraction of T is equal to the trace of the matrix T^i_j.

Notice that the contraction of $T \in V^* \otimes V$ is well defined without reference to any basis and hence so too is the contraction of a linear map $T : V \to V$. The lemma says that if we write T in terms of a basis, then the trace of this matrix is equal to the contraction of T. Therefore we recover immediately, the well known result of linear algebra that the trace of a matrix A is invariant under similarity transformation, $A \mapsto P^{-1}AP$ for P a change of basis.

Remark 3.20. The lemma is an example of a more general phenomena in linear algebra. Using the canonical isomorphisms, between tensor products and Hom , we may give *basis independent* definitions of all notions in linear algebra that are usually defined in terms of matrices and then shown to be independent of the basis chosen. That is, with the right definitions in place, we need never check independence of the chosen basis used when computing explicitly with respect to a basis. **This is an advance that should not be underestimated!**

3.2 Week 03, Lecture 02: Tensor Bundles

3.2.1 The Tensor Algebra

Definition 3.21. A tensor T, of degree p is an element

$$T \in T^p(V) = V^p = \bigotimes^p V = V \bigotimes \cdots \bigotimes V.$$

A tensor of degree 0 is defined to be an element of $\mathbb{R}$.

A tensor T, of *contravariant degree* p and *covariant degree* q is an element

$$T \in T^p_q(V) = V^p \bigotimes (V^*)^q,$$

also referred to as a tensor of type (p, q).

There is a multi-liner map (i.e. linear in each slot when holding all other arguments fixed)

$$\pi^p : \prod^p V \to \bigotimes^p V = T^p(V)$$

mapping

$$(v_1, \cdots, v_p) \mapsto v_1 \otimes \cdots \otimes v_p.$$

Recall that the tensor product is associative up to isomorphism and so $\bigotimes^p V$ is well defined up to isomorphism.

Lemma 3.22. *The pair, $(T^p(V), \pi^p)$ is characterised uniquely by the following universal property: given any multi-linear map $\phi : \prod^p V \to Z$, there is a unique linear map $\psi : T^p(V) \to Z$ such that $\phi = \psi \circ \pi^n$.*

The proof is left as an exercise. Alternatively, one may define the p-fold tensor product $\bigotimes^p V$ by a construction similar to the construction for the tensor product and then prove the universal property. Note that *any $p^n : \prod^n V \to U$ satisfying the universal property is isomorphic to $\pi^p : \prod^p V \to T^p(V)$*.

For reference, though it won't play a major role, here's the definition of the tensor algebra.

Definition 3.23. The *tensor algebra* is the vector space,

$$T(V) = T^\infty(V) = \bigoplus_{p \geq 0} T^p V = \mathbb{R} \bigoplus V \bigoplus (V \bigotimes V) \bigoplus \cdots .$$

Elements $T \in T(V)$ are finite, linear combinations,

$$T = \sum_{p \geq 0} T^p$$

with each $T^p \in T^p(V)$ and all but finitely many $T^p = 0$.

The *bi-graded tensor algebra* is the vector space,

$$T_\infty^\infty(V) = \bigoplus_{p,q \geq 0} T_q^p V = \mathbb{R} \bigoplus V \bigoplus V^* \bigoplus (V \bigotimes V^*) \bigoplus \cdots .$$

Elements $T \in T_\infty^\infty(V)$ are finite, linear combinations,

$$T = \sum_{p,q \geq 0} T_q^p$$

with each $T_q^p \in T_q^p(V)$ and all but finitely many $T_q^p = 0$.

The vector space structure is the direct-sum of vector spaces. The algebra structure (i.e. the vector space structure along with a multiplication) comes from the tensor product $\otimes$: given a tensor $T \in T^p$ of degree p and a tensor $S \in T^q$ of degree q, the tensor product,

$$T \otimes S \in T^p \bigotimes T^q \simeq T^{p+q}$$

is a tensor of degree $p + q$.

Given a tensor $T \in T_q^p$ of type (p, q) and a tensor $S \in T_s^r$ of type (r, s) the tensor product

$$T \otimes S \in T_q^p \bigotimes T_s^r \simeq T_{q+s}^{p+r}$$

is a tensor of type $(p + r, q + s)$. Note here that

$$T_q^p \bigotimes T_s^r = V^p \bigotimes (V^*)^q \bigotimes V^r \bigotimes (V^*)^s$$

while

$$T_{q+s}^{p+r} = V^{p+r} \bigotimes (V^*)^{q+s}.$$

The isomorphism between the two is given by the canonical isomorphism reordering the factors. For example,

$$T_1^0 \bigotimes T_1^1 = V^* \otimes V \otimes V^* \simeq V \otimes V^* \otimes V^*.$$

Remark 3.24. Often the notation $T^*(V)$ and $T_*^*(V)$ is used to denote the tensor algebras defined above, but we are already using the symbol $*$ to denote the dual and so I have avoided this notation and opted for $T^\infty(V)$ and $T_\infty^\infty(V)$ instead. Other authors get around this issue by using # to denote the dual space.

3.2.2 New Bundles from Old

Throughout this section, let $\pi : E \to M$ and $\rho : F \to M$ denote vector bundles over M. Our aim is to construct various new vector bundles over M from existing bundles. In particular, we will construct the dual bundle $\pi^* : E^* \to M$, the Hom bundle Hom (E, F) and the various tensor bundles, $\pi^p_q : T^p_q(E) \to M$.

There is a very general method of constructing such bundles that we will employ, known as *gluing*. To motivate the construction recall that the vector bundle E is defined by specifying local trivialisations,

$$\phi_\alpha : \pi^{-1}[U_\alpha] \to U_\alpha \times \mathbb{R}^k$$

such that the transition maps

$$\tau_{\alpha\beta} = \phi_\beta \circ \phi_\alpha^{-1} : U_\alpha \cap U_\beta \times \mathbb{R}^k \to U_\alpha \cap U_\beta \times \mathbb{R}^k$$

are diffeomorphisms of the form

$$\tau_{\alpha\beta}(x, v) = (x, A_{\alpha\beta}(x) \cdot v)$$

for $A_{\alpha\beta} : U_\alpha \cap U_\beta \to GL_k(\mathbb{R})$ a smooth map.

The maps $A_{\alpha\beta}$ satisfy the following conditions:

1. $A_{\alpha\alpha} = \mathrm{Id}$ (since $\tau_{\alpha\alpha} = \phi_\alpha \circ \phi_\alpha^{-1} = \mathrm{Id}$)
2. $A_{\alpha\beta}A_{\beta\alpha} = \mathrm{Id}$ (since $\tau_{\alpha\beta} \circ \tau_{\beta\alpha} = \phi_\beta \circ \phi_\alpha^{-1} \circ \phi_\alpha \circ \phi_\beta^{-1} = \mathrm{Id}$)
3. $A_{\alpha\beta}A_{\beta\gamma}A_{\gamma\alpha} = \mathrm{Id}$ (since $\tau_{\alpha\beta} \circ \tau_{\beta\gamma}\tau_{\gamma\alpha} = \phi_\beta \circ \phi_\alpha^{-1} \circ \phi_\gamma \circ \phi_\alpha^{-1}\phi_\alpha \circ \phi_\gamma^{-1} = \mathrm{Id}$)

The last condition is known as the *co-cycle condition*. You might like to try to see why it does not follow from the second condition. The second condition is however redundant since it follows from 1 and 3 by taking $\gamma = \beta$.

Lemma 3.25 (Vector Bundle Gluing Lemma)**.** *Let U_α be an open cover of M and $A_{\alpha\beta} : U_\alpha \cap U_\beta \to GL_k(\mathbb{R})$ a collection of smooth maps satisfying conditions 1-3 above. Then there exists a unique (up to isomorphism) vector bundle $\pi : E \to M$ of rank k, locally trivial over U_α and with transition maps $\tau_{\alpha\beta}(x, v) = (x, A_{\alpha\beta}(x) \cdot v)$.*

The proof is a little technical and only a sketch is given. All the relevant constructions are there and if you are so inclined, with a little persistence you should be able to fill in the details. **On a first reading, and indeed for the purposes of this course, only a very rudimentary knowledge of this proof is required**.

(Sketch of proof). As a topological space,

$$E = \sqcup_\alpha U_\alpha \times \mathbb{R}^k / \sim$$

where $(x, v) \in U_\alpha \times \mathbb{R}^k \sim (y, u) \in U_\beta \times \mathbb{R}^k$ if

$$x = y \in U_\alpha \cap U_\beta \text{ and } u = A_{\alpha\beta}(x) \cdot v.$$

Each $U_\alpha \times \mathbb{R}^k$ is a topological space, and the disjoint union carries the disjoint union topology for which a set $U \subseteq \sqcup_\alpha U_\alpha \times \mathbb{R}^k$ is open set if and only if $\iota_\alpha^{-1}(U)$ is open for each α where $\iota_\alpha : U_\alpha \times \mathbb{R}^k \to \sqcup_\beta U_\beta \times \mathbb{R}^k$ is the natural inclusion. The topology on E is then the quotient topology.

Notice that $\sim$ is an equivalence relation:

- reflexive: from condition 1,
- symmetric: from condition 2,
- transitive: from conditions 2 and 3

and so the quotient E is a well defined topological space.

Denote by $[x, v]$ the equivalence class of (x, v). The projection $\pi : E \to M$ is then just $\pi([x, v]) = x$ which is well defined independent of the chosen representative by the definition of the equivalence relation $\sim$.

Let $V_\alpha = \pi^{-1}[U_\alpha] = \{[x, v] : x \in U_\alpha\}$ and define

$$\begin{aligned} \phi_\alpha : \pi^{-1}[U_\alpha] &\to U_\alpha \times \mathbb{R}^k \\ z &\mapsto (x, v) \end{aligned}$$

where $(x, v) \in U_\alpha \times \mathbb{R}^k \subseteq \sqcup_\beta U_\beta \times \mathbb{R}^k$ is in the unique representative of the equivalence class z with $x \in U_\alpha$ (if also $(x, u) \sim (x, v)$ with $x \in U_\alpha$, then $u = A_{\alpha\alpha}(x) \cdot v = v$). This defines a local trivialisation on E and may also be used to define a smooth structure since it exhibits a homeomorphism of the open set $\pi^{-1}[U_\alpha]$ with the open subset $U_\alpha \times \mathbb{R}^k \subseteq M \times \mathbb{R}^k$.

One may now check directly the remaining conditions, such as that π is smooth and $\pi = p \circ \phi_\alpha$ for $p : U_\alpha \times \mathbb{R}^k \to U_\alpha$ the projection. This latter condition also automatically implies π is a surjective submersion.

To see that the transition functions are $A_{\alpha\beta}$, note that for $z \in \pi^{-1}[U_\alpha \cap U_\beta]$,

$$\phi_\alpha(z) = (x, v)$$

and

$$\phi_\beta(z) = (y, u)$$

with $[x, v] = z = [y, u]$ so that $(x, v) \sim (y, u)$, which by definition says that

$$x = y, \quad u = A_{\alpha\beta}(x) \cdot v.$$

Therefore,

$$\tau_{\alpha\beta}(x, v) = \phi_\beta \circ \phi_\alpha^{-1}(x, v) = \phi_\beta(z) = (y, u) = (x, A_{\alpha\beta}(x) \cdot v).$$

□

Remark 3.26. Given a vector bundle, one obtains transition maps $A_{\alpha\beta} : U_\alpha \cap U_\beta \to GL_n(\mathbb{R})$ satisfying the conditions 1-3 above. The lemma says the converse is true, given transition maps satisfying 1-3, one obtains a unique bundle with the given transition maps. In other words, to specify a vector bundle it is equivalent to specify the local trivialisations and the transition maps between them. Often you will encounter statements like "the dual bundle E^* is the bundle whose fibres are the duals E_x^* of the fibres E_x. Strictly speaking this **does not define a vector bundle**. Why? Because the transition maps have not been given. Compare for instance the fact that any two bundles E, F, of rank k has isomorphic fibres $E_x \simeq F_x$ since the fibres are isomorphic to $mathbbR^k$. To give the fibres is clearly not enough (e.g. every bundle has fibres isomorphic to the fibres of a trivial bundle!), one must also specify how the fibres *vary locally* via the local trivialisations and how these local descriptions glue together via the transition maps.

It would be more correct to say that E^* is the vector bundle that is locally the dual to E. We make this precise using the gluing lemma as follows.

Definition 3.27. Let $\pi : E \to M$ be a vector bundle with local trivialisations $\phi_\alpha : \pi^{-1}[U_\alpha] \to U_\alpha \times \mathbb{R}^k$ and transition maps $\tau_{\alpha\beta} = (\mathrm{Id}_M, A_{\alpha\beta})$. The dual bundle is the unique bundle E^* with local trivialisations,

$$A^*_{\alpha\beta} = (A^{-1}_{\alpha\beta})^T$$

To check this definition makes sense, one needs to check conditions 1-3 hold (as remarked above we don't need to check 2).

1. $A^*_{\alpha\beta}A^*_{\beta\alpha} = (A^{-1}_{\alpha\beta})^T \cdot (A^{-1}_{\beta\alpha})^T = \left[(A_{\alpha\beta}A_{\beta\alpha})^{-1}\right]^T = \mathrm{Id}$.
2. Exercise!

Remark 3.28. The vector bundle gluing lemma tells us this bundle is well defined, but why do we call it the dual? That is, how do the transition maps $(A^{-1}_{\alpha\beta})^T$ relate to the dual? To see how, we can think of local trivialisations giving a local basis $\{e_i(x) = (x, (0, \cdots, 0, 1, 0, \cdots, 0,)\}$ for each fibre E_x, and changing trivialisation gives a new local basis $\{f_i(x) = (x, A_{\alpha\beta(x)}e_i\}$. Let $\{\theta_i(x)\}$ and $\{\phi_i(x)\}$ be the respective dual bases. Then as an exercise, show that

$$\phi_i = (A^{-1}_{\alpha\beta})^T\theta_i$$

and hence that if we write $\alpha = \alpha_i\theta^i = \beta_i\phi^i$, that

$$\beta_i = \sum_{j=1}^{n} B^j_i\alpha^i$$

where B^j_i is the change of basis matrix $(A^{-1}_{\alpha\beta})^T$.

The most important dual bundle for us is the co-tangent bundle.

Definition 3.29. Let M be a smooth manifold. The *co-tangent bundle*, denoted T^*M is the dual of the bundle TM. Elements of T^*M are called *co-vectors*, or *1-forms* (more on forms later).

Definition 3.30. Let $\pi : E \to M$, $\rho : F \to M$ be vector bundles and with local trivialisations $\phi_\alpha : \pi^{-1}[U_\alpha] \to U_\alpha \times \mathbb{R}^k$ and $psi_\mu : \pi^{-1}[V_\mu] \to V_\mu \times \mathbb{R}^\ell$ respectively. Also let $\tau_{\alpha\beta} = (\mathrm{Id}_M, A_{\alpha\beta})$ and $\sigma_{\mu\nu} = (\mathrm{Id}_M, B_{\mu\nu})$. The bundle $\mathrm{Hom}(E, F)$ is the unique bundle with local trivialisations given by the transition matrices $C_{(\alpha,\mu)(\beta,\nu)}$ determine by

$$f_i \otimes \phi^j = C_{(\alpha,\mu)(\beta,\nu)}e_i \otimes \theta^j$$

over the open sets

$$U_{\alpha,\mu} = U_\alpha \cap V_\mu$$

Remark 3.31. Note that here E and F may not have local trivialisations over the same open cover, but both bundles are locally trivial over the doubly indexed sets $U_{\alpha,\mu} = U_\alpha \cap V_\mu$. The transition maps $C_{(\alpha,\mu)(\beta,\nu)}$ are then defined on the overlaps $(U_\alpha \cap V_\mu) \cap (U_\alpha \cap V_\mu)$ and take values in $GL_{k\ell}(\mathbb{R})$.

Also observe that it is easier to define the transition maps on $E^* \otimes F$ rather than directly on $\mathrm{Hom}(E, F)$. Explicitly,

$$C_{(\alpha,\mu)(\beta,\nu)} = (A_{\alpha\beta}^{-1})T \otimes B_{\mu\nu}.$$

You might like to check that under the identification $\theta^i_j = \theta^i \otimes f_j$, the transition maps obtained directly on $\mathrm{Hom}(V, W)$ are exactly the same transition maps obtained here.

Check this definition is well defined as with the dual bundle, and that these really are the transition maps for Hom.

Very important to us are the various tensor bundles. The remainder of this course will deal with defining and studying various sections of these bundles (i.e. with *tensor fields*).

Definition 3.32. The tensor bundles, $T^p_q(M)$ on M are the tensor bundles with transition maps

$$\bigotimes^p A_{\alpha\beta} \otimes \bigotimes^q (A_{\alpha\beta}^{-1})^T$$

where $A_{\alpha\beta}$ are the transition maps of TM.

Note that here, as opposed to the case of $\mathrm{Hom}(E, F)$, we don't have to worry about intersecting different local trivialisations since the tensor products and duals are formed from the single bundle TM. Therefore, all these bundles are locally trivial over any open set on which TM is locally trivial.

Example 3.33. Given a section $T \in \Gamma(\mathrm{End}(TM)) = \Gamma(T^*M \otimes TM)$ and a section $X \in \mathfrak{X}(M) = \Gamma(TM$, we may define a new section,

$$x \in M \mapsto Y_x = T_x(X_x) \in T_xM.$$

We will generally simply write $Y = T(X)$ for this vector field. Thinking of T as a section of $T^*M \otimes TM$ we can also think of Y being given by first forming the section $X \otimes T$ and then contracting the T^*M part of T with X.

This is a very important operation to understand. The way to think of this is as follows. Let us give the vector space construction and leave it as an exercise to see that this also works at the level of vector bundles. The contraction in question is the map

$$\begin{aligned} V \otimes V^* \otimes V &\to V \\ v \otimes \theta \otimes w &\mapsto \theta(v)w \end{aligned}$$

on indecomposable elements and extended to all of $V \otimes V^* \otimes V$ by using the universal property for tensors. So this map contracts the first two entries together to get a real number and then multiplies the last entry by this real number. To define this map on sections, just perform the map pointwise in each fibre.

Remark 3.34. In the exercises, you will be asked to prove the so-called *test for tensorality*. There you will see that extending the map from vector spaces to vector bundles by pointwise maps in the fibres shows that Y is a well defined section of the tangent bundle.

Remark 3.35. Given two linear maps $T : V \to W$ and $S : U \to Z$, we may define a new linear map

$$\begin{aligned} T \otimes S : V \otimes U &\to W \otimes Z \\ v \otimes u &\mapsto T(v) \otimes S(u) \end{aligned}$$

The example is then the map

$$\begin{aligned} \mathrm{Tr} \otimes \mathrm{Id} : (V \otimes V^*) \otimes V &\to \mathbb{R} \otimes V \simeq V \\ (v \otimes \theta) \otimes w &\mapsto \mathrm{Tr}(v \otimes \theta) \otimes \mathrm{Id}(w) \end{aligned}$$

More generally we can contract any pair of entries in $T^{p+1}_{q+1}M$ $p, q \geq 0$ to obtain maps

$$\mathrm{Tr}^i_j : T^{p+1}_{q+1}M \to T^p_q$$

where the i'th upper index is contracted with the j'th lower index. For example

$$\mathrm{Tr}^2_1(u \otimes v \otimes w \otimes \theta \otimes \phi) = \theta(v)u \otimes w \otimes \phi.$$

3.3 Week 03, Lecture 03: Bundle Metrics

3.3.1 Tensors and Change of Basis

We first covered the material from the end of the previous lecture.

3.3.2 Inner Product on Vector Spaces

Definition 3.36. An inner product on a vector space V is a positive definite, symmetric bilinear form. That is, an inner product is a map

$$g(\cdot,\cdot) : V \times V \to \mathbb{R}$$

that is linear in each slot,

$$g(c_1v_1 + c_2v_2, w) = c_1g(v_1, w) + c_2g(v_2, w)$$

and similar for the second slot. This defines a bilinear form. The form must also satisfy,

- $g(v_1, v_2) = g(v_2, v_1)$ (symmetry),
- $g(v, v) \geq 0$ with $g(v, v) = 0 \Rightarrow v = 0$ (positive definite).

Remark 3.37. Note that for V finite dimensional, since $\dim V^* = \dim V$, the two spaces are isomorphic. In fact there are infinitely many (what is the cardinality?) isomorphisms between the two spaces corresponding to each choice of basis on each space. But none are canonical. However, for each inner-product there exists an isomorphism,

$$v \in V \mapsto [\alpha_v : u \mapsto g(v, u)] \in V^*.$$

One can easily check directly that this map is injective, hence in the finite dimensional case it must be an isomorphism. In the infinite dimensional case, how does one check surjectivity?

This metric isomorphism is very important in differential geometry. It is sometimes referred to as a *musical isomorphism*, or the *raising and lowering indices*. The terminology comes from writing this isomorphism in terms of a basis $\{e_i\}$ for V with corresponding dual basis θ^i for V^*. Then $v = v^ie_i$ is mapped to α_v which acts as

$$\alpha_v(u^ie_i) = g(v^ie_i, u^je_j) = v^iu^jg(e_i, e_j).$$

If we write $g_{ij} = g(e_i, e_j)$, then

$$\alpha_v(u) = g_{ij}v^i u^j = g_{ij}v^i\theta^j(u)$$

and this is commonly written as z[$\alpha_{\mathrm{v}} = \mathrm{v_j}\ \theta^{\mathrm{j}}$] with $v_j = g_{ij}v^i$. We've used the metric g to *lower* the index of v. The musical part comes from writing this as

$$v^\flat = \alpha_v.$$

Analogously, given $\alpha = \alpha_i\theta^i \in V^*$, we define a vector,

$$V^\sharp = V_\alpha = V^j e_j = g^{ij}\alpha_i e_j$$

by *raising* the index. Here g^{ij} denotes the inverse matrix to g_{ij} (remember g is positive definite), $g^{ik}g_{kj} = \delta^i_j$. In the exercises, you will see precisely why we use the inverse matrix, though it should a reasonable thing to try given that the assignment $\alpha \mapsto \alpha^\sharp$ should be the inverse of $V \mapsto V^\flat$.

The inner-product can be extended to tensors of type \((p,q)).

Definition 3.38. Let g be an inner product on V. Define the dual metric g^* by

$$g^*(\alpha, \beta) = g(\alpha^\sharp, \beta^\sharp.$$

This defines a metric (check it!) because $\sharp$ is an isomorphism. On tensors of type (p, q) use the universal property to define for indecomposable tensors

$$g^p_q(v_1\otimes\cdots\otimes v_p\otimes\theta^1\otimes\cdots\otimes\theta^q, w_1\otimes\cdots\otimes w_p\otimes\phi^1\otimes\cdots\otimes\phi^q) = g(v_1, w_1)\cdots g(v_p, w_p)g^*(\theta^1, \phi^1)\cdots g^*(\theta^q, \phi^q)$$

where the right hand side is the product of real numbers. In other words,

$$g^p_q = \otimes^p g \otimes \otimes^q g^*.$$

Typically, the letter g is used to denote all these inner-products which can admittedly be a little confusing, but is at least unambiguously defined.

3.3.3 Inner Product on Vector Bundles

Remark 3.39. From the canonical isomorphism, $B(V) \to V^* \otimes V^*$ we may equivalently think of an inner product as an element of $V^* \otimes V^*$ that is symmetric and positive definite.

Definition 3.40. A smooth Riemannian bundle is vector bundle E equipped with a smooth section g of the bundle $B(E) \simeq V^* \otimes V^*$, such that $g_x \in V_x^* \otimes V_x$ is symmetric and positive definite (i.e. it is an inner-product on E_x) for each $x \in M$. The section g is called a *smooth metric* on E.

A smooth Riemannian manifold is a manifold M, such that the tangent bundle is equipped with a smooth metric.

Definition 3.41. Two Riemannian bundles E, F are *isometric* if there exists an *isometry* between them, namely a bundle isomorphism $\phi : E \to F$ such that

$$g_x^F(\phi(u), \phi(v)) = g_x^E(u, v)$$

for every $u, v \in E_x$. Equivalently, for any local sections $u, v \in \Gamma(U, E)$, defining local sections $\phi(u), \phi(u) in \Gamma(U, F)$ by $\phi(u)(x) = \phi_x(u_x)$ (and like wise for $\phi(v)$), the smooth functions

$$x \mapsto g_x^E(u, v), \quad \text{and} \quad x \mapsto g_x^F(\phi(u)(x), \phi(v)(x))$$

agree for all $x \in U$.

Example 3.42. Any regular surface with local parametrisation $\phi : U \to \mathbb{R}^3$ is a Riemannian manifold with

$$g(u, v) = \langle d\phi \cdot u, d\phi \cdot v \rangle_{\mathbb{R}^3}.$$

This definition is smooth in any local parametrisation (since ϕ is smooth) and is independent of the choice of local parametrisation by identifying $u \in TU$ with $d(\psi \circ \phi) \cdot u \in TV$ where $\psi : V \to \mathbb{R}^3$ is another local parametrisation.

4 Week 04

4.1 Week 04, Lecture 01: Integration and the Anti-Symmetric Algebra

4.1.1 Area of a Regular Surface

Let $\phi : U \to \mathbb{R}^3$ be a local parametrisation of a regular surface, S. The tangent plane is spanned by $\{\partial_{x^1}\phi, \partial_{x^2}\phi\}$ for (x^1, x^2) coordinates on $\mathbb{R}^2$. A small rectangle $R \subseteq U$ is mapped to a parallelogram with area,

$$A = |\partial_{x^1}\phi \times \partial_{x^2}\phi|^2 \text{Area}(R)$$

and so the infinitesimal area element,

$$dA(x^1, x^2) = |\partial_{x^1}\phi(x^1, x^2) \times \partial_{x^2}\phi(x^1, x^2)| dx^1 dx^2.$$

The area of $\phi(U) \subset \mathbb{R}^3$ is then

$$\int_U dA(x^1, x^2) = \int_U |\partial_{x^1}\phi(x^1, x^2) \times \partial_{x^2}\phi(x^1, x^2)| dx^1 dx^2.$$

To compute the entire area of S, let $\{\phi_\alpha : U_\alpha \to S\}$ be a covering of S by local parametrisations and ρ_α a partition of unit subordinate to this cover. Then we define

$$A(S) = \sum_\alpha \rho_\alpha \int_{U_\alpha} dA_\alpha.$$

4.1.2 Volume on a Riemannian Manifolds

Next we want to do the same thing on a Riemannian manifold, but we don't have local parametrisations into an ambient space in which to define the infinitesimal volume element, referred to as the *Riemannian volume form.* To proceed, observe that on a surface, we defined the metric, $g_{ij} = \langle \partial_{x^1}\phi, \partial_{x^2}\phi \rangle$ and one can check that

$$g = \lambda^T \lambda$$

where λ is the 3×2 matrix $\lambda_{ai} = \partial_{x^i}\phi^a$, $1 \leq i \leq 2$ and $1 \leq a \leq 3$. Therefore (again write it out!),

$$\det g = \det(\lambda^T \lambda) = |\partial_{x^1}\phi(x^1, x^2) \times \partial_{x^2}\phi(x^1, x^2)|^2$$

so that on a regular surface, we may write

$$A(U) = \int_U \sqrt{\det g} dx^1 dx^2.$$

More generally, if $\phi : U \subseteq \mathbb{R}^n \to \mathbb{R}^{n+k}$ is an immersion (this ensures $g_{ij} = \langle d\phi \cdot \partial_i, d\phi \cdot \partial_j \rangle$ is positive-definite), then a n-cube, C is mapped to a n-parallelotope which volume,

$$V = \sqrt{\det g_{ij}} \mathrm{Vol}(C).$$

Therefore, once again the infinitesimal element of volume is $\sqrt{\det g_{ij}}$.

That is, we have expressed the volume element entirely in terms of the metric and we may now generalise this to a Riemannian manifold.

Definition 4.1. Let (M, g) be a Riemannian manifold. The *Riemannian volume form* is

$$\mu_g = \sqrt{\det g}.$$

That is, let U_α be a chart on M and define

$$\mathrm{Vol}(U_\alpha) = \int_{U_\alpha} \mu_g = \int_{U_\alpha} \sqrt{\det g_{ij}^\alpha} dx^1 \cdots dx^n.$$

Let $U \subset M$ be an open set. To define the volume of U, let $\{U_\alpha\}$ be a cover by charts and ρ_α a partition of unity subordinate to the cover. Then define,

$$\mathrm{Vol}(U) = \sum_\alpha \rho_\alpha \int_{U_\alpha} \sqrt{\det g_{ij}^\alpha} dx^1 \cdots dx^n.$$

Remark 4.2. In fact we can work U a Borel set above, or more generally we can define measurable sets, but we won't need such generality here. For reference however, the assignment,

$$U \mapsto \mathrm{Vol}(U)$$

defines an *outer-measure* on M called the *Riemannian measure* also denoted μ_g. The pair (M, μ_g) is measure space with many properties similar to the Lebesgue measure on Euclidean space. For example, Borel sets are measurable.

Example 4.3. Compute the measure on the two-sphere in polar coordinates and verify the usual formula for the area of the sphere. The higher dimensional area may be computed in a number of ways, induction on dimension being one of them.

4.1.3 Anti-Symmetric Algebra

The discussion above fits into a more general framework, and can be expressed well in terms of alternating forms. As usual, we begin with the vector space version and then move on to the vector bundle version.

Definition 4.4. The wedge product of V with itself, denoted $V \wedge V$ is the unique vector space and projection map $\pi : V \times V \to V \wedge V$ satisfying the following universal property: given any alternating bilinear map $\phi : V \times V \to \mathbb{R}$ ($\phi(u, v) = -\phi(v, u)$), there exists a unique linear map $\psi : V \wedge V \to \mathbb{R}$ such that $\phi = \psi \circ \pi$. The image $\pi(u, v)$ is denoted $u \wedge v$.

Remark 4.5. Here are three ways to define the wedge product. To see they give the same vector space (up to isomorphism), one merely has to verify the universal property in each case.

First, the wedge product can be defined in a similar manner to the tensor product, i.e. as a quotient of the free vector space on $V \times V$. One simply adds the relation $(u, v) + (v, u)$ to $R(U, V)$.

Alternatively, it may be realised as the quotient

$$V \otimes V/U$$

where U is the subspace of $V \otimes V$ generated by elements of the form $v \otimes u + u \otimes v$.

By the rank-nullity theorem of linear algebra, the quotient $V \otimes V/U \simeq V \wedge V$ is actually isomorphic to a subspace of $V \otimes V$. Explicitly, define the map

$$\begin{aligned} \mathrm{Alt} : V \otimes V &\to V \otimes V \\ u \otimes v &\mapsto u \otimes v - v \otimes u. \end{aligned}$$

Then the image is isomorphic to $V \otimes V/U$ since the kernel is precisely U.

Just as with the tensor product, we may take repeated wedge products, and this operation is associative.

Definition 4.6. An alternating tensor of degree p is an element of

$$\Lambda^p(V) = \underbrace{V \wedge \cdots \wedge V}_{\text{p times}}.$$

Then we have a projection,

$$\pi : \prod^{p} V \to \Lambda^p(V)$$

satisfying a universal property analogous to the definition above for $V \wedge V$.

There is also the important alternating map $\otimes^p V \to \otimes^p V$,

$$\mathrm{Alt} : v_1 \otimes \cdots \otimes v_p \mapsto \frac{1}{p!} \sum_{\sigma \in S_p} \mathrm{sgn}(\sigma) v_{\sigma_1} \otimes \cdots \otimes v_{\sigma_p}$$

where S_p denotes the symmetric group on p-letters. This image of this map is isomorphic to $\Lambda^p V$. This exhibits $\Lambda^p V$ as a subspace of $\otimes^p V$ and (by the rank-nullity theorem) as also as a quotient of $\otimes^p V$. Note for $p = 2$ this map is exactly the map defined above realising $V \wedge V$ as a quotient of $V \otimes V$.

Definition 4.7. Let $\alpha \in \Lambda^p$, $\beta \in \Lambda^q$. Define the wedge product

$$\alpha \wedge \beta = \frac{(p+q)!}{p!q!} \mathrm{Alt}(\alpha \otimes \beta) \in \Lambda^{p+q}.$$

Remark 4.8. The constants here are chosen so that,

$$(v_1 \wedge \cdots \wedge v_p) \wedge (w_1 \wedge \cdots \wedge w_q) = v_1 \wedge \cdots \wedge v_p \wedge w_1 \wedge \cdots \wedge w_q.$$

To see this, note that the left hand side is

$$\alpha \wedge \beta = \frac{(p+q)!}{p!q!} \mathrm{Alt}(\alpha \otimes \beta).$$

with $\alpha = v_1 \wedge \cdots \wedge v_p$ and $\beta = w_1 \wedge \cdots \wedge w_q$. Now as an exercise, prove the formula,

$$v_1 \wedge \cdots \wedge v_p = p! \mathrm{Alt}(v_1 \otimes \cdots \otimes v_p) \tag{1}$$

where the left hand side is iterated wedge products as defined above (which is associative - check this!). For example, $v \wedge w = 2\mathrm{Alt}(v \otimes w) = v \otimes w - w \otimes v$.

Recalling that Alt is linear, we then have

$$(v_1\wedge\cdots\wedge v_p)\wedge(w_1\wedge\cdots\wedge w_q) = \frac{(p+q)!}{p!q!}\mathrm{Alt}(p!\mathrm{Alt}(v_1\otimes\cdots\otimes v_p)\otimes(q!\mathrm{Alt}(w_1\otimes\cdots\otimes w_q))$$
$$= (p+q)!\mathrm{Alt}(\mathrm{Alt}(v_1\otimes\cdots\otimes v_p)\otimes(\mathrm{Alt}(w_1\otimes\cdots\otimes w_q))$$
$$= (p+q)!\mathrm{Alt}\left(\frac{1}{p!}\sum_{\sigma\in S_p}\mathrm{sgn}(\sigma)v_{\sigma_1}\otimes\cdots\otimes v_{\sigma_p}\bigotimes\frac{1}{q!}\sum_{\rho\in S_q}\mathrm{sgn}(\rho)v_{\rho_q}\otimes\cdots\otimes v_{\rho_q}\right)$$
$$= (p+q)!\mathrm{Alt}\left(\frac{1}{p!}\frac{1}{q!}\sum_{(\sigma,\rho)\in S_p\times S_q}\mathrm{sgn}(\sigma)\mathrm{sgn}(\rho)v_{\sigma_1}\otimes\cdots\otimes v_{\sigma_p}\otimes v_{\rho_q}\otimes\cdots\otimes v_{\rho_q}\right).$$

There are exactly $p!q!$ terms in the sum and they are all permutations of $v_1\otimes\cdots\otimes v_p\otimes w_1\otimes\cdots\otimes w_q$, hence we obtain

$$(v_1\wedge\cdots\wedge v_p)\wedge(w_1\wedge\cdots\wedge w_q) = (p+q)!\mathrm{Alt}(v_1\otimes\cdots\otimes v_p\otimes w_1\otimes\cdots\otimes w_q)$$
$$= v_1\wedge\cdots\wedge v_p\wedge w_1\wedge\cdots\wedge w_q$$

by equation (1). Note that we may think of $S_p\times S_q\subset S_{p+q}$ where S_p acts on the first p letters and S_q acts on the last q letters. Then $\mathrm{sgn}(\sigma,\rho) = \mathrm{sng}(\sigma)\mathrm{sgn}(\rho)$ and

$$\mathrm{Alt}(v_{\sigma_1}\otimes\cdots\otimes v_{\sigma_p}\otimes w_{\rho_1}\otimes\cdots\otimes w_{\rho_q}) = \mathrm{sgn}(\sigma)\mathrm{sgn}(\rho)\mathrm{Alt}(v_1\otimes\cdots\otimes v_p\otimes w_1\otimes\cdots\otimes w_q).$$

With respect to a basis $\{e_i\}$ for V, a basis for $\otimes^p V$ is $e_{i_1}\otimes\cdots\otimes e_{i_p}$ for any choice (with repeats allowed!) of p-tuples of indices $(i_1,\cdots,i_p)\in\prod^p\{1,\cdots,n\}$. A basis for $\Lambda^p(V)$ is $\{e_{i_1}\wedge\cdots\wedge e_{i_p}\}$ with $1\le i_1<i_2<\cdots<i_p\le n$. Notice that no repeats are allowed and that the indices are strictly increasing! In particular,

$$\dim\Lambda^1 = n,\quad \dim\Lambda^n = 1,\quad \dim\Lambda^p = 0, p\ge n.$$

The general formula is,

$$\dim\Lambda^p = \binom{n}{p}.$$

Last of all, we can express the notion of determinant quite nicely (and in a basis independent way) using the wedge product. If $T: V\to V$ is a linear transformation, then there is an induced linear transformation,

$$\wedge^n T: \Lambda^n\to\Lambda^n$$
$$v_1\wedge\cdots\wedge v_n\mapsto T(v_1)\wedge\cdots\wedge T(v_n).$$

Since $\dim \Lambda^n = 1$, this map is given by scalar multiplication! We then define the determinant of the linear map T (*there are no matrices here!*) by

$$T(v_1) \wedge \cdots \wedge T(v_n) = \det T(v_1 \wedge \cdots \wedge v_n).$$

That is $\det T$ is the scalar.

Remark 4.9. If $\{e_i\}$ is a basis for V and T^i_j is the matrix representing T with respect to that basis, then $\det T^i_j = \det T$ where the left hand side is the determinant of a matrix and the right hand side is the determinant just defined. Therefore, we automatically obtain the well known fact that the determinant is invariant under similarity transformation. Compare this with our definition of the trace via tensor contractions!

4.2 Week 04, Lecture 02: Orientation, Integration and Differential Forms

4.2.1 Orientation

To connect the anti-symmetric algebra with integration and volume, we need first to introduce the notion of orientation.

Definition 4.10. Let $\omega, \nu \in \Lambda^n(V)$. Define the equivalence relation, $\omega \sim \nu$ if and only if $\nu = \lambda\omega$ with $\lambda > 0$. This equivalence relation partitions $\Lambda^n(V)$ into three equivalence classes (one of which is the equivalence class of 0 containing only the 0 n-vector). An *orientation* on V is choice of one or the other non-zero equivalence classes. An *oriented vector space* is a vector space with a choice of orientation.

An non-zero n-vector, ω determines an orientation on V denoted by $[\omega]$. The other non-zero orientation is referred to as the *opposite orientation* and is equal to $[-\omega]$.

We can use an inner-product, g on V to define a n-vector μ_g representing an orientation on V. First, recall that g induces an inner product on $\otimes^n V$, on indecomposable elements,

$$g(v_1 \otimes \cdots \otimes v_p, w_1 \otimes \cdots \otimes w_p) = g(v_1, w_1) \cdots g(v_p, w_p).$$

Since $\Lambda^p(V) \subseteq \otimes^p V$, by restriction we also have an inner product on $\Lambda^p(V)$. However, for convenience later, let us add a normalising factor, and define

$$g(v_1 \wedge \cdots \wedge v_p, w_1 \wedge \cdots \wedge w_p) = \frac{1}{p!} g(v_1, w_1) \cdots g(v_p, w_p).$$

Also define,

$$g(c_1, c_2) = c_1 c_2$$

for $c_1, c_2 \in \mathbb{R} \simeq \Lambda^0(V)$. Before moving on to orientation, let us record a rather useful expression for this inner-product.

Lemma 4.11. *The inner-product g on $\Lambda^p(V)$ satisfies,*

$$g(v_1 \wedge \cdots \wedge v_p, w_1 \wedge \cdots \wedge w_p) = \det g(v_i, w_j)$$

Proof. We make use of a useful formula for the determinant,

$$\det A_{ij} = \sum_{\sigma \in \Sigma_k} \mathrm{sgn}(\sigma) A_{1\sigma_1} \cdots A_{k\sigma_k}$$

for any $k \times k$ matrix A. We also need the formula $e_1 \wedge \cdots \wedge e_k = k!\mathrm{Alt}(e_1 \otimes e_k)$. Then we may write

$$\begin{aligned} &g(v_1 \wedge \cdots \wedge v_p, w_1 \wedge \cdots \wedge w_p) \\ &\qquad = \frac{1}{p!} g\left(\sum_{\sigma \in S_p} \mathrm{sgn}(\sigma) v_{\sigma_1} \otimes \cdots \otimes v_{\sigma_p}, \sum_{\tau \in S_p} \mathrm{sgn}(\tau) w_{\tau_1} \otimes \cdots \otimes w_{\tau_p} \right) \\ &\qquad = \frac{1}{p!} \sum_{\sigma \in S_p} \mathrm{sgn}(\sigma) \sum_{\tau \in S_p} \mathrm{sgn}(\tau) g(v_{\sigma_1}, w_{\tau_1}) \cdots g(v_{\sigma_p}, w_{\tau_p}) \\ &\qquad = \frac{1}{p!} \sum_{\sigma \in S_p} \mathrm{sgn}(\sigma) \det g(v_{\sigma_i}, w_j). \end{aligned}$$

Now observe that there are exactly $p!$ terms in the sum, and that they are all equal to $\det g(v_i, w_j)$! The latter because the determinant is multiplied by $\mathrm{sgn}(\sigma)$ when σ permutes the rows (or columns). □

Note the appearance of the normalising constant $\frac{1}{p!}$ we used in the definition of the inner-product on $\Lambda^p(V)$.

Definition 4.12. Let V be an oriented vector space with orientation $[\omega]$ and let g be an inner product on V. The orientation form of g is the unique form $\mu_g \in [\omega] \subseteq \Lambda^n$, and such that $g(\mu_g, \mu_g) = 1$ where g is the inner-product on $\Lambda^n(V)$.

That μ_g is unique follows since $\Lambda^n(V)$ is one dimensional and V is oriented: any element $\eta \in \Lambda^n(V)$ may be written $\eta = \lambda \mu_g$ and so $g(\eta, \eta) = \lambda^2 g(\mu_g, \mu_g)$ implies $\eta = \pm \mu_g$ are the only two n-vectors such that $g(\eta, \eta) = 1$. However, $-\mu_g$ also represents the opposite orientation, $[-\mu_g] = [-\omega]$.

With respect to a basis $\{e_i\}$ for V, a basis for $\Lambda^n(V)$ is $e_1 \wedge \cdots \wedge e_n$. Therefore, $\mu_g = \lambda e_1 \wedge \cdots \wedge e_n$ for some real number, $\lambda \neq 0$. Lemma 4.11 implies that,

$$1 = g(\mu_g, \mu_g) = \lambda^2 g(e^1 \wedge \cdots \wedge e_n, e^1 \wedge \cdots \wedge e_n) = \lambda^2 \det g_{ij}$$

and hence

$$\mu_g = \pm \frac{1}{\sqrt{\det g_{ij}}} e_1 \wedge \cdots \wedge e_n.$$

This looks familiar to the definition of the volume form on a Riemannian manifold defined in the previous lecture, but something funny is going on with the determinant in the denominator. More on this in the next section! Note also that a positive definite matrix always has positive determinant and so the square root is well defined.

A basis $\{e_i\}$, for V is *orthonormal* if $g_{ij} = g(e_i, e_j) = \delta_{ij}$. An *ordered basis* $\{e_i\}$ is *positively oriented* (with respect to μ_g) if $e_1 \wedge \cdots \wedge e_n = \lambda \mu_g$ with $\lambda > 0$, in which case $\mu_g = +\frac{1}{\sqrt{\det g_{ij}}} e_1 \wedge \cdots \wedge e_n$. What happens to λ under reordering of the basis? With respect to a positively oriented, orthonormal basis, the orientation form is simply, $\mu_g = e_1 \wedge \cdots \wedge e_n$.

Given an inner-product space (V, g) with orientation form μ_g we can define a very useful operation, known as the Hodge star operator.

Definition 4.13. Let (V, g) be an inner-product space with orientation form μ_g and define the *Hodge star* operator,

$$* : \Lambda^k(V) \to \Lambda^{n-k}(V)$$

which given $\beta \in \Lambda^k(V)$ is the unique element $*\beta \in \Lambda^{n-k}(V)$ such that for every $\alpha \in \Lambda^k(V)$,

$$\alpha \wedge *\beta = g(\alpha, \beta)\mu_g \in \Lambda^n(V).$$

Remark 4.14. Observe that $\dim \Lambda^k = \dim \Lambda^{n-k}$ (the binomial coefficients are equal!). As an exercise, try to show $*$ is an isomorphism. *Hint*: for fixed β, the map $\alpha \mapsto g(\alpha, \beta)$ defines an element $\beta^\flat \in (\Lambda^k(V))^*$. This is the musical isomorphism described in a previous lecture for a vector space with an inner-product, which applies in particular to $\Lambda^k(V)$. On the other hand, for fixed $\gamma \in \Lambda^{n-k}(V)$, the map $\alpha \mapsto \alpha \wedge \gamma \in \Lambda^n(V)$ also defines an isomorphism $\Lambda^{n-k}(V) \simeq (\Lambda^k(V))^*$ where we consider $\Lambda^n(V) \simeq \mathbb{R}$ via $\mu_g \mapsto 1$. The Hodge star is the composition of these two isomorphisms.

In the particular case $k = 0$, $* : \Lambda^0 \simeq \mathbb{R} \to \Lambda^n$ and we can characterise the orientation form by

$$\mu_g = *1$$

since by definition,

$$1 \wedge *1 = g(1, 1)\mu_g = \mu_g = 1 \wedge \mu_g$$

implies that $*1 = \mu_g$ by uniqueness (which follows from the remark).

With respect to a positively oriented, orthonormal basis,

$$*(e_1 \wedge \cdots \wedge e_k) = e_{k+1} \wedge \cdots \wedge e_n.$$

With respect to an arbitrary positively oriented basis,

$$*(e_1 \wedge \cdots \wedge e_k) = \sqrt{\det g(e_i, e_j)} e_{k+1} \wedge \cdots \wedge e_n.$$

This latter follows from Lemma 4.11 and $\mu_g = \frac{1}{\sqrt{\det g_{ij}}} e_1 \wedge \cdots \wedge e_n$, while the former follows from the latter. What if the basis is not positively oriented?

4.2.2 Differential Forms

Definition 4.15. A differential k-form on M is a section of the bundle

$$\Lambda^k(M) = \Lambda^k(T^*M).$$

The set of differential forms is written,

$$\Omega^k(M) = \Gamma(\Lambda^k(M), M).$$

As with the tensor bundles, we may define the bundle $\wedge^k(T^*M)$ by specifying transition maps and it is well defined (though a little messy to write down explicitly).

In local coordinates we write dx^i for the basis of T^*M dual to the basis ∂_{x^i} for TM. A local section $\alpha \in \Omega^k(M)$ may thus be written

$$\alpha = \alpha_{i_1, \cdots, i_k} dx^{i_1} \wedge \cdots \wedge dx^{i_k}$$

where we use the summation convention to sum over all indices $1 \leq i_1 < \cdots < i_k \leq n$. More briefly, adopting multi-index notation $I = (i_1, \cdots, i_k)$ we write,

$$\alpha = \alpha_I dx^I$$

with $\alpha_I = \alpha_{i_1, \cdots, i_k}$ and $dx^I = dx^{i_1} \wedge \cdots \wedge dx^{i_k}$. Again the summation convention implies a sum over all multi-indices $(i_1, \cdots, i_k)$ such that $1 \leq i_1 < \cdots < i_k \leq n$.

Remark 4.16. Note that differential forms are sections of $\Lambda^k(T^*M)$ and not sections of $\Lambda^k(TM)$. In particular, given a metric g on $\Lambda^k(TM)$ the dual metric $g = g^*$ (remember we denote this by g as well) is a metric on T^*M which induces a metric g on $\Lambda^k(T^*M)$. Given local coordinates, with local frame ∂_i the metric may be written $g_{ij} = g(\partial_i, \partial_j)$. The dual metric is written $g^{ij} = g(dx^i, dx^j) = g((dx^i)^\flat, (dx^j)^\flat)$. As an exercise show that g^{ij} is the inverse matrix to g_{ij}.

Definition 4.17. A *volume form*, ω on M is any non-vanishing section of $\Lambda^n(M)$ and this defines an orientation on M as the set of volume forms ν such that $\nu = f\mu$ with $f \in C^\infty(M)$ and $f(x) > 0$ for all $x \in M$. A manifold is *orientable* if there exist volume forms and is *oriented* if a choice of equivalence class of volume forms is made. As with vector spaces, $\nu \sim \mu$ if $\nu = f\mu$ for a smooth positive function and this relation partitions volume forms into exactly two equivalence classes (note here we exclude n-forms that are zero anywhere).

In particular, a Riemannian metric determines a volume form on M provided that M is orientable. The *Riemannian volume form* is

$$\mu_g = *1$$

where here $*$ is the Hodge dual, which extends from a map of vector spaces to a bundle map $\Lambda^k(M) \to \Lambda^{n-k}(M)$ and $1 \in C^\infty(M) \simeq \Gamma(\Lambda^0(M))$ denotes the constant function $x \mapsto 1$. With respect to local coordinates,

$$\mu_g = \frac{1}{\sqrt{\det g^{ij}}} dx^1 \wedge \cdots \wedge dx^n = \sqrt{\det g_{ij}} dx^1 \wedge \cdots \wedge dx^n,$$

since g^{ij} is the inverse of g_{ij}. This looks very much like the expression we used for the Riemannian measure in the previous lecture. The following definition identifies the Riemannian volume form with the Riemannian measure.

Definition 4.18. Let ω be a volume form on M. On a compact local chart, $\phi_\alpha : U_\alpha \to V_\alpha$, we may write

$$\omega = \rho dx^1 \wedge \cdots \wedge dx^n$$

for $\rho > 0$ a smooth function (where we orient the x_i so that ρ is positive). Define

$$\int_{U_\alpha} \omega = \int_{V_\alpha} \rho dx^1 \cdots dx^n$$

Now using a partition of unity $\{\rho_\alpha\}$ subordinate to an open cover by charts $\{U_\alpha\}$, we define

$$\int_M \omega = \sum_\alpha \rho_\alpha \int_{U_\alpha} \omega.$$

Given an open set $U \subset M$, define

$$\int_U \omega = \sum_\alpha \rho_\alpha \int_{U \cap U_\alpha} \omega.$$

We may also integrate smooth functions with respect to ω,

$$\int_M f\omega = \sum_\alpha \rho_\alpha \int_{V_\alpha} f \circ \phi_\alpha^{-1} \rho dx^1 \cdots dx^n.$$

4.2.3 Exterior Derivative

Now that we have differential forms, let's differentiate them! First observe that if $f \in C^\infty(M)$, the differential is a map,

$$df_x : T_x M \to T_{f(x)}\mathbb{R} \simeq \mathbb{R},$$

or in other words, $df_x \in T_x^* M$. The assignment, $x \mapsto df_x$ defines a section $df \in \Gamma(T^*M) = \Lambda^1(M)$, i.e a 1-form. In local coordinates,

$$df = \frac{\partial f}{\partial x^i} dx^i$$

and so df is indeed a smooth section since the components, $\frac{\partial f}{\partial x^i}$ are smooth. Now let's differentiate higher order forms. Recall that $\Lambda^0(M) \simeq C^\infty(M)$.

Proposition 4.19. *For each k, there exists a unique map,*

$$d : \Lambda^k(M) \to \Lambda^{k+1}(M)$$

such that

1. *$d^0 f = df$ for any smooth function f (agrees with the differential on functions),*

2. *$d^{k+l}(\alpha \wedge \beta) = d^k \alpha \wedge \beta + (-1)^k \alpha \wedge d^l \beta$ (Leibniz product rule),*

3. *$d^{k+1} \circ d^k = 0$ (co-chain)*

Typically we drop the superscript from d^k and simply write d for these operators.

Proof. Assuming d exists, in local coordinates $\{x^i\}$ an arbitrary k-form may be written

$$\alpha = \alpha_I dx^I = \alpha_{i_1,\cdots,i_k} dx^{i_1} \wedge \cdots \wedge dx^{i_k}$$

with the summation convention in force and each α_I a smooth function. Then using 1-3, we have

$$\begin{aligned} d\alpha &= d(\alpha_I dx^{i_1} \wedge \cdots \wedge dx^{i_k}) \\ &= d\alpha_I \wedge (dx^{i_1} \wedge \cdots \wedge dx^{i_k}) + (-1)^0 \alpha_I \wedge d(dx^{i_1} \wedge \cdots \wedge dx^{i_k}) \quad (2) \\ &= d\alpha_I \wedge dx^I + \alpha_I \wedge \sum_{p=1}^{k} (-1)^{p-1} dx^{i_1} \wedge \cdots dx^{i_{p-1}} \wedge d^2 x^{i_p} \wedge dx^{i_{p+1}} \wedge \cdots dx^{i_k} \quad (2) \\ &= d\alpha_I \wedge dx^I \quad (3) \\ &= \frac{\partial \alpha_I}{\partial x^i} dx^i \wedge dx^I \quad (1). \end{aligned}$$

Note that the summation convention applies to both i and I here so that written in full,

$$d\alpha = \sum_{i=1}^{k} \sum_{(i_1,\cdots,i_p)} \frac{\partial \alpha_I}{\partial x^i} dx^i \wedge dx^{i_1} \wedge \cdots \wedge dx^{i_k}.$$

Therefore, if d exists it is given explicitly by this local formula hence is unique.

To prove existence, define $d\alpha$ locally by this formula. It is an exercise to show that $d\alpha$ transforms correctly under change of coordinates, hence gives a well defined section. This latter calculation may be simplified by using the notion of pull back, which we will define shortly, so if you're worried about the calculation, wait until then! □

Example 4.20. Let $\alpha = f_1 dx^1 + f_2 dx^2$ be the local expression of a one form on a two-dimensional manifold. Then

$$\begin{aligned} d(f_1 dx^1) &= \sum_{i=1}^{2} \frac{\partial f_1}{\partial x^i} dx^i \wedge dx^1 \\ &= \frac{\partial f_1}{\partial x^1} dx^1 \wedge dx^1 + \frac{\partial f_1}{\partial x^2} dx^2 \wedge dx^1 \\ &= -\frac{\partial f_1}{\partial x^2} dx^1 \wedge dx^2 \end{aligned}$$

and similarly for $d(f_2 dx^2)$. Therefore,

$$d\alpha = \left(\frac{\partial f_2}{\partial x^1} - \frac{\partial f_1}{\partial x^2} \right) dx^1 \wedge dx^2.$$

In the plane you should recognise this as the *scalar curl* of the vector field $f_1 \partial_{x^1} + f_2 \partial_{x^2}$. More generally, here are some formulae relating classic quantities to the exterior derivative d, the Hodge star $*$ and the musical isomorphisms.

- gradient: $\nabla f = df^\sharp$ (i.e. $df(X) = g(\nabla f, X)$ for all vectors X),
- divergence: $\operatorname{div} X = *d * X^\flat$,
- curl: $\text{curl}X = (*dX^\flat)^\sharp$,
- Laplacian: $\Delta f = \operatorname{div} \nabla f = *d * df$.

Note that all of these definitions make sense on a Riemannian manifold and may serve as definitions! As an exercise, verify these formulae in Euclidean space. This is actually not as hard as it first appears since $g_{ij} = \delta_{ij} \Rightarrow g^{ij} = \delta^{ij}$ and so for example,

$$df = \frac{\partial f}{\partial x^i} dx^i$$

from which one deduces,

$$\nabla f = g^{ij} \frac{\partial f}{\partial x^j} \partial_{x^i} = \frac{\partial f}{\partial x^i} \partial_{x^i}.$$

You'll need an expression for the Hodge star for the others. Note that the Laplacian follows immediately from the gradient and divergence.

4.3 Week 04, Lecture 03: Connections

In this lecture we would like to define a way to differentiate vector fields. First observe that in Euclidean space, we may define the *directional derivative* $D_X Y$, of a vector field $Y = Y^j \partial_j$ in the direction $X = X^i \partial_x$:

$$D_X Y = X^i \partial_i (Y^j \partial_j) = X^i \frac{\partial Y^j}{\partial x_i} \partial_j.$$

That is we just differentiate the component functions Y^j (note the summation convention!). Writing,

$$D_X Y = (D_X Y)^j \partial_j,$$

we have

$$(D_X Y)^j = X^i \frac{\partial Y^j}{\partial x^i}$$

with a sum over i. Or in other words, the j'th component of the directional derivative is just the directional derivative of the j'th component. **This does not work in general on a manifold!**

Example 4.21. Let $\phi : U \to V, \bar{\phi} : \bar{U} \to \bar{V}$ be two local coordinate charts for M, and write $X = X^i \partial_i$ and $\bar{X} = \bar{X}^k \bar{\partial}_k$ for the coordinate representations of the same vector field on M. That is,

$$\bar{X}(\tau(x)) = d\tau_x \cdot X(x)$$

where $\tau = \bar{\phi} \circ \phi^{-1}$ is the transition map. In particular,

$$d\tau \cdot \partial_i = \tau_i^k \bar{\partial}_k$$

where $\tau_i^k = \frac{\partial}{\partial x^i} (\bar{\phi} \circ \phi^{-1})^k$. Hence,

$$\bar{X}^k \bar{\partial}_k = d\tau \cdot (X^i \partial_i) = X^i d\tau \cdot \partial_i = \tau_i^k X^i \bar{\partial}_k,$$

so that

$$\bar{X}^k(\bar{x}) = \tau_i^k(x) X^i(x)$$

where $\bar{x} = \tau(x)$. We may also use the chain rule to relate partial derivatives. For a smooth function $f : V \to \mathbb{R}$, define $\bar{f} : \bar{V} \to \mathbb{R}$ by $\bar{f}(\bar{x}) = f(\tau^{-1}(\bar{x})) = f(x)$. Then

$$\bar{\partial}_l \bar{f}(\bar{x}) = \partial_j f(x) \bar{\partial}_l (\tau^{-1})^j(\bar{x}) = \sigma_l^j(\bar{x}) \partial_j f(x)$$

where $\sigma = d(\tau^{-1}) = (d\tau)^{-1}$. In particular,

$$\bar{\partial}_l \bar{X}^k = \sigma_l^j \partial_j (\tau_i^k X^i) = \sigma_l^j \left(\tau_i^k \partial_j X^i + X^i \partial_j \tau_i^k \right). \tag{2}$$

You might already be able to see a problem with partial derivatives from this equation.

Pressing on heedlessly, for a vector field Y, let $\bar{Y} = d\tau \cdot Y$ and define the derivatives,

$$D_X Y = X^j \partial_j (Y^i) \partial_i$$

and

$$\bar{D}_{\bar{X}} \bar{Y} = \bar{X}^l \bar{\partial}_l (\bar{Y}^k) \bar{\partial}_k$$

in the charts. These are perfectly well defined vector fields on the charts, but if we want these to be the local coordinate representations of a vector field on M, then they need to be identified by the transition map,

$$\bar{D}_{\bar{X}} \bar{Y} = d\tau \cdot (D_X Y).$$

Expressed component-wise, this requires that

$$\bar{X}^l \bar{\partial}_l (\bar{Y}^k) \bar{\partial}_k = \tau_i^k X^j \partial_j (Y^i) \bar{\partial}_k. \tag{3}$$

On the other hand, since $\bar{X} = d\tau \cdot X$ and $\bar{Y} = d\tau \cdot Y$, we must have

$$\bar{D}_{\bar{X}} \bar{Y} = \bar{D}_{d\tau \cdot X} d\tau \cdot Y.$$

Expressed in components this gives

$$\begin{aligned} \bar{X}^l \bar{\partial}_l (\bar{Y}^k) \bar{\partial}_k &= \tau_j^l X^j \bar{\partial}_l (\bar{Y}^k) \bar{\partial}_k \\ &= \tau_j^l X^j \sigma_l^m \left(\tau_i^k \partial_m Y^i + Y^i \partial_m \tau_i^k \right) \bar{\partial}_k \\ &= \delta_m^j \left(\tau_i^k X^j \partial_m Y^i + X^j Y^i \partial_m \tau_i^k \right) \bar{\partial}_k \\ &= \left(\tau_i^k X^j \partial_j Y^i + X^j Y^i \partial_j \tau_i^k \right) \bar{\partial}_k \\ &= d\tau \cdot D_X Y + \left(X^j Y^i \partial_j \tau_i^k \right) \bar{\partial}_k. \end{aligned} \tag{4}$$

where in the second line we used equation (2), in the third line we used $\sigma = d\tau^{-1}$ so that $\tau_j^l \sigma_l^m = \delta_j^m$, and in the final line we used equation (3).

The conclusion is that,

$$\bar{D}_{\bar{X}} \bar{Y} = \bar{D}_{d\tau \cdot X} (d\tau \cdot Y) \neq d\tau \cdot D_X Y$$

in general. In fact, equality will hold for all vector fields X, Y if and only if $d\tau$ is constant so that $\partial_j \tau_i^k = 0$ in the last line of equation (4). This is most certainly not true in general! In fact, if we perform a non-constant change of variables (i.e. τ is not an affine transformation) in Euclidean space (e.g. polar coordinates), the directional derivative expressed in the new coordinates ($d\tau \cdot D_X Y$ is not the directional derivative of those coordinates, $\bar{D}_{\bar{X}} \bar{Y}$. If you don't know the expression for the directional derivative in polar coordinates, check it as an exercise (or look it up).

Example 4.22. Let $X, Y \in \mathfrak{X}(S)$ be vector fields on a regular surface $S \subset \mathbb{R}^3$. Let $U \subset \mathbb{R}^3$ be an open set on which X and Y extend to $\bar{X}$ and $\bar{Y}$, vector fields on U. That is $\bar{X} : U \mapsto \mathbb{R}^3$ is a smooth function such that $\bar{X}|_S = X$ and likewise for $\bar{Y}$. That such extensions exist can be first proven locally by using the implicit function theorem to obtain a diffeomorphism of an open set $V \subset \mathbb{R}^3 \to W \subset \mathbb{R}^3$ such that $S \cap V \mapsto \{x^3 = 0\} \cap W$ and then $X = X(x^1, x^2)$ and we may define $\bar{X}(x^1, x^2, x^3) = X(x^1, x^2)$. Now, as usual, a partition of unity may be used to patch together the locally defined $\bar{X}$.

Then we would like to define,

$$\nabla_X Y = (D_{\bar{X}} \bar{Y})|_S$$

where D is the directional derivative on $\mathbb{R}^3$. Now certainly $D_{\bar{X}} \bar{Y}$ is a well defined vector field on U and we may restrict to S to obtain a smooth map $S \to \mathbb{R}^3$. The problem here is that in general, the so defined $\nabla_X Y$ will not be tangent to S!

As a concrete exercise, check that for the vector fields $X = \partial_\theta$ and $Y = \partial_\phi$ on the sphere in polar coordinates, $D_X Y$ is not tangent to the sphere in general. *Hint*: Write X and Y as vector fields on $\mathbb{R}^3$ as functions of (x, y, z) and then it's easy to calculate.

Remark 4.23. There is a well defined way to differentiate vector fields on a manifold, namely the Lie derivative,

$$\mathcal{L}_X Y = [X, Y].$$

This may be thought of as a derivative since it obeys the Leibniz rule,

$$\mathcal{L}_X (fY) = (df \cdot X) Y + f \mathcal{L}_X Y.$$

This is not however, a generalisation of directional derivative on Euclidean space. In particular (even on Euclidean space), the Lie derivative does not satisfy the rule,

$$\mathcal{L}_{fX+gY}Z = f\mathcal{L}_X Z + g\mathcal{L}_Y Z$$

which **is** satisfied by the directional derivative:

$$D_{fX+gY}Z = fD_X Z + gD_Y Z.$$

Thus the Lie derivative does not generalise the notion of directional derivative.

The preceding discussion suggests that we need to work a bit harder to differentiate vector fields on manifolds. The appropriate notion turns out to be that of a connection. Essentially we define a connection to be any operation that behaves like the directional derivative. We will also be able to largely avoid these messy change of local coordinate calculations as in the first example. This is because our definitions will automatically be well defined independently of any choice of local coordinates (as with all our definition so far) and we need never check that they transform correctly under change of coordinates. This should be a great relief! The calculation above is perhaps one of the easiest such calculations and that approach only gets worse so we really want to avoid it wherever possible.

Definition 4.24. A *connection* on a vector bundle $\pi : E \to M$ is an $\mathbb{R}$-linear map

$$\nabla : \Gamma(E) \to \Gamma(T^*M \otimes E)$$

satisfying the Leibniz rule,

$$\nabla(fs) = df \otimes s + f\nabla s$$

for any $f \in C^\infty(M)$ and $s \in \Gamma(E)$.

Given a vector field X on M, the *covariant derivative* of $s \in \Gamma(E)$ in the direction X, is the contraction,

$$\nabla_X s = \mathrm{Tr}(X \otimes \nabla s)$$

of the first slots of $X \otimes \nabla S \in TM \otimes T^*M \otimes E$.

Remark 4.25. Noting that $C^\infty(M) \otimes E \simeq E$ via $f \otimes s \mapsto fs$ we may also write the Leibniz rule more suggestively as

$$\nabla(f \otimes s) = df \otimes s + f \otimes \nabla s.$$

Remark 4.26. To understand the covariant derivative, recall that the contraction of $X \otimes \alpha \in TM \otimes T^*M$ is the map

$$\mathrm{Tr}(X \otimes \alpha) = \alpha(X)$$

from $TM \otimes T^*M \to M \times \mathbb{R}$. Then we define the contraction,

$$\begin{aligned} \mathrm{Tr}^E = \mathrm{Tr} \otimes \mathrm{Id} : TM \otimes T^*M \otimes E &\to E \\ X \otimes \alpha \otimes s &\mapsto \alpha(X)s. \end{aligned}$$

This operation is also defined on sections by acting pointwise yielding the covariant derivative,

$$\begin{aligned} \nabla : \Gamma(TM \otimes E) &\to \Gamma(E) \\ X \otimes s &\mapsto \nabla_X s = \mathrm{Tr}(X \otimes \nabla s). \end{aligned}$$

If we fix X, and let s vary we obtain the covariant derivative in the direction X,

$$\nabla_X(s) = \mathrm{Tr}(X \otimes \nabla s).$$

Note that $\nabla_X : \Gamma(E) \to \Gamma(E)$. That is the derivative of a section $s \in \Gamma(E)$ is again a section $\nabla_X s \in \Gamma(E)$. Also note that we may think of ∇s as a map $TM \to E$ via,

$$\nabla s(X) = \nabla_X s.$$

Another way of saying this is that the map

$$\alpha \otimes s \in T^*M \otimes E \mapsto (X \mapsto \alpha(X)s) \in \mathrm{Hom}(TM, E)$$

is an isomorphism.

The most important bundle E is the tangent bundle TM and a connection on TM is a map,

$$\nabla : \Gamma(TM) \to \Gamma(T^*M \otimes TM)$$

satisfying the Liebniz rule. Fixing X, the corresponding covariant derivative is the map

$$\begin{aligned} \nabla_X : TM &\to TM \\ Y &\mapsto \nabla X(Y) \end{aligned}$$

where $\nabla X \in T^*M \otimes M \simeq \mathrm{Hom}(TM, TM)$.

From the properties of the connection, we have the following properties of the covariant derivative,

- $\nabla_X(c_1Y_1 + c_2Y_2) = c_1\nabla_X Y_1 + c_2\nabla_X Y_2$ ($\mathbb{R}$-linearity),
- $\nabla_{f_1X_1+f_2X_2}Y = f_1\nabla_{X_1}Y + f_2\nabla_{X_2}Y$ ($C^\infty(M)$-linearity in X),
- $\nabla_X(fY) = f\nabla_X Y + X(f)Y$ (Leibniz rule)

These properties in particular are shared by the directional derivative in Euclidean space.

It might be helpful to see this in local coordinates. Let $X = X^i\partial_i$ and $Y = Y^j\partial_j$. Then using the properties of covariant derivatives,

$$\nabla_{X^i\partial_i}(Y^j\partial_j) = X^i\nabla_{\partial_i}(Y^j\partial_j) = X^i\partial_i(Y^j)\partial_j + X^iY^j\nabla_{\partial_i}\partial_j.$$

The first term is the partial derivatives we saw before that cause so much difficulty under change of coordinates. **In general, the other term** $\nabla_{\partial_i}\partial_j \neq 0$. We know this must be true since if this term were zero, we'd be right back where we started with partial derivatives and we know that doesn't work!

More concretely, if x^i are the standard coordinates on Euclidean space, and $\nabla = D$ the directional derivative, then $\nabla_{\partial_i}\partial_j = 0$. However, if we change to polar coordinates, then the vector fields $\partial_r, \partial_\theta$ are non-constant and hence $\nabla_{\partial_i}\partial_j \neq 0$ in these coordinates. Thus in the general case, for each pair of indices (i,j), $\nabla_{\partial_i}\partial_j$ is a vector field and we may write it in terms of the basis $\{\partial_k\}$:

$$\nabla_{\partial_i}\partial_j = \Gamma^k_{ij}\partial_k$$

for smooth functions Γ^k_{ij}. These functions are known as the *Christoffel Symbols*. As an exercise, calculate these in spherical polar coordinates in the plane.

Remark 4.27. As an exercise, compute how the Christoffel symbols transform under a change of coordinates. You'll find out that $\Gamma^k_{ij}\partial_k$ *does not transform like a vector field and hence the Christoffel symbols are coordinate dependent.* You may feel this implies the connection is not well defined on the manifold, but only locally. This is actually not true, since we defined a connection as a map on the manifold satisfying certain properties.

So what is going on? Recall that the problem with partial derivatives is that extra, unwanted terms pop up when changing coordinates: the derivatives of the transition maps. **The Christoffel symbols transform in such a way to cancel out these unwanted terms**. Since the unwanted terms are coordinate dependent so are the Christoffel symbols. What is true however, is that for any connection,

$$d\tau \cdot \nabla_X Y = \nabla_{d\tau \cdot X} d\tau \cdot Y$$

under a change of coordinates. That is, given a connection defined on the manifold, if we express this well defined object in local coordinates,

$$\nabla_{X^i \partial_i}(Y^j \partial_j) = X^i \partial_i (Y^j) \partial_j + \Gamma^k_{ij} X^i Y^j \partial_k,$$

and then apply $d\tau$ we get the same answer if instead we first applied $d\tau$ to X and Y, then took the covariant derivative in the new coordinates using the new Christoffel symbols in these coordinates.

The moral of the story here is that the Christoffel symbols express the connection in local coordinates, but neither the partial derivatives, nor the Christoffel symbols alone transform correctly. It is only together that they transform correctly. Another way of saying this is that the connection $\nabla_X Y$ is tensorial only in X ($C^\infty(M)$ linearity, a.k.a. the test for tensorality in the exercises) and not in Y because of the Liebniz rule. That is the map $X \otimes Y \mapsto \nabla_X Y$ is not $C^\infty(M)$ linear in the second slot and hence does not define a section of $T^*M \otimes T^*M \otimes TM \simeq \mathrm{Hom}(TM \otimes TM, TM)$.

Using the test for tensorality however, it is possible to show that $(\nabla_X Y)(x)$ depends only on the value $X(x)$ of X at x. Therefore for a tangent vector $v \in T_x M$ we may define $\nabla_v Y \in T_x M$ for any vector field Y by extending v arbitrarily to a vector field X with $X(x) = v$ and defining

$$(\nabla_v Y) = (\nabla_X Y)(x)$$

since the result will be independent of the choice of extension X.

Show that there are infinitely many vector fields Y with the same value at x and that in general, $Y_1(x) = Y_2(x)$ does not imply that $\nabla_v Y_1 = \nabla_v Y_2$. *Hint*: The question is local. The reason for this behaviour is that to differentiate a vector field (which after all is just a smooth function $M \to TM$), we need to know the value of Y in a neighbourhood of x, e.g. $f(x) = x$ and $g(x) = x^2$ satisfy $f(0) = g(0)$ but $f'(0) \neq g'(0)$. Soon we will see that actually, all we need to know is the value of Y along the *integral curve* of X through x.

5 Week 05

5.1 Week 05, Lecture 01: Dependence of Connections on Vector Fields

Let us look a little more closely at the dependence of the covariant derivative $\nabla_X Y$ on the vector fields X and Y. More generally, we can ask how $\nabla_X s$ depends on the vector field X and the section $s \in \Gamma(E)$ for a vector bundle E with a connection ∇. The tensorality of ∇s tells us that the dependence of X is pointwise, whereas the Leibniz rule implies that the dependence on s is not pointwise. Since ∇s is a derivative of s, it depends on s in an open neighbourhood of a point. We will see however, that $\nabla_X s$ depends only on s along the integral curves of X.

5.1.1 Differentiation Along a Curve

Lemma 5.1. *Given vector fields, X, Y, the tangent vector $(\nabla_X Y)(x)$ is independent of $X(y)$ for $y \neq x$. That is, the dependence on X is only on the value of X at x. Therefore, given any tangent vector $v \in T_xM$, we may define $\nabla_v Y = (\nabla_X Y)(x) \in T_xM$ where X is any vector field such that $X(x) = v$.*

Proof. We simply write in local coordinates,

$$(\nabla_X Y)(x) = X^i(x)(\partial_i Y^j)(x)\partial_j + X^i(x)Y^j(x)\Gamma_{ij}^k(x)\partial_k.$$

The X dependence is only on the value of $X^i(x)$.

For the second part, we need to check that given a tangent vector $v \in T_xM$ that there exists a vector field X such that $X(x) = v$. Then $\nabla_v Y$ may be defined as $(\nabla_X Y)(x)$ and the first part shows $\nabla_v Y$ is independent of the choice of such X. For this, write $v = v^i\partial_i \in T_xM$ and define,

$$X(x) = v^i\partial_i$$

for x in the chart, $\phi : U \to V$. Then choose any open set W with $x \in W$ and $\bar{W} \subset U$ and let ρ be a smooth bump function, supported in U and such that $\rho|_W \equiv 1$. Then ρX is a smooth vector field on M such that $(\rho X)(x) = \rho(x)X(x) = 1 \cdot v = v$. □

Remark 5.2. The fact that $\nabla Y \cdot X = \nabla_X Y$ depends only on the value of X at x is another way (by the test for tensorality) of saying that $\nabla Y \in \Gamma(T^*M \otimes$

TM, i.e. that ∇Y is a homomorphism $TM \to TM$ and $(\nabla Y(X))(x) = (\nabla Y)_x(X(x))$ and so depends only on the value of X at x. Note however, that $(\nabla Y)_x$ depends on Y in a neighbourhood of x, below we will see that it actually depends on far less. Compare for instance, with the differential, df of a map f. We may compute this as $df_x \cdot v = (f \circ \gamma)'(0)$ where $\gamma'(0) = v$ and $\gamma(0) = v$. Thinking of X as the vector field $\gamma'(t)$, we see that $df_x \cdot v$ depends only on $X(0) = v$ and not on $X(t)$ for $t \neq 0$.

Remark 5.3. The same proof works with connections on arbitrary vector bundles, E. Namely, choose an open set U for which we have a chart on M and for which E is locally trivial (e.g. intersect a chart with a local trivialisation). The we may write $X = X^i(x)\partial_i$, $1 \leq i \leq n$, and $s = s^p(x)e_p$, $1 \leq p \leq k$ where e_p is the standard basis element on $\mathbb{R}^k$ which is a well defined local section of E via the vector bundle isomorphism, $E|_U \simeq U \times \mathbb{R}^k$. Here $E|_U = \pi^{-1}[U]$ is just a convenient shorthand. Define Christoffel symbols for the bundle by

$$\nabla_{\partial_i} e_p = (\Gamma^E)^q_{ip} e_q.$$

Then just as when $E = TM$, using the properties of ∇, we have the formula,

$$(\nabla_X Y)(x) = X^i(x)(\partial_i s^p)(x)e_p + X^i(x)Y^p(x)\Gamma^q_{ip}(x)e_q.$$

The X dependence is only via $X^i(x)$ again.

Definition 5.4. Let $\gamma : I \to M$ be a smooth curve (with I an open interval). A vector field X, along γ is a smooth map $X : I \to TM$ such that $\pi \circ X = \gamma(t)$, or in other words, $X(t) \in T_{\gamma(t)}M$. The set of tangent vectors along γ is denoted $\mathfrak{X}_\gamma(M)$.

Proposition 5.5. *Let M be a smooth manifold and ∇ a connection on TM. Then, for any curve γ, there exists a unique map*

$$\nabla_t : \mathfrak{X}_\gamma(M) \to \mathfrak{X}_\gamma(M)$$

such that for any $Y, Y_1, Y_2 \in \mathfrak{X}$, $c_1, c_2 \in \mathbb{R}$ and $f \in C^\infty(I)$, we have

1. *$\nabla_t(c_1 Y + c_2 Y_2) = c_1 \nabla_t(Y_1) + c_2 \nabla_t(Y_2)$ ($\mathbb{R}$-linearity),*
2. *$\nabla_t(fY) = \frac{df}{dt} Y + f\nabla_t Y$ (Liebniz rule),*
3. *If $Y \in \mathfrak{X}(M)$, then $\nabla_{\gamma'} Y = \nabla_t(Y \circ \gamma)$.*

Remark 5.6. Some remarks on point 3 are in order. Given $Y \in \mathfrak{X}(M)$, $\bar{Y} = Y \circ \gamma$ defines a vector field along γ and so $\nabla_t(Y \circ \gamma)$ makes sense.

As to $\nabla_{\gamma'(t)}Y$, differentiation on all of M, we have two ways to think about this.

The first method, which avoids any technical issues, is to simply note that by the lemma, for each $t \in I$, $\nabla_{\gamma'}Y$ depends only on $\gamma'(t)$ and so is well defined independently of any extension of γ'.

The second method, which suffers from some technicalities, is similar to the proof of the lemma above. We may extend $\gamma'(t)$ to a vector field on all of M. The key step here is to note, that we can choose coordinates $\phi : U \to V$ around a point $x \in \gamma(I)$ so that $\phi(\gamma(I) \cap U) = \{x^2 = \cdots = x^n = 0\} \cap V$. In these coordinates, $\gamma' = \gamma'(x^1)$ and we simply extend it to all of V by $(x^1, \cdots, x^n) \mapsto \gamma'(x^1)$. This extends γ' in a chart, and now we cover $\gamma(I)$ by charts and use a partition of unity to patch it together and then finally a bump function to extend to all of M.

Some technical issues arise in that irregular points of γ ($\gamma' = 0$) do not necessarily have a coordinate neighbourhood of the desired form (e.g. $\gamma(t) = (t, \sqrt{|t|}) \in \mathbb{R}^2$ has a cusp at the origin). But we avoid this problem by noting that $\nabla_{\gamma'}Y = 0$ at such points anyway and the set of irregular points is closed, so we can still cover the complement by open charts. We also have to take some care with self-intersections of γ, but by working on a open set of points $t \in I$ where $\gamma'(t) \neq 0$, we may further restrict to an open set where γ is diffeomorphic with it's image and ignore any other arcs that may intersect the image. Another issue, is that we may have to shrink V a little so that if $(x^1, \cdots, x^n) \in V$, then $(x^1, 0, \cdots, 0) \in V$ also, but this causes no harm to the argument.

Proof. Because, we have a connection, we have Christoffel symbols, Γ^k_{ij} defined locally. We may also write $Y = Y^i(t)\partial_i$ for $Y \in \mathfrak{X}_\gamma$ and $\gamma' = \dot{\gamma}^i\partial_i$. Then define,

$$\nabla_t Y = \frac{dY^i}{dt}\partial_i + Y^i\dot{\gamma}^j\Gamma^k_{ij}(\gamma(t))\partial_k.$$

It's easy to verify the first two assertions by direction calculation: the right hand side is clearly $\mathbb{R}$-linear in Y and the Liebniz rule follows from the product rule applied to $\frac{d(fY^i)}{dt}$.

For the third assertion, for $Y = Y^i(x)\partial_i \in \mathfrak{X}(M)$, we then have $\bar{Y} =$

$Y^i(\gamma(t))\partial_i$ so that,

$$\frac{d\bar{Y}^i}{dt} = \frac{dY^i(\gamma(t))}{dt} = \partial_j(Y^i)\dot{\gamma}^j,$$

and hence,

$$\begin{aligned}\nabla_t\bar{Y} &= \frac{d\bar{Y}^i}{dt}\partial_i + \bar{Y}^i\dot{\gamma}^j\Gamma^k_{ij}(\gamma(t))\partial_k \\ &= \dot{\gamma}^j\partial_j(Y^i)\partial_i + \bar{Y}^i\dot{\gamma}^j\Gamma^k_{ij}(\gamma(t))\partial_k \\ &= \nabla_{\gamma'}Y.\end{aligned}$$

□

Remark 5.7. Similar to the extension of γ' in the remark before the proof, given $\bar{Y} \in \mathfrak{X}_\gamma(M)$, we may also extend $\bar{Y}$ to $Y \in \mathfrak{X}(M)$ such that $Y \circ \gamma = \bar{Y}$. Well, we may do this at least by restricting γ to an open sub-interval of I on which γ is an embedding. Since everything is local, this causes no harm. Later we will develop a systematic way of handling such situations.

Now, by the proposition, $\nabla_{\gamma'}Y = \nabla_t\bar{Y}$, the latter depending only on the values of Y along γ. Thus, for any $X \in \mathfrak{X}(M)$, if we let γ be the integral curve of X though x (i.e. $\gamma(0) = x$ and $\gamma'(t) = X(\gamma(t))$, then at $x = \gamma(0)$,

$$\nabla_X Y = \nabla_t\bar{Y}.$$

That is $(\nabla_X Y)(x)$, rather than depending on Y in a full open neighbourhood of x, in fact only depends on Y along the integral curve of X through x! Essentially, this is just the chain rule:

$$(\partial_t Y^j \circ \gamma)(t) = (\partial_i Y^j)(\gamma(t))\partial_t\gamma^j(t)$$

used in the proof.

Remark 5.8. Just as with the lemma, the same proof applies to sections of a vector bundle E, resulting in the fact that $\nabla_X s$ depends only on s along the integral curves of X. Again, this is just the chain rule and the properties of connections.

5.2 Week 05, Lecture 02: Parallel Transport and Holonomy

We begin this lecture by reviewing the notion of integral curves. Then we look a little more closely at the notion of tensorality.By direct calculation, in the previous lecture we saw that $(\nabla XY)(x)$ depends only on $X(x)$. Here we argue that this is in fact a direct consequence of the definition of the connection and no checking is strictly required!

We shall also see how a connection allows us to identify the fibres of a vector bundle E, at different points by *connecting* the points with a smooth path, which is known as *parallel transport*. Such a construction is dependent on the choice of path and this dependence is measured by the notion of *Holonomy*. We won't treat Holonomy in any detail in this course, but it is important to be aware of it, because it tells us to what extent we may identify different fibres of a bundle. In Euclidean space, we do this all the time, by considering vectors based at a points $x \neq y$ as being based at the origin and then comparing them. In the language of this lecture, we parallel translate all tangent vectors to the origin in Euclidean space. It turns out however, that there are connections on Euclidean space (but not the usual directional derivative!) for which parallel transport has non-trivial Holonomy.

5.2.1 Integral Curves

In the previous lecture we required the existence of *integral curves* of a vector field X. Let us define exactly what is meant by this, and then prove that such curves exist, and are unique.

Definition 5.9. Given a vector field X on M, an integral curve for X, is a smooth curve $\gamma : I \to M$ such that

$$\gamma'(t) = X(\gamma(t)).$$

Lemma 5.10. *For any smooth vector field X, and a point $x \in M$, there exists a unique integral curve for X passing through the point x.*

Proof. The required curve satisfies, $\gamma(0) = x$ and $\gamma'(t) = X(\gamma(t))$. In local coordinates about x this may be expressed as,

$$\begin{cases} (\gamma^i)'(t) & = X^i(\gamma^1(t), \cdots, \gamma^n(t)) \\ \gamma^i(0) & = x^i \end{cases}$$

This is a system of first order ODE's for the functions $\gamma^1, \cdots, \gamma^n$, and since X is smooth, the theory of ODE's asserts the existence and uniqueness of γ in the chart for $t \in (-\epsilon, \epsilon)$ for some $\epsilon > 0$. In an overlapping chart, take a point $y = \gamma(t)$ in the intersection. In the new chart, we obtain a unique solution to the system with initial data based at y and thus may extend γ uniquely into this chart. By uniqueness, this new integral curve agrees with the old one on the overlap.

Proceeding like this over all overlapping charts, we obtain a unique, integral curve through x and defined on a maximal interval I. In general $I \neq \mathbb{R}$. □

5.2.2 Tensorality

In the exercise sheets, you were asked to prove that tensor fields may be identified with $C^\infty(M)$-mulit-linear maps. Lets take a look at exactly what that means and how this relates to the assertion that $(\nabla_X Y)(x)$ depends on $X(x)$ only. By definition, $\nabla Y \in \Gamma(T^*M \otimes TM)$ and so we may think of, ∇_Y as endomorphism of sections of the tangent bundle,

$$\nabla Y : \mathfrak{X}(M) \to \mathfrak{X}(T)$$

that is $C^\infty(M)$-linear since by definition ∇Y is a tensor field. This gives us a covariant derivative,

$$\nabla_X Y = \nabla Y \cdot X \in \mathfrak{X}(M)$$

with the property that

$$\nabla_{fX} Y = f \nabla_X Y$$

for any $f \in C^\infty(M)$ by tensorality of ∇Y.

Why does this imply that $\nabla_X Y$ depends only on $X(x)$? In the previous lecture, we computed in local coordinates and saw from the resulting expression that this was true. But I claim that in fact, **it had to be true**. The reason is very simple, given a section $T \in T^*M \otimes TM$, and a vector field $Z \in \mathfrak{X}(M)$, *by definition*,

$$T(X)(x) = T_x(X(x))!$$

That is, $T_x : T_xM \to T_xM$ and $X(x) \in T_xM$ and the definition of $T(X)(x)$ is exactly obtained by first taking the endomorphism T_x of the fibre $T_x^*M \otimes T_xM$

and applying it to the tangent vector $X(x)$ in the fibre T_xM. This is the definition of how tensors fields act on other tensor fields. Writing in local coordinates shows that this is smooth in x.

What about Y? By the Liebniz rule,

$$\nabla_X(fY) = f\nabla_X Y + df \otimes Y$$

and the second term means that after fixing X, them map

$$Y \mapsto \nabla_X Y$$

is not $C^\infty(M)$-linear in Y! Thus for fixed X, this map does not define a section of $T^*M \otimes TM$.

All this is another way of restating that $(\nabla_X Y)(x)$ *depends only on* $X(x)$*, but depends on* $Y(y)$ *for points* y *near* \(x).

Now recall that a vector v may be though of a double-dual element by $v(\alpha) = \alpha(v)$ for α in the dual. Thus a section, T of $T^*M \otimes TM$ may be thought of as a multi-linear map,

$$T : \Gamma(TM) \otimes \Gamma(T^*M) \to C^\infty(M)$$

In summary then, given an $\mathbb{R}$ multi-linear map $T : TM \times TM \to C^\infty(M)$, if it satisfies,

$$T(fX, Y) = fT(X, Y)$$

the test for tensorality says this is tensorial in X and hence, $T(fX, Y)(x)$ only depends on $X(x)$. If however,

$$T(X, fY) \neq fT(X, Y)$$

then $T(fX, Y)(x)$ depends on Y at points $y \neq x$ (otherwise T would be a tensor and hence $C^\infty(M)$ linear in Y).

How does this work with ∇Y then? This is higher order tensor than above, if we allow Y to vary. That is, we have an $\mathbb{R}$-linear map $Y \in \mathfrak{X}(M) \mapsto \Gamma(T^*M \otimes TM)$ and we can form the $\mathbb{R}$-multi-linear map,

$$(Y, X, \alpha) \mapsto \alpha(\nabla Y \cdot X).$$

This is $C^{infty}(M)$-multi-linear in X, α, but not Y hence does not define a tensor field. *For each* Y*, nablaY is a tensor field certainly, but allowing* Y *to vary does not produce a tensor field because of the Leibniz rule.*

5.2.3 Parallel Transport

On Euclidean space, we have constant vector fields $X(x) = X_0$ where X_0 is a fixed vector. These vector fields satisfy,

$$D_Y X = DX \cdot Y = 0$$

for any other vector field Y. In other words, $DX \equiv 0$ as a section of $T^*\mathbb{R}^n \otimes T\mathbb{R}^n$. The generalisation of the notion of constant vector field to manifolds is given in the following definition, the terminology presumably coming from the fact that a constant vector field in Euclidean space determines a set of parallel lines.

Definition 5.11. Let E be a vector bundle with connection ∇. A section $s \in \Gamma(E)$ is *parallel* if $\nabla s \equiv 0$. Given a curve γ, s is *parallel along* γ if $\nabla_{\gamma'} s = \nabla_t s \equiv 0$.

A parallel section is automatically parallel along every curve. Conversely, if a section is parallel along every curve, it is parallel: just note that for every $v \in T_x M$, there is a curve, γ passing through x with tangent vector v and s is parallel along γ, hence $\nabla s \cdot v = \nabla_t s = 0$. A section may be parallel along some curves, but not parallel along others. A simple example being the vector field $X(x, y) = x\partial_x$ in $\mathbb{R}^2$ which is parallel along any curve of the form $\gamma(t) = (x_0, y(t))$ but not along $\gamma(t) = (t, 0)$.

Another way of looking at this in Euclidean space, is that given any tangent vector $(x_0, v) \in \mathbb{R}^n \times \mathbb{R}^n \simeq T_{x_0}\mathbb{R}^n$, there is a unique parallel vector field X on $\mathbb{R}^n$ such that $X(x_0) = (x_0, v)$, namely $X(x) = (x, v)$. In general we cannot hope to extend a vector field to a parallel vector field on all of a manifold M (for one thing this would give a non-vanishing vector field and we know no such vector fields exist on $\mathbb{S}^2$). However, given a curve γ we can extend a vector $v \in T_{\gamma(t_0)}M$ to a unique, parallel vector field X along γ.

First, it will be convenient to enlarge the range of allowable curves γ to *piecewise-smooth curves.*

Definition 5.12. A continuous curve $\gamma : I \to M$ is *piecewise-smooth* if it is smooth except for a discrete set of points $\{t_k\}$ (i.e. around each point there is an interval $I_k = (t_k - \epsilon_k, t_k + \epsilon_k)$ not containing any of the other t_j and such that $\lim_{t \to t_k} \gamma'(t)$ is defined. In particular, if $I = [a, b]$ a closed, bounded interval, then the set $\{t_k\}$ is finite. These conditions allow corners, where the tangents from either side don't meet up, but not cusps where the tangent becomes infinite.

Proposition 5.13. *Let $\gamma : I \to M$ be a piecewise-smooth curve. Then for any $t_1, t_2 \in I$, there is a unique isomorphism, $P_{t_1,t_2} : T_{\gamma(t_1)}M \to T_{\gamma(t_2)}M$ such that the vector field,*

$$X(t) = P_{t_1,t}(X_1)$$

is the unique vector field parallel along γ and such that $X(t_1) = X_1 \in T_{\gamma(t_1)}M$.

Proof. This result is a result on the existence and uniqueness of solutions to ODE's. Let $x_1 = \gamma(t_1)$ and choose any vector $V \in T_xM$. Set $V_1 = V$.

In a chart containing x, the condition that $X(t) = X^i(t)\partial_i$ is parallel along γ in the chart, and that $X^i(t_0)\partial_i = X(t_0) = V_1 = V_1^i\partial_i$ is the following system of first order ODE's:

$$\begin{cases} \left(\frac{dX^i}{dt} + \Gamma^i_{kl}\dot{\gamma}^k X^l\right)\partial_i & = 0 \\ X^i(t_0) & = V_1^i \end{cases} \tag{5}$$

for the n-functions X^i. It might be easier to understand this system (which contains several instances of the summation convention), if for each fixed i, l we let $\Gamma^i_l(t) = \Gamma^i_{lk}\dot{\gamma}^k$ so that we have the coupled system

$$\frac{d}{dt}X^i = -\Gamma^i_1(t)X^1(t) - \cdots - \Gamma^i_n(t)X^n(t)$$

where the coefficients Γ^i_l depend only on the Christoffel symbols and the curve γ. For each different connection and curve, we get a different system of ODE's.

In any case, the coefficient functions $\Gamma^i_{kl}\dot{\gamma}^k$ are piecewise-continuous and bounded, and the theory of ODE's ensures the existence of a unique solution, $X(t)$ using which, we then define

$$P_{t_1,t}(V_1) = X(t).$$

Now we may cover $\gamma(I)$ by an at most countable number of charts, $\phi_i : U_i \to V_i$, $i = 1, 2, \cdots$, with $U_i \cap U_{i+1} \neq \emptyset$. By the above construction, we have a solution $X_1(t)$ defined in the chart ϕ_1. Choose any point $x_2 = \gamma(t_2) \in U_1 \cap U_2$, and let $V_2 = X_1(t_2)$. Let $X_2(t)$ be the unique parallel vector field along γ in the chart U_2 such that $X_2(t_2) = X_1(t_2)$.

On the overlap, $U_1 \cap U_2$ we have two parallel vector fields along γ (note this is independent of any choice of coordinates, we just *expressed* this requirement above in coordinates but it makes sense without reference to any coordinates) X_1 and X_2 such that $X_1(t_2) = X_2(t_2)$. Hence by uniqueness, $X_1(t) = X_2(t)$ on the overlap and thus we obtain a unique, parallel vector field, X on $U_1 \cup U_2$ with $X(t_1) = V$.

Now proceed by induction to complete the existence and uniqueness of a parallel vector field X, along γ such that $X(t_1) = V$.

To check that the map $P_{t_1,t}$ is is linear, one needs to verify that

$$P_{t_1,t}(aV + bW) = aX(t) + bY(t)$$

where X, Y are the unique, parallel vector fields along γ satisfying $X(t_1) = V$ and $Y(t_1) = W$ (or in other words, $X = P_{t_1,t}V$ and $Y = P_{t_1,t}W$). Certainly, the initial condition is satisfied:

$$(aX + bY)(t_1) = aX(t_1) + bY(t_1) = aV + bW.$$

Moreover, the ODE (5) is linear in X and hence $aX + bY$ also satisfies $\nabla_t(aX + bY) = 0$. By uniqueness of solutions, $aX + bY$ is the unique, parallel vector field along γ such that $(aX + bY)(t_1) = aV + bW$. □

Remark 5.14. The previous proposition implies that a connection determines parallel translation along curves. The converse is also true. As an exercise, prove the formula,

$$\nabla_t X = P_{t_1,t}\left[\frac{d}{dt}\left(P_{t_1,t}^{-1}X(t)\right)\right]$$

for $X \in \mathfrak{X}_\gamma(M)$. *Hint*: For $\{e_i\}$ a basis for γ, define $E_i(t) = P_{t_1,t}e_i$ which are now parallel vector fields along γ that form a basis for the tangent space at each point of γ (since $\{e_1\}$ is a basis and parallel transport is an isomorphism). The vector fields $\{E_i(t)\}$ are called a *parallel frame* along γ. Now we may write $X(t) = X^i(t)E_i(t)$. Differentiate this expression using the left hand side and right hand side above to see they are equal.

5.2.4 Holonomy

Parallel transport as defined above gives a way to identify T_xM with T_yM by joining x to y by a smooth path and parallel transporting along it. That is parallel transport furnishes us with an isomorphism of vector spaces,

$$P_\gamma : T_xM \to T_yM.$$

However, if we use a different path σ to join x to y, then we get another map P_σ and in general,

$$P_\gamma \neq P_\sigma.$$

As remarked earlier, if it were the case that parallel transport is independent of the chosen curve, we could fix a point $x_0 \in M$ and a non-zero vector $v \in T_{x_0}M$ and unambiguously define a vector field

$$X(y) = P_{x,y}(v)$$

where $P_{x,y}$ is parallel transport along any curve joining x to y. Since each parallel transport is an isomorphism, the resulting vector field would be a non-vanishing vector field on all of M. In fact, it's possible to arrange the curves to vary in a smooth way and by smooth dependence on coefficients of ODE's, X will in fact be smooth. But we know that such a construction is not possible on the two-sphere, hence on the two-sphere, **parallel transport depends on the choice of curve**.

We can see this explicitly through the following example.

Example 5.15. Let $M = \mathbb{S}^2$, the two-sphere embedded into $\mathbb{R}^3$ via the *standard embedding* as the unit sphere. We may define a connection on M by,

$$\nabla_X Y = \pi^{TM} D_X Y$$

where $D_X Y$ denotes the directional derivative in $\mathbb{R}^3$, and π^{TM} denotes the projection of a tangent vector in Euclidean space along M to the tangent space of M. Explicitly, if ν denotes the unit, outer normal to the sphere $\nu(x) = (x, x) \in T\mathbb{R}^3 \simeq \mathbb{R}^3 \times \mathbb{R}^3$, then

$$\nabla_X Y = D_X Y - \langle D_X Y, \nu \rangle \nu.$$

It is an exercise to check this defines a connection.

Consider the loop, formed by the three arcs:

$$\begin{aligned}
\gamma_1(t) &= (\sin(t), 0, \cos(t)), 0 \leq t \leq \pi/2 \\
\gamma_2(t) &= (\cos(t - \pi/2), \sin(t - \pi/2), 0), \pi/2 \leq t \leq \pi \\
\gamma_3(t) &= (0, \cos(t - \pi), \sin(t - \pi)), \pi \leq t \leq 3\pi/2.
\end{aligned}$$

The first arc starts at the north pole $(0, 0, 1)$ and follows a great circle to $(1, 0, 0)$. The next follows a great circle to $(0, 1, 0)$ and the last follows a great circle back to the start point.

Notice that this loop is only piecewise smooth, but our constructions above allow for this.

Consider the vector $(0,1,0)$ based at the north pole $(0,0,1)$. The tangent plane at the north pole is parallel to the $z=0$ pane and so this is indeed tangent to the sphere. I claim the vector $X(t)=(0,1,0)$ is in fact tangent to the sphere all along γ_1. In fact at any point $\gamma_1(t)=(\sin(t),0,\cos(t))$ the curve

$$\sigma_t(\theta)=(\sin(t)\cos(\theta\sin^{-1}(t)),\sin(t)\sin(\theta\sin^{-1}(t)),\cos(t))$$

has $\sigma_t(0)=\gamma_1(t)$ and $\sigma_t'(0)=(0,1,0)$. Certainly the vector field $X(t)=(0,1,0)$ satisfies, $D_{\gamma_1'}X=0$ and hence $\nabla_{\gamma'}X=0$ so that X is parallel along γ_1.

Now at the point $(1,0,0)$ we have $X=(0,1,0)$ and we parallel translate it along γ_2. I'll leave it as an exercise to check that along γ_2, $X(t)=(-\sin(t-\pi/2),\cos(t-\pi/2),0)=\gamma_2'(t)$. This follows since the derivative $D_{\gamma_2}X$ is perpendicular to the sphere so that $\nabla'_{\gamma_2}X=0$.

In particular, $X(\pi)=(-1,0,0)$ and now in a similar fashion to γ_1, parallel translation along γ_3 back to the north pole gives $X(3\pi/2)=(-1,0,0)$.

Thus parallel translating around the loop formed by the three curves takes $(0,1,0)$ to $(-1,0,0)$. It's rotation by $\pi/2$ and the result is not what we started with!

Remark 5.16. The above example could be rephrase by parallel translating $(0,1,0)$ at the north pole along γ_1 and then γ_2 to obtain the tangent vector $(-1,0,0)$ at the point $(0,0,1)$. On the other hand, we could parallel translate from the north pole along γ_3 to obtain the tangent vector, $(0,1,0)$ at the point $(0,0,1)$ thus exhibiting dependence of parallel translation on the chosen curve.

The above discussion motivates the definition of Holonomy. Given a loop, γ based at a point $x\in M$, parallel transportation defines an isomorphism $P_\gamma: T_xM\to T_xM$, or in other words $P_\gamma\subset GL_n(\mathbb{R})$.

Definition 5.17. Let E be a vector bundle with connection ∇. For each fixed point $x\in M$ define the holonomy group

$$\mathrm{Hol}_x(\nabla)=\{P_\gamma:\gamma \text{ is a loop based at } x\subset GL_n(\mathbb{R})\}.$$

The group operation is composition of linear maps $T_xM\to T_xM$. The composite,

$$P_\gamma\cdot P_\sigma=P_{\sigma*\gamma}$$

where $\sigma * \gamma$ is the loop based at x obtained by first following σ and then γ. Inverses are,

$$P_\gamma^{-1} = P_{-\gamma}$$

where $-\gamma$ is following γ in the opposite direction. Thus $\mathrm{Hol}_x(\nabla)$ is closed under the group operation of $GL_n(\mathbb{R})$.

Example 5.18. Euclidean space with directional derivative has zero holonomy.

The two-sphere with the connection from the example above has non-zero holonomy as shown in the previous example.

Consider steriographic projection, ϕ mapping the two-sphere minus the point $(0, -1, 0)$ diffeomorphicly onto the plane. The loop $\gamma_1 * \gamma_2 * \gamma_3$ from the example above is contained in the two-sphere minus the point $(0, -1, 0)$ and thus is mapped to a loop in the plane. We can define a connection, ∇^ϕ on the plane by,

$$\nabla_X^\phi Y = d\phi^{-1} \cdot \nabla_{d\phi \cdot X} d\phi \cdot Y$$

where ∇ is the connection on the two-sphere. Since ∇ has non-zero holonomy, so does ∇^ϕ. Thus we obtain a connection on the plane with non-zero holonomy. The usual "constant" vector fields on Euclidean space are not parallel with respect to this connection!

A similar construction may be made by working in polar coordinates on the plane minus the origin. The connection in (r, θ) coordinates is not the directional derivative with respect to these coordinates and hence constant vector fields are not parallel. Does this connection have holonomy?

5.3 Week 05, Lecture 03: Levi-Civita Connection

So far, we have not dealt with the question of whether connections even exist on arbitrary manifolds. Of course, we wouldn't have gone through all the trouble of studying connections if we couldn't produce a sufficiently large set of useful examples of connections on manifolds. This lecture deals with showing the existence and uniqueness of canonical connection on an arbitrary Riemannian manifold with certain desirable properties. This connection is known as the Levi-Civita connection (or sometimes the Riemannian connection). Recalling that every smooth manifold, possess a smooth metric, we see that every smooth manifold possesses a connection.

5.3.1 The Levi-Civita Connection

Theorem 5.19 (The Fundamental Theorem of Riemannian Geometry)**.** *Let (M, g) be a Riemannian manifold. Then there exists a unique connection on TM, called the* Levi-Civita connection *satisfying the following properties,*

- $X(g(Y, Z)) = g(\nabla_X Y, Z) + g(Y, \nabla_X Z)$ *(metric compatability),*
- $\nabla_X Y - \nabla_Y X - [X, Y] = 0$ *(torsion free),*

for all smooth vector fields X, Y, Z.

The term $X(g(Y, Z))$ should be interpreted as X acting as a derivation on the smooth function $x \mapsto g_x(Y(x), Z(x))$.

Proof. For each X, Y define a one-form,

$$\alpha_{X,Y}(Z) = \frac{1}{2}\left[Xg(Y, Z) + Yg(X, Z) - Xg(X, Y)\right.$$
$$\left. - g(X, [Y, Z]) - g(Y, [X, Z]) + g([X, Y], Z)\right). \quad (6)$$

Then we define,

$$\nabla_X Y = \sharp\alpha_{X,Y},$$

that is, for every Z,

$$g(\nabla_X Y, Z) = \alpha_{X,Y}(Z).$$

We must check four points:

1. $\alpha_{X,Y}$ is indeed a one-form,

2. $\nabla_X Y$ is a connection,

3. $\nabla_X Y$ is metric compatible and torsion free,

4. If $\bar{\nabla}$ is a metric compatible, torsion free connection, then $\bar{\nabla} = \nabla$.

There's a repeating theme to the proof so we won't give all the details.

1. We must check that $\alpha_{X,Y}(fZ) = f\alpha_{X,Y}(Z)$ and that $\alpha_{X,Y}(Z_1 + Z_2) = \alpha_{X,Y}(Z_1) + \alpha_{X,Y}(Z_2)$.

2. We need to check that $\nabla_X Y$ is tensorial in X, additive in Y and satisfies the Leibniz rule in Y.

3. We use the formula $g(\nabla_X Y, Z) = \alpha_{X,Y}(Z)$ to check the required properties.

   - Metric compatability:
   - Torsion free:

4. If $\bar{\nabla}$ is metric compatible and torsion free, we'll show that $g(\bar{\nabla}_X Y, Z) = \alpha_{X,Y}(Z)$ for all Z and hence since g is non-degenerate, $\bar{\nabla}_X Y = \nabla_X Y$.

□

5.3.2 Higher Derivatives

Recall that given a connection ∇ and a vector field X, $\nabla X \in T^*M \otimes TM$ and therefore we take it's trace.

Definition 5.20. Let ∇ be a connection on TM and X be a vector field. Define the divergence by,

$$\operatorname{div} X = \mathrm{Tr} \nabla X$$

Definition 5.21. More generally, let $T \in \Gamma(T_q^{p+1}M$ with $p, q \geq 0$ and let ∇ be a connection on $T_q^p M$. Then $\nabla T \in \Gamma(T_{q+1}^{p+1}M$ and define

$$\operatorname{div} T = \mathrm{Tr} \nabla T \in \Gamma(T_q^p M).$$

For $T \in \Gamma(T^p_{q+1}M$ we have two ways to define the divergence, depending on whether we have a connection on $\Gamma(T^{p+1}_q M$ or on $\Gamma(T^p_{q+1}$. In the former case, we define

$$\operatorname{div} T = \operatorname{Tr}(\nabla \sharp T),$$

while in the latter case we define,

$$\operatorname{div} T = \operatorname{Tr}(\sharp \nabla T).$$

Remark 5.22. Note here that for $T \in T^{p+1}_q$, the trace is defined by contracting one of the $p+1$ TM slots with the T^*M slot obtained by applying ∇: $\nabla T \in T^*M \otimes T^{p+1}_q$. By writing T^{p+1}_q we are implicitly suggesting that the trace is over the first TM slot, i.e. by writing $T \in TM \otimes T^p_1$ and $\nabla T \in T^*M \otimes TM \otimes T^p_q$ and contracting on the first two slots. In general we need to specify with TM slot to contract on.

Remark 5.23. In general for $T \in \Gamma^p_{q+1}M$, $\sharp T \in \Gamma^{p+1}_q$ and the first definition above defines $\operatorname{div} T$ to be the divergence of the $(p+1, q)$-tensor $\sharp T$. The second divergence, first applies ∇ to obtain $\nabla T \in T^p_{q+2}$ and then applies $\sharp$ to the second T^*M index:

$$\nabla T \in T^*M \otimes T^p_{q+1}M \simeq T^*M \otimes T^p_0 M \otimes T^*M \otimes T^0_q M$$

and we apply $\sharp$ to the second T^*M occurring and then take a trace.

This can get quite confusing, and it is important to clearly distinguish whether one is using musical isomorphisms or not, and whether one applies the musical isomorphism before taking ∇ or after. Below we will see how one can extend a connection on TM to all of the tensor spaces $T^p_q M$. *The order in which one takes derivatives, musical isomorphisms and contractions is important.* However, as we will see, if ∇ is the Levi-Civita connection (or in fact any metric-compatible connection) then we always get the same answer.

There's a rather useful operation in play behind the scenes here, that of *metric contraction*.

Definition 5.24. Let $T \in T^p_{q+2}$. Then $\sharp T \in T^{p+1}_{q+1}$ and we define the *metric contraction*, or *trace with respect to the metric* of T, by,

$$\operatorname{Tr}_g T = \operatorname{Tr} \sharp T$$

On the other hand, if $T \in T^{p+2}_q$, define,

$$\operatorname{Tr}_g T = \operatorname{Tr} \flat T$$

In this notation, the second approach to the divergence of a T^p_{q+1} tensor may be written,

$$\operatorname{div} T = \operatorname{Tr}\sharp\nabla T = \operatorname{Tr}_g \nabla T.$$

Another rather useful application of musical isomorphisms is to define the *gradient*.

Definition 5.25. Let $T \in T^p_q M$ and define the *gradient*,

$$\operatorname{grad} T = \sharp\nabla T$$

where $\sharp$ is applied to the T^*M factor arising from the connection.

Example 5.26. If f is a function, then $\operatorname{grad} f = \sharp df$ is true since this is equivalent to,

$$df \cdot X = g(\operatorname{grad} f, X)$$

for all vectors X.

Definition 5.27. Let f be a smooth function. The *Hessian* of f is defined to be,

$$\operatorname{Hess} f = \nabla\sharp df.$$

It is a section of $T^*M \otimes TM$.

Define the *Laplacian*,

$$\Delta f = \operatorname{div} \nabla f = \operatorname{Tr}(\nabla\sharp df) = \operatorname{Tr} \operatorname{Hess} f$$

The result is a smooth function.

Remark 5.28. Notice that we could instead, considered ∇df which is a $(0,2)$ tensor, and we could instead take the trace with respect to the metric,

$$\operatorname{Tr}_g \nabla df = \operatorname{Tr}\sharp\nabla df.$$

The result is again a smooth function, but in general not equal to Δf. Again, the order in which we apply $\sharp$ and ∇ is important! The way to remember it, is to remember that $\Delta = \operatorname{div} \operatorname{grad}$. We remarked on this phenomena above, and noted that for the Levi-Civita connection, in fact the order doesn't matter. A little more patience is requested before we prove this claim.

Notice that we have taken two derivatives of f and to extend this operation to tensor fields, we'll need to be able to differentiate more that once. That is, we need connections on all the bundles $T^p_q M$.

Theorem 5.29. *Let ∇ be a connection on a vector bundle E. Then there exists unique connections, ∇_q^p on $E_q^p = \otimes^p E \otimes \otimes^q E^*$ such that,*

1. *$nabla_0^0 f = df$ for all $f \in E_0^0 = C^\infty(M)$,*
2. *$\nabla_0^1 s = \nabla s$ for all $s \in \Gamma(E_0^1) = \Gamma(E)$,*
3. *$\nabla_{q+s}^{p+r}(T \otimes S) = \nabla_q^p T \otimes S + T \times \nabla_s^r$,*
4. *$dT(\alpha_1, \cdots, \alpha_p, s^1, \cdots, s^q) =$*

Remark 5.30. In other words, we start with a connection on E, use df as our connection on E_0^0 and then extend it to arbitrary tensor products of E by requiring that it obeys the Leibniz rule with respect to tensor products (3). To extend it to E^*, we require that it commutes with traces:

$$\nabla \mathrm{Tr} \alpha \otimes s = \mathrm{Tr}\left(\nabla \alpha \otimes s + \alpha \otimes \nabla s\right).$$

This is part (4), and defines $\nabla \alpha$ by the formula,

$$(\nabla \alpha)(s) = \nabla\left(\alpha(s)\right) - \alpha(\nabla s).$$

It's a little hard to understand what's going on here, and is probably easiest seen by using the covariant derivative. Remember $\nabla \alpha \in T^*M \otimes E^*$, $\nabla_X \alpha = \nabla \alpha(X) \in E^*$ and $\nabla(\alpha(S)) \in T^*M$. Then we define,

$$\nabla \alpha(X, s) = (\nabla_X \alpha)(s) = (\nabla(\alpha(s)))(X) - \alpha(\nabla_X s).$$

The formula in 4. is obtained by repeated application of this construction and by 3. Or, by requiring that ∇ commutes with all traces.

Proof. □

Theorem 5.31. *Let ∇ be a connection on TM with (M, g) a Riemannian manifold. The connections ∇_q^p are metric compatible with the metric induced on $T_q^p M$ by g.*

Proof. □

Remark 5.32. The notion of torsion doesn't make sense for the bundles T_q^p so there is no requirement that ∇ be torsion-free here.

6 Week 06

6.1 Week 06, Lecture 01: Sub-Manifolds

In this lecture, we consider maps between manifolds, $f : M^k \to N^n$, and how we may relate the differential and geometric structures of the domain and range via a smooth map. The simplest case is the inclusion of an open set into the ambient manifold. Then, since most of what we have done is local (i.e. it depends only on open sets) the differential structure and geometry is the same as the ambient space. More interesting, but still reasonably simple (from a certain point of view!) is the case of an embedding. The geometry and differential structure are induced from the ambient space, but notice that the image, $f(M)$ won't be an open set of N in general. So a-priori, it's perhaps not clear how the structures relate. The key observation is that there is a coordinate neighbourhood of N such that the image is mapped to $\mathbb{R}^k \subset \mathbb{R}^n$. Relaxing the requirement of embedding to immersion, locally on M this is still true and any local statement for an embedding holds true for an immersion.

6.1.1 Submanifolds

Definition 6.1. A k-dimensional submanifold of Euclidean space is a subset $M \subseteq \mathbb{R}^n$ such that M may be covered by open sets $\{U_\alpha\}$ of $\mathbb{R}^n$ for which there diffeomorphisms $\phi_\alpha : U_\alpha \to V_\alpha$ of U_α onto an open set $V_\alpha \subseteq \mathbb{R}^n$ and such that,

$$\phi_\alpha(M \cap U_\alpha) = V_\alpha \cap \{x^{k+1} = \cdots = x^n = 0\}.$$

We may now use this definition to define general submanifolds.

Definition 6.2. Let N^n be a smooth manifold. A k-dimensional submanifold of N is a subset $M \subset N$ such that for every chart in an atlas for N, $\{\phi_\alpha : U_\alpha \to V_\alpha\}$ of N, we have,

$$\phi_\alpha(U_\alpha) \subset \mathbb{R}^{\ltimes}$$

is a submanifold of Euclidean space.

Remark 6.3. The basic idea is that a subset M should be a submanifold of N precisely when $\{U_\alpha \cap M\}$ gives an atlas for M where $\{U_\alpha\}$ is an atlas for N. The difficulty is what? I don't think there is any difficulty!

Now let $f : M \to N$ be a smooth map.

Lemma 6.4. *The smooth map, $f : M \to N$ is an embedding if and only if $f(M)$ is a submanifold, and f is a diffeomorphism with respect to this smooth structure.*

Proof. Essentially this is the definition. Details left as an exercise. □

Lemma 6.5. *Any smooth immersion $f : M \to N$ is a local embedding, in the sense that there is a cover of M by open sets U_α such that for each α, $f|_{U_\alpha}$ is an embedding. In particular, this implies that $f(U_\alpha)$ is a submanifold.*

Proof. Implicit function theorem. □

Remark 6.6. Sometimes the terminology, *embedded submanifold* and *immersed submanifold* is used, to mean the image of an embedding or an immersion respectively. Any submanifold defined above is an embedded submanifold, with embedding given by the inclusion as a subset $\iota : M \to N$.

Let us establish some notation to be used throughout this lecture. Let M be a smooth manifold, let $(\bar{M}, \bar{g}$ be a Riemannian manifold, and let $F : M \to \bar{M}$ an immersion. For convenience, let us also write $\bar{X} = dF \cdot X$ for any $X \in \mathfrak{X}(M)$. We define a *vector field along F* to be an element of the set

$$\mathfrak{X}_F(M) = \{\bar{X} : M \to T\bar{M} | \pi \circ \bar{X} = F(x)\}.$$

Then we have that, $\bar{X} = dF \cdot X \in \mathfrak{X}_F(\bar{M})$ for any $X \in \mathfrak{X}(\mathfrak{M})$.

Definition 6.7. The *pull-back* (or *induced*) metric $g = F^*\bar{g}$ on TM, is defined by,

$$g(X, Y) = \bar{g}(\bar{X}, \bar{Y}) = \bar{g}(dF \cdot X, dF \cdot Y)$$

for any vector fields $X, Y \in \mathfrak{X}(M)$.

Proposition 6.8. *The pull-back metric is a metric.*

Proof. Symmetry follows immediately from the symmetry of $\bar{g}$. The local coordinate formula,

$$dF \cdot X(x) = (\partial_i F^j(x) X^i(x))\bar{\partial}_j$$

shows that

$$dF \cdot fX = f(x) dF \cdot X.$$

That is, dF is tensorial (more on this below). Then $C^\infty(M)$-bilinearity of g follows immediately from $C^\infty(\bar{M})$-bilinearity of $\bar{g}$.

We obtain smoothness of g from smoothness of $\bar{g}$ and F, but also can see this from the local coordinate expression,

$$g_{kl} = g(\partial_k, \partial_l) = g((\partial_i F^j \delta^i_k)\bar{\partial}_j, (\partial_m F^n \delta^m_l)\bar{\partial}_n = \partial_k F^j \partial_l F^n \bar{g}_{jn}$$

where here $\bar{\partial}_i$ denotes a coordinate vector field on $\bar{M}$ and not $dF \cdot \partial_i$.

Positive-definiteness of g follows from positive-definiteness of $\bar{g}$ and the fact that F is an immersion so that $dF \cdot X = 0$ if and only if $X = 0$. □

Remark 6.9. Note that in general, for $X \in \mathfrak{X}(M)$, while it is true that $\bar{X} \in \mathfrak{X}_F(M)$ is a vector field along F, it may not be possible to extend $\bar{X}$ to a vector field on all of $\bar{M}$ (or even locally on $\bar{M}$). For example, this happens if $F(x) = F(y)$ for $x \neq y$ but $dF \cdot X(x) \neq dF \cdot X(y)$. Can you think of an example of this behaviour?

Suppose now that $\bar{M}$ come equipped also with a connection $\bar{\nabla}$. Can we use this to induce a connection on M even though we may not be able to extend vector fields along F to $\bar{M}$? The answer to this question is the same as for differentiating vector fields along curves. There is another problem, in that $\bar{\nabla}_{\bar{X}}\bar{Y}$ will not in general be tangent to M, i.e. will not be equal to $dF \cdot Z$ for some vector field $Z \in \mathfrak{X}(M)$. We deal with this in the same way as we dealt with this issue for regular surfaces, by projecting onto the tangent space.

Since F is an immersion, by definition, dF is injection and hence for each $x \in M$, $dF_x(T_xM) \subset T_{F(x)}\bar{M}$ is a k-dimensional vector subspace. Let $\pi^{TM} : T_{F(x)}\bar{M} \to T_{F(x)}\bar{M}$ denote the orthogonal (with respect to $\bar{g}$) projection onto this subspace.

If $\{X_i\}$ is a basis for T_xM, then $\bar{X}_i$ is a basis for $dF_x(T_xM$ and we may complete it by adding $\{\nu_1, \cdots, \nu_{n-k}\}$ to a basis for $T_{F(x)}\bar{M}$. Then we may write,

$$\pi^{TM}(\bar{Y}) = \bar{Y} - \sum_{i=1}^{n-k} \bar{g}(\bar{Y}, nu_i)\nu_i.$$

Definition 6.10. Let $\bar{\nabla}$ be a connection on $T\bar{M}$. Define the *pull-back*, or *induced* connection *nabla*, on TM by

$$\nabla_X Y = \pi^{TM}\left(\bar{\nabla}_{\bar{X}}\bar{Y}\right)$$

Proposition 6.11. *∇ is a connection on TM. Moreover, if $\bar{\nabla}$ is the Levi-Civita connection for $\bar{g}$, then ∇ is the Levi-Civita connection for g.*

Proof. The proof that ∇ is a connection is essentially the same as for regular surfaces, and is left as an exercise.

For metric compatibility, we have,

$$\begin{aligned} Xg(Y,Z) &= \bar{X}\bar{g}(\bar{X},\bar{Z}) \\ &= \bar{g}(\bar{\nabla}_{\bar{X}}\bar{Y},\bar{Z}) + \bar{g}(\bar{Y},\bar{\nabla}_{\bar{X}}\bar{Z}) \\ &= \bar{g}(\pi^{TM}(\bar{\nabla}_{\bar{X}}\bar{Y}),\bar{Z}) + \bar{g}(\bar{Y},\pi^{TM}(\bar{\nabla}_{\bar{X}}\bar{Z})) \\ &= g(\nabla_X Y),Z) + g(Y,\nabla_X Z)). \end{aligned}$$

The second to last line is since $\bar{Z}$ lies in $dF(TM)$ and so is orthogonal to any normal component of $\bar{\nabla}_{\bar{X}}\bar{Y}$.

Another exercise is to show that,

$$[dF \cdot X, dF \cdot Y](F(x)) = dF_x \cdot ([X,Y](x)).$$

This can be shown using the local expression for $dF \cdot X$ to show that for $\bar{X} = dF \cdot X$, $\bar{Y} = dF \cdot Y$, we have

$$[\bar{X}(\bar{f})](F(x)) = [X(\bar{f} \circ F)](x).$$

With this in hand, we may prove the connection is torsion-free:

$$\nabla_X Y - \nabla_Y X = \pi^{TM}\left(\bar{\nabla}_{\bar{X}}\bar{Y} - \bar{\nabla}_{\bar{Y}}\bar{X}\right) = \pi^{TM}\left([\bar{X},\bar{Y}]\right) = [X,Y].$$

□

Let us record here a rather useful notion.

Definition 6.12. A diffeomorphism, $F : (M,g) \to (\bar{M},\bar{g})$ is called an *isometry* (or *Riemannian isomorphism*) if $F^*\bar{g} = g$. In particular, if we are given a metric on M, the unique choice of $g = F^*\bar{g}$ makes F an isometry. If F is just an immersion such that $F^*\bar{g} = g$, then we say F is an *isometric immersion.* Finally, a *local isometry* is a map $F : U \to \bar{M}$ for $U \subseteq M$ and open set that is isometric with it's image $F(U)$.

Isometries are the isomorphisms of differential geometry. That is we may identify isometric manifolds by means of the equivalence relation $(M,g) \sim (\bar{M},\bar{g})$ if there exists an isometry, $F : (M,g) \to (\bar{M},\bar{g})$.

6.2 Week 06, Lecture 02: Pull-backs Bundles

The sub-manifold geometry of the previous lecture relied on the fact that F is an immersion. In the lecture, we relax the assumption that F be an immersion and study arbitrary smooth maps $F : M \to N$. In general, things are not so nice, but still we may develop a good theory framed in terms of pull-back bundles. Not everything we did with immersion extends to this level of generality however. In particular, the pull back metric, well defined though it is as a symmetric bilinear form, is not positive-definite. Consider the example of an irregular curve, e.g. $\gamma(t) = (t^2, t^3)$ which suffers from the drawback that $\gamma'(0) = 0$, at which point it has a cusp. In higher dimensions, the map $(r, \theta) \mapsto (r\cos\theta, r\sin\theta, r)$ which has non-injective differential at $(r, \theta) = (0, 0)$. The image is a cone, suffering this time from a vertex at the origin. Keep these examples in mind as we develop the theory in this lecture - in particular out theory applies to such irregular curves and surfaces.

The pull-back bundle can be thought of as a way to view $\bar{M}$ from the perspective of M. For example, us terrestrial beings may look out into space, off the earth and into the surrounding three dimensional space. Certain aspects of the previous section carry over to the more general situation of an arbitrary smooth map $F : M \to \bar{M}$, i.e. not necessarily an immersion. There are two crucial points to note: 1. that $dF \cdot X$ will not in general extend to a vector field on $\bar{M}$, but it does define a section of the pull-back bundle, and 2. even at points where dF does not have maximal rank, and in particular if the rank of dF changes from point to point, the pull-back bundle is perfectly well defined. The degeneracy is recorded in dF and not in the bundle structure. Furthermore, it turns out that we do not even require a metric to define the normal bundle, but that a metric allows us to canonically choose a structure for the normal bundle.

6.2.1 Pull Back Bundles

Definition 6.13. Let $F : M \to \bar{M}$ be a smooth map, and let $\bar{\pi} : \bar{E}^k \to \bar{M}$ be a smooth vector bundle. Let $\{\bar{U}_\alpha\}$ be an open cover by trivialisations with transition maps, $\bar{\tau}_{\alpha\beta} = \mathrm{Id}_\alpha \times \bar{A}_{\alpha\beta}$ with $\bar{A}_{\alpha\beta} \in C^\infty(\bar{U}_\alpha \cap \bar{U}_\beta) \to GL_k(\mathbb{R})$. Define the *pull-back bundle* by applying the vector bundle gluing lemma as follows:

Let $U_\alpha = F^{-1}[\bar{U}_\alpha]$ and define transition maps,

$$A_{\alpha\beta} = \bar{A}_{\alpha\beta} \circ F : U_\alpha \cap U_\beta \to GL_k(\mathbb{R}).$$

Then all the required conditions on $A_{\alpha\beta}$ (e.g. co-cycle condition) are automatically satisfied because they are satisfied for $\bar{A}_{\alpha\beta}$! Hence there is a well defined, rank k vector bundle $\pi : E \to M$, called the *pull-back bundle*, with $A_{\alpha\beta}$ as transition maps. Typically this bundle is written,

$$F^*\pi : F^*\bar{E} \to M.$$

Definition 6.14. Let $\bar{s} \in \Gamma(\bar{E}$ be a section of $\bar{E}$, and define the *pull back section*,

$$s(x) = F^*\bar{s}(x) = \bar{s} \circ F(x).$$

This defines $F^* \in \Gamma(F^*\bar{E})$ as section of the pull-back bundle.

Remark 6.15. In general, not all sections of the pull back are pulled-back sections! For example, if $F(x) = F(y)$, but $x \neq y$, then there is a section s of the pull back bundle with $s(x) \neq s(y)$. To construct such a section, use bump functions in disjoint open neighbourhoods of x and y! Similarly, let $z = F(x)$ and suppose there is an open neighbourhood U of x, such for all open sets $\bar{U} \subset \bar{M}$ containing z, there exists $y \in U$ with $F(y) \in \bar{U}$. Then a similar issue arises. Both of these examples may occur even in the case of an immersion.

Note that in general, we have $\Gamma(F^*\bar{E}) = \Gamma_F(\bar{E}) = \{s : M \to \bar{E} | \pi \circ s = F\}$, and we refer to the latter set as the set of sections of $\bar{E}$ along F. In particular, in an induced trivialisation over $U_\alpha = F^{-1}[\bar{U}_\alpha]$ we have sections $\bar{e}_i \in \Gamma(\bar{E}, \bar{U}_\alpha)$ and these pull back to sections $e_i = F^*\bar{s}_i \in \Gamma(F^*\bar{E}, U_\alpha)$ by

$$e_i(x) = \bar{e}_i \circ F(x).$$

By definition, the sections e_i are a local frame for $F^*\bar{E}$, i.e. $\{e_i(x)\}$ is a basis for the fibre $(F^*E)_x$ at each $x \in U_\alpha$. Thus any local section $s \in \Gamma(F^*\bar{E}, U_\alpha)$ may be written,

$$s(x) = s^i(x)e_i(x)$$

But note that if there are points $x \neq y$ with $z = F(x) = F(y)$, then in general, $s^i(x) \neq s^i(y)$ and in this case, s cannot be extended to a section on $\bar{E}$ by $\bar{s}(z) = s^i(F(x))\bar{e}_i(z)$ where $z = F(x)$.

Remark 6.16. In the case of an embedding we may be rather more direct with our definition. In that case,

$$E = \bar{\pi}^{-1}(F(M))$$

describing E as a submanifold of $\bar{E}$. To check it's a submanifold one needs to use the fact that $F(M)$ is a submanifold of $\bar{M}$ and that $\bar{\pi}$ is a submersion. The projection is then,

$$\pi(z) = F^{-1}(\bar{\pi}(z)).$$

Local trivialisations $\phi_\alpha : \pi^{-1}[U_\alpha] \to U_\alpha \times \mathbb{R}^k$ are,

$$\phi_\alpha(z) = \bar{\phi}_\alpha(z).$$

Noting that F is an embedding, $U_\alpha = F^{-1}[\bar{U}_\alpha]$, and that $\pi = F^{-1} \circ \bar{\pi}$, it's not hard to verify the axioms of a vector bundle and that E is isomorphic to $F^*T\bar{M}$.

In the case of an immersion, we must take a little more care, because either F is not one-to-one with it's image, or F is not diffeomorphic to it's image. Essentially $U_\alpha = F^{-1}[\bar{U}_\alpha]$ cannot be guaranteed to be connected even if $\bar{U}_\alpha$ is connected. However, if we perform essentially the same construction as for embeddings above on each of the connected components of U_α then glue together everything works out and again E so constructed is isomorphic to $F^*T\bar{M}$.

As an exercise, prove that if $F : M \to \bar{M}$ is any smooth map (not necessarily an immersion) and $X \in \mathfrak{X}(M)$, then $\bar{X}(x) = dF_x \cdot X(x)$ is a well defined, smooth section of $F^*T\bar{M}$. This result may be phrased by saying that $dF : TM \to F^*T\bar{M}$ is bundle morphism. Convince yourself of this fact! In the case of an immersion, show also that the bundle morphism $TM \to F^*T\bar{M}$ is injective on each of the fibres. Typically however, **there is not** a morphism $F^*T\bar{M} \to T\bar{M}$.

In the case that F is an embedding however, show that for $\bar{X} \in \Gamma(F^*T\bar{M}$, $Y(y) = \bar{X}(F^{-1}(y))$ defines a vector field along $F(M)$ in $T\bar{M}$, and moreover there exists a vector field $\bar{Y} \in \mathfrak{X}(\bar{M})$ such that $\bar{Y}(y) = Y(y)$ for all $y \in F(M)$. *Hint*: for this last part, reduce to the case of $M = \mathbb{R}^k \subset \mathbb{R}^n = \bar{M}$. Thus for an embedding, there **is** a morphism $F^*T\bar{M} \to T\bar{M}$ and composing with the morphism $F^*T\bar{M} \to T\bar{M}$ we obtain an injective morphism, $TM \to T\bar{M}$.

Definition 6.17. Let $F : M \to \bar{M}$ be an immersion. Then by the exercise, there is an injective bundle map $TM \to F^*T\bar{M}$ and hence we may identify TM with the sub-bundle, $dF \cdot TM$ and then form the quotient bundle, $NM = F^*T\bar{M}/TM$. This bundle may be constructed by the vector bundle gluing lemma, where the transition maps are the induced quotient maps $\mathbb{R}^{\bar{n}}/\mathbb{R}^n \to \mathbb{R}^{\bar{n}}/\mathbb{R}^n$ where the numerator is a trivialisation on $\bar{M}$ and then denominator is a trivialisation of M thought of as a subspace of the numerator via the map dF. Thus the transition maps are elements of $GL_{\bar{n}-n}(\mathbb{R})$ and so NM has rank $\bar{n}-n$.

Remark 6.18. There are infinitely many ways of realising the quotient, these correspond precisely to the choices of complementary subspace in $T\bar{M}$ to TM. For example, in $\mathbb{R}^2$ consider the subspace $V = \{y = 0\}$. Choose any vector $v = (x, y))$ with $y \neq 0$, then $\mathbb{R}^2 \simeq V \oplus \mathbb{R}v$. But if $T\bar{M}$ is equipped with a metric, then we can make a canonical choice of complementary subspace; the orthogonal complement of TM. In the example, this corresponds to choosing $v = (0, 1)$ when the metric is the standard Euclidean inner-product.

Lastly, in an effort to generalise the notion of the induced connection of an immersion, we define the pull-back connection of an arbitrary smooth map. This is a connection on the pull-back bundle however and not on the tangent bundle of M.

Definition 6.19. Let $\pi : E \to \bar{M}$ be a smooth vector bundle equipped with a smooth connection, $\bar{\nabla}$. Let $F : M \to \bar{M}$ be a smooth map. For a section of the form, $s(x) = \bar{s}(F(x)) \in F^*\bar{E}$ where $\bar{s} \in \Gamma(\bar{E}$, and a vector field $X \in \mathfrak{X}(M)$, we define

$$\left(F^*\nabla_X s\right)(x) = \left(\bar{\nabla}_{dF\cdot X}\bar{s}\right)(F(x)).$$

Let us also write $\nabla_X s$ for $F^*\bar{\nabla}_X s$. Since, locally any section may be written as a linear combination of $e_i(x) = \bar{e}_i \circ F(x)$, we may apply the Liebniz rule and define for arbitrary sections $s = s^i e_i$,

$$\nabla_X s = \nabla_X(s^i e_i) = X(s^i)e_i + s^i \nabla_X e_i.$$

This definition is independent of the choice of trivialisations and local frame, since $\bar{\nabla}$ is defined independently of any local trivialisations, hence $\bar{\nabla}$ tranforms correctly. Then notice that $\nabla_X e_i = \bar{\nabla}_{dF\cdot X}\bar{e}_i$ also transforms correctly.

Remark 6.20. Using the pull-back metric, and given an immersion $F : M \to \bar{M}$, the induced connection defined above is just the projection of the pull-back connection onto $dF(TM)$.

6.3 Week 06, Lecture 03: The Second Fundamental Form

6.3.1 The Second Fundamental Form

In this section, we assume that F is an immersion. We have a metric $\bar{g}$ on $F^*T\bar{M}$ equipped with the pull-back connection, $\bar{\nabla}$ whose tangential projection is the Levi-Civita connection for g the induced metric on TM. We also have $\bar{g}$-orthogonal sub-bundles, TM and NM, so that we may write,

$$F^*T\bar{M} = TM \oplus NM.$$

The normal part of $\bar{\nabla}_{\bar{X}}\bar{Y}$ also has a name.

Definition 6.21. The second fundamental form, h of an immersion F is,

$$h(X, Y) = \bar{\nabla}_{\bar{X}}\bar{Y} - \pi^{TM}(\bar{\nabla}_{\bar{X}}\bar{Y}).$$

Let us write N_xM for the orthogonal complement of $dF_x(T_xM)$ with respect to the metric $\bar{g}$ on $\bar{M}$. Thus we may write,

$$T_{F(x)}\bar{M} \simeq T_xM \oplus N_xM.$$

In fact, working in local trivialisations, one can see that there is a natural bundle NM on M known as the normal bundle. It is just the orthogonal complement of TM in $T\bar{M}$

Then the second fundamental form is

$$h(X, Y) = \pi^{NM}(\bar{\nabla}_{\bar{X}}\bar{Y}).$$

Definition 6.22. For any normal vector $\nu \in N_xM$, define a bilinear form,

$$A_\nu(X, Y) = \bar{g}(h(X, Y), \nu)$$

In the particular, case of an oriented hypersurface ($M^n \to \bar{M}^{n+1}$, we may choose a single unit normal vector field ν, and thus simply define,

$$A(X, Y) = \bar{g}(h(X, Y), \nu).$$

So A_ν projects h onto a one-dimensional subspace of NM. In the particular case of a hypersurface, NM is already one dimensional and so via A we can think of h as being a scalar. How does this relate to metric duality?

Definition 6.23. For each section $\nu \in \Gamma(NM)$, define the *Weingarten map* $\mathcal{W}_\nu : TM \to TM$ in the direction ν by

$$\mathcal{W}(X) = -\pi^{TM}\left(\bar{\nabla}_{\bar{X}}\nu\right).$$

Again, for hypersurfaces, there is only one Weingarten map.

Remark 6.24. Note that the definition makes sense because to differentiate ν using $\bar{\nabla}$, we only need to know it's values along the integral curves of $\bar{X}$. But since $\bar{X}$ is tangential to M, the integral curves lie entirely on M and so we do indeed know the value of ν along the integral curves of $\bar{X}$ (since we actually know the values of ν on all of M!).

For an oriented hypersurface, choosing a unit length normal vector field ν ($\bar{g}(\nu,\nu) = 1$, then automatically $\bar{\nabla}_{\bar{X}}\nu \in TM$ because, by metric compatability,

$$0 = \bar{\nabla}_{\bar{X}}\bar{g}(\nu,\nu) = 2\bar{g}(\bar{\nabla}_{\bar{X}}\nu,\nu).$$

Proposition 6.25. *The second fundamental form and the Wiengarten map are related by,*

$$A_\nu(X,Y) = \bar{g}(\bar{\mathcal{W}_\nu}\cdot X,\bar{Y}).$$

Proof. Since $\bar{g}(\nu,\bar{Y}) = 0$ for any $\bar{Y}$, using metric compatability we have,

$$0 = \bar{\nabla}_{\bar{X}}\bar{g}(\nu,\bar{Y}) = \bar{g}(\bar{\nabla}_{\bar{X}}\nu\cdot X,\bar{Y}) + \bar{g}(\nu,\bar{\nabla}_{\bar{X}}\bar{Y}) = -\bar{g}(\bar{\mathcal{W}_\nu}\cdot X,\bar{Y}) + A_\nu(X,Y).$$

□

Remark 6.26. Another way of phrasing the previous proposition, is to note that A_ν is a symmetric bilinear form, and hence there exists a self-adjoint map $T : TM \to TM$ such that $g(T\cdot X,Y) = A_\nu(X,Y)$. The proposition says that this self adjoint map is in fact the Weingarten map $\mathcal{W}_\nu$.

Proposition 6.27. *The second fundamental form is symmetric: $h(X,Y) = h(Y,X)$.*

Proof. From the previous proposition, and metric compatability,

$$A_\nu(X,Y) = -\bar{g}(\bar{\nabla}_{\bar{X}}\nu,\bar{Y}) = -\bar{\nabla}_{\bar{X}}\bar{g}(\nu,\bar{Y}) + \bar{g}(\nu,\bar{\nabla}_{\bar{X}}\bar{Y}) = \bar{g}(\nu,\bar{\nabla}_{\bar{X}}\bar{Y}).$$

Since $\bar{\nabla}$ is torsion free, and since $[\bar{X},\bar{Y}]$ is tangent to M,

$$\bar{g}(\nu,\bar{\nabla}_{\bar{X}}\bar{Y}) = \bar{g}(\nu,\bar{\nabla}_{\bar{Y}}\bar{X}) + \bar{g}(\nu,[\bar{X},\bar{Y}]) = \bar{g}(\nu,\bar{\nabla}_{\bar{Y}}\bar{X}) = A_\nu(Y,X).$$

□

Proposition 6.28. *The second fundamental form is tensorial in both arguments.*

Proof. Tensorality in the first argument follows since $\bar{\nabla}_{\bar{X}}\bar{Y}$ is tensorial in $\bar{X}$. Tensorality in the second argument follows by symmetry! $\square$

Thus we may think of h as a bilinear map $TM \times TM \to NM$, and because of tensorality, we realise h is as a section of $\Gamma(T^*M \otimes T^*M \otimes NM)$. You may encounter the terminology, *NM valued bilinear forms* to describe such sections.

7 Week 07

7.1 Week 07, Lecture 01: The Metric Structure of a Riemannian Manifold

In this lecture, we finally get to the metric space structure of a Riemannian manifold, (M, g). We deal with how to define the length of path, and how this induces a metric, $d = d_g$ (in the sense of metric spaces!) on M. Now that (M, d) is a metric space, we have a topology. But M already had a topology. We'll see that these topologies are in fact the same which is reassuring. We will also encounter the notion of geodesics. We will define such curves in terms of parallel vector fields, generalising the straight lines of Euclidean space. Later will see that in fact these curves minimise length (at least locally), and hence generalise the notion of straight lines in this way too.

Throughout this lecture, (M, g) will denote a smooth, Riemannian manifold.

7.1.1 Length of paths

Let $\gamma : I \to M$ be a smooth path with I an interval. Here smooth means γ extends to a smooth map on an open interval J containing I. If I is open, then we may take $J = I$ and so this definition is about dealing with closed and half closed intervals. It says that the derivative \(γ') is defined at the end points.

Definition 7.1. The length of γ is defined to be,

$$L[\gamma] = L_g[\gamma] = \int_I \sqrt{g(\gamma', \gamma')}.$$

For the integrand, we will also use the notation,

$$|\gamma'| = |\gamma'|_g = \sqrt{g(\gamma', \gamma')}.$$

It will be of great utility to generalise the notion of length to piece-wise smooth paths. Recall, these are continuous maps $\gamma : [a, b] \to M$ for which there exists a partition, $a = t_0 < t_1 \cdots < t_N = b$ such that $\gamma|_{[t_k, t_{k+1}]}$ is smooth as defined above.

Definition 7.2. Let γ be a piecewise smooth curve and let $\gamma_k = \gamma|_{[t_{k-1},t_k]}$, for $k = 1, \cdots, N$. Define the length of γ by,

$$L[\gamma] = \sum_{k=1}^{N} L[\gamma_k].$$

Remark 7.3. Let $\phi : [c,d] \to [a,b]$ be a diffeomorphism. Then define $\bar{\gamma} = \gamma \circ \phi : [c,d] \to M$. Then it is an exercise to show that $L[\bar{\gamma}] = L[\gamma]$. This is essentially the same calculation as for curves in Euclidean space, but now we have a (possibly varying) metric g. If you know the implicit function theorem, then it's possible to show that for $\gamma : [a,b] \to M$ smooth and regular ($\gamma' \neq 0$), given any other regular parametrisation, $\bar{\gamma} : [c,d] \to M$ with $\gamma([a,b] = \bar{\gamma}([c,d])$, there exists such a ϕ.

7.1.2 Riemannian Distance

Definition 7.4. Let (M, g) be a connected, Riemannian manifold. Then for any $x, y \in M$, define,

$$d(x,y) = d_g(x,y) = \inf\{L[\gamma] : \gamma(a) = x, \gamma(b) = y\}$$

where the infimum is taken over all piece-wise smooth curves.

Remark 7.5. The set over which the infimum is taken, is non empty. Fix a point x and consider the set of all points y such that there exists a piecewise smooth curve joining x to y. This set is open since in any chart about y, we can extend any curve joining x to y to any other point in the chart. It is closed, because the complement is open, which can be shown by similar reasoning applied to a point y that cannot be connected to x by a piece-wise smooth curve.

The g in d_g signifies the fact that length of paths depends on the metric, and hence the distance also depends on the metric.

Proposition 7.6. *The pair, (M, d) is a metric space.*

Proof. We need to verify the axioms for a metric space,

1. $d(x,y) \geq 0$,
2. $d(x,y) = 0 \Rightarrow x = y$,

3. $d(x, y) = d(y, x)$,

4. $d(x, z) \leq d(x, y) + d(y, z)$.

Let us take these in turn.

1. This is clear since, $L[\gamma] \geq 0$ for any curve γ, the integrand in the definition of L being non-negative.

2. Suppose that $x \neq y$, and let γ be any curve joining x to y. We need to show that $L[\gamma] \geq C > 0$ for a constant independent of γ. Working in $V \subset \mathbb{R}^n$ for a chart $\phi : U \to V$ around x, let $z = \phi(x) \in V$ and choose $\epsilon > 0$ so that the closed ball of radius ϵ, $\bar{B}_\epsilon(z) \subseteq V$ and $y \notin \bar{B}_\epsilon$. Now, we may think of $g : V \to GL_n(\mathbb{R})$ with g a positive definite, symmetric matrix, hence diagonalisable with strictly positive eigenvalues, λ. The function $\lambda_{\min}(p)$ on V giving the minimum eigenvalue of g at p is continuous, and strictly positive, hence has a positive lower bound, $m > 0$ on $\bar{B}_\epsilon$. Therefore,

$$g(u, u) \geq m|u|^2$$

where $|u|$ denotes the Euclidean norm. Then any curve, γ joining x to y must contain an arc, σ from z to a point in the boundary ∂B_ϵ. Then,

$$L[\gamma] \geq L[\sigma] = \int_a^b \sqrt{g(\sigma', \sigma')} \geq m \int_a^b |\sigma'| \geq m d_{\mathbb{R}^n}(\sigma(a), \sigma(b)) \geq m\epsilon.$$

Letting $C = m\epsilon > 0$ does the job.

3. Any curve joining x to y may traversed backwards to obtain a curve joining y to x.

4. Let $x, y, z \in M$ and consider any pair of curves, $\mu : [a, b] \to M$, $\sigma : [c, d] \to M$ with $\mu(a) = x$, $\mu(b) = \sigma(c) = y$ and $\sigma(d) = z$. Define the concatenation, $\gamma = \mu * \sigma$, by

$$\gamma(t) = \begin{cases} \mu(t), & t \in [a, b] \\ \sigma(t - b + c), & t \in [b, b + d - c] \end{cases}$$

Then γ is piece-wise smooth, joining x to z, and $L[\gamma] = L[\mu] + L[\sigma]$. Now for any $\epsilon > 0$ choose μ, σ so that,

$$L[\mu] \leq d(x, y) + \epsilon, \quad L[\sigma] \leq d(y, z) + \epsilon.$$

Then we have for any $\epsilon > 0$,

$$d(x, z) \leq L[\mu * \sigma] = L[\mu] + L[\sigma] \leq d(x, y) + d(y, z) + 2\epsilon$$

and the result follows by taking $\epsilon \to 0$.

□

Theorem 7.7. *The metric topology induced by d_g is the same as the original topology on M.*

Proof. It suffices to prove that for any open set $U \subseteq M$, and any $x \in U$, there exists a metric ball,

$$B_r(x) = \{y \in M : d_g(x, y) < r\}$$

contained in U and conversely, that for any metric ball $B_r(y)$, and any $x \in B_r(y)$, there exists an open set $U \subseteq B_r(y)$ with $x \in U$. Actually, we can get away with proving less than this, but it doesn't substantially make the proof any easier.

So, let U be an open set and $x \in U$. Consider a chart $\phi : W \to V$ about x. By restricting the chart, we assume that $W \subseteq U$ and that V is a bounded open set in Euclidean space. Arguing as previously, but this time using the supremum over V, of the maximum eigenvalue of g_x,

$$M = \sup\{g_x(u, u) : x \in V, |u| = 1\}$$

we have that $g(u, u) \leq M|u|^2$. Choose $\epsilon > 0$ sufficiently small so that $\{y \in \mathbb{R}^n : d_{\mathbb{R}^n}(x, y) < \epsilon\} \subseteq V$. Then,

$$B_\epsilon^{\mathbb{R}^n} = \phi^{-1}[\{y \in \mathbb{R}^n : d_{\mathbb{R}^n}(x, y) < \epsilon\}] \subseteq W \subseteq U.$$

Choosing $r = \epsilon/M$ we obtain,

$$B_r(x) = \{y \in M : d_g(x, y) < r\} \subseteq B_\epsilon^{\mathbb{R}^n}(x) \subseteq U.$$

For the other inclusion, let $B = B_r(y)$ be a metric ball and let $xinB$. This time in the chart we define

$$m = \inf\{g_x(u, u) : x \in V, |u| = 1\}$$

to be the smallest eigenvalue so that $g(u, u) \geq m|u|^2$ to obtain,

$$U = B^{\mathbb{R}^n}_{\epsilon/m}(x) \subseteq B_\epsilon(x)$$

where $\epsilon > 0$ is chosen small enough so that $B_\epsilon(x) \subseteq B_r(y)$. Remember we can always do this in a metric space by the triangle inequality. Note that U is indeed open being the inverse image under the continuous map ϕ of an open ball in Euclidean space, and certainly $x \in U$. $\square$

7.2 Week 07, Lecture 02: Geodesics

Definition 7.8. A *smooth variation*, γ_s of a curve $\gamma : [a, b] \to M$ is a smooth map,

$$F : [a, b] \times (-\epsilon, \epsilon) \to M$$

for some $\epsilon > 0$ and such that,

$$F(t, 0) = \gamma(t)$$

for all $t \in [a, b]$.

The variation, F *fixes the endpoints*, if

$$F(a, s) = F(a, 0), \quad F(b, s) = F(b, 0)$$

for all $s \in (-\epsilon, \epsilon)$.

Write $\gamma_s(t) = F(t, s)$ so that for each s, γ_s is a curve, and $\gamma_0 = \gamma$. Let us write,

$$T = F_* \partial_t = \gamma_s'$$

for the tangent vector to γ_s, and

$$V = F_* \partial_s$$

for the *variation vector*. Let us also write,

$$T_0(t) = T(t, 0), \quad V_0(t) = V(t, 0)$$

for the restriction of T, V to γ.

Finally, let us assume (for simplicity not necessity) that γ is parametrised by arc-length, so that $|T|_g \equiv 1$.

We want now to compute the infinitesimal variation of length. That is, let $L(s) = L[\gamma_s]$ and compute $L'(0)$.

Proposition 7.9 (First Variation Formula)**.**

$$L'(0) = \left.\frac{d}{ds}\right|_{s=0} L[\gamma_s] = g(T_0, V_0)_{t=a}^{t=b} - \int_a^b g(V_0, \nabla_{T_0} T_0) dt.$$

In particular, if F fixes the endpoints, then

$$L'(0) = -\int_a^b g(V_0, \nabla_{T_0} T_0) dt.$$

Proof. We use that,

$$[T,V] = [F_*\partial_t, F_*\partial_s] = F_*[\partial_t, \partial_s] = 0$$

so that (using the pull back connection!),

$$\nabla_T V - \nabla_V T = [T,V] = 0.$$

Therefore, using metric compatability we have,

$$\begin{aligned}
\frac{\partial}{\partial s}\sqrt{g(T,T)} &= \frac{1}{2}\frac{1}{\sqrt{g(T,T)}} 2g(\nabla_V T, T) \\
&= \frac{1}{\sqrt{g(T,T)}} g(\nabla_T V, T) \\
&= \frac{1}{\sqrt{g(T,T)}}\left(\frac{\partial}{\partial t} g(V,T) - g(V, \nabla_T T)\right).
\end{aligned}$$

Evaluating at $s = 0$ where we have $g(T,T) \equiv 1$ we obtain,

$$\left.\frac{\partial}{\partial s}\right|_{s=0} \sqrt{g(T,T)} = \frac{\partial}{\partial t} g(V_0, T_0) - g(V_0, \nabla_{T_0} T_0).$$

Now we integrate from a to b to obtain,

$$\begin{aligned}
\left.\frac{d}{ds}\right|_{s=0} L[\gamma_s] = \left.\frac{d}{ds}\right|_{s=0} \int_a^b \sqrt{g(T,T)}dt &= \int_a^b \left.\frac{\partial}{\partial s}\right|_{s=0} \sqrt{g(T,T)}dt \\
&= \int_a^b \left[\frac{\partial}{\partial t} g(V_0, T_0) - g(V_0, \nabla_{T_0} T_0)\right] dt \\
&= g(T_0, V_0)_{t=a}^{t=b} - \int_a^b g(V_0, \nabla_{T_0} T_0) dt
\end{aligned}$$

If F fixes the endpoints, then $F(a,s) \equiv F(a,0)$ implies,

$$V_0(a) = 0$$

and similarly,

$$V_0(b) = 0$$

and hence

$$g(T_0, V_0)_{t=a}^{t=b} = 0.$$

□

Definition 7.10. A smooth curve, γ is a *geodesic* if,

$$\nabla_{\gamma'}\gamma' = 0.$$

That is, if γ' is parallel along γ.

Remark 7.11. Note that the first variation formula says that geodesics are precisely the critical points of the length functional among all smooth variations fixing the endpoints: Clearly, if γ is a geodesic, then for any variation fixing the endpoints,

$$L'(0) = -\int_a^b g(V_0, \nabla_{\gamma'}\gamma')dt = 0.$$

Conversely, suppose that for all variations F fixing the endpoints of a curve γ, we have

$$L'(0) = 0.$$

Let $N = \nabla_{\gamma'}\gamma'$ and for any smooth function $\phi : [a,b] \to \mathbb{R}$, let $V = \phi N$. Define a variation,

$$F(t,s) = f_s(V(t))$$

where f_s is the flow of V (so that $\partial_s F|_{t=0} = V$). Then, by the first variation formula

$$0 = L'(0) = -\int_a^b \phi g(\nabla_{T_0}T_0, \nabla_{T_0}T_0)dt.$$

Since this is true for any smooth function ϕ, by choosing bump functions supported in small neighbourhoods of points $t_0 \in (a,b)$, we have for every $t_0 \in (a,b)$,

$$(\nabla_{T_0}T_0)(t_0) = 0.$$

Later we will prove that in fact, geodesics are precisely the *local* minimisers of length. Here local means on any sufficiently small sub-interval of $[a,b]$, a geodesic minimises length between it's endpoints. Note that since geodesics are precisely the critical points of the length functional, these are the only possible candidates for length minimising curves.

Let us now examine the question of whether geodesics exist, and whether they are unique in any way.

Proposition 7.12. *Let $x \in M$ and $v \in T_xM$. Then there exists a unique geodesic γ such that,*

$$\gamma(0) = x, \text{ and } \gamma'(0) = v.$$

Remark 7.13. Here uniqueness means that for any other geodesic σ with $\sigma(0) = \gamma(0)$ and $\sigma'(0) = \gamma'(0)$, we must have σ equals γ on their common domain of definition. As usual, there exists a *maximal geodesic* though x and with tangent vector v at x, meaning γ is defined on a maximal interval $(-t_1, t_2)$ with $t_1, t_2 > 0$ and cannot be extended to a geodesic defined on a larger interval.

Proof. The proof consists in writing down the ODE satisfied by geodesics in local coordinates. Then as usual, since the geodesic equation is defined on M without reference to coordinates, the unique solutions on overlapping charts must agree (on M).

Let us write, in local coordinates,

$$\gamma(t) = (\gamma^1(t), \cdots, \gamma^n(t)),$$

and

$$\gamma'(t) = \dot{\gamma}^i(t)\partial_i(\gamma(t))$$

where $\partial_i(\gamma(t))$ means the vector field ∂_i evaluated at the point $\gamma(t)$. Thus

$$\dot{\gamma}^i = \partial_t \gamma^i$$

Let us also write

$$\ddot{\gamma}^i = \partial_t^2 \gamma^i.$$

From the chain rule,

$$\partial_i \dot{\gamma}^j = \partial_t \gamma^i \partial_t \dot{\gamma}^j = \dot{\gamma}^i \ddot{\gamma}^j.$$

Using the pull-back connection, and since $\gamma' = \gamma_* \partial_t$, we have

$$\nabla_{\gamma'} \dot{\gamma}^j \partial_j = \partial_t(\dot{\gamma}^j)\partial_j + \dot{\gamma}^j \nabla_{\gamma'} \partial_j.$$

Thus we may conclude that γ is a geodesic if and only if,

$$\begin{aligned} 0 = \nabla_{\gamma'}\gamma' &= \ddot{\gamma}^j \partial_j + \dot{\gamma}^j \nabla_{\dot{\gamma}^i \partial_i} \partial_j \\ &= \ddot{\gamma}^j \partial_j + \dot{\gamma}^j \dot{\gamma}^i \Gamma_{ij}^k \partial_k \\ &= \left(\ddot{\gamma}^j + \dot{\gamma}^i \dot{\gamma}^k \Gamma_{ik}^j \right) \partial_j \end{aligned}$$

This is a second order system of ODE's with smooth coefficients and hence has a unique solution on some interval $(-t_1, t_2)$. □

7.3 Week 07, Lecture 03: The Exponential Map

Given any $v \in TM$ let $x = \pi(v)$. We know that there is a unique, maximally defined geodesic, γ_v with $\gamma_v(0) = x$ and $\gamma_v'(0) = v$. This geodesic will not in general, be defined for all $t \in \mathbb{R}$. Let $DM \subset TM$ be the set of tangent vectors v such that $\gamma_v(1)$ is defined.

Definition 7.14. The exponential map is the map,

$$\exp : v \in DM \mapsto \gamma_v(1) \in M.$$

Restricting to the fibre, T_xM we write

$$\exp_x : D_xM = DM \cap T_xM \to M$$

We call DM the *domain of the exponential map.*

The exponential map turns our to have some rather remarkable properties and is fundamental in Riemannian geometry. Let us begin with a basic property, namely *homogeneity.*

Lemma 7.15. *Let $v \in DM$ and $\lambda \in \mathbb{R}$ such that $\lambda v \in DM$. Then*

$$\exp(\lambda v) = \gamma_{\lambda v}(1) = \gamma_v(\lambda)$$

Proof. Define the curve,

$$\tilde{\gamma}(t) = \gamma_v(\lambda t)$$

We need to show that $\tilde{\gamma}(1) = \gamma_{\lambda v}(1)$.

First observe that $\tilde{\gamma}(0) = \gamma_v(0) = x = \gamma_{\lambda v}(0)$ and that,

$$\tilde{\gamma}'(0) = \partial_t|_{t=0} (\gamma_v(\lambda t)) = \lambda \gamma_v'(0) = \lambda v$$

That is, $\tilde{\gamma}$ satisfies the same initial conditions as $\gamma_{\tilde{v}}$.

Next,

$$\nabla_{\tilde{\gamma}'} \tilde{\gamma}' = \nabla_{\lambda \gamma_v'} \lambda \gamma_v' = \lambda^2 \nabla_{\gamma_v'} \gamma_v' = 0$$

since γ_v is a geodesic.

Thus $\tilde{\gamma}$ and $\gamma_{\lambda v}$ are both geodesics with the same initial conditions, hence are equal. In particular, they are equal at $t = 1$. ☐

We will also need the following result, whose proof we delay until the next lecture.

Proposition 7.16. *There exists a unique, smooth vector field $G \in \mathfrak{X}(TM)$ such that the integral curves of G are precisely the maps*

$$t \mapsto \gamma'(t) \in TM$$

for γ a geodesic.

Note that integral curves are always smooth and that ϕ_t is defined on an open set $\mathcal{U} \subseteq TM \times \mathbb{R}$. The map ϕ_t is called the *geodesic flow.*

Theorem 7.17.

1. *The exponential map is smooth,*
2. *The domain DM is an open neighbourhood of the zero section in TM,*
3. *For each $x \in M$, there is an open neighbourhood U_x of the origin $O \in T_xM$ such that $\exp_x : U \to M$ is a diffeomorphism onto the open set $V_x = \exp_x(U_x)$,*
4. *For each x, $D_xM \subset T_xM$ is star-shaped. That is, for all $v \in D_xM$, $tv \in D_xM$ for all $t \in [0, 1]$.*
5. *The each $v \in DM$, curve $t \mapsto \exp(tv)$ is defined for all $t \in [0, 1]$ (and maybe other t as well) and is equal to the geodesic $\gamma_v(t)$.*

Proof of Theorem.

1. The exponential map is smooth.

   The geodesic flow is a smooth map and has the property that $\exp(v) = \gamma_v(1) = \pi \circ \phi_1(v)$ where ϕ_t is the geodesic flow. This is the composition of smooth maps, hence smooth.

2. The domain DM is an open neighbourhood of the zero section in TM.

   For any point $v_0 \in DM$, $\phi_1(v_0)$ is defined and so $\phi_t(v)$ is defined on an open set $U \times (1 - \delta, 1 + \delta) \subset \mathcal{U}$, with $U \subset TM$ open. Hence, $\phi_1(v)$ is defined for all v in the open neighbourhood U of v_0. That is $U \subset DM$ and hence DM is open. Now of course, $\exp_x(0_x) = x$ where $0_x \in T_xM$ is the zero vector. Thus the zero section, $Z = \{0_x \in T_xM : x \in M\} \subseteq DM$.

3. For each $x \in M$, there is an open neighbourhood U_x of the origin $O \in T_xM$ such that $\exp_x : U \to M$ is a diffeomorphism onto the open set $V_x = \exp_x(U_x)$.

Let $v \in T_xM$ and consider the curve $\alpha(t) = tv \in T_xM$ This curve has the property that $\alpha(0) = O$. Let $V = [\alpha] \in T_O(T_xM)$ be the tangent vector represented by α. Note that every tangent vector in $T_O(T_xM)$ may be represented this way, and so we have a map $v \in T_xM \mapsto V = [tv] \in T_O(T_xM)$ that is in fact a vector space isomorphism.

Now, since $\exp_x : T_xM \to M$, the differential $d\exp_x : T_O(T_xM) \to T_{\exp_x(O)}M = T_xM$.

Using homogeneity, we then compute

$$\begin{aligned} d\exp_x \cdot V &= \frac{d}{dt}\bigg|_{t=0} \exp_x(tv) = \frac{d}{dt}\bigg|_{t=0} \gamma_{tv}(1) \\ &= \frac{d}{dt}\bigg|_{t=0} \gamma_v(t) \\ &= \gamma_v'(0). \end{aligned}$$

But γ_v is the unique geodesic such that $\gamma_v'(0) = v$ and hence,

$$d\exp_x \cdot V = v.$$

Or in other words, the composition,

$$v \in T_xM \mapsto V \in T_O(T_xM) \mapsto d\exp_x \cdot V in T_xM$$

is the identity and hence $d\exp_x$ is an isomorphism. The inverse function theorem furnishes us with a local, smooth inverse to $\exp_x$.

4. The each $v \in DM$, curve $t \mapsto \exp(tv)$ is defined for all $t \in [0,1]$ (and maybe other t as well) and is equal to the geodesic $\gamma_v(t)$.

   This is more or less by definition. If $v \in DM$, then the unique geodesic γ_v is defined at least for $t \in [0,1]$. Therefore by homogeneity,

   $$\gamma_v(t) = \gamma_{tv}(1) = \exp(tv).$$

5. For each x, $D_xM \subset T_xM$ is star-shaped. That is, for all $v \in D_xM$, $tv \in D_xM$ for all $t \in [0,1]$.

   Let $v \in D_xM$ and $t \in [0,1]$. Then, by the previous part, $\exp(tv)$ is defined and hence $tv \in D_xM$.

$\square$

8 Week 08

8.1 Week 08, Lecture 01: Geodesic Flow and Gauss Lemma

A fundamental tool in studying geodesics is the Gauss lemma. It tells us how to interpret *geodesic polar coordinates* on a Riemannian manifold. Our first application of the Gauss lemma is to show that geodesics are locally length minimising. Recall that if a curve minimises length amongst all curves with the same end points, then that curve must be a geodesic. In Euclidean space, the converse is true; geodesics (straight lines) minimise length between their end points. The converse is actually false in general. For example on $\mathbb{S}^n$, an arc of a great circle (intersection of a 2-plane through the origin with the sphere) only minimises length between end-points provided those points are with distance π of each other. For greater distances, traversing the great circle in the opposite direction provides a shorter path. The best we can hope to do therefore, is that a geodesic minimises length between points sufficiently close to each other.

8.1.1 Geodesic Flow

Fundamental to the proof is the *geodesic flow*, which is of independent interest.

Proposition 8.1. *There exists a unique vector field $G \in \mathfrak{X}(TM)$ such that the integral curves of G are precisely the maps*

$$t \mapsto \gamma'(t) \in TM$$

for γ a geodesic.

This might seem a little odd. We are asserting the existence of a vector field on TM. This makes perfect sense since TM is itself a manifold, but may be a little difficult to understand. Essentially the idea here is smooth dependence on parameters, and the way we phrase this here is to turn the second order geodesic system of ODE's on M into a system of first order ODE's on TM. What then are the parameters upon which the solutions of the geodesic ODE should smoothly depend? They are precisely the tangent vectors v. So in other words, for each v we obtain a unique geodesic and as we vary v, the geodesics should smoothly vary. Namely, the map,

$$(v,t) \mapsto \gamma_v(t)$$

should be smooth in both t and v. We already have smoothness in t which is just saying that γ_v is a smooth curve, and smooth dependence on parameters ensures smoothness in v. The proof makes this precise and in so doing furnishes us with the geodesic flow.

Since we're working on the tangent bundle, let us recall the differential structure. For $(x^1, \cdots, x^n)$ local coordinates $\phi : U \to V$, writing $v = v^i \partial_i$, we have coordinates on $\pi^{-1}[U]$,

$$Phi(v) = (x^1, \cdots, x^n, v^1, \cdots, v^n)$$

where $(x^1, \cdots, x^n) = \phi(\pi(v))$. Let ∂_i denote the coordinate vector field corresponding to x^i and let $\dot{\partial}_i$ denote the coordinate vector field corresponding to v^i. Thus a tangent vector $V \in T(TM)$ to the tangent bundle may be written,

$$V = V^i \partial_i + \dot{V}^i \dot{\partial}_i.$$

Proof. We do the usual thing. Write the equation for G in coordinates (on TM!) and prove existence and uniqueness in coordinates, thereby allowing us to patch together the local expressions on the coordinate overlaps.

For $v \in TM$, let $(x^1, \cdots, x^n) = \phi \circ \pi(v)$ denote the coordinates of the base point of v and let $(v^1, \cdots, v_n)$ denote the coordinates of the vector part, and then define

$$G(v) = v^i \partial_i - v^k v^l \Gamma^i_{kl}((x^1, \cdots, x^n))\dot{\partial}_i.$$

Before verifying this is well defined and unique, let us check that $\phi_t(v) = \gamma_v'(t)$ is the unique integral curve of G through v where γ_v is the unique geodesic such that $\gamma_v(0) = \pi(v)$ and$\gamma_v'(0) = v$. Note that we only need to check that,

$$\partial_t \phi_t(v) = G(\phi_t(v))$$

since we already know that integral curves are unique and $\phi_0(v) = \gamma_v'(0) = v$ so that ϕ_t satisfies the required initial condition.

Now, in coordinates,

$$\phi_t(x^1, \cdots, x^n, v^1, \ldots, v^n) = (\gamma^1(t), \cdots, \gamma^n(t), \dot{\gamma}^1(t), \cdots, \dot{\gamma}^n(t)),$$

and so

$$\partial_t \phi_t(x^1, \cdots, x^n, v^1, \ldots, v^n) = \dot{\gamma}^i \partial_i + \ddot{\gamma}^i \dot{\partial}_i.$$

On the other hand,

$$\begin{aligned} G(\phi_t(x^1, \cdots, x^n, v^1, \ldots, v^n)) &= G((\gamma^1(t), \cdots, \gamma^n(t), \dot{\gamma}^1(t), \cdots, \dot{\gamma}^n(t))) \\ &= \dot{\gamma}^i \partial_i - \dot{\gamma}^k \dot{\gamma}^i \Gamma^i_{kl} \dot{\partial}_i. \end{aligned}$$

Therefore $\partial_t \phi_t(v) = G(\phi_t(v))$ provided,

$$\ddot{\gamma}^i = -\dot{\gamma}^k \dot{\gamma}^i \Gamma^i_{kl}$$

and this precisely the geodesic equation. Marvellous!

Now, existence and uniqueness of G. Both are quite easy, since any vector field whose flow is ϕ_t must satisfy

$$G(v) = G(\phi_0(v)) = \partial_t|_{t=0} \phi_t(v)$$

so G is determined by ϕ_t (this is in fact generally true - if two vector fields have the same flow then they are equal). This expression also defines G and so existence is assured. □

Remark 8.2. You may be wondering why we didn't simply define, $G(v) = \partial_t|_{t=0} \phi_t(v)$. The reason is that a priori, we know neither whether ϕ_t is a smooth map, nor whether in fact ϕ_t is the flow of any vector field! Assuming smoothness, the theory of integral curves gives a necessary condition for ϕ_t to be the flow of a vector field: $\phi_{t+s} = \phi_t \circ \phi_s$ and $\phi_0 = \mathrm{Id}$. You might like to think about how to prove this statement directly for $\phi_t(v) = \gamma'_v(t)$. Of course, this doesn't help to show ϕ_t is smooth. The difficulty here is that $\phi_t(v)$ is defined as the solution of an ODE with v fixed. The proof gave us the required smooth dependence on parameters.

Remark 8.3. The geodesic flow is a particular case of a general construction in Hamiltonian dynamics. There one treats the derivatives as independently varying variables thereby reducing a second order system of equations to a first order system. This general approach is equivalent to finding a vector field on TM whose integral curves are the solutions of the second order equation.

For example, consider the second order equation,

$$y''(x) = y(x)$$

for $x \in \mathbb{R}$. Letting $p = y'$, this becomes the first order system,

$$p' = y, \quad y' = p$$

for two functions $y(x), p(x)$. In other words, if we let $(p, y) \in \mathbb{R}^2 \simeq T\mathbb{R}$ be coordinates on $T\mathbb{R}$ we wish to find a vector field G on $T\mathbb{R}^2$, or equivalently a function,

$$G : \mathbb{R}^2 \to \mathbb{R}^2$$

such that for any (y_0, p_0), the integral curves, $\phi_t(y_0, p_0) = (y(t), p(t))$ satisfy

$$p = y', \quad y'' = y.$$

Thus the unique solutions of $y'' = y$ satisfying the initial conditions $y(0) = y_0$, $y'(0) = p_0$ is precisely $y(t) = \pi \circ \phi_t(y_0, p_0)$ where $\pi(y, p) = y$ is the projection $T\mathbb{R} \to \mathbb{R}$.

The vector field G is easily solved to be,

$$G(y, p) = (p, y).$$

The corresponding flow is,

$$\phi_t(y_0, p_0) = (y(t), p(t)) = \frac{1}{2}((y_0+p_0)e^t+(y_0-p_0)e^{-t}, (y_0+p_0)e^t+(p_0-y_0)e^{-t}).$$

Thus the unique solution of $y'' = y$ subject to $y(0) = y_0$, $y'(0) = p_0$ is

$$y(t) = \pi \circ \phi_t(y_0, p_0) = \frac{1}{2}(y_0 + p_0)e^t + \frac{1}{2}(y_0 - p_0)e^{-t}$$

which you may easily verify by direct calculation.

8.1.2 Gauss' Lemma

Recall that for any $w \in T_xM$ we can form the curve $t \mapsto tw \in T_xM$ and thus obtain a tangent vector $W = [tw] \in T_O(T_xM)$.

Lemma 8.4 (Gauss' Lemma)**.** *Let $x \in M$, fix $v \in T_xM$ and let γ_v be the corresponding geodesic. For any $w \in T_xM$, let $W(t) = d\exp_x|_{tv} \cdot W$. Then*

$$g(\gamma_v'(t), W(t)) = g(v, w).$$

Proof. First observe that since $\gamma'_{\lambda v} = \lambda\gamma'_v$, we may assume that $g_x(v,v) = 1$. Let us write, $|v|_g = \sqrt{g_x(v,v)}$. Next notice that the map $(w,t) \mapsto g(\gamma'_v(t), W(t))$ is linear in w and so it suffices to prove the result for $w = v$ and also for $w \perp v$ with $|w|_g = 1$.

First, if $w = v$, then

$$W(t) = d\exp_x|_{tv} \cdot w = \partial_s|_{s=0} \exp_x(tv + sv) = \partial_s|_{s=0}\gamma_v(t+s) = \gamma'_v(t).$$

From metric compatibility,

$$\partial_t g(\gamma'_v(t), W(t)) = \partial_t g(\gamma'_v(t), \gamma'_v(t)) = 2g(\nabla'_{\gamma_v}\gamma'_v, \gamma'_v) = 0$$

since γ_v is a geodesic. Therefore $g(\gamma'_v, W)$ is constant in the case $w = v$.

Now suppose that $w \in T_xM$ such that $|w|_g = 1$ and $g_x(w,v) = 0$. Define,

$$F(t,\theta) = \exp_x(t(\cos\theta v + \sin\theta w)) = \gamma_{\cos\theta v + \sin\theta w}(t),$$

for θ sufficiently small so that $\exp_x(t(\cos\theta v + \sin\theta w))$ is defined. For each fixed t, $F(t,\theta)$ traces out an arc of a curve in T_xM such that

$$|F(t,\theta)|^2 = t^2\left(\cos^2\theta|v|_g^2 + \sin^2\theta|w|_g^2\right) = t^2$$

since v, w are unit length, orthogonal vectors. Let us write $u(\theta) = \cos\theta v + \sin\theta w$ and so we also have $|u|_g = 1$ by the same reasoning and $F = tu(\theta)$.

Define the vectors,

$$V = F_*\partial_t, U = F_*\partial_\theta.$$

Then we have,

$$V = \partial_t\gamma_{\cos\theta v + \sin\theta w}(t) = \gamma'_{u(\theta)}(t).$$

is the tangent to a geodesic. Then

$$\partial_t g(V,V) = \partial_t g(\gamma'_u, \gamma'_u) = 2g(\nabla_{\gamma'_u}\gamma'_u, \gamma'_u) = 0$$

Thus $g(V,V) = g_x(u,u) = 1$ is constant in both t and θ.

Now using $[V,U] = 0$:

$$\begin{aligned}
\partial_t g(V,U) &= g(\nabla_V V, U) + g(V, \nabla_V U) \\
&= g(V, \nabla_U V) \\
&= \frac{1}{2}\partial_\theta g(V,V) = 0.
\end{aligned}$$

Thus $g(V,U)$ is also constant. Now observe that when $\theta = 0$ and $t > 0$,

$$U = F_*\partial_\theta = d\exp_x|_{tv} \cdot (\cos\theta v + \sin\theta w)'|_{\theta=0} = d\exp_x|_{tv} \cdot W = W(t).$$

□

8.2 Week 08, Lecture 02: Length Minimisation

We can use Gauss' lemma to define *geodesic polar coordinates.* First define,

$$Q : v \in T_xM \mapsto g_x(v, v) \in \mathbb{R}$$

Then we have, writing $V = [(1+t)v] \in T_v(T_xM)$ that,

$$dQ_v \cdot V = \partial_t|_{t=0} g_x((1+t)v, (1+t)v) = \partial_t|_{t=0}(1+t)^2 g(v, v) = 2g(v, v) \neq 0$$

Letting $U_xM = Q^{-1}(\{1\})$, we find that 1 is a regular value of Q and hence U_x is a smooth submanifold by the implicit function theorem. The set $U_x = \{v \in T_xM : g_x(v, v) = 1\}$ is the *unit tangent space.* More generally, let $\mathbb{S}_r(x) = \{\exp_x(v) : |v|_g = r\}$ which is defined for r sufficiently small. We will see that in fact,

$$\mathbb{S}_r(x) = \{y \in M : d(x, y) = r\}$$

and that,

$$B_r(x) = \{y \in M : d(x, y) < r\} = \{\exp_x(v) : |v|_g < r\}.$$

That is, the exponential map identifies the ball centred on the origin of radius r with respect to g_x with the metric ball of radius r centred on x in M.

Definition 8.5 (Geodesic Polar Coordinates)**.** Define the map,

$$\begin{aligned} \phi : (0, \infty) \times U_xM &\to T_xM - \{O\} \\ (r, v) &\mapsto rv. \end{aligned}$$

This map is a diffeomorphism and induces a diffeomorphism,

$$\begin{aligned} \bar{\phi} : (0, R) \times U_xM &\to T_xM - \{O\} \\ (r, v) &\mapsto \exp(rv) \end{aligned}$$

for $r \in (0, R)$ where R is chosen sufficiently small, so that $\exp(rv)$ is defined for all $v \in U_xM$ and all $r \in (0, R)$. We consider $(O, R) \times U_xM$ as geodesic polar coordinates on M.

To connect this notion with more familiar polar coordinates, note that the diffeomorphism $\bar{\phi}$ induces coordinates on $M - \{x\}$ by taking a g_x-orthonormal basis $\{e_i\}$ for T_xM and defining the diffeomorphism,

$$(x^1, \cdots, x^n) \mapsto x^i e_i$$

of $\mathbb{R}^n$ with T_xM. In particular, this map takes the unit sphere, $\mathbb{S}^{n-1} = \{x \in \mathbb{R}^n : |x| = 1\}$ onto U_x. Taking polar coordinates $(r, \theta) \in (0, \infty) \times \mathbb{S}^{n-1}$ on $\mathbb{S}^{n-1}$ induces geodesic polar coordinates in M,

$$(r, \theta) \mapsto \exp_x(r\theta).$$

As an exercise, using Gauss' lemma prove that in the coordinates provided by the diffeomorphism,

$$(x^1, \cdots, x^n) \mapsto \exp_x(x^i e_i)$$

with $\{e_i\}$ a g_x-orthonormal basis, we have,

$$g_{ij}(0) = \delta_{ij}, \quad \Gamma_{ij}^k(0) = 0.$$

Such coordinates are called *normal coordinates*.

Similarly, in geodesic polar coordinates, Gauss' lemma implies that the pull-back metric takes the form,

$$g = dr^2 + \bar{g}(r) = dr \otimes dr + \bar{g}(r)$$

where $\bar{g}(r)$ is a smoothly varying (in r) metric on U_xM (or $\mathbb{S}^{n-1}$) after inducing coordinates via an orthonormal frame $\{e_i\}$) for each r. In other words,

$$g(\partial_r, \partial_r) = 1$$

and

$$g(\partial_r, w) = 0$$

for any w tangent to U_x (or $\mathbb{S}^{n-1}$).

Thus in particular, any tangent vector $v \in T_xM$ may be written uniquely as

$$v = v^r \partial_r + w$$

where $g_x(\partial_r, w) = 0$ and then,

$$g_x(v, v) = (v^r)^2 + \bar{g}_r(w, w).$$

Definition 8.6. With the notation of the previous definition, we may also define the radial function,

$$r(y) = g_x(\exp_x^{-1} y, \exp_x^{-1} y)$$

for $y \in \exp_x(U)$ where $\exp_x$ is a diffeomorphism with it's image when restricted to U. Thus $r(y) = L[\gamma_v]$ where $v = \exp^{-1}(y)$ is the unique tangent vector such that, $y = \exp(v)$.

Then in fact (exercise!), $\partial_r = \nabla r$. That is for any vector field v,

$$dr(v) = g(\partial_r, v).$$

The following proposition, shows that in particular, $r(y) = d(x, y)$ provided y is sufficiently close to x. That is, r is the distance along the radial geodesic $t \mapsto \exp_(tv)$ joining x to y. This proves the claim about the sets $B_r(x)$ and $\mathbb{S}_r(x)$ above.

Proposition 8.7. *Geodesics locally minimise length. That is, if $\gamma : I \to M$ is a geodesic for the Levi-Civita connection, then there exists a $\delta > 0$ such that for all $t_0 \in I$,*

$$L[\gamma|_{[t_1,t_2]}] = d(\gamma(t_1), \gamma(t_2))$$

whenever $|t_2 - t_1| < \delta$. Moreover, given any $x \in M$, there exists an open neighbourhood, U of x such that for all $y \in U$ there exists a unique*, length minimising geodesic joining x to y.*

Proof. Given $\delta > 0$, choose $t_1, t_2 \in I$ with $|t_2 - t_1| < \delta$ and write $x = \gamma(t_1)$ and $y = \gamma(t_2)$. Now choose $\delta > 0$ small enough so that $\exp_x$ is diffeomorphism on the ball $B_r(O) = \{v \in T_xM : \sqrt{g_x(v, v)} < r\}$ for some $r > \delta$.

Normalise, γ so that $g(\gamma', \gamma') = 1$ and let $v = \gamma'(t_1) \in T_xM$ so that,

$$y = \gamma(t_2) = \exp_x((t_2 - t_1)v).$$

This follows since γ_v and γ are the same geodesics, just with time shifted. Note that $g_x((t_2 - t_1)v, (t_2 - t_1)v) = (t_2 - t_1)^2 g_x(v, v) < \delta^2 < r^2$ so that $y \in B_r(x) = \exp_x(B_r(O))$ and also that $L[\gamma_{[t_1,t_2]}] = |t_2 - t_1| < \delta < r$ since $g(\gamma', \gamma') = 1$.

Now let μ be any curve joining x and y. The connected component of $\mathrm{Image}(\mu) \cap B_r(x)$ containing x is a curve entirely contained within $B_r(x)$ joining x to either y or a point $x \in \partial B_r(x)$.

From the Gauss lemma, we may write

$$\mu'(t) = r(t)\partial_r + w(t)$$

for some function $r(t)$ and a vector field $w(t)$ along μ with $g(\partial_r, w) = 0$. Then,

$$|\mu'|_g^2 = r^2 + |w|_g^2 \geq r^2$$

whence,

$$|\mu'| = |r| \geq r = dr(\mu') = \frac{d}{dt} r(\mu(t)),$$

for $t > 0$. Therefore,

$$L[\mu] = \int_0^1 |\mu'| dt \geq \int_0^1 \frac{d}{dt} r(\mu(t)) dt = r(z) - r(x) = r(z) = L[\gamma].$$

Note that we do need to take care at $t = 0$ since r is not differentiable there, and we should really take,

$$L[\mu] = \lim_{\epsilon \to 0} \int_\epsilon^1 |\mu'| dt \geq r(z) - \lim_{\epsilon \to 0} r(\gamma(\epsilon)) = r(z).$$

Now suppose that σ is any other length minimising curve connecting x to y. Let $w = \sigma'(0) \in T_x M$. Then in fact, $\sigma = \gamma_w$. We can see this either, by recalling that the first variation formula implies any length minimising curve is a geodesic, or by noting that $L[\sigma] = r(y)$ implies that $\sigma' = \sigma^r \partial_r$ so that $\sigma(t) = \exp(tw) = \gamma_w(t)$. Now we have that

$$\exp_x(w) = y = \exp_x(v)$$

and hence $w = v$ so that $\sigma = \gamma_v$ is the unique, length minimising geodesic joining x to y. $\square$

8.3 Week 08, Lecture 03: Completeness

So we know that geodesics locally minimise length, but given $x, y \in M$ does there exist a length minimising geodesic joining x to y? If so, is it unique? In general the answer is no to both questions. For example on the sphere, two antipodal points are connected by infinitely many geodesics all of equal length, so uniqueness fails in general. The existence question is a little more subtle and we'll return to that in a moment.

First, let us note that if x and y are sufficiently close, in particular if $y \in \exp_x(U)$ where $U \subset T_xM$ is an open neighbourhood of the origin on which $\exp_x$ is a diffeomorphism, letting $v = \exp_x^{-1}(y)$ we have that $y = \gamma_v(1)$ and $x = \gamma_v(0)$ so existence of a geodesic is always assured locally, and the previous lemma states that γ_v is length minimising for $t < \delta$ for some $\delta > 0$. In fact, it's possible to prove a stronger version of this statement:

Proposition 8.8. *Let $x \in M$. Then there exists an open neighbourhood $U \subset T_xM$ such that,*

$$d(x, \exp_x(v)) = L[\gamma_v]$$

for all $v \in U$.

We won't prove this proposition here, but the idea is to use the geodesic flow to obtain a uniform δ for which γ_v is length minimising on $(0, \delta)$ for all $v \in U_xM$ and then use homogeneity of geodesics. Thus in particular, for all $y \in \{\exp_x(tv) : t \in (0, \delta), v \in U_xM\}$, there exists a length minimising geodesic joining x to y. In fact, an even strong result holds by allowing x to vary:

Proposition 8.9. *There exists an open set $U \subset M \times M$ containing the diagonal $\Delta = \{(x, x) \in M \times M\}$ such that for all $(x, y) \in U$ there exists a length minimising geodesic joining x to y.*

Lastly, the general question of length minimising geodesics has to be dealt with by *defining the problem away.*

Theorem 8.10 (Hopf-Rinow)**.** *Let (M, g) be a Riemannian manifold and d_g the induced distance function. The following are equivalent:*

1. *(M, d_g) is a complete metric space,*
2. *$DM = TM$, i.e. $\exp$ is defined on all of TM,*

3. There exists $x \in M$ such that $\exp_x$ is defined on all of T_xM.

Any of the three conditions implies that,

- *$(*)$ For all $x, y \in M$ there exists a length minimising geodesic joining x to y.*

Remark 8.11. If condition 2, is satisfied, then (M, g) is said to be *geodesically complete* since it implies by homogeneity that that every geodesic is complete, i.e. exists on the maximal interval $\mathbb{R}$.

Remark 8.12. The last condition, $(*)$ is sometimes mistakenly claimed to be equivalent to the first three. This is false in general. A simple example is the open unit disc equipped with the Euclidean metric. Any two points may be connected by a straight line, entirely contained within the disc but the other three conditions are not satisfied. Another example to keep in mind is that of $\mathbb{R}^2 - \{O\}$ equipped with the Euclidean metric. Again the three equivalent conditions fail to be satisfied, and in this case the final condition is also not satisfied. Why?

Sketch.

- 2. implies 3. trivially.
- 1. implies 2.

  Suppose 1. is true, that (M, d_g) is a complete metric space, but 2. is false so that there exists a $v \in TM$ such that γ_v is defined on (a, b) with either $a > -\infty$ or $b < \infty$. Assume the latter (the former is the same argument). Let $t_n \nearrow b$. Then (t_n) is a Cauchy sequence and,

  $$d(\gamma_v(t_n), \gamma_v(t_m)) \leq L[(\gamma_v|_{[t_n,t_m]})] = |v||t_n - t_m|$$

  and so $\gamma(t_n)$ is also Cauchy hence by assumption, convergent to some $x \in M$. But now for a relatively compact open neighbourhood U of x, $K = \{u \in \pi^{-1}[\bar{U}] : |u| = |v|\}$ is compact in TM and $\gamma_v'(t_n) \in K$ for all n large enough. Passing to a sub-sequence we then have

  $$\gamma_v'(t_n) \to u \in T_xM$$

  as $n \to \infty$. That is,

  $$\phi_{t_n}(v) \to u$$

  where ϕ_t is the geodesic flow. Now the Flow-Box Theorem (Picard-Lindel\"of) implies $\phi_t(v)$ may be extended past b contradicting maximality.

- 2. implies 1.

  Let $C \subset M$ be a closed and bounded set. By showing that C is compact we obtain the result. To this end, pick any $x \in M$ and let $r = \sup\{d(x,y) : y \in C\} < \infty$ and let $\overline{B_r(x)}$ denote the closed metric ball of radius r with center x. Then $C \subset \overline{B_r(x)}$ and the result will follow by showing that $\overline{B_r(x)} = \exp_x(\{vinT_xM : |v| \leq r\})$, the latter being the continuous image of a compact set, hence continuous. Thus C is a closed subset of a compact set, hence itself compact. To prove that $\overline{B_r(x)} = \exp_x(\{vinT_xM : |v| \leq r\})$ note that if $|v| \leq r$ then $L[\gamma_v]|_{[0,1]} = |v| \leq r$ so that the inclusion $\supseteq$ follows. The reverse inclusion follows if we can prove that 3. implies the last condition, $(*)$.

- 3. implies the last condition, $(*)$.

  This is the part that requires the most work, and we refer to Andrews notes, chapter 11. or Theorem 1.7.1 of Chavel Riemannian Geometry. The idea is that we know that there is a $\delta > 0$ such that if $d(x,y) < \delta$ then there is a length minimising geodesic joining x to y. If $d(x,y) \geq \delta$, we choose z such that $d(x,z) = \delta$ and $d(z,y) \leq d(w,y)$ for all other w such that $d(x,w) = \delta$ which can be done by compactness of the set of all such δ. Then it turns out that the geodesic joining x to z also joins x to y and minimises length.

- 3. implies 1.

  Essentially, if $\exp_x$ is defined on all of T_xM, then $(*)$ holds and for any y we may join x to y by a length minimising geodesic and then a similar argument to 2. implies 1. shows that M is complete.

$\square$

9 Week 09

9.1 Week 08, Lecture 01: Higher Derivatives

9.1.1 Induced Connections on Tensor Bundles

Throughout, let ∇ be a connection on TM. Let us also write $\nabla f = df$. With this notation, ∇ is a connection on T^0M (the sections of which area smooth functions $f \in C^\infty(M)$) and the Leibniz rule reads,

$$\nabla(f \otimes X) = \nabla f \otimes X + f \otimes \nabla X$$

Recall that $\mathbb{R} \otimes V \simeq V$ for any vector space V, with the isomorphism $\lambda \otimes v \mapsto \lambda v$. At the sections of bundles level $T^0M \otimes TM \simeq TM$ via the map $f \otimes X \mapsto fX$ for $f \in C^\infty(M)$ and $X \in \mathfrak{X}(M)$. This rephrasing of the Leibniz rule suggests a way to extend ∇ to tensor fields; by requiring the Leibniz rule to hold for tensor products.

First, let us see how to extend the connection to the dual bundle as this will help with the general construction.

Lemma 9.1. *Let $\alpha \in \Gamma(T^*M)$ and define $\nabla^*\alpha \in \Gamma(T^*M \otimes T^*M$ by the formula*

$$(\nabla^*\alpha)(X, Y) = \nabla_X(\alpha(Y)) - \alpha(\nabla_X Y).$$

Then ∇^* is a connection on $\Gamma(T^*M)$.

Moreover, if ∇ is the Levi-Civita connection for a metric g on TM, then ∇^* is the Levi-Civita connection for the dual metric g^*.

Proof. The right hand side of

$$(\nabla^*\alpha)(X, Y) = \nabla_X(\alpha(Y)) - \alpha(\nabla_X Y),$$

consists of the first term, the derivative of the smooth function $\alpha(Y)$ in the direction X and the second term, the one-form α evaluated at (contracted with) the vector field $\nabla_X Y$. Thus at least the right hand side is well defined.

Let's check that $\nabla^*\alpha$ is indeed tensorial in X and Y. Tensorality in X

follows from tensorality of ∇_X and of α. For Y, we compute,

$$\begin{aligned}\nabla^*\alpha(X, fY) &= \nabla_X\alpha(fY) - \alpha(\nabla_X fY)\\ &= \nabla_X(f\alpha(Y)) - \alpha((\nabla_X f)Y) + f\nabla_X Y)\\ &= (\nabla_X f)\alpha(Y) + f\nabla_X\alpha(Y) - (\nabla_X f)\alpha(Y) - f\alpha(\nabla_X Y)\\ &= f(\nabla_X\alpha(Y) - \alpha(\nabla_X Y))\\ &= f\nabla^*\alpha(X, Y).\end{aligned}$$

Additivity is easy to check.

For the Leibniz rule,

$$\begin{aligned}\nabla^* f\alpha(X, Y) &= \nabla_X(f\alpha(Y)) - f\alpha(\nabla_X Y)\\ &= (\nabla_X f)\alpha(Y) + f(\nabla_X(\alpha(Y)) - \alpha(\nabla_X Y))\\ &= \nabla_X f\alpha(Y) + f\nabla^*\alpha(X, Y).\end{aligned}$$

In other words,

$$\nabla^*_X(f\alpha) = (\nabla_X f)\alpha + f\nabla^*_X\alpha.$$

□

By the usual abuse of notation, we write ∇ for the connection ∇^*.

We can rephrase the lemma as follows: For any vector field, X and one-form α we have

$$\nabla(\alpha(Y)) = \nabla\mathrm{Tr}(\alpha\otimes Y) = \mathrm{Tr}\left(\nabla\alpha\otimes Y\right) + \mathrm{Tr}\left(\alpha\otimes\nabla X\right).$$

Thus if define,

$$\nabla(\alpha\otimes X) = \nabla\alpha\otimes X + \alpha\otimes\nabla X$$

we obtain,

$$\begin{aligned}\nabla\mathrm{Tr}(\alpha\otimes Y) &= \mathrm{Tr}\nabla(\alpha\otimes Y) \text{ and}\\ \nabla(\alpha\otimes X) &= \nabla\alpha\otimes X + \alpha\otimes\nabla X.\end{aligned}$$

Thus ∇ commutes with traces and obeys the Leibniz rule for tensor products. One can check directly that the definition of $\nabla(\alpha\otimes X)$ defines a connection T^1_1M.

What have we achieved here? First we extended ∇ from T^1_0 to T^0_1M and then to T^1_1M by requiring ∇ commutes with traces and obeys the Leibniz rule with respect to tensor products. Now we inductively define ∇ on any tensor field in $T^p_q M$. Note that in this setting, a connection on $T^p_q M$ is precisely a map,

$$\nabla : \Gamma(T^p_q M) \to \Gamma(T^* M \otimes T^p_q M) = \Gamma(T^p_{q+1} M)$$

satisfying the Leibniz rule with respect to tensor products with $f \in \Gamma(T^0_0 M)$. But what about general tensor products? Notice that the connections on T^1_0, T^0_1, T^1_1 satisfy the Liebniz rule with respect to tensor products of T^1_0 and T^0_1. In general, this is how we proceed.

Proposition 9.2. *Let E be a vector bundle with connection ∇. Then there exists a unique connection ∇^* on E^* such that,*

$$\nabla_X [(\omega(s))] = (\nabla^*_X \omega)(s) + \omega(\nabla_X s).$$

Equivalently,

$$\nabla_X Tr(\omega \otimes s) = Tr\left(\nabla^*_X \omega \otimes s + \omega \otimes \nabla_X s\right).$$

Let E_1 and E_2 be vector bundles with connections ∇_1 and ∇_2. Then there exists a unique connection, ∇ on $E_1 \otimes E_2$ such that

$$\nabla(s_1 \otimes s_2) = \nabla_1 s_1 \otimes s_2 + s_1 \otimes \nabla s_2.$$

Putting together the two parts of the proposition we can now define a connection on $E \otimes E^*$ by

$$\nabla^1_1(\omega \otimes s) = \nabla^* \omega \otimes s + \omega \otimes \nabla s$$

so that

$$\nabla_X \mathrm{Tr}(\omega \otimes s) = \mathrm{Tr}\left(\nabla^*_X \omega \otimes s + \omega \otimes \nabla_X s\right) = \mathrm{Tr}(\nabla^1_1)_X(\omega \otimes s).$$

Inductively repeating this construction in the particular case that $E = TM$ we obtain a connection on $T^* M$ and then taking $E_1 = T^p_q M$, $E_2 = TM$ (or $E_2 = T^* M$) we obtain a connection on $T^{p+1}_q M$ (or $T^p_{q+1} M$).

Corollary 9.3. *Let ∇ be a connection on TM. Then there exists unique connections ∇^p_q on T^p_q such that,*

1. $\nabla^0_0 = d$,
2. $\nabla^1_0 = \nabla$,

3. $\nabla^{p+k}_{q+l}(s \otimes t) = \nabla^p_q s \otimes t + s \otimes \nabla^k_l t$ *for* $s \in T^p_q$, $t \in T^k_l$

4. $Tr \nabla s = \nabla Trs$ *if* $s \in T^{p+1}_{q+1}$ *for any trace.*

As usual, we usually simply write ∇ for all these connections. An important result is the following:

Corollary 9.4. *If* ∇ *is the Levi-Civita connection on* TM *for a metric* g*, then the connections on* $T^p_q M$ *are metric compatible with the metrics induced by* g*.*

The proof of these results is left as an exercise, but let's look at an example to see how the construction works in practice.

Example 9.5. Let $s \in \Gamma(T^1_2(M))$ which we may think of as a $C^\infty(M)$-multilinear map,

$$\Gamma(T^*M) \times \Gamma(TM) \times \Gamma(TM) \to C^\infty(M).$$

Then, $\nabla s \in \Gamma(T^1_3(M))$ and is defined by

$$\nabla s(X, \alpha, Y_1, Y_2) = \nabla_X s(\alpha, Y_1, Y_2) - s(\nabla_X \alpha, Y_1, Y_2) - s(\alpha, \nabla_X Y_1, Y_2) - s(\alpha, Y_1, \nabla_X Y_2)$$

where

$$(\nabla_X \alpha)(Y) = \nabla_X \alpha(Y) - \alpha(\nabla_X Y).$$

For say, $s \in \Gamma^2_0(M)$, we have

$$\nabla s(X, \alpha_1, \alpha_2) \nabla_X (s(\alpha_1, \alpha_2)) - s(\nabla_X \alpha_1, \alpha_2) - s(\alpha_1, \nabla_X \alpha_2).$$

You should be able to see the general idea now which is a little tedious to write out, but straight forward.

9.1.2 Second Derivatives

The constructions of the previous section allow us now to take second derivatives.

Definition 9.6. Let $f \in C^\infty(M)$ and define the *Hessian* of f,

$$\nabla^2 f = \nabla(\nabla f).$$

This is a section of $T^0_2(M)$, i.e. a bilinear form. As a bilinear form,

$$\nabla^2 f(X, Y) = \nabla(\nabla f)(X, Y) = \nabla_X(\nabla f(Y)) - \nabla f(\nabla_X Y) = \nabla_X(\nabla Y) - \nabla_{\nabla_X Y} f.$$

In fact this definition makes sense for any tensor field T,

$$\nabla^2 T = \nabla(\nabla T)$$

and we have,

$$\nabla^2 T(X,Y) = \nabla_X(\nabla_Y T) - \nabla_{\nabla_X Y} T.$$

It's also common to write,

$$\nabla^2_{X,Y} T = \nabla^2 T(X,Y).$$

Using index notation,

$$(\nabla^2 T)_{ij} = (\nabla^2 T)(\partial_i, \partial_j) = \nabla_{\partial_i}(\nabla_{\partial_j} T) - \nabla_{\nabla_{\partial_i}\partial_j} T = \nabla_{\partial_i}(\nabla_{\partial_j} T) - \Gamma^k_{ij} \nabla_{\partial_k} T.$$

Note here that $\nabla_{\partial_j} T$ is a tensor of the same type as T and then we can differentiate this new tensor also to obtain $\nabla_{\partial_i}(\nabla_{\partial_j} T)$ also a tensor of the same type. But the map,

$$(X,Y) \mapsto \nabla_X(\nabla_Y T)$$

is tensorial only in X and not in Y. Then extra term $\nabla_{\nabla_X Y} T$ may be thought of as correction term that ensures $\nabla^2 T$ is tensorial in both arguments. I prefer to think of this correction term as simply arising from the fact that ∇ commutes with tensor products and traces. Refer back to the original construction of ∇ on T^0_1 in this lecture where it was shown that $\nabla\alpha$ is tensorial in both arguments. This followed in precisely the same way, the extra term in $\nabla_X(\nabla_{fY} T)$ arising from the Leibniz rule is cancelled by the extra term in $\nabla_{\nabla_X fY} T$ arising from the Leibniz rule!

A word of caution here: Often $(\nabla^2 T)_{ij}$ is written as $\nabla_i \nabla_j T$. There are certain advantages (mostly notational) to writing it this way, but personally I find it too confusing since in this notation,

$$\nabla_i \nabla_j T \neq \nabla_{\partial_i}(\nabla_{\partial_j} T).$$

This confusion is ameliorated somewhat, by the general maxim that *all expressions should be tensorial* and so $\nabla_i \nabla_j T$ must stand for $(\nabla^2 T)_{ij}$ and not the iterated derivative. In this course, we will never use the notation $\nabla_i \nabla_j T$, but it is *extremely common in the literature* and you would do well to keep that in mind.

9.2 Week 09, Lecture 02: Curvature

9.2.1 The Curvature Tensor

Now that we can differentiate twice, we can define the curvature tensor.

Definition 9.7. Let ∇ be a connection on a vector bundle E. The *Riemann curvature tensor* is defined by

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

where X, Y are vector fields and T is a section of E. If E also posses a metric g, we may define the metric contraction as,

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

where S is also a section of E.

Remark 9.8. There are basically two common conventions in use; the one given here and the opposition convention,

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

That is X and Y are reversed, which simply introduces a minus sign. I prefer the convention given above since the order of X, Y, T are preserved in the first and last terms. Whereas, using the convention in this remark one has to remember to swap the order of X and Y (or introduce a minus sign).

In the particular case that $E = T^p_q M$, immediately we have the Ricci identities.

Lemma 9.9 (Ricci Identities)**.** *Let ∇ be a torsion-free connection on TM and denote also by ∇ the induced connection on $T^p_q M$. Then,*

$$R(X,Y)T = \nabla^2 T(X,Y) - \nabla^2 T(Y,X) = \nabla^2_{X,Y}T - \nabla^2_{Y,X}T$$

for any $X, Y \in \mathfrak{X}(M)$, $T \in \Gamma(T^p_q M)$.

Proof. Left as an exercise. *Hint*: The key thing here is that ∇ is torsion free. The result follows easily now from the definition of ∇^2. □

Thus one interpretation of the curvature tensor is that it measures the failure of second derivatives to commute (when we take tensorial second derivatives).

Remark 9.10. For a connection on an arbitrary vector bundle, the concept of torsion is not well defined. Recall that the torsion tensor of a connection on TM is defined by,

$$T(X,Y) = \nabla_X Y - \nabla_Y X - [X,Y]$$

and clearly this only makes sense for $X, Y \in \mathfrak{X}(M)$. We can however, as in the lemma, speak of a connection on $T^p_q M$ induced by a torsion free connection on TM and then the Ricci identities hold. In fact, the converse is also true.

Let us now consider the Levi-Civita connection on a Riemannian manifold (M, g). The curvature tensor possess several symmetries in this situation. Some of these hold for general connections, some depend on either metric compatability or the connection being torsion free. See if you can spot which are which.

Proposition 9.11. *1. $R(X,Y)Z = -R(Y,X)Z$,*

2. $R(X,Y,Z,W) = -R(X,Y,W,Z)$,

*3. $R(X,Y)Z + R(Y,Z)X + R(Z,X)Y = 0$ (*Bianchi Identity*),*

4. $R(X,Y,Z,W) = R(Z,W,X,Y)$.

Proof. 1. Immediate.

2. We use metric compatability.

$$\begin{aligned} g(\nabla_X \nabla_Y Z, W) &= \nabla_X g(\nabla_Y Z, W) - g(\nabla_Y Z, \nabla_X W) \\ &= \nabla_X \nabla_Y g(Z, W) - \nabla_X g(Z, \nabla_Y W) \\ &\quad - \nabla_Y g(Z, \nabla_X W) + g(Z, \nabla_Y \nabla_X W). \end{aligned}$$

Similarly,

$$\begin{aligned} g(\nabla_Y \nabla_X Z, W) &= \nabla_Y g(\nabla_X Z, W) - g(\nabla_X Z, \nabla_Y W) \\ &= \nabla_Y \nabla_X g(Z, W) - \nabla_Y g(Z, \nabla_X W) \\ &\quad - \nabla_X g(Z, \nabla_Y W) + g(Z, \nabla_X \nabla_Y W). \end{aligned}$$

Thus,

$$g(\nabla_X \nabla_Y Z - \nabla_Y \nabla_X, W) = (\nabla_X \nabla_Y - \nabla_Y \nabla_X) g(Z, W) - g(Z, \nabla_X \nabla_Y W - \nabla_Y \nabla_X W)$$

Subtracting,

$$g(\nabla_{[X,Y]}Z, W) = \nabla_{[X,Y]}g(Z,W) - g(Z, \nabla_{[X,Y]}W)$$

we obtain,

$$\begin{aligned} R(X,Y,Z,W) &= g(\nabla_X\nabla_Y Z - \nabla_Y\nabla_X Z - \nabla_{[X,Y]}Z, W) \\ &= (\nabla_X\nabla_Y - \nabla_Y\nabla_X - \nabla_{[X,Y]})g(Z,W) \\ &\quad - g(Z, \nabla_X\nabla_Y W - \nabla_Y\nabla_X W - \nabla_{[X,Y]}W) \\ &= R(X,Y)g(Z,W) - R(X,Y,W,Z). \end{aligned}$$

The result now follows by showing that for any smooth function, $R(X,Y)f = 0$ applied to the smooth function $f = g(Z,W)$.

3. This follows from the Jacobi identity,

$$[X,[Y,Z]] + [Y,[Z,X]] + [Z,[X,Y]] = 0$$

and that ∇ is torsion free. Notice that both what we are trying to prove and the Jacobi identity satisfy the sum of cyclic permutations is zero, which is at least an indication that this might work. The details of the Jacobi identity and the curvature tensor identity are left as an exercise. Alternatively, one may prove the required identity in coordinates, since this has the added bonus that $[\partial_i, \partial_k] = 0$.

4. This is purely algebraic. It follows from the previous symmetries. From the Bianchi identity,

$$\begin{aligned} \mathrm{Rm}(X,Y,W,Z) + \mathrm{Rm}(Y,W,X,Z) + \mathrm{Rm}(W,X,Y,Z) &= 0, \\ \mathrm{Rm}(Y,W,Z,X) + \mathrm{Rm}(W,Z,Y,X) + \mathrm{Rm}(Z,Y,W,X) &= 0, \\ \mathrm{Rm}(W,Z,X,Y) + \mathrm{Rm}(Z,X,W,Y) + \mathrm{Rm}(X,W,Z,Y) &= 0, \\ \mathrm{Rm}(Z,X,Y,W) + \mathrm{Rm}(X,Y,Z,W) + \mathrm{Rm}(Y,Z,X,W) &= 0. \end{aligned}$$

Adding these and applying symmetry 2., we find that the first two column's cancel (each term in the second column cancels with term from the subsequent line of the first column). Applying symmetry 1. and symmetry 2. to the last column yeilds,

$$\begin{aligned} 0 &= \mathrm{Rm}(W,X,Y,Z) + \mathrm{Rm}(Z,Y,W,X) + \mathrm{Rm}(X,W,Z,Y) + \mathrm{Rm}(Y,Z,X,W) \\ &= \mathrm{Rm}(W,X,Y,Z) - \mathrm{Rm}(Y,Z,W,X) + \mathrm{Rm}(W,X,Y,Z) - \mathrm{Rm}(Y,Z,W,X) \\ &= 2\left[\mathrm{Rm}(W,X,Y,Z) - \mathrm{Rm}(Y,Z,W,X)\right]. \end{aligned}$$

□

Corollary 9.12. *The curvature tensor is a tensor.*

Proof. Consider the map,

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

We need to show this is $C^\infty(M)$ linear in each slot. In fact, since $R(X, Y)Z = -R(X, Y)Z$, we need only show this is for X and Z.

Now, just write it out. For example,

$$\begin{aligned} R(fX, Y)Z &= \nabla_{fX}(\nabla_Y Z) - \nabla_Y(\nabla_{fX}Z) - \nabla_{[fX,Y]}Z \\ &= f\nabla_X\nabla_Y Z - Y(f)\nabla_X Z - f\nabla_Y\nabla_X Z - \nabla_{-Y(f)X+f[X,Y]}Z \\ &= f\left(\nabla_X\nabla_Y Z - \nabla_Y\nabla_X Z - \nabla_{[X,Y]}Z\right) - Y(f)\nabla_X Z + Y(f)\nabla_X Z \\ &= fR(X, Y)Z. \end{aligned}$$

The proof for tensorality in Z is left as an exercise. □

Remark 9.13. Given a connection, ∇ on a vector bundle E, the curvature tensor is always a tensor. One must check that if Z is replaced by $s \in \Gamma(E)$ in the proof that tensorality holds for s. The conclusion then is that,

$$(X, Y, s) \in \mathfrak{X}(M) \times \mathfrak{X}(M) \times \Gamma(E) \mapsto R(X, Y)s \in \Gamma(E)$$

is $C^\infty(M)$ tri-linear and hence determines a section $R \in \Gamma(T^*M \otimes T^*M \otimes E^* \otimes E)$. The proof of this fact, uses the Liebniz rule for a connection only.

9.2.2 Ricci and Scalar Curvature

The curvature tensor is rather difficult to deal with, and may be quite complicated. Simpler curvature tensors may be obtained by averaging. That is, by taking traces.

Definition 9.14. The *Ricci curvature* Ric is the trace

$$\mathrm{Ric}(X, Y) = \mathrm{Tr}(Z \mapsto R(Z, X)Y).$$

The curvature identities imply that Ric is symmetric:

$$\mathrm{Ric}(X, Y) = \mathrm{Ric}(Y, X).$$

While we're at it, why don't we take yet another trace?

Definition 9.15. The *scalar curvature* is the trace of the Ricci curvature with respect to the metric,

$$S = \mathrm{Tr}\mathrm{Ric}^\flat.$$

Since Ric is symmetric, we may take the trace by raising in either slot to produce an endomorphism and then take the trace obtaining the same value for S.

In coordinates, if we write $\mathrm{Ric} = \mathrm{Ric}_{ij} dx^i \otimes dx^j$ with

$$\mathrm{Ric}_{ij} = \mathrm{Ric}(\partial_i, \partial_j)$$

then raising in the first slot gives,

$$g^{ik}\mathrm{Ric}_{kj},$$

while raising in the second slot gives, by symmetry, the same thing,

$$g^{ki}\mathrm{Ric}_{jk} = g^{ik}\mathrm{Ric}_{kj}.$$

Thus we may unambiguously write,

$$S = \mathrm{Ric}^i_i$$

for either metric contraction. That is, explicitly,

$$\mathrm{Ric}^\flat = \mathrm{Ric}^i_j \partial_i \otimes dx^j = g^{ik}\mathrm{Ric}_{kj}\partial_i \otimes dx^j$$

and then

$$\begin{aligned} S &= \mathrm{Tr}\left(g^{ik}\mathrm{Ric}_{kj}\partial_i \otimes dx^j\right) \\ &= g^{ik}\mathrm{Ric}_{kj}\mathrm{Tr}(\partial_i \otimes dx^j) \\ &= g^{ik}\mathrm{Ric}_{kj} dx^j(\partial_i) \\ &= g^{ik}\mathrm{Ric}_{kj}\delta^j_i \\ &= g^{ik}\mathrm{Ric}_{ki} \\ &= \mathrm{Ric}^i_i. \end{aligned}$$

Many results in Riemannian geometry pertain to upper or lower bounds on the Ricci or scalar curvature. We will see some of these results in the final week.

9.3 Week 09, Lecture 03: Curvature Operators and Sectional Curvature

9.3.1 The Curvature Operator

Definition 9.16. For each X, Y, define the curvature operator, $R_{X,Y} \in \Gamma(E^* \otimes E)$ by,

$$R_{X,Y} : s \mapsto R(X,Y)s$$

The assignment $(X,Y) \mapsto R_{X,Y}$ is antisymmetric by the curvature tensor symmetries, and hence induces a map

$$\Lambda^2 TM = TM \wedge TM \to E^* \otimes E$$

sending indecomposable elements,

$$X \wedge Y \mapsto R_{X,Y}.$$

Using more symmetries, we can actually use the operator to induce a symmetric, bi-linear form on $TM \wedge TM = \Lambda^2 TM$.

Proposition 9.17. *The assignment,*

$$Q(X \wedge Y, Z \wedge W) = R(X,Y,Z,W)$$

defines a symmetric, bilinear form on $\Lambda^2 TM$. The metric dual of Q is a self-adjoint, endomorphism $\mathcal{R} : \Lambda^2 TM \to \Lambda^2 TM$.

Proof. We already observed that the right hand side of the definition of Q is anti-symmetric in X and Y. It is also anti-symmetric in Z and W hence is well defined. The fact that Q is symmetric follows from the curvature identity,

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

Since Q is symmetric, there exists unique, self-adjoint endomorphism $\mathcal{R} : \Lambda^2 TM \to \Lambda^2 TM$ such that,

$$g(\mathcal{R}(X \wedge Y), Z \wedge W) = g(X \wedge Y, \mathcal{R}(Z \wedge W))$$

where g is the induced metric on $\Lambda^2 TM$ from the metric g on TM. But either side is precisely the definition of the metric dual. $\square$

9.3.2 Sectional Curvature

Definition 9.18. A 2-plane Π at $x \in M$ is a two dimensional subspace of T_xM.

Definition 9.19. The *sectional curvature* of a two-plane Π is defined to be,

$$K_\Pi = \frac{Q(X \wedge Y, Y \wedge X)}{|X \wedge Y|_g^2}.$$

where $\Pi = \text{span} X, Y$ for $X, Y \in T_xM$.

Remark 9.20. Pay close to attention to the fact that we have reversed the orientation in the second argument to Q. This convention ensures that we obtain the correct sign for the sectional curvatures. For example, spheres have positive sectional curvature while hyperbolic space has negative sectional curvature. If we chose the opposite convention for the definition of the curvature tensor, then we would not need to reverse the orientation in the second argument. But then we would have to remember to switch the order of X and Y in the definition of the curvature tensor, and that seems a far greater sin.

We will see shortly that there is a close connection between sectional curvature and the curvature operators, which provides us a means of unambiguously choosing the appropriate sign for sectional curvatures as in the definition so that positive sectional curvature correspond to positive curvature operator.

Proposition 9.21. *The sectional curvature K is well defined, independent of the chosen oriented basis X, Y.*

Proof. This follows since if X', Y' is any other basis, so that $X' = AY$ and $Y' = AY$ for $A : \Pi \to \Pi$ and invertible, linear transformation, we have,

$$X' \wedge Y' = \det A X \wedge Y, \quad \text{and} \quad |X' \wedge Y'|^2 = (\det A)^2 |X \wedge Y|,$$

and the fact that Q is bilinear. □

Remark 9.22. Notice that the orientation of Π is irrelevant since Q is symmetric:

$$Q(Y \wedge X, X \wedge Y) = Q(X \wedge Y, Y \wedge X)$$

and $|X \wedge Y| = |Y \wedge X|$.

Also, another way to see that $K = K_\Pi$ does not depend on the chosen basis X, Y is that

$$K = Q(\frac{X \wedge Y}{|X \wedge Y|}, \frac{Y \wedge X}{|Y \wedge X|})$$

Thinking of

$$\frac{X \wedge Y}{|X \wedge Y|}$$

as a unit length *two-vector*, the change of basis computation in the proof shows that (up to orientation), there is a unique *two-vector* representing Π.

A rather useful concept here is the notion of a *curvature-like* function.

Definition 9.23. A *curvature-like function* is a section F of $T^*M \otimes T^*M \otimes T^*M \otimes TM$ satisfying the symmetries of the curvature tensor:

1. $F(X, Y)Z = -F(Y, X)Z$,
2. $g(F(X, Y)Z, W) = -g(F(X, Y)W, Z)$,
3. $F(X, Y)Z + F(Y, Z)X + F(Z, X)Y = 0$,
4. $g(F(X, Y)Z, W) = g(F(Z, W)X, Y)$.

Proposition 9.24. *The sectional curvatures determine the curvature tensor. That is, if two curvature-like functions, F_1, F_2 have the same sectional curvatures for all 2-planes, then $F_1 = F_2$.*

Proof. It suffices to show that $K_F = 0$ for all two-planes implies $F = 0$. For then, by multi-linearity, $K_{F_1} = K_{F_2}$ if and only if $K_{F_1 - F_2} = 0$ and hence $K_{F_1 - F_2} = 0$ implies $F_1 = F_2$. Notice here that we used that fact that if F_1, F_2 are curvature-like functions, then so too is $F_1 - F_2$.

By the assumption then,

$$F(X, Y, Y, X) = 0$$

for all X, Y. Using this equation, multi-linearity and the symmetries,

$$\begin{aligned} 0 &= F(W + X, Y, Y, W + X) \\ &= F(W, Y, Y, W) + F(X, Y, Y, W) + F(W, Y, Y, X) + F(X, Y, Y, X) \\ &= 2F(X, Y, Y, W). \end{aligned}$$

Then

$$F(W, Y+Z, Y+Z, X) = 0$$

and a similar computation gives

$$F(W, Z, Y, X) = -F(W, Y, Z, X)$$

But symmetry 1. implies

$$F(W, Z, Y, X) = F(Y, W, Z, X).$$

In other words, F is unchanged by a 2 cyclic permutation of the first three slots. Since 2 and 3 are co-prime, F is unchanged by any cyclic permutation of the first three slots. But the Bianchi identity (symmetry 3.) now shows that

$$0 = F(X, Y, Z, W) + F(Y, Z, X, W) = F(Z, X, Y, W) = 3F(X, Y, Z, W)$$

for any X, Y, Z, W. □

10 Week 10

10.1 Week 10, Lecture 01: Positive and Constant Curvature

10.1.1 Positive Curvature Operator

Definition 10.1. The *curvature operator* in direction X is the endomorphism,

$$R_X : Y \in TM \mapsto R(Y, X)X \in TM.$$

The *Ricci curvature* in direction X is the trace:

$$\text{Ric}(X) = \text{Trace} R_X.$$

Notice that $\text{Ric}(X) = \text{Ric}(X, X)$. This is just the usual relationship between a bilinear form and a quadratic form.

Definition 10.2. We say that R_X is a positive (negative) operator if for every Y linearly independent of X,

$$g(R_X(Y), Y) > (<) 0.$$

Non-negative and Non-positive operators are those where the inequality need not be strict.

Note that $R_X(X) = 0$ by the curvature symmetries, and this is why required Y linearly independent of X.

Remark 10.3. Using the symmetries of the curvature tensor, show that R_X is self-adjoint. That is,

$$g(R_X(Y), Z) = g(Y, R_X(Z))$$

for all Y, Z. Therefore, R_X is diagaonalisable and hence has a full set of n-eigenvalues (some of which may be repeated, or even 0). In particular, if λ is an eigenvalue of R_X with eigenvector Y (which we may assume is unit length), then

$$g(R_X(Y), Y) = \lambda$$

and we see that $R_X > 0$ if and only if it has all positive eigenvalues, $R \geq 0$ if and only if all eigenvalues are non-negative. Analogous statements also apply to $R_X < 0$ and $R_X \leq 0$ with negative, and non-positive eigenvalues respectively.

Lemma 10.4. *For a 2-plane Π, $K_\pi > 0$ if and only if,*

$$R_X(Y) > 0$$

for all X, Y spanning Π. In particular, $K_\Pi > 0$ for every Π if and only if $R_X > 0$ for every X.

Proof. We have the following formula then for the sectional curvature,

$$\begin{aligned} K_\Pi &= \frac{Q(X \wedge Y, Y \wedge X)}{|X \wedge Y|_g^2} \\ &= \frac{1}{|X \wedge Y|_g^2} R(X, Y, Y, X) \\ &= \frac{1}{|X \wedge Y|_g^2} g(R(X, Y)Y, X) \\ &= \frac{1}{|X \wedge Y|_g^2} g(R_Y(X), X). \end{aligned}$$

Thus we see that $K_\Pi > 0$ if and only if $R_Y > 0$.

The last result follows simply by applying this result to arbitrary X, Y linearly independent. □

Remark 10.5. We can now obtain bounds for K_Π in terms of the curvature operator R_X. Choose any unit length X and Y, with Y orthogonal to X for which we then have,

$$g(R_X(Y), Y) = K_{X \wedge Y}.$$

Writing Y as a linear combination of eigen-vectors of R_X, we obtain

$$\lambda_- \leq K_{X \wedge Y} \leq \lambda_+$$

where λ^- denotes the minimum eigenvalue of R_X and λ^+ denotes the maximum eigenvalue of R_X.

Whilst on the topic of eigen-values, we also have that,

$$\operatorname{Ric}(X) = \lambda_1 + \cdots + \lambda_n$$

is the (unscaled by $1/n$) average curvature in the X direction.

10.1.2 Constant Curvature Spaces

Proposition 10.6 (Gauss' Formula)**.** *Let $F : M^n \to (\bar{M}^{n+1}, \bar{g})$ be an immersion. Let $\bar{\nabla}$ be the Levi-Civita connection on $\bar{M}$, ∇ the induced Levi-Civita connection on M, and A the second fundamental form. Then*

$$\bar{R}(\bar{X}, \bar{Y}, \bar{Z}, \bar{W}) = R(X, Y, Z, W) + A(X, Z)A(Y, W) - A(X, W)A(Y, Z)$$

for any $X, Y, Z, W \in \Gamma(TM)$ and where $\bar{X} = dF \cdot X$.

Proof. It follows from the decomposition,

$$\bar{\nabla}_{\bar{X}}\bar{Y} = \nabla_X Y + h(X, Y)$$

where h is the second fundamental form, and

$$A(X, Y) = \bar{g}(h(X, Y), \nu)$$

where ν is a unit, normal vector field. □

Corollary 10.7. *Let ∇ be the Levi-Civita connection on the sphere $\mathbb{S}^n$ equipped with the "round" metric, i.e. the pull-back metric of the Euclidean metric under the standard embedding of $\mathbb{S}^n \to \mathbb{R}^{n+1}$. Then*

$$R(X, Y)Z = g(Y, Z)X - g(X, Z)Y.$$

Thus every sectional curvature $K_\Pi = 1$. For the sphere or radius R, we have

$$R(X, Y)Z = \tfrac{1}{R^2}\left(g(Y, Z)X - g(X, Z)Y\right)$$

and $K_\Pi = 1/R^2$.

Proof. We have,

$$A(X, Y) = -g(X, Y)$$

since the unit normal $\mathbf{n}(x) = x$ and so $\mathcal{W}(X) = -\nabla_X \mathbf{n} = -X$. Now use the fact that $Rm \equiv 0$ on Euclidean space and apply the Gauss equation. For radius R, one may find the result by scaling. □

Definition 10.8. *Minkowski space*, written $\mathbb{R}^{n,1}$ is the smooth manifold $\mathbb{R}^n \times \mathbb{R}$ equipped with the metric,

$$g = \sum_{i=1}^{n} dx^i \otimes dx^i - dt \otimes dt.$$

Hyperbolic space, written $\mathbb{H}^n$ is the smooth manifold,

$$\{z = (x,t) \in \mathbb{R}^n \times \mathbb{R} : g(z,z) = |x|^2 - t^2 = -1, t > 0\}$$

equipped with the pull-back of the Minkowski metric. Here $|x|$ is the computed with the usual Euclidean metric, $\sum_{i=1}^n dx^i \otimes dx^i$.

We may also consider Hyperbolic space as $\mathbb{S}^n(i) = \{\sqrt{g(z,z)} = i\} \cap \{t > 0\}$ the sphere with radius i. Note that the set $\{\sqrt{g(z,z)} = i\}$ has two connected components, and we are choosing the one with $t > 0$. The other component is the reflection across the hyperplane $\{t = 0\}$.

Proposition 10.9. *Hyperbolic space is a Riemannian manifold with curvature tensor,*

$$R(X,Y)Z = -g(Y,Z)X + g(X,Z)Y.$$

and hence has constant sectional curvature -1*.*

Proof. That $\mathbb{H}^n$ is Riemannian follows since the null-cone $t = \pm|x|$ divides tangent vectors into those which have positive or negative length and $\mathbb{H}^n$ lies in the region with positive length and asymptotes to the null cone at infinity.

The curvature tensor computation is similar to the sphere, but taking care with the mixed-signature Minkowski metric. One also needs to check directly that the Levi-Civita connection is just the Euclidean directional derivative. □

Theorem 10.10. *Let* (M,g) *be a complete, smooth Riemannian manifold of constant sectional curvature* K*. Let* $\tilde{M}$ *denote the universal cover with covering map* $\pi : \tilde{M} \to M$ *and equip* $\tilde{M}$ *with the pull back metric,* $\tilde{g} = \pi^* g$*. Then*

$$(\tilde{M}, \tilde{g}) = \begin{cases} \mathbb{S}^n(1/\sqrt{K}), & K > 0 \\ \mathbb{R}^n, & K = 0 \\ \mathbb{S}^n(i/\sqrt{-K}), & K < 0. \end{cases}$$

Definition 10.11. A complete manifold with constant sectional curvature is called a *space form.* By the theorem above, it is a quotient of one of the three model spaces from the theorem.

Theorem 10.12 (B\"ohm-Wilking, 2006)**.** *Let* (M,g) *be compact with positive curvature operator. Then* M *admits a metric of constant positive sectional curvature. Therefore* M *is diffeomorphic to a space-form and admits a metric making it a space-form.*

10.2 Week 10, Lecture 02: Classical Results

10.2.1 Second Variation of Length and Bonnet-Myers

Let $\gamma : [a, b] \to M$ be a geodesic and F a smooth variation,

$$F : [a, b] \times (-\epsilon, \epsilon) \to M$$

with $F(t, 0) = \gamma(t)$ for all $t \in [a, b]$. Let us also assume that F fixes the endpoints:

$$F(a, s) = F(a, 0), \quad F(b, s) = F(b, 0)$$

for all $s \in (-\epsilon, \epsilon)$.

For ease of notation, define $\gamma_s(t) = F(t, s)$ so that $\gamma_0 = \gamma$ and each γ_s is a curve defined on $[a, b]$ with fixed endpoints. Let us also write

$$T = F_* \partial_t = \gamma_s'$$

for the tangent vector to γ_s and,

$$V = F_* \partial_s$$

for the variation vector. Finally, let

$$V^T = g(V, T)T, \quad V^\perp = V - g(V, T)T$$

be the components of V tangent to γ_s and orthogonal to γ_s respectively. In particular,

$$V = V^T + V^\perp \quad \text{and} \quad g(V^T, V^\perp) = 0.$$

Theorem 10.13 (Second Variation Formula)**.**

$$\frac{d}{ds}\bigg|_{s=0} L[\gamma_s] = \int_a^b |\nabla_T V^\perp|^2 + Rm(V^\perp, T, V^\perp, T) dt.$$

Proof. Since $\gamma_0 = \gamma$ is a geodesic we may assume that $g(T, T) = 1$ when $s = 0$. But note that we cannot assume $g(T, T)$ is constant. By the computation from the first variation formula (essentially just since $[V, T] = 0$),

$$\partial_s \sqrt{g(T, T)} = \frac{1}{\sqrt{g(T, T)}} g(\nabla_T V, T).$$

Differentiating again,

$$\begin{aligned}\partial_s^2 \sqrt{g(T,T)} &= \partial_s \left(\frac{1}{\sqrt{g(T,T)}} \right) g(\nabla_T V, T) + \frac{1}{\sqrt{g(T,T)}} \partial_s \left(g(\nabla_T V, T) \right) \\ &= -\frac{1}{\sqrt{g(T,T)}^{3/2}} \frac{1}{\sqrt{g(T,T)}} g(\nabla_T V, T) g(\nabla_T V, T) \\ &\quad + \frac{1}{\sqrt{g(T,T)}} \left(g(\nabla_V \nabla_T V, T) + g(\nabla_T V, \nabla_V T) \right).\end{aligned}$$

Evaluating at $s = 0$, where $g(T,T) = 1$, we have

$$\begin{aligned}\frac{d}{ds}\bigg|_{s=0} \sqrt{g(T,T)} &= -g(\nabla_T V, T)^2 + g(\nabla_V \nabla_T V, T) + g(\nabla_T V, \nabla_V T) &(7)\\ &= -g(\nabla_T V, V)^2 + g(\nabla_T V, \nabla_V T) + g(\nabla_T \nabla_V V, T) &(8)\\ &\quad + \mathrm{Rm}(V, T, V, T). &(9)\end{aligned}$$

Note that in the last line we used $[V, T] = 0$ for the curvature tensor.

For the second term in equation (7), we have

$$g(\nabla_T V, \nabla_V T) = g(\nabla_T V, \nabla_T V) = |\nabla_T V|^2,$$

again using $[V, T] = 0$. Now I claim that the first two terms of equation (7) give

$$-g(\nabla_T V, V)^2 + |\nabla_T V|^2 = |\nabla_T V^{\perp}|^2.$$

This follows since γ is a geodesic, hence

$$\begin{aligned}\nabla_T V^T &= \nabla_T \left[g(V,T)T \right] \\ &= g(\nabla_T V, T)T + g(V, \nabla_T T)T + g(V,T)\nabla_T T \\ &= g(\nabla_T V, T)T \\ &= (\nabla_T V)^T.\end{aligned}$$

Therefore, also we have

$$(\nabla_T V)^{\perp} = \nabla_T V - (\nabla_T V)^T = \nabla_T V - \nabla_T V^T = \nabla_T (V - V^T) = \nabla_T V^{\perp}.$$

Thus,

$$|\nabla_T V|^2 = |\nabla_T V^{\perp}|^2 + |\nabla_T V^T|^2 = |\nabla_T V^{\perp}|^2 + g(\nabla_T V, T)^2$$

as required.

For the third term in equation (7),

$$g(\nabla_T \nabla_V V, T) = \partial_t g(\nabla_V V, T) - g(\nabla_V V, \nabla_T T) = \partial_t g(\nabla_V V, T)$$

since γ is geodesic (and we've already set $s = 0$).

For the last term of equation (7),

$$\mathrm{Rm}(V, T, V, T) = \mathrm{Rm}(V^\perp + V^T, T, V^\perp + V^T, T) = \mathrm{Rm}(V^\perp, T, V^\perp, T).$$

Putting all this together, equation (7) becomes we have

$$\begin{aligned} \frac{d}{ds}\bigg|_{s=0} \sqrt{g(T,T)} &= -g(\nabla_T V, V)^2 + |\nabla_T V|^2 + g(\nabla_T \nabla_V V, T) + \mathrm{Rm}(V, T, V, T) \\ &= |\nabla_T V^\perp|^2 + \partial_t g(\nabla_V V, T) + \mathrm{Rm}(V^\perp, T, V^\perp, T). \end{aligned}$$

Since $F(a, s) = F(a, 0)$ and $F(b, s) = F(b, 0)$ we have $V(a, 0) = V(b, 0) = 0$. Then integrating from a to b we obtain,

$$\begin{aligned} \frac{d}{ds}\bigg|_{s=0} L[\gamma_s] &= \int_a^b |\nabla_T V^\perp|^2 + \partial_t g(\nabla_V V, T) + \mathrm{Rm}(V^\perp, T, V^\perp, T) dt \\ &= \int_a^b |\nabla_T V^\perp|^2 + \mathrm{Rm}(V^\perp, T, V^\perp, T) dt + g(\nabla_V V, T)|_a^b \\ &= \int_a^b |\nabla_T V^\perp|^2 + \mathrm{Rm}(V^\perp, T, V^\perp, T) dt. \end{aligned}$$

□

Corollary 10.14. *Let γ be a length minimising geodesic. Then for all $V \in \mathfrak{X}_\gamma(M)$ with $V(a) = V(b) = 0$ and $g(V, \gamma') = 0$*

$$Q(V) = \int_a^b |\nabla_{\gamma'} V|^2 + Rm(V, \gamma', V, \gamma') dt \geq 0.$$

Remark 10.15. The integrand is called the *index form*. It arises from the bilinear form,

$$I(V_1, V_2) = \int_a^b g(\nabla_{\gamma'} V_1^\perp, \nabla_{\gamma'} V_2^\perp) + \mathrm{Rm}(V_1^\perp, \gamma', V_2^\perp, \gamma') dt.$$

Notice that I is bilinear and $Q(V) = I(V, V)$ is the associated quadratic form. In particular, if either V_1 or V_2 is proportional to γ', then $Q(V) = 0$ so we really think of I (and hence Q) acting on $(\gamma')^\perp$.

The corollary says that a length minimising geodesic has non-negative index form.

Proof. If V is given as above, then let

$$F(t, s) = \exp_{\gamma(t)}(sV(t)).$$

Then F is a variation of γ fixing the end points and such that $F_*\partial_s = V$ when $s = 0$. By assumption $s = 0$ is a local minimum for $s \mapsto L[\gamma_s]$ and now just apply the second derivative test and the second variation formula. □

Theorem 10.16 (Bonnet-Myers). *Let (M^n, g) be a complete Riemannian manifold such that*

$$Ric \geq \frac{n-1}{r^2} g$$

(i.e. for all vector fields X, $Ric(X, X) \geq\geq \frac{n-1}{r^2} g(X, X)$ for some $r > 0$. Then M is compact with diameter bounded above by πr.

Remark 10.17. An n dimensional sphere of radius r equipped with the round metric satisfies

$$\mathrm{Ric} = \frac{n-1}{r^2} g$$

and has diameter πr. Thus the theorem is *sharp*.

If we relax the assumption from a positive lower bound for Ric to just Ric > 0, then the theorem is false. A counter example is given by the paraboloid of revolution $(r\cos\theta, r\sin\theta, r^2)$ which is not compact. The study of non-compact manifolds with non-negative Ricci is an active area of study.

Proof. First, suppose we can prove the diameter bound $D < \infty$ where the diameter is defined to be

$$D = \sup_M \{d(x, y) : x, yinM\}.$$

Then at point $x \in M$,

$$M = \exp_x(\bar{B}_D(0))$$

is the image of the compact set $\bar{B}_D(0) \subset T_xM$ under the continuous function $\exp_x$ hence is compact. Note that we use completeness here to ensure

$\exp_x$ is defined on all of $\bar{B}_D(0)$. Without completeness, we need not have compactness. For example, any bounded open subset of Euclidean space has bounded diameter, but is not compact.

Now for the diameter bound. Let $x, y \in M$ and we need to show that $d(x, y) \leq \pi r$. By the Hopf-Rinow theorem, there exists a (not necessarily unique) length minimising geodesic γ joining x to y so that $L[\gamma] = d(x, y)$ and we the theorem will follow by showing that

$$L = L[\gamma] \leq \pi r.$$

To this end, assume that γ is parmetrised with unit speed on $[0, L]$. Let $\{e_i\}$ be an orthonormal basis for $T_x M$ with $e_n = \gamma'$. Let $E_i(t)$ be the parallel transport of e_i along γ. Then since γ' is parallel along γ we have $\gamma' = E_n$. Moreover, since $e_i \perp e_n$ for $1 \leq i \leq n-1$, and parallel transport is an isometry, $\{E_i\}$ is an orthonormal frame along γ. In particular, $E_i \perp \gamma'$ for $1 \leq i \leq n-1$. Define

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

for $1 \leq i \leq n-1$ and we also have $Y_i \perp \gamma'$.

From the corollary above, for each i we have,

$$0 \leq I(Y_i) = \int_0^L |\nabla_T Y_i|^2 + \mathrm{Rm}(Y_i, \gamma', Y_i, \gamma')dt.$$

Since $\nabla_T E_i = 0$ and E_i is unit length,

$$|\nabla_T Y_i|^2 = \left|\nabla_T \left[\sin\left(\frac{\pi t}{L}\right)\right] E_i\right|^2 = \frac{\pi^2}{L^2}\cos^2\left(\frac{\pi t}{L}\right)$$

We also have

$$\mathrm{Rm}(Y_i, \gamma', Y_i, \gamma') = \sin^2\left(\frac{\pi t}{L}\right)\mathrm{Rm}(E_i, \gamma', E_i, \gamma').$$

Thus we have

$$0 \leq I(Y_i) = \int_0^L \frac{\pi^2}{L^2}\cos^2\left(\frac{\pi t}{L}\right) + \sin^2\left(\frac{\pi t}{L}\right)\mathrm{Rm}(E_i, \gamma', E_i, \gamma')dt.$$

Now observe that, since E_i is orthonormal, $E_n = \gamma'$, and using the assumed lower bound on Ric,

$$\begin{aligned}\frac{n-1}{r^2} = \frac{n-1}{r^2} g(\gamma', \gamma') \leq \operatorname{Ric}(\gamma', \gamma') &= \sum_{i=1}^{n} \operatorname{Rm}(E_i, \gamma', \gamma', E_i) \\ &= -\sum_{i=1}^{n-1} \operatorname{Rm}(E_i, \gamma', E_i, \gamma').\end{aligned}$$

Thus summing $I(Y_i)$ over over i we obtain

$$\begin{aligned}0 &\leq \sum_{i=1}^{n-1} I(Y_i) = \int_0^L (n-1)\frac{\pi^2}{L^2} \cos^2\left(\frac{\pi t}{L}\right) + \sin^2\left(\frac{\pi t}{L}\right) \sum_{i=1}^{n-1} \operatorname{Rm}(E_i, \gamma', E_i, \gamma') \\ &\leq \int_0^L (n-1)\frac{\pi^2}{L^2} \cos^2\left(\frac{\pi t}{L}\right) - \sin^2\left(\frac{\pi t}{L}\right) \frac{n-1}{r^2} \\ &= \frac{1}{2}\left(\frac{(n-1)\pi^2}{L^2} - \frac{n-1}{r^2}\right).\end{aligned}$$

Thus $L \leq \pi r$. □

One question that may arise, is what is the significance of the choice $Y_i = \sin(\frac{\pi t}{L})E_i(t)$? It is simply that it works, or is there some deeper reason for the choice. The answer is the latter.

Definition 10.18. A *Jacobi field* along a geodesic γ is a section $J \in \mathfrak{X}_\gamma(M)$ such that,

$$\nabla_{\gamma'} \nabla_{\gamma'} J + \operatorname{Ric}_{\gamma'}(J) = 0.$$

Definition 10.19. A *geodesic variation* is a map $F : [a, b] \times (-\epsilon, \epsilon) \to M$ such that $\gamma_s = F(\cdot, s)$ is a geodesic for each $s \in (-\epsilon, \epsilon)$.

The connection between Jacobi fields and geodesic variations is given by the next proposition.

Proposition 10.20. *Let γ_s be a geodesic variation. Then the variation vector field, J is a Jacobi field. Conversely, if J is a Jacobi field along γ, then there exists a geodesic variation γ_s of γ with variation vector field J.*

The proof can be found in most texts on Riemannian Geometry. Returning to the vector fields Y_i, we have the following lemma.

Lemma 10.21. *Let γ be a geodesic on the round sphere. Then Y_i is a Jacobi field.*

Thus we may interpret the Bonnet-Myers theorem as a sort of comparison of geodesic variations for Riemannian manifolds with positive lower bounds for Ricci, with an appropriate sphere. This then yields a comparison on the *geometric* quantity, diameter. One might expect that there are other geometric comparisons, such as volume and this is indeed the case and such questions lead to the rich area of comparison geometry.