Basic Riemannian Geometry

F.E. Burstall
Department of Mathematical Sciences
University of Bath

Introduction

My mission was to describe the basics of Riemannian geometry in just three hours of lectures, starting from scratch. The lectures were to provide background for the analytic matters covered elsewhere during the conference and, in particular, to underpin the more detailed (and much more professional) lectures of Isaac Chavel. My strategy was to get to the point where I could state and prove a Real Live Theorem: the Bishop Volume Comparison Theorem and Gromov's improvement thereof and, by appalling abuse of OHP technology, I managed this task in the time alloted. In writing up my notes for this volume, I have tried to retain the breathless quality of the original lectures while correcting the mistakes and excising the out-right lies.

I have given very few references to the literature in these notes so a few remarks on sources is appropriate here. The first part of the notes deals with analysis on differentiable manifolds. The two canonical texts here are Spivak [5] and Warner [6] and I have leaned on Warner's book in particular. For Riemannian geometry, I have stolen shamelessly from the excellent books of Chavel [1] and Gallot–Hulin–Lafontaine [3]. In particular, the proof given here of Bishop's theorem is one of those provided in [3].

1 What is a manifold?

What ingredients do we need to do Differential Calculus? Consider first the notion of a continuous function: during the long process of abstraction and generalisation that leads from Real Analysis through Metric Spaces to Topology, we learn that continuity of a function requires no more structure on the domain and co-domain than the idea of an open set.

By contrast, the notion of differentiability requires much more: to talk about the difference quotients whose limits are partial derivatives, we seem to require that the (co-)domain have a linear (or, at least, affine) structure.

However, a moment's thought reveals that differentiability is a completely *local* matter so that all that is really required is that the domain and co-domain be *locally* linear, that is, each point has a neighbourhood which is homeomorphic to an open subset of some linear space. These ideas lead us to the notion of a *manifold*: a topological space which is locally Euclidean and on which there is a well-defined differential calculus.

We begin by setting out the basic theory of these spaces and how to do Analysis on them.

1.1 Manifolds

Let M be a Hausdorff, second countable[1], connected topological space.

M is a C^r *manifold* of dimension n if there is an open cover $\{U_\alpha\}_{\alpha\in I}$ of M and homeomorphisms $x_\alpha : U_\alpha \to x_\alpha(U_\alpha)$ onto open subsets of $\mathbb{R}^n$ such that, whenever $U_\alpha \cap U_\beta \neq \emptyset$,

$$x_\alpha \circ x_\beta^{-1} : x_\beta(U_\alpha \cap U_\beta) \to x_\alpha(U_\alpha \cap U_\beta)$$

is a C^r diffeomorphism.

Each pair (U_α, x_α) called a *chart*.

Write $x_\alpha = (x^1, \dots, x^n)$. The $x^i : U_\alpha \to \mathbb{R}$ are *coordinates*.

1.1.1 Examples

1. Any open subset $U \subset \mathbb{R}^n$ is a C^∞ manifold with a single chart $(U, 1_U)$.
2. Contemplate the unit sphere $S^n = \{v \in \mathbb{R}^{n+1} : \|v\| = 1\}$ in $\mathbb{R}^{n+1}$. Orthogonal projection provides a homeomorphism of any open hemisphere onto the open unit ball in some hyperplane $\mathbb{R}^n \subset \mathbb{R}^{n+1}$. The sphere is covered by the $(2n + 2)$ hemispheres lying on either side of the coordinate hyperplanes and in this way becomes a C^∞ manifold (exercise!).
3. A good supply of manifolds is provided by the following version of the Implicit Function Theorem [6]:

   Theorem. *Let $f : \Omega \subset \mathbb{R}^n \to \mathbb{R}$ be a C^r function $(r \geq 1)$ and $c \in \mathbb{R}$* a regular value, *that is, $\nabla f(x) \neq 0$, for all $x \in f^{-1}\{c\}$.*

   Then $f^{-1}\{c\}$ is a C^r manifold.

   Exercise. Apply this to $f(x) = \|x\|^2$ to get a less tedious proof that S^n is a manifold.

[1] This means that there is a countable base for the topology of M.

4. An open subset of a manifold is a manifold in its own right with charts $(U_\alpha \cap U, x_\alpha|_{U_\alpha \cap U})$.

1.1.2 Functions and maps

A continuous function $f : M \to \mathbb{R}$ is C^r if each $f \circ x_\alpha^{-1} : x_\alpha(U_\alpha) \to \mathbb{R}$ is a C^r function of the open set $x_\alpha(U_\alpha) \subset \mathbb{R}^n$.

We denote the vector space of all such functions by $C^r(M)$.

Example. Any coordinate function $x^i : U_\alpha \to \mathbb{R}$ is C^r on U_α.

Exercise. The restriction of any C^r function on $\mathbb{R}^{n+1}$ to the sphere S^n is C^r on S^n.

In the same way, a continuous map $\phi : M \to N$ of C^r manifolds is C^r if, for all charts (U, x), (V, y) of M and N respectively, $y \circ \phi \circ x^{-1}$ is C^r on its domain of definition.

A slicker formulation[2] is that $h \circ \phi \in C^r(M)$, for all $h \in C^r(M)$.

At this point, having made all the definitions, we shall stop pretending to be anything other than Differential Geometers and henceforth take $r = \infty$.

1.2 Tangent vectors and derivatives

We now know what functions on a manifold are and it is our task to differentiate them. This requires some less than intuitive definitions so let us step back and remind ourselves of what differentiation involves.

Let $f : \Omega \subset \mathbb{R}^n \to \mathbb{R}$ and contemplate the derivative of f at some $x \in \Omega$. This is a linear map $\mathrm{d}f_x : \mathbb{R}^n \to \mathbb{R}$. However, it is better for us to take a dual point of view and think of $\mathbf{v} \in \mathbb{R}^n$ is a linear map $\mathbf{v} : C^\infty(M) \to \mathbb{R}$ by

$$\mathbf{v}f \stackrel{\text{def}}{=} \mathrm{d}f_x(\mathbf{v}).$$

The Leibniz rule gives us

$$\mathbf{v}(fg) = f(x)\mathbf{v}(g) + \mathbf{v}(f)g(x). \tag{1.1}$$

Fact. *Any linear* $\mathbf{v} : C^\infty(\Omega) \to \mathbb{R}$ *satisfying* (1.1) *arises this way.*

Now let M be a manifold. The preceding analysis may give some motivation to the following

[2]It requires a little machinery, in the shape of bump functions, to see that this *is* an equivalent formulation.

Definition. A *tangent vector* at $m \in M$ is a linear map $\xi : C^\infty(M) \to \mathbb{R}$ such that

$$\xi(fg) = f(m)\xi(g) + \xi(f)g(m)$$

for all $f, g \in C^\infty(M)$.

Denote by M_m the vector space of all tangent vectors at m.

Here are some examples

1. For $\gamma : I \to M$ a (smooth) path with $\gamma(t) = m$, define $\gamma'(t) \in M_m$ by

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

   Fact. *All $\xi \in M_m$ are of the form $\gamma'(t)$ for some path γ.*

2. Let (U, x) be a chart with coordinates $x^1, \dots, x^n$ and $x(m) = p \in \mathbb{R}^n$. Define $\partial_{i|m} \in M_m$ by

   $$\partial_{i|m} f = \left.\frac{\partial(f \circ x^{-1})}{\partial x^i}\right|_p$$

   Fact. *$\partial_{1|m}, \dots, \partial_{n|m}$ is a basis for M_m.*

3. For $p \in U \subset \mathbb{R}^n$ open, we know that U_p is canonically isomorphic to $\mathbb{R}^n$ via

   $$\mathbf{v}f = \mathrm{d}f_p(\mathbf{v})$$

   for $\mathbf{v} \in \mathbb{R}^n$.

4. Let $M = f^{-1}\{c\}$ be a regular level set of $f : \Omega \subset \mathbb{R}^n \to \mathbb{R}$. One can show that M_m is a linear subspace of $\Omega_m \cong \mathbb{R}^n$. Indeed, under this identification,

   $$M_m = \{\mathbf{v} \in \mathbb{R}^n : \mathbf{v} \perp \nabla f_m\}.$$

Now that we have got our hands on tangent vectors, the definition of the derivative of a function as a linear map on tangent vectors is almost tautological:

Definition. For $f \in C^\infty(M)$, the *derivative* $\mathrm{d}f_m : M_m \to \mathbb{R}$ *of f at $m \in M$* is defined by

$$\mathrm{d}f_m(\xi) = \xi f.$$

CAMBRIDGE

Cambridge University Press
0521777496 - Spectral Theory and Geometry: ICMS Instructional Conference, Edinburgh 1998
Edited by Brian Davies and Yuri Safarov
Excerpt
More information

We note:

1. Each $\mathrm{d}f_m$ is a linear map and the Leibniz Rule holds:

$$\mathrm{d}(fg)_m = g(m)\mathrm{d}f_m + f(m)\mathrm{d}g_m.$$

2. By construction, this definition coincides with the usual one when M is an open subset of $\mathbb{R}^n$.

Exercise. If f is a constant map on a manifold M, show that each $\mathrm{d}f_m = 0$.

The same circle of ideas enable us to differentiate maps between manifolds:

Definition. For $\phi : M \to N$ a smooth map of manifolds, the *tangent map* $\mathrm{d}\phi_m : M_m \to N_{\phi(m)}$ *at* $m \in M$ is the linear map defined by

$$\mathrm{d}\phi_m(\xi)f = \xi(f \circ \phi),$$

for $\xi \in M_m$ and $f \in C^\infty(N)$.

Exercise. Prove the chain rule: for $\phi : M \to N$ and $\psi : N \to Z$ and $m \in M$,

$$\mathrm{d}(\psi \circ \phi)_m = \mathrm{d}\psi_{\phi(m)} \circ \mathrm{d}\phi_m.$$

Exercise. View $\mathbb{R}$ as a manifold (with a single chart!) and let $f : M \to \mathbb{R}$. We now have two competing definitions of $\mathrm{d}f_m$. Show that they coincide.

The *tangent bundle of* M is the disjoint union of the tangent spaces:

$$TM = \coprod_{m\in M} M_m.$$

1.3 Vector fields

Definition. A *vector field* is a linear map $X : C^\infty(M) \to C^\infty(M)$ such that

$$X(fg) = f(Xg) + g(Xf).$$

Let $\Gamma(TM)$ denote the vector space of all vector fields on M.

We can view a vector field as a map $X : M \to TM$ with $X(m) \in M_m$: indeed, we have

$$X_{|m} \in M_p$$

where

$$X_{|m} f = (Xf)(m).$$

In fact, vector fields can be shown to be exactly those maps $X : M \to TM$ with $X(m) \in M_m$ which satisfy the additional smoothness constraint that for each $f \in C^\infty(M)$, the function $m \mapsto X(m)f$ is also C^∞.

The *Lie bracket* of $X, Y \in \Gamma(TM)$ is $[X,Y] : C^\infty(M) \to C^\infty(M)$ given by

$$[X,Y]f = X(Yf) - Y(Xf).$$

The point of this definition is contained in the following

Exercise. Show that $[X,Y] \in \Gamma(TM)$ also.

The Lie bracket is interesting for several reasons. Firstly it equips $\Gamma(TM)$ with the structure of a Lie algebra; secondly, it, and operators derived from it, are the only differential operators that can be defined on an arbitrary manifold without imposing additional structures such as special coordinates, a Riemannian metric, a complex structure or a symplectic form.

There is an extension of the notion of vector field that we shall need later on:

Definition. Let $\phi : M \to N$ be a map. A *vector field along* ϕ is a map $X : M \to TN$ with

$$X(m) \in N_{\phi(m)},$$

for all $m \in M$, which additionally satisfies a smoothness assumption that we shall gloss over.

Denote by $\Gamma(\phi^{-1}TN)$ the vector space of all vector fields along ϕ.

Here are some examples:

1. If $c : I \to N$ is a smooth path then $c' \in \Gamma(\phi^{-1}TN)$.
2. More generally, for $\phi : M \to N$ and $X \in \Gamma(TM)$, $\mathrm{d}\phi(X) \in \Gamma(\phi^{-1}TN)$. Here, of course,

$$\mathrm{d}\phi(X)(m) = \mathrm{d}\phi_m(X_{|m}).$$

3. For $Y \in \Gamma(TN)$, $Y \circ \phi \in \Gamma(\phi^{-1}TN)$.

1.4 Connections

We would like to differentiate vector fields but as they take values in different vector spaces at different points, it is not so clear how to make difference quotients and so derivatives. What is needed is some extra structure: a *connection* which should be thought of as a "directional derivative" for vector fields.

Definition. A *connection on* TM is a bilinear map

$$\begin{aligned} TM \times \Gamma(TM) &\to TM \\ (\xi, X) &\mapsto \nabla_\xi X \end{aligned}$$

such that, for $\xi \in M_m$, $X, Y \in \Gamma(TM)$ and $f \in C^\infty(M)$,

1. $\nabla_\xi X \in M_m$;
2. $\nabla_\xi(fX) = (\xi f)X_{|m} + f(m)\nabla_\xi X$;
3. $\nabla_X Y \in \Gamma(TM)$.

A connection on TM comes with some additional baggage in the shape of two multilinear maps:

$$\begin{aligned} T_m &: M_m \times M_m \to M_m \\ R_m &: M_m \times M_m \times M_m \to M_m \end{aligned}$$

given by

$$\begin{aligned} T_m(\xi, \eta) &= \nabla_\xi Y - \nabla_\eta X - [X, Y]_{|m} \\ R_m(\xi, \eta)\zeta &= \nabla_\eta \nabla_X Z - \nabla_\xi \nabla_Y Z - \nabla_{[Y,X]_{|m}} \end{aligned}$$

where X, Y, $Z \in \Gamma(TM)$ with $X_{|m} = \xi$, $Y_{|m} = \eta$ and $Z_{|m} = \zeta$.

T_m and R_m are, respectively, the *torsion* and *curvature* at m of ∇.

Fact. *R and T are well-defined—they do not depend of the choice of vector fields X, Y and Z extending ξ, η and ζ.*

We have some trivial identities:

$$\begin{aligned} T(\xi, \eta) &= -T(\eta, \xi) \\ R(\xi, \eta)\zeta &= -R(\eta, \xi)\zeta. \end{aligned}$$

and, if each $T_m = 0$, we have the less trivial *First Bianchi Identity*:

$$R(\xi, \eta)\zeta + R(\zeta, \xi)\eta + R(\eta, \zeta)\xi = 0.$$

A connection ∇ on TN induces a similar operator on vector fields along a map $\phi : M \to N$. To be precise, there is a unique bilinear map

$$TM \times \Gamma(\phi^{-1}TN) \to TN$$
$$(\xi, X) \mapsto \phi^{-1}\nabla_\xi X$$

such that, for $\xi \in M_m$, $X \in \Gamma(TM)$, $Y \in \Gamma(\phi^{-1}TN)$ and $f \in C^\infty(M)$,

1. $\phi^{-1}\nabla_\xi Y \in N_{\phi(m)}$;
2. $\phi^{-1}\nabla_\xi(fY) = (\xi f)Y_{|\phi(m)} + f(m)\phi^{-1}\nabla_\xi Y$;
3. $\phi^{-1}\nabla_X Y \in \Gamma(\phi^{-1}TN)$ (this is a smoothness assertion);
4. If $Z \in \Gamma(TN)$ then $Z \circ \phi \in \Gamma(\phi^{-1}TN)$ and
$$\phi^{-1}\nabla_\xi(Z \circ \phi) = \nabla_{\mathrm{d}\phi_m(\xi)} Z.$$

$\phi^{-1}\nabla$ is the *pull-back of* ∇ *by* ϕ. The first three properties just say that $\phi^{-1}\nabla$ behaves like ∇, it is the last that essentially defines it in a unique way.

2 Analysis on Riemannian manifolds

2.1 Riemannian manifolds

A rich and useful geometry arises if we equip each M_m with an inner product:

Definition. A *Riemannian metric* g on M is an inner product g_m on each M_m such that, for all vector fields X and Y, the function

$$m \mapsto g_m(X_{|m}, Y_{|m})$$

is smooth.

A *Riemannian manifold* is a pair (M, g) with M a manifold and g a metric on M.

Here are some (canonical) examples:

1. Let $(\ ,\ )$ denote the inner product on $\mathbb{R}^n$.

   An open $U \subset \mathbb{R}^n$ gets a Riemannian metric via $U_m \cong \mathbb{R}^n$:
$$g_m(v, w) = (v, w).$$

2. Let $S^n \subset \mathbb{R}^{n+1}$ be the unit sphere. Then $S^n_m \cong m^\perp \subset \mathbb{R}^{n+1}$ and so gets a metric from the inner product on $\mathbb{R}^{n+1}$.

3. Let $D^n \subset \mathbb{R}^n$ be the open unit disc but define a metric by

$$g_z(v,w) = \frac{4(v,w)}{(1-|z|^2)^2}$$

(D^n, g) is *hyperbolic space*.

Much of the power of Riemannian geometry comes from the fact that there is a *canonical* choice of connection. Consider the following two desirable properties for a connection ∇ on (M,g):

1. ∇ is *metric*: $Xg(Y,Z) = g(\nabla_X Y, Z) + g(Y, \nabla_X Z)$.

2. ∇ is *torsion-free*: $\nabla_X Y - \nabla_Y X = [X,Y]$

Theorem. *There is a unique torsion-free metric connection on any Riemannian manifold.*

Proof. Assume that g is metric and torsion-free. Then

$$\begin{aligned} g(\nabla_X Y, Z) &= Xg(Y,Z) - g(Y, \nabla_X Z) \\ &= Xg(Y,Z) - g(Y,[X,Z]) - g(Y,\nabla_Z X)\ldots \end{aligned}$$

and eventually we get

$$\begin{aligned} 2g(\nabla_X Y, Z) = {} & Xg(Y,Z) + Yg(Z,Y) - Zg(X,Y) \\ & - g(X,[Y,Z]) + g(Y,[Z,X]) + g(Z,[X,Y]). \end{aligned} \tag{2.1}$$

This formula shows uniqueness and, moreover, *defines* the desired connection. □

This connection is the *Levi–Civita connection* of (M,g).

For detailed computations, it is sometimes necessary to express the metric and Levi–Civita connection in terms of local coordinates. So let (U,x) be a chart and $\partial_1, \ldots, \partial_n$ be the corresponding vector fields on U. We now define $g_{ij} \in C^\infty(U)$ by

$$g_{ij} = g(\partial_i, \partial_j)$$

and *Christoffel symbols* $\Gamma^k_{ij} \in C^\infty(U)$ by

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

(Recall that $\partial_{1|m}, \dots, \partial_{n|m}$ form a basis for M_m.)

Now let (g^{ij}) be the matrix inverse to (g_{ij}). Then the formula (2.1) for ∇ reads:

$$\Gamma_{ij}^k = \tfrac{1}{2} \sum_l g^{kl}(\partial_i g_{jl} + \partial_j g_{li} - \partial_l g_{ij}) \tag{2.2}$$

since the bracket terms $[\partial_i, \partial_j]$ vanish (exercise!).

2.2 Differential operators

The metric and Levi–Civita connection of a Riemannian manifold are precisely the ingredients one needs to generalise the familiar operators of vector calculus:

The *gradient of* $f \in C^\infty(M)$ is the vector field $\operatorname{grad} f$ such that, for $Y \in \Gamma(TM)$,

$$g(\operatorname{grad} f, Y) = Yf.$$

Similarly, the *divergence* of $X \in \Gamma(TM)$ is the function $\operatorname{div} f \in C^\infty(M)$ defined by:

$$(\operatorname{div} f)(m) = \operatorname{trace}(\xi \to \nabla_\xi X)$$

Finally, we put these together to introduce the hero of this volume: the *Laplacian* of $f \in C^\infty(M)$ is the function

$$\Delta f = \operatorname{div} \operatorname{grad} f.$$

In a chart (U, x), set $\mathbf{g} = \det(g_{ij})$. Then

$$\operatorname{grad} f = \sum_{i,j} g^{ij}(\partial_i f)\partial_j$$

and, for $X = \sum_i X_i \partial_i$,

$$\begin{aligned} \operatorname{div} X &= \sum_i (\partial_i X_i + \sum_j \Gamma_{ij}^i X_j) \\ &= \frac{1}{\sqrt{\mathbf{g}}} \sum_j \partial_j(\sqrt{\mathbf{g}} X_j). \end{aligned}$$

Here we have used $\sum_i \Gamma_{ij}^i = (\partial_j \sqrt{\mathbf{g}})/\sqrt{\mathbf{g}}$ which the Reader is invited to deduce from (2.2) together with the well-known formula for a matrix-valued function A:

$$\mathrm{d} \ln \det A = \operatorname{trace} A^{-1} \mathrm{d}A.$$