Differential Geometry *

Zhiyuan Bai

Compiled on November 29, 2022

This document serves as a set of revision materials for the Cambridge Mathematical Tripos Part III course *Differential Geometry* in Michaelmas 2022. However, despite its primary focus, readers should note that it is NOT a verbatim recall of the lectures, since the author might have made further amendments in the content. Therefore, there should always be provisions for errors and typos while this material is being used.

Contents

*Based on the lectures under the same name taught by Dr. J. Smith in Michaelmas 2022.

0 Introduction

Differential Geometry is the study of smooth manifolds, i.e. spaces that locally looks like $\mathbb{R}^n$ (in the smooth sense). We happen to have a good notion of smooth functions on these manifolds, so we can do calculus and be happy (or not).
There are two approaches to differential geometry: The first is that of embedded manifolds, which are spaces which sit nicely inside $\mathbb{R}^n$ for some n. They usually have a very rich extrinsic picture. Examples of this are manifolds defined by equations on the coordinates in $\mathbb{R}^n$.
Another way to study differnetial geometry would be to look at abstract manifolds, which are topological spaces equipped with reasonable local coordinates around each point, such that the coordinate transformations are smooth. This gives a way to talk about smooth manifolds in an intrinsic way.
It turns out that these two approaches are equivalent in the following sense: Every embedded manifold inherits local coordinates from the ambient space, and every abstract manifold can be embedded into $\mathbb{R}^N$ for some N (the "Whitney embedding theorem").

Once we know what a smooth manifold is, there are many basic constructions we can do, e.g. tangent spaces, smooth maps, vector fields and flows, submanifolds, etc.. It is also often useful to consider manifolds with some extra structure, like those with a compatible group structure (Lie groups). We'll briefly discuss Lie groups, and how its tangent space at the identity (its associated Lie algebra) lets us recover much information about them through the so-called exponential map and some other things. For example, the Lie group $\mathrm{GL}(n) \subset \mathbb{R}^{n^2}$ has Lie algebra $\mathbb{R}^{n^2}$ with the "Lie bracket" given by $[A, B] = AB - BA$ and exponential map $\exp A = I + A + A^2/2! + A^3/3! + \cdots$.
There is also the question about how one might differentiate a vector field. Given a surface $\Sigma \subset \mathbb{R}^3$ and a vector field V on Σ. What could the derivative of V possibly mean? One might try to differentiate in $\mathbb{R}^3$, but V is not necessarily defined outside the surface, so there is no way we can differentiate in a direction pointing "out" of the surface. When we restrict our attention to directions tangent to the surface, we find that the derivatives could still point "out" of the surface, which is kinda annoying since it doesn't seem to mean anything in the intrinsic picture. In a desperate grasp of faith, we decided to then project the result back on the surface, which does give something but it does rely on the embedding you chose.
In a way, these problems are eventually surround the question of "what kind of object are the derivatives". Even more so when we try to understand very simple questions like the interpretation of the derivative being zero.
We'll develop a language one can use to answer these questions. More precisely, we'll discuss the notions of differential forms, tensors, connections, parallel transport, curvature.
Most of these are motivated by embedded manifolds, but we can develop all these tools in a more general setting building from vector bundles and principal bundles. We'll soon meet the complex projective space $\mathbb{CP}^n$ which parameterises lines in $\mathbb{C}^{n+1}$ passing through the origin. Associated to each $p \in \mathbb{CP}^n$ is a one-dimensional (complex) vector space E_p, namely the line in $\mathbb{C}^{n+1}$ represented by p. This family of vector spaces form what's called a complex line bundle. This particular bundle, for various reasons, is called the tautological bundle $\mathcal{O}(-1)$.
In a more abstract setting, a vector field is a section of a vector bundle which, loosely speaking, is the choice of a family $s(p) \in E_p$ varying smoothly in p. What's good about doing this? For $z, w \in E_p$, we get to talk about a lot of things, like the distance between them and their relative phase. These encode geometrical structures to a vector bundle, which is very convenient for studying them.
For a physical example, suppose we have a quantum particle in a space-time X described by a wavefunction Ψ. Ψ is (approximately) a function on X, it is in fact the section of a line bundle on X. The physically meaningful quantities about this particle are $|\psi|^2$ and its relative phase compared to a different particle. The issue is, to write down the Schrödinger equation and solve for the motion of our particle, we must make sense of how Ψ can be differentiated. For this, we use the notion of connections. This connection turns out to be the electromagnetic potential, and the curvature is the field strength! This is the start of Gauge theory.

1 Manifolds and Smooth Maps

1.1 Manifolds

Definition 1.1. A topological n-manifold is a Hausdorff second-countable topological space X locally homeomorphic to $\mathbb{R}^n$, in the sense that every $p \in X$ there is an open neighbourhood $U \ni p$ ("coordinate patch") and an open set $V \in \mathbb{R}^n$ such that there is a homeomorphism $\phi : U \to V$ ("chart").

A collection of charts covering a topological manifold is called an atlas. If $x_1, \dots, x_n$ is the standard coordinates on $\mathbb{R}^n$, then the functions $x_1 \circ \phi, \dots, x_n \circ \phi$ are called the local coordinates on U (or at p), which we usually just denote as $x_1, \dots, x_n$. ϕ^{-1} is called a parameterisation.

Example 1.1. $\mathbb{R}^n$ is a topological manifold (duh).

Remark. One can show with non-trivial argument that for spaces locally homeomorphic to $\mathbb{R}^n$, being Hausdorff and second countable is the same as being metrisable and has countably many components.

Example 1.2. If X is a topological n-manifold, so is any open subset of X (with the subspace topology, of course).

If $\phi_\alpha : U_\alpha \to V_\alpha, \phi_\beta : U_\beta \to V_\beta$ are charts with $U_\alpha \cap U_\beta \neq \varnothing$, then $\phi_\beta \circ \phi_\alpha^{-1} : \phi_\alpha(U_\alpha \cap U_\beta) \to \phi_\beta(U_\alpha \cap U_\beta)$ is called the transition map from U_α to U_β. It describes the ϕ_β coordinates in terms of the ϕ_α coordinates.

Definition 1.2. Given an atlas $\mathcal{A} = \{\phi_\alpha : U_\alpha \to V_\alpha : \alpha \in A\}$ for a topological n-manifold X and a function $f : W \to \mathbb{R}^m$ on some open $W \subset X$, we say f is smooth with respect to $\mathcal{A}$ if for every α the map $f \circ \phi_\alpha^{-1} : \phi_\alpha(U_\alpha \cap W) \to \mathbb{R}^m$ is smooth.
The atlas $\mathcal{A}$ is smooth if every transition map is smooth (with the definition above, this is to say that every chart is smooth).

Lemma 1.1. *If $\mathcal{A}$ is a smooth atlas, then a function $f : W \to \mathbb{R}^m$ is smooth with respect to $\mathcal{A}$ iff for all $p \in W$, there is a chart $\phi_\alpha : U_\alpha \to V_\alpha$ such that $f \circ \phi_\alpha^{-1}$ is smooth.*

Proof. The "only if" direction is obvious. For the "if" direction, suppose f satisfies the hypothesis in the lemma and suppose $\phi_\beta : U_\beta \to V_\beta$ is a chart on X. For $p \in W \cap U_\beta$, there is a chart $\phi_\alpha : U_\alpha \to V_\alpha$ around p such that $f \circ \phi_\alpha^{-1}$ is smooth. Then $f \circ \phi_\beta^{-1} = f \circ \phi_\alpha^{-1} \circ \phi_\alpha \circ \phi_\beta^{-1}$ which is smooth. □

Corollary 1.2. *If $\mathcal{A}$ is smooth, then all local coordinates are smooth.*

Definition 1.3. Two atlases $\mathcal{A}, \mathcal{B}$ are smoothly equivalent if $\mathcal{A} \cup \mathcal{B}$ is smooth.

One can check that this is indeed an equivalence relation.

Lemma 1.3. *Two atlases are smoothly equivalent iff they have the same smooth functions, in the sense that for any open set $W \subset X$ and any function $f : W \to \mathbb{R}$, f is smooth with respect to $\mathcal{A}$ iff it is smooth with respect to $\mathcal{B}$.*

Proof. Example sheet. □

Definition 1.4. A smooth equivalence class of smooth atlases on a topological n-manifold X is called a smooth structure on X. A smooth n-manifold is a topological n-manifold equipped with a smooth structure.
Given a smooth n-manifold X and a function $f : W \to \mathbb{R}^m$ (with $W \subset X$ open), we say that f is smooth if it's smooth with respect to some, equivalently all, atlases in the smooth structure of X.

Example 1.3. 1. $\mathbb{R}^n$ is a smooth manifolds with the usual smooth structure given by the atlas $\{\mathrm{id}_{\mathbb{R}^n}\}$.
2. Open subsets of smooth manifolds inherit sturctures of smooth manifolds.
3. If X, Y are smooth n- and m-manifolds, then $X \times Y$ with the product topology is a smooth $n \times m$ manifolds by taking the products of charts.

Remark. 1. Being a topological n-manifold is a property, whereas being a smooth n-manifold is a choice, since one need to admit a smooth structure.
2. In dimension at most 3, every topological manifold has an essentially unique smooth structure. However, in dimensions at least 4, there exists topological manifolds with no smooth structure, e.g. the E_8 manifold, and those with many different smooth structures, e.g. exotic $\mathbb{R}^4$'s, exotic 7-sphere.

Definition 1.5. Given an n-manifold X, either topological or smooth, n is the dimension $\dim X$ of X.

We'll prove later that this is well-defined, i.e. X cannot be both an n- and an m-manifold if $n \neq m$.

Example 1.4. The n-sphere S^n is the smooth n-manifold defined by $S^n = \{y \in \mathbb{R}^{n+1} : |y|^2 = 1\}$ with the subspace topology and smooth structure given by stereographic projection, i.e. the charts on $U_\pm = S^n \backslash \{(0, \dots, 0, \pm 1)\}$ via $\phi_\pm : U_\pm \to \mathbb{R}^n, y \mapsto (1 \mp y_{n+1})^{-1}(y_1, \dots, y_n)$. The local coordinates $x^\pm$ associated to these charts are given by $x^\pm = (1 \mp y_{n+1})^{-1}(y_1, \dots, y_n)$ (lol). The height function $y_{n+1} : S^n \to \mathbb{R}$ is smooth since in local coordinates it's given by $y_{n+1} = \pm(|x^\pm|^2 - 1)/(|x^\pm|^2 + 1)$ on $U^\pm$.

From now on, when we say a manifold we automatically mean a smooth manifold.

1.2 Manifolds from Sets

Suppose we're given a set X and a collection of subsets $(U_\alpha)_\alpha \subset X$ that covers X. For each α, we are also given a bijection $\phi_\alpha : U_\alpha \to V_\alpha \subset \mathbb{R}^n$ for V_α open such that $\phi_\alpha(U_\alpha \cap U_\beta)$ is always open in $\mathbb{R}^n$ and each $\phi_\beta \circ \phi_\alpha^{-1} : \phi_\alpha(U_\alpha \cap U_\beta) \to \phi_\beta(U_\alpha \cap U_\beta)$.

Definition 1.6 (Non-standard)**.** The data as described above is called a smooth pseudo-atlas on X, and each ϕ_α is called a pseudo-chart. Two smooth pseudo-atlases are equivalent if their union is a smooth pseudo-atlas.

Say $W \subset X$ is open iff $\phi_\alpha(W \cap U_\alpha)$ is open for all α.

Lemma 1.4. *This defines a topology on X that makes it locally Euclidean, and the pseudo-atlas becomes a smooth atlas under this topology. Furthermore, the topology and smooth structure produced this way depends only on the equivalence class of the pseudo-atlas.*

Proof. Check. □

Example 1.5. The real projective space $\mathbb{RP}^n$ is constructed as follows: The underlying set is the set of lines through the origin in $\mathbb{R}^{n+1}$. Any $x \in \mathbb{R}^{n+1} \setminus \{0\}$ gives rise to a point in this set, namely the line $\mathbb{R}x$ through x. Moreover, every line through origin arises in this way. Two points x, y defines the same line iff they differ by a rescaling. So we can describe points in $\mathbb{RP}^n$ by the ratios $[x_0 : \cdots : x_n]$ ("homogeneous coordinates"). Explicitly, $[x_0 : \cdots : x_n] = [y_0 : \cdots : y_n]$ iff there is some $\lambda \in \mathbb{R}^\times$ such that $(x_0, \ldots, x_n) = (\lambda y_0, \ldots, \lambda y_n)$.
To define the pseudo-charts, we consider $\phi_i : U_i \to \mathbb{R}^n$ where $U_i = \{[x_0 : \cdots : x_n] : x_i \neq 0\}$ and $\phi_i([x_0 : \cdots : x_n]) = x_i^{-1}(x_0, \ldots, \hat{x}_i, \ldots, x_n)$. One easily checks that this gives a pseudo-atlas, and is Hausdorff and second-countable, hence defines a smooth (n-)manifold.
One can similarly define $\mathbb{CP}^n$ and show that it's a $2n$-manifold.

1.3 Smooth Maps

Fix n-manifold X and m-manifold Y with atlases $\{\phi_\alpha : U_\alpha \to V_\alpha\}$ and $\{\psi_\beta : S_\beta \to T_\beta\}$ respectively.

Definition 1.7. A map $F : X \to Y$ is smooth if it is continuous and for any α, β, the map $\psi_\beta \circ F \circ \phi_\alpha^{-1} : \phi_\alpha(U_\alpha \cap F^{-1}(S_\beta)) \to T_\beta$ is smooth.

The standard argument shows that this depends only on the smooth structure on the manifolds.

Remark. We do need to require a priori that F is continuous to ensure the openness of $F^{-1}(S_\beta)$ and hence the openness of $\phi_\alpha(U_\alpha \cap F^{-1}(S_\beta))$.

Example 1.6. 1. The identity map on any manifold, constant maps, projections from a product of manifolds are all smooth.
2. The inclusion $S^n \to \mathbb{R}^{n+1}$ is smooth.
3. Equip $\mathbb{R}^m$ with the usual manifold structure. Then $f : X \to \mathbb{R}^m$ is a smooth map iff it is a smooth function as defined earlier.
4. A map $X \to Y$ with $X \subset \mathbb{R}^n, Y \subset \mathbb{R}^m$ open is a smooth map iff it is smooth.

Lemma 1.5. *(i) Smoothness is local at the source. More precisely, $f : X \to Y$ is smooth iff for every $p \in X$, there is an open neighbourhood U of p such that $f|_U$ is smooth.*
(ii) A composition of smooth maps is smooth.

Proof. Trivial. □

Example 1.7. Consider $S^{2n+1} \subset \mathbb{C}^{2n+2}$. The map $h : S^{2n+1} \to \mathbb{CP}^n, x \mapsto [x_0 : \cdots : x_n]$ is called the Hopf map. This is smooth (example sheet).

Definition 1.8. A diffeomorphism between manifolds X, Y is a smooth map $X \to Y$ with a smooth two-sided inverse. If such a diffeomorphism exists, we say X, Y are diffeomorphic, written $X \cong Y$.

Remark. Smooth bijections need not be diffeomorphisms.

Example 1.8. $\mathbb{CP}^1$ is diffeomorphic to S^2 (example sheet), so $\mathbb{CP}^1$ is often referred to as the Riemann sphere. We have the Hopf map $S^3 \to S^2$ which has some funny properties.

Lemma 1.6. *If $X \cong Y$ are diffeomorphic (nonempty) manifolds, then* $\dim X = \dim Y$.

Proof. Let $F : X \to Y$ be a diffeomorphism and $G : Y \to X$ its inverse. Fix $p \in X$ and pick charts $\phi : U \to V$ about p and $\psi : S \to T$ about $q = F(p)$. Translating the charts if necessary, let's assume $\phi(p) = 0, \psi(q) = 0$. Now $\psi \circ F \circ \phi^{-1}, \phi \circ G \circ \psi^{-1}$ are inverse smooth maps between V and T, sending 0 to 0, so their derivatives are inverse linear maps, hence square matrices. □

1.4 Tangent Spaces

Fix an n-manifold X and $p \in X$.

Definition 1.9. A curve based at p is a smooth map $\gamma : I \to X$ where I is an open interval containing 0 and $\gamma(0) = p$. We say two curves γ_1, γ_2 based at p agree to first order if there exists a chart ϕ about p such that $(\phi \circ \gamma_1)'(0) = (\gamma \circ \gamma_2)'(0)$.

Write π_p^ϕ to denote the map $\gamma \mapsto (\phi \circ \gamma)'(0)$.

Lemma 1.7. *If γ_1, γ_2 have $\pi_p^\phi(\gamma_1) = \pi_p^\phi(\gamma_2)$ for some chart ϕ around p, then $\pi_p^\psi(\gamma_1) = \pi_p^\psi(\gamma_2)$ for any chart ψ around p.*

Proof. $\pi_p^\psi(\gamma) = (\psi \circ \gamma)'(0) = (\psi \circ \phi^{-1} \circ \phi \circ \gamma)'(0) = \mathrm{D}_{\phi(p)}(\psi \circ \phi^{-1})(\phi \circ \gamma)'(0) = A\pi_p^\phi(\gamma)$ where $A = \mathrm{D}_{\phi(p)}(\psi \circ \phi^{-1})$. □

Corollary 1.8. *Agreement to first order is an equivalence relation.*

Definition 1.10. The tangent space T_pX of X at p is the set of curves based at p modulo agreement to first order.

Elements of T_pX are called tangent vectors at p. We write $[\gamma]$ to denote the tangent vector represented by γ.

Proposition 1.9. *T_pX is naturally an n-dimensional real vector space.*

Proof. Given a chart ϕ about p, π_p^ϕ factors exactly through the relation of agreement to first order. So it induces a natural injection $T_pX \to \mathbb{R}^n$. This is also surjective. Indeed, given $v \in \mathbb{R}^n$, let γ be defined by $\gamma(t) = \phi^{-1}(\phi(p) + tv)$ which is a curve based at p and $\pi_p^\phi(\gamma) = v$.
Different charts are differed by an linear automorphism given by the derivative of the transition map, so we do indeed have a natural structure of an n-dimensional real vector space. □

Definition 1.11. Given a chart ϕ about p with local coordinates $x_1, \ldots, x_n$, define $\partial/\partial x_i \in T_pX$ to be $(\pi_p^\phi)^{-1}(e_i)$, where $e_i \in \mathbb{R}^n$ is the i^{th} standard basis vector. We'll often abbreviate $\partial_i = \partial_{x_i} = \partial/\partial x_i$.

Guess, what, these vectors form a basis for T_pX.

Example 1.9. In $\mathbb{R}^2$, ∂_x and ∂_y are what you think they are. If we instead use the polar coordinates (on $\mathbb{R}^2 \setminus \{0\}$), ∂_r are the unit vectors pointing away from the origin, and $\partial\theta$ are the vectors perpendicular to ∂_r to the anticlockwise direction, whose magnitude has unit angular length.

Remark. ∂_{x_i} depends very much on the choice of the whole system $x_1, \dots, x_n$, not just the choice of one coordinate x_i. If $y_1, \dots, y_n$ is another set of local coordinates with $x_1 = y_1$, we need not have $\partial_{y_1} = \partial_{x_1}$.

Lemma 1.10. *On overlapping charts with local coordinates $(x_j), (y_i)$ respectively, we have*

$$\frac{\partial}{\partial y_i} = \sum_{j=1}^{n} \frac{\partial x_j}{\partial y_i} \frac{\partial}{\partial x_j}$$

Proof. Let ϕ, ψ be the charts associated to $(x_j), (y_i)$ respectively. We want $(\pi_p^\psi)^{-1}(e_i)$ in terms of $(\pi_p^\phi)^{-1}(e_j)$. Let $A = \mathrm{D}_{\phi(p)}(\psi \circ \phi^{-1})$, then $(\psi_p^\psi)^{-1}(e_i) = (\pi_p^\phi)^{-1}(A^{-1}e_i)$. Now we have

$$A^{-1} = \mathrm{D}_{\psi(p)}(\phi \circ \psi^{-1}) = \left(\frac{\partial x_j}{\partial y_i}\right)\bigg|_p, A^{-1}e_i = \begin{pmatrix} \partial x_1/\partial y_i \\ \vdots \\ \partial x_n/\partial y_i \end{pmatrix} = \sum_j \frac{\partial x_j}{\partial y_i} e_j$$

Hence the result. □

Example 1.10. Suppose $[\gamma] = \sum_i a_i \partial_{x_i}$. Then $(\phi \circ \gamma)'(0) = \pi_p^\phi(\gamma) = \sum_i a_i e_i$, so a_i is the "speed of the curve to the direction e_i". For example, in $\mathbb{R}^2$, we can take $\gamma = (t^2 + 2t, -3t)$ based at 0. Then $[\gamma] = 2\partial_x - 3\partial_y$.

1.5 Derivatives

Fix manifolds X, Y and a smooth map $F : X \to Y$.

Definition 1.12. The derivative of F at $p \in X$ is the map $\mathrm{D}_p F : T_p X \to T_{F(p)} Y$ defined by $[\gamma] \mapsto [f \circ \gamma]$.

Sometimes we'll also write F^* for this ("the pushforward of F on tangent vectors").

Lemma 1.11. *The map $\mathrm{D}_p F$ is well-defined and linear.*

Proof. Fix charts ϕ, ψ around $p, q = F(p)$ respectively. We have $\pi_q^\psi(F \circ \gamma) = (\psi \circ F \circ \gamma)'(0) = (\psi \circ F \circ \phi^{-1} \circ \phi \circ \gamma)'(0) = A\pi_p^\phi(\gamma)$ where A is the derivative of $\psi \circ F \circ \phi^{-1}$ at $\phi(p)$. So $\pi_q^\psi(F \circ \gamma)$ depends on γ only through $\pi_p^\phi(\gamma)$ and does so linearly, hence we're done. □

Note that $\psi \circ F \circ \phi^{-1}$ is essentially F viewed under local coordinates on both sides. And on these local coordinates (say $(x_i), (y_j)$), we have

$$\begin{aligned} \mathrm{D}_p F(\partial_{x_i}) &= \mathrm{D}_p F((\pi_p^\phi)^{-1}(e_i)) = (\pi_{F(p)}^\psi)^{-1}(Ae_i) \\ &= \sum_j (\pi_{F(p)}^\psi)^{-1}\left(\frac{\partial y_j}{\partial x_i} e_j\right) = \sum_j \frac{\partial y_j}{\partial x_i} \partial y_j \end{aligned}$$

Remark. 1. For $X = \mathbb{R}^n, Y = \mathbb{R}^m$, this notion of derivative is the same as the one we used to have.
2. Given $f : X \to \mathbb{R}$, we have $\mathrm{D}_p f(\partial_{x_i}) = \partial f/\partial x_i|_p$.
3. We can write $[\gamma] = \mathrm{D}_0 \gamma(\partial_t)$ where t is the parameter of γ.

Lemma 1.12 (Chain Rule). *Suppose we have smooth maps $F : X \to Y$ and $G : Y \to Z$, then $\mathrm{D}_p(G \circ F) = \mathrm{D}_{F(p)}G \circ \mathrm{D}_pF$.*

Proof. For $[\gamma] \in T_pX$, we have $\mathrm{D}_p(G \circ F)([\gamma]) = [G \circ F \circ \gamma] = \mathrm{D}_{F(p)}G([F \circ \gamma]) = \mathrm{D}_{F(p)}G \circ \mathrm{D}_pF([\gamma])$. □

1.6 Immersions, Submersions and Local Diffeomorphisms

Definition 1.13. A smooth map $F : X \to Y$ is a(n) immersion (resp. submersion, local diffeomorphism) at $p \in X$ if D_pF is a(n) injection (resp. surjection, isomorphism). It is called a(n) immersion (resp. submersion, local diffeomorphism) if it is a(n) immersion (resp. submersion, local diffeomorphism) at every $p \in X$.

Definition 1.14. We say $p \in X$ is a regular point of F if F is a submersion at p. Otherwise, it is called a critical point of F.
We say $q \in Y$ is a regular value of F if every $p \in F^{-1}(\{q\})$ is a regular point. Otherwise, it is called a critical value of F.

Lemma 1.13. *F is a local diffeomorphism at p iff there are open neighbourhoods U of p and V of $F(p)$ such that F restricts to a diffeomorphism $U \to V$.*

Proof. The "if" part follows from the chain rule.
For the "only if" part, we make use of the inverse function theorem.
Pick charts $\phi : U \to B$ around p and $\psi : V \to D$ around $q = F(p)$. WLOG $F(U) \subset V$ by shrinking ϕ. Now consider $f = \psi \circ F \circ \phi^{-1} : B \to D$, which is a smooth map whose derivative is an isomorphism at $\phi(p)$ by the chain rule. By inverse function theorem, after some further shrinking, f has a smooth inverse $g : D \to B$. Then $G = \psi^{-1} \circ g \circ \phi$ is a smooth inverse to $F|_U : U \to V$. □

Remark. So $\psi \circ F$ is also a chart around p. If we actually take the charts $\psi \circ F, \psi$, then F becomes the identity in local coordinates. Similarly, $\phi \circ G$ is a chart around q under which the same happens.

Proposition 1.14. *Suppose F is an immersion at p and $x_1, \ldots, x_n$ are local coordinates on X around p, then there is a system of local coordinates $y_1, \ldots, y_m$ on Y about $F(p)$ with respect to which F has the form $(x_1, \ldots, x_n) \mapsto (x_1, \ldots, x_n, 0, \ldots, 0)$.*
Similarly, if F is a submersion at p, then we can choose such local coordinates under which F looks like $(x_1, \ldots, x_n) \mapsto (x_1, \ldots, x_m)$.

Proof. Example sheet. □

1.7 Submanifolds

Fix an n-manifold X.

Definition 1.15. A subset $Z \subset X$ is a submanifold of codimension k if for every $p \in Z$, there is a system of local coordinates $x_1, \ldots, x_n$ around p such that Z is locally given by the vanishing locus of $x_1, \ldots, x_k$.
Z is a properly embedded submanifold if the same thing holds but for all $p \in X$.

Remark. Technically the empty set is always a submanifold of codimension k for every k, I guess.

Example 1.11. $\{0\} \times \mathbb{R} \subset \mathbb{R}^2$ is properly embedded but $\{0\} \times \mathbb{R}^\times \subset \mathbb{R}^2$ is a submanifold that's not properly embedded.

Given a codimension k submanifold $Z \subset X$, we can equip it with the subspace topology, which automatically makes it a topological manifold. Indeed, it's immediately Hausdorff and second countable, and near any point $p \in Z$, if Z is locally the vanishing locus of $x_1, \ldots, x_k$, then $x_{k+1}, \ldots, x_n$ define a chart of Z around p. The transition functions between these charts are smooth, so we actually have the structure of a smooth $(n-k)$-manifold on Z (which, as one easily checks, depends only on the smooth structure on X).

Proposition 1.15. *A codimension k submanifold $Z \subset X$ has the natural structure of an $(n-k)$-manifold. The inclusion map $\iota : Z \to X$ is a smooth immersion and a homeomorphism onto its image, and composition with ι gives a bijection between smooth maps to Z and smooth maps to X with image contained in Z.*

Proof. Follows from our discussion above. □

Definition 1.16. A map $F : Y \to X$ is an embedding if it's a smooth immersion and a homeomorphism onto its image.

Lemma 1.16. *The image of an embedding $F : Y \to X$ is a submanifold Z of X. Moreover, F induces a diffeomorphism $Y \to Z$.*

So submanifolds are exactly the same as images of embeddings.

Example 1.12. The inclusion $S^n \to \mathbb{R}^{n+1}$ is an embedding. So S^n is a submanifold of $\mathbb{R}^{n+1}$ and the induced manifold structure coincides with the standard one.

Proposition 1.17. *If $F : X \to Y$ is a smooth map and $q \in Y$ is a regular value of F, then $F^{-1}(q)$ is a submanifold of X of codimension* $\dim Y$.

Proof. Write $n = \dim X, m = \dim Y$.
Take a point $p \in F^{-1}(q)$. Since q is a regular value, F is a submersion at p, so there are local coordinates $(x_1, \ldots, x_n)$ about p and $(y_1, \ldots, y_m)$ about q (WLOG $y(q) = 0$) under which $y \circ F = (x_1, \ldots, x_m)$.
Let U be a small open neighbourhood of p on which the x-coordinates and $y \circ F$ are defined. Then $U \cap F^{-1}(q) = (y \circ F)^{-1}(0) = \{x_1 = \cdots = x_m = 0\}$, which is what we wanted. □

Example 1.13. Consider $F : \mathbb{R}^{n+1} \to \mathbb{R}, y \mapsto \|y\|^2 = \sum_j y_j^2$. Then $\mathrm{D}F = (2y_1, \ldots, 2y_{n+1})$. So $F^{-1}(\lambda)$ is a codimension 1 submanifold of $\mathbb{R}^{n+1}$ for all $\lambda \neq 0$ (for $\lambda < 0$ it'll be empty, fun). In particular, $S^n = F^{-1}(\{1\})$ is a submanifold.

Theorem 1.18 (Sard's Theorem)**.** *The set of critial values of a smooth map $F : X \to Y$ has measure 0 in Y. More precisely, for any chart $\psi : S \to T$ on Y, $\psi(\{$critical values of $F\} \cap S) \subset T \subset \mathbb{R}^{\dim Y}$ has measure 0 with respect to the Lesbegue measure on $\mathbb{R}^{\dim Y}$.*

Proof. Omitted. □

Corollary 1.19. *The set of regular values for $F : X \to Y$ is dense in Y.*

In particular, if Y is nonempty then a regular value must exist.

Remark. This however does not mean that regular points must exist. Indeed, if $\dim X < \dim Y$, then there cannot be any regular points. This is not a contradiction because the image would have measure 0.

Definition 1.17. Submanifolds $Y, Z \subset X$ are transverse at $p \in Y \cap Z$ if $T_pY + T_pZ = T_pX$. We say Y, Z are transverse if they are transverse at any $p \in Y \cap Z$. Equivalently, $\operatorname{ann}(T_pY) \cap \operatorname{ann}(T_pZ) = 0$ in $(T_pX)^\vee$.

Remark. Here, we mean by T_pY the image of T_pY under the derivative of the inclusion $Y \to X$. Same for T_pZ.

Example 1.14. In $\mathbb{R}^2$, the axes are transverse but the x-axis is not transverse to the submanifold $\{y = x^2\}$.

Proposition 1.20. *Suppose Y, Z are submanifolds of X of codimensions k, l. If they intersect transversely, then $Y \cap Z$ is a submanifold of X with codimension $k + l$.*

Proof. Pick $p \in Y \cap Z$. Since Y, Z are submanifolds, there are coordinates $y_1, \dots, y_n$ and $z_1, \dots, z_n$ about p such that we locally have $Y = \{y_1 = \dots = y_k = 0\}, Z = \{z_1 = \dots = z_l = 0\}$. Let U be an open neighbourhood of p on which y, z are defined. Consider $F : U \to \mathbb{R}^{k+l}$ given by $(y_1, \dots, y_k, z_1, \dots, z_l)$. This is a submersion at p since Y, Z are transverse at p. So there exists local coordinates $(x_1, \dots, x_n)$ on U about p in which F is the projection to the first $k + l$ coordinates. In this local system of coordinates, $Y \cap Z = \{x_1 = \dots = x_{k+l} = 0\}$. □

2 Vector Bundles and Tensors

2.1 Vector Bundles

A vector bundle over a manifold X is, loosely speaking, a "family of vector spaces" parameterised by X.

Definition 2.1. A (real) vector bundle of rank k over a manifold B is the data of a manifold E, a smooth map $\pi : E \to B$, an open cover $\{U_\alpha\}_{\alpha \in A}$ of B, and a diffeomorphism $\Phi_\alpha : \pi^{-1}(U_\alpha) \to U_\alpha \times \mathbb{R}^k$ for each $\alpha \in A$ such that:
(i) $\mathrm{pr}_1 \circ \Phi_\alpha = \pi$ on $\pi^{-1}(U_\alpha)$.
(ii) For all α, β, $\Phi_\beta \circ \Phi_\alpha^{-1} : (U_\alpha \cap U_\beta) \times \mathbb{R}^k \to (U_\alpha \cap U_\beta) \times \mathbb{R}^k$ has the form $(b, v) \mapsto (b, g_{\beta\alpha}(b)v)$ for some (necessarily smooth) map $g_{\beta\alpha} : U_\alpha \cap U_\beta \to \mathrm{GL}(k, \mathbb{R})$.

So I can't be bothered to draw a picture, but you probably should.
E is called the total space of the vector bundle, π the projection, Φ_α the local trivialisations, and $g_{\beta\alpha}$ the transition funcitons. The fibres $\pi^{-1} = (b), b \in B$ are denoted E_b.
We sometimes just write $\pi : E \to B$ to denote the data of such a vector bundle.

Remark. 1. Each E_b has the natural structure of a vector space via $\Phi_\alpha : E_b \to \{b\} \times \mathbb{R}^k$.
2. If we just replace $\mathbb{R}$ by $\mathbb{C}$ everywhere, we get the definition of a complex vector bundle.

Example 2.1. 1. The trivial bundle $\underline{\mathbb{R}^k}$ of rank k over B is given by taking $E = B \times \mathbb{R}^k$, π the projection onto B, $\{U_\alpha\}_{\alpha \in A} = \{B\}$ and $\Phi_\alpha = \mathrm{id}_{B \times \mathbb{R}^k}$.
2. The tautological bundle $\mathcal{O}_{\mathbb{RP}^n}(-1)$ on $\mathbb{RP}^n$ is the line bundle (i.e. a vector bundle of rank 1) over $\mathbb{RP}^n$ given by

$$E = \{(p, v) \in \mathbb{RP}^n \times \mathbb{R}^{n+1} : v \text{ lies on the line described by } p\}$$

which is a submanifold of $\mathbb{RP}^n \times \mathbb{R}^{n+1}$. We take π to be the projection onto the first factor. The open cover we'll take is the standard cover $U_i = \{[x_0 : \cdots : x_n] : x_i \neq 0\}$ and $\Phi_i : \pi^{-1}(U_i) \to U_i \times \mathbb{R}$ is given by $([x_0 : \cdots : x_n], (y_0, \ldots, y_n)) \mapsto ([x_0 : \cdots : x_n], y_i)$. On $U_i \cap U_j$, we have

$$\Phi_j \circ \Phi_i^{-1}([x_0 : \cdots : x_n], t) = ([[x_0 : \cdots : x_n]], (x_j/x_i)t)$$

So this is indeed a line bundle, with $g_{ji}([x_0 : \cdots : x_n])(t) = (x_j/x_i)t$.
3. Replace every $\mathbb{R}$ in the above example by $\mathbb{C}$ gives the tautological complex line bundle on $\mathbb{CP}^n$.
4. Let's construct the tangent bundle of an n-manifold X, which is supposed to be a rank n vector bundle.
The total space is given by $E = TX = \coprod_{p \in X} T_pX = \{(p, v) : p \in X, v \in T_pX\}$. Given a coordinate patch U on X with coordinates x_i, we get a pseudochart $\coprod_{p \in U} T_pX \to U \times \mathbb{R}^n$ via $(p, \sum_i a_i \partial_{x_i}) \mapsto (p, (a_1, \ldots, a_n))$. This makes TX a manifold, and we have a smooth map $\pi : TX \to X$ by $(p, v) \mapsto p$. The pseudocharts we took already give a satisfactory local trivialisation.

Definition 2.2. A (global) section of a vector bundle (or indeed any smooth map) $\pi : E \to B$ is a smooth map $s : B \to E$ such that $\pi \circ s = \mathrm{id}_B$. A local section on an open set $s : U \subset B$ is a smooth map $s : U \to E$ such that $\pi \circ s = \mathrm{id}_U$.

Example 2.2. 1. Every vector bundle $\pi : E \to B$ has a zero section given by $s(b) = 0 \in E_b$ for all b.
2. A vector field on X is the same as a section of TX.

Definition 2.3. Given vector bundles $\pi_1 : E_1 \to B_1, \pi_2 : E_2 \to B_2$ and a smooth map $F : B_1 \to B_2$, a morphism of vector bundles $E_1 \to E_2$ covering F is a smooth map $G : E_1 \to E_2$ such that $\pi_2 \circ G = F \circ \pi_1$ and for any $p \in B$, the induced map $G_p : (E_1)_p \to (E_2)_{F(p)}$ is linear.

Remark. We usually only care about vector bundles over a fixed base space B, in which case we'll always take $B = B_1 = B_2$ and $F = \mathrm{id}_B$.

Definition 2.4. An isomorphism of vector bundles over B is a morphism of vector bundles covering id_B with a two-sided inverse. Equivalently, it's a diffeomorphism on the total spaces that's linear on the fibres.

Example 2.3. Consider $B = S^1 = \{e^{i\theta} : \theta \in \mathbb{R}\} \subset \mathbb{C}$. The vector field $\partial\theta$ on S^1 is well-defined and nonzero in each fibre. So we get an isomorphism of vector bundles $S \times \mathbb{R} \to TS^1, (e^{i\theta}, a) \mapsto (e^{i\theta}, a\partial_\theta)$ over S^1. Hence we say S^1 has trivial tangent bundle. On the other hand, TS^n is not trivial for $n > 1$ (example sheet).

Remark. A morphism $G : \underline{\mathbb{R}} \to E$ is the same as a section s of E. Indeed, any such morphism G gives a section $s(b) = G(b, 1)$, and any section s gives $G(b, t) = ts(b)$. More generally, a morphsim $\underline{\mathbb{R}^k} \to E$ is the same as a k-tuple of sections. This is an isomorphism iff the sections form a basis in each fibre.

Definition 2.5. Given a vector bundle $\pi : E \to B$ of rank k, a subbundle of rank $l \leq k$ is a subset F of E such that B can be covered by local trivialisations $\Phi_\alpha : \pi^{-1}(U_\alpha) \to U_\alpha \times \mathbb{R}^k$ under which $F = \Phi_\alpha^{-1}(U_\alpha \times (\mathbb{R}^l \times \{(0, \dots, 0)\}))$. This is naturally a vector bundle of rank l, and we certainly have quotient bundle E/F of E by F, which is a vector bundle over B of rank $k - l$, and we have morphisms $F \to E \to E/F$.

Example 2.4. $\mathcal{O}_{\mathbb{RP}^n}(-1)$ is a subbundle of $\underline{\mathbb{R}^{n+1}}$ over $\mathbb{RP}^n$. We have the Euler sequence $\mathcal{O}_{\mathbb{RP}^n}(-1) \to \underline{\mathbb{R}^{n+1}} \to \underline{\mathbb{R}^{n+1}}/\mathcal{O}_{\mathbb{RP}^n}(-1) \cong T\mathbb{RP}^n(-1)$, except we haven't defined the last thing yet.

2.2 Gluing

To define a vector bundle over a base B of rank k, it suffices to give an open cover $\{U_\alpha\}_{\alpha \in A}$ of B and a smooth map $g_{\beta\alpha} : U_\alpha \cap U_\beta \to \mathrm{GL}(k, \mathbb{R})$ such that:
(i) (the identity condition) $g_{\alpha\alpha} = \mathrm{id}_{\mathbb{R}^k}$.
(ii) (the cocycle condition) $g_{\gamma\alpha} = g_{\gamma\beta} \circ g_{\beta\alpha}$. Note that we thus have $g_{\alpha\beta} = g_{\beta\alpha}^{-1}$.
Given this collection of data, we define

$$E = \left(\coprod_{\alpha \in A} U_\alpha \times \mathbb{R}^k \right) / (\forall b \in U_\alpha \cap U_\beta, (b, v) \sim (b, g_{\beta\alpha} v))$$

and let π be the obvious map to B.
This comes with identifications $\pi^{-1}(U_\alpha) \cong U_\alpha \times \mathbb{R}^k$, which defines a pseudochart on E that also gives local trivialisations. It's easy but slightly tedious to check that this does define a vector bundle.

Proposition 2.1. *If $\pi : E \to B$ is a vector bundle of rank k trivialised over $\{U_\alpha\}_{\alpha \in A}$ with transition functions $g_{\beta\alpha}$. Then these transition functions satisfy the identity condition and the cocycle condition, and the vector bundle they glue to is isomorphic to $\pi : E \to B$.*

Proof. The conditions follow from the corresponding conditions for Φ_α's. The trivialisations of E and their inverses define diffeomorphisms between E and the gluing construction which are compatible with projections and linear on fibres, hence give an isomorphism. □

Corollary 2.2. *Two vector bundles are isomorphic iff they can be trivialised over a common cover with the same transition functions.*

Example 2.5. 1. For $r \in \mathbb{Z}$, we can define a line bundle on $\mathbb{RP}^n$ to be the trivial bundle over the standard open sets $U_i = \{[x_0 : \cdots : x_n] : x_i \neq 0\}$ and $g_{ji} = (x_j/x_i)^{-r}$. We call this line bundle $\mathcal{O}_{\mathbb{RP}^n}(r)$. When $r = -1$, this is just the tautological bundle.
2. The Möbius bundle $M \to \mathbb{RP}^1$ is a line bundle which is defined for it to be trivialised over U_0, U_1 with $g_{10} = \mathrm{sgn}(x_1/x_0)$. Of course M is basically the (open) Möbius band.
We claim that $M \cong \mathcal{O}_{\mathbb{RP}^1}(-1)$. It suffices to show that we can modify trivialisation of M to make the transition functions become $g_{10} = x_1/x_0$.
Let's rescale the local trivialisation Φ_i of M by a smooth map $\psi_i : U_i \to \mathbb{R}^\times = \mathrm{GL}(1, \mathbb{R})$. Explicitly, consider the trivialisation $\Phi^{-1}(U_i) \ni (p, v) \mapsto (p, \psi_i(p) \, \mathrm{pr}_2 \circ \Phi_i(p, v))$. This modifies g_{10} to $\psi_1/\psi_0 g_{10}$. We're left to choose

ψ_0, ψ_1 such that $\psi_1/\psi_0 = |x_1|/|x_0|$. One way to do this would be to set $\psi_1 = \sqrt{x_1^2/(x_0^2+x_1^2)}, \psi_0 = \sqrt{x_0^2/(x_0^2+x_1^2)}$.

Definition 2.6. Given a vector bundle $\pi : E \to B$ and a smooth map $F : B' \to B$, the pullback bundle F^*E is a vector bundle $\pi : F^*E \to B'$ defined as follows: The total space of F^*E is $\coprod_{p \in B'} E_{F(p)}$. Suppose E is trivialised over $\{U_\alpha\}_{\alpha \in A}$ with transition functions $g_{\beta\alpha}$, then F^*E is trivialised over $\{f^{-1}U_\alpha\}_{\alpha \in A}$ with transition functions $F^*g_{\beta\alpha} = g_{\beta\alpha} \circ F$.

The fibre $(F^*E)_p$ is then naturally identified with $E_{F(p)}$.

Example 2.6. Consider the Hopf map $H : S^{2n+1} \to \mathbb{CP}^n$. $H^*\mathcal{O}_{\mathbb{CP}^1}(-1)$ is trivial. Indeed, it is trivialised by the section $S^{2n+1} \ni p \mapsto (p, \text{line through } p)$.

Definition 2.7. Given a vector bundle $E \to B$, the dual bundle $E^\vee \to B$ has total space $\coprod_{p \in B}(E_p)^\vee$. Suppose E is trivialised over $\{U_\alpha\}_{\alpha \in A}$ with transition functions $g_{\beta\alpha}$, then $E^\vee$ is trivialised over $\{U_\alpha\}_\alpha$ with transition functions $(g_{\beta\alpha}^\vee)^{-1}$.
If E is trivialised over $U \subset B$ by a fibrewise basis of sections $\sigma_1, \dots, \sigma_k$, then the fibrewise dual basis $\sigma_1^\vee, \dots, \sigma_k^\vee$ are smooth sections of $E^\vee$ which trivialises it over U.

2.3 Cotangent Bundle

Fix an n-manifold X.

Definition 2.8. The cotangent bundle T^*X of X is the dual of the tangent bundle TX of X. Its fibre over $p \in X$ is denoted T_p^*X (the "cotangent space at p"). We then have $T_p^*X = (T_pX)^\vee$.

We might describe T_p^*X alternatively as certain classes of functions $X \to \mathbb{R}$. Consider the set G of germs of smooth functions around p. We say $f, g \in G$ agree to first order if $\mathrm{D}_p f = \mathrm{D}_p g$.

Proposition 2.3. *The vector space of equivalence classes in G given by first order agreement is naturally isomorphic to T_p^*X.*

Proof. Consider the map $e : G \to T_p^*X$ via $f \mapsto ([\gamma] \mapsto (f \circ \gamma)'(0))$. In local coordinates, this reads

$$f \mapsto \left(\sum_i a_i \partial_{x_i} \mapsto \sum_i a_i \frac{\partial f}{\partial x_i} \right)$$

This is certainly surjective since $\{e(x_j)\}_j$ is the dual basis to ∂x_i. The kernel, on the other hand, is the set of germs $f \in G$ with $\mathrm{D}_p f = 0$. □

So for any smooth functions $f : U \to \mathbb{R}$, we get an element of T_p^*X at each $p \in U$ via e, which is often denoted $\mathrm{d}_p f$.

Lemma 2.4. *These $\mathrm{d}_p f$ define a smooth section of T^*X on U, denoted $\mathrm{d}f$, which is called the differential of f.*

Proof. In local coordinates, we simply have $\mathrm{d}f = \sum_i (\partial f/\partial x_i)\mathrm{d}x_i$ (more precisely, $e(f) = \sum_i (\partial f/\partial x_i) e(x_i)$). But $\mathrm{d}x_i$ are fibrewise duals to ∂_{x_i}, so they are smooth. Consequently, as a smooth linear combination of them, $\mathrm{d}f$ must too be smooth. □

Note, by construction, that $\mathrm{d}f(v)$ is exactly the directional derivative of f in the direction of v.

Definition 2.9. The sections of T^*X are called 1-forms.

Remark. Unlike ∂_{x_i}, each $\mathrm{d}x_i$ depends only on x_i but not on the whole system.

Definition 2.10. Given a smooth map $F : X \to Y$, we have a pullback map $(\mathrm{D}_p F)^\vee : T^*_{F(p)}Y \to T^*_p X$, often denoted F^* when p is understood.

Lemma 2.5. *Given $F : X \to Y$ and a smooth function g on Y, then $F^*\mathrm{d}g = \mathrm{d}(F^*g)$ (recall $F^*g = g \circ F$).*

Proof. Tautologies. Given $[\gamma] \in T_pX$, we have $(F^*\mathrm{d}g)([\gamma]) = (\mathrm{d}g)(\mathrm{D}_pF([\gamma])) = \mathrm{d}g([F \circ \gamma]) = (g \circ F \circ \gamma)'(0) = (\mathrm{d}(F^*g))([\gamma])$. □

2.4 Multilinear Algebra

So, go read some multilinear algebra.

2.5 Tensors and Forms

We can apply any functorial operations to the transition functions on any existing bundles, and by functoriality they guarentee to give another bundle.

Definition 2.11. Given vector bundles $E, F \to B$ trivialised over a common cover $\{U_\alpha\}_\alpha$ of B with transition functions $g_{\beta\alpha}, h_{\beta\alpha}$. The direct sum of E, F is the bundle $E \oplus F \to B$ with fibres $E_p \oplus F_p$, trivialised over $\{U_\alpha\}_\alpha$ with transition functions $g_{\beta\alpha} \oplus h_{\beta\alpha}$.
Doing exactly the same thing but with the tensor product $\otimes$ allows the definition of $E \otimes F$, the tensor product of E and F. Similar for exterior powers and symmetric powers of a vector bundle.

Example 2.7. Given a smooth map $F : X \to Y$, $\mathrm{D}F$ is a section of $(T^*X) \otimes (F^*TY)$. Indeed, for each $p \in X$, $\mathrm{D}_pF \in \mathrm{Hom}_{\mathbb{R}}(T_pX, T_{F(p)}(Y)) \cong (T^*_pX) \otimes (T_{F(p)}(Y)) = ((T^*X) \otimes (F^*TY))_p$.

Definition 2.12. A tensor (or tensor field) of type (p, q) is a section of $(TX)^{\otimes p} \otimes (T^*X)^{\otimes q}$.
An r-form is a section of $\bigwedge^r T^*X$. The space of r-forms on an open $U \subset X$ is denoted $\Omega^r(U)$.

Remark. Note that $\Lambda^1 T^*X = T^*X$, so the $r = 1$ case agrees with our eariler definition of a 1-form.

Example 2.8. 1. A tensor of type $(0, 0)$ is a section of $\underline{\mathbb{R}}$, i.e. a smooth function (or, as some call it, a scalar field).
2. A tensor of type $(1, 0)$ is a section of TX, i.e. a vector field.
3. A tensor of type $(0, 1)$ is a section of T^*X, i.e. a 1-form.
4. A tensor of type $(0, q)$ is something which "eats q vectors multilinearly and spits out a number (at each point)". By contrast, an r-form "eats q vectors multilinearly and antisymmetrically and spits out a number" (more details later).

2.6 Index Notation

For the rest of the course, we're gonna put indices on local coordinates as superscripts, i.e. the coordinates should be $x^1, \ldots, x^n$. A section T of $TX \otimes T^*X \otimes TX$ (a tensor of type $(2,1)$) can be written in local coordinates (x^i) uniquely as

$$T = \sum_{i,j,k} T^i{}_j{}^k \, \partial_i \otimes \mathrm{d}x^j \otimes \partial_k$$

where the horizontal positions of indices refer to ordering of tensor factors, and vertical position refer to whether it's from the tangent (superscript) or cotangent (lowerscript) bundle.

We'll often use summation convention, where repeated indices once up and once down are summed over, e.g. $T = T^i{}_j{}^k \, \partial_i \otimes \mathrm{d}x^j \otimes \partial_k$. Even more lazily, we might just write $(T^i{}_j{}^k)$ or just $T^i{}_j{}^k$ for T.

Tensor products corresponds to juxtaposition, e.g.

$$(T^i{}_j{}^k \, \partial_i \otimes \mathrm{d}x^j \otimes \partial_k) \times (S_{lm} \, \mathrm{d}x^l \otimes \mathrm{d}x^m) = T^i{}_j{}^k S_{lm} \, \partial_i \otimes \mathrm{d}x^j \otimes \partial_k \otimes \mathrm{d}x^l \otimes \mathrm{d}x^m$$

On the other hand, contraction corresponds to summation, e.g. contraction of the third facor of $T = (T^i{}_j{}^k)$ with the second factor of $S = (S_{lm})$ is the tensor

$$\begin{aligned} T^i{}_j{}^k S_{lm}(\mathrm{d}x^m(\partial_k)) \, \partial_i \otimes \mathrm{d}x^j \otimes \mathrm{d}x^l &= T^i{}_j{}^k S_{lm} \delta^m{}_k \, \partial_i \otimes \mathrm{d}x^j \otimes \mathrm{d}x^l \\ &= T^i{}_j{}^k S_{lk} \, \partial_i \otimes \mathrm{d}x^j \otimes \mathrm{d}x^l \end{aligned}$$

Similarly, in a system of local coordinates (x^i), an r-form α can be writen uniquely as

$$\sum_{I \text{ multi-index}} \alpha_I \, \mathrm{d}x^I = \sum_{I \text{ multi-index}} \alpha_I \, \mathrm{d}x^{i_1} \wedge \cdots \wedge \mathrm{d}x^{i_r}$$

summing over multi-indices $I = (i_1 < \cdots < i_r)$.

Given r vectors $v_{(1)}, \ldots, v_{(r)}$, we can feed them to our beloved r-form α to get

$$\sum_{I \text{ multi-index}, \sigma \in S_r} \epsilon(\sigma) \alpha_I v_{(1)}^{i_{\sigma(1)}} \cdots v_{(r)}^{i_{\sigma(r)}}$$

where $\epsilon : S_r \to \{\pm 1\}$ is the sign function for permutations. This is equivalent to viewing α as the tensor

$$\sum_{I \text{ multi-index}, \sigma \in S_r} \epsilon(\sigma) \alpha_I \, \mathrm{d}x^{i_{\sigma(1)}} \otimes \cdots \otimes \mathrm{d}x^{i_{\sigma(r)}}$$

of type $(0, r)$, and evaluation at $v_{(1)}, \ldots, v_{(r)}$ is the same as contracting with them as tensors.

Remark. Some people include $1/r!$ in the formula. We don't.

We refer to the components of this tensor as $\alpha_{i_1 \cdots i_r}$, the coefficient of $\mathrm{d}x^{i_1} \otimes \cdots \otimes \mathrm{d}x^{i_r}$. When the i_j's form a multi-index I, this agrees with α_I.

Example 2.9. On $\mathbb{R}^2$, we view $\mathrm{d}x^1 \wedge \mathrm{d}x^2$ as $\mathrm{d}x^1 \otimes \mathrm{d}x^2 - \mathrm{d}x^2 \otimes \mathrm{d}x^1$. A general 2-form then looks like $\alpha_{12} \, \mathrm{d}x^1 \wedge \mathrm{d}x^2 = \alpha_{ij} \, \mathrm{d}x^i \otimes \mathrm{d}x^j$ where $\alpha_{21} = -\alpha_{12}, \alpha_{11} = \alpha_{22} = 0$.

In general, $\alpha = \alpha_{i_1 \cdots i_r} \, \mathrm{d}x^{i_1} \otimes \cdots \otimes \mathrm{d}x^{i_r}$ for a similar definition of $\alpha_{i_1 \cdots i_r}$ as in the example above. However, one should note that $\alpha_{i_1 \cdots i_r} \, \mathrm{d}x^{i_1} \wedge \cdots \wedge \mathrm{d}x^{i_r} = (r!)\alpha$ instead of α.

2.7 Pushforward and Pullback

Fix manifolds X, Y and a smooth map $f : X \to Y$. Given $p \in X$ and a tensor T of type $(r, 0)$ at p (i.e. a section of $(T_pX)^{\otimes r}$). We can push this forward to F_*T which is a section of $(T_{F(p)}Y^{\otimes r}$ by applying D_pF on each tensor factor. Similarly, given $p \in X$ and a tensor T of type $(0, r)$ at $F(p)$, we can pull this back to a tensor F^*T of type $(0, r)$ at p by using $(\mathrm{D}_pF)^\vee$ to every factor. We can do the same for r-forms at $F(p)$ using $\bigwedge^r(\mathrm{D}_pF)^\vee$.
Given a tensor of T type $(0, r)$ on Y, we can pull it back to a tensor F^*T on X by $(F^*T)_p = F^*(T_{F(p)})$. Similar for r-forms.
To summarise, we can push "up" tensors forward at a point and pull "down" tensors or forms at a point or over an open set.
If F is a diffeomorphism, then one can pushforward or pullback any tensor at a point or over an open set, and the pushforward is the same as the pullback along its inverse.

Example 2.10. Suppose we have a tensor of type $(1, 1)$ on X, then $(F_*T)_q = F_*(T_{F^{-1}(q)})$, where the second F_* simply means applying $\mathrm{D}_{F^{-1}(q)}F$ on the TX factor, and $(\mathrm{D}_{F^{-1}(q)}F)^\vee$ on the T^*X factor.

3 Differential Forms

3.1 Exterior Derivative

Given a 1-form $\alpha = \alpha_i\,\mathrm{d}x^i$ on X, we might naïvely want its derivative to be $\sum_{i,j}(\partial\alpha_i/\partial x_j)\,\mathrm{d}x^j \otimes \mathrm{d}x^i$. But this is bad, since a different local coordinates (y^i) (say with $\alpha = \alpha'_i\,\mathrm{d}y^i = \sum_{i,j}\alpha'_i(\partial y^i/\partial x^j)\,\mathrm{d}x^j)$ would give the "derivative"

$$\frac{\partial\alpha_i}{\partial x^j}\,\mathrm{d}x^j \otimes \mathrm{d}x^i = \frac{\partial}{\partial x^j}\left(\alpha'_k\frac{\partial y^k}{\partial x^i}\right)\mathrm{d}x^j \otimes \mathrm{d}x^i = \frac{\partial\alpha'_k}{\partial y^j}\,\mathrm{d}y^j \otimes \mathrm{d}y^k + \alpha'_k\frac{\partial^2 y^k}{\partial x^j\partial x^i}$$

where the summation convention applies and the indices appearing in the denominators (e.g. ∂x^j's) are treated as down-indices. This is very clearly coordinate independent, so how do we fix this?

Definition 3.1. The exterior derivative $\mathrm{d}\alpha$ of a 1-form α is

$$\mathrm{d}\alpha = \partial\frac{\partial\alpha_i}{\partial x^j}\mathrm{d}x^j \wedge \mathrm{d}x^i = \mathrm{d}\alpha_i \wedge \mathrm{d}x^i$$

By the calculation above, this is independent of the choice of local coordinates.

Definition 3.2. In general, given a p-form $\alpha = \alpha_I\,\mathrm{d}x^I$ (summing over multi-indices I), its exterior derivative is

$$\mathrm{d}\alpha = \mathrm{d}\alpha_I \wedge \mathrm{d}x^I = \frac{\partial\alpha_I}{\partial x^j}\,\mathrm{d}x^j \wedge \mathrm{d}x^I$$

Proposition 3.1. *(i) The exterior derivative is* $\mathbb{R}$*-linear*
(ii) It agrees with the differential on 0*-forms.*
(iii) $\mathrm{d}^2 = 0$*, i.e.* $\mathrm{d}(\mathrm{d}\alpha) = 0$ *for all differential forms* α*.*
(iv) Given p*-form* α *and* q*-form* β*, we have* $\mathrm{d}(\alpha\wedge\beta) = (\mathrm{d}\alpha)\wedge\beta + (-1)^p\alpha\wedge(\mathrm{d}\beta)$*.*

Proof. (i), (ii) Immediate from definition.
(iii) Consider $\alpha = \alpha_I \, \mathrm{d}x^I$, then we have

$$\mathrm{d}^2\alpha = \mathrm{d}\left(\frac{\partial \alpha_I}{\partial x^j}\, \mathrm{d}x^j \wedge \mathrm{d}x^I\right) = \frac{\partial^2 \alpha_I}{\partial x^k \partial x^j}\, \mathrm{d}x^k \wedge \mathrm{d}x^j \wedge \mathrm{d}x^I = 0$$

since the summand is anti-symmetric in j, k.
(iv) Write $\alpha = \alpha_I \, \mathrm{d}\alpha^I, \beta = \beta_J \, \mathrm{d}x^J$, then

$$\begin{aligned}
\mathrm{d}(\alpha \wedge \beta) &= \mathrm{d}(\alpha_I \beta_J \, \mathrm{d}x^I \wedge \mathrm{d}x^J) = \mathrm{d}(\alpha_I \beta_J) \wedge \mathrm{d}x^I \wedge \mathrm{d}x^J \\
&= (\mathrm{d}\alpha_I)\beta_J \wedge \mathrm{d}x^I \wedge \mathrm{d}x^J + \alpha_I \, \mathrm{d}\beta_J \wedge \mathrm{d}x^I \wedge \mathrm{d}x^J \\
&= (\mathrm{d}\alpha_I \wedge \mathrm{d}x^I) \wedge (\beta_J \, \mathrm{d}x^J) + (-1)^p (\alpha_I \, \mathrm{d}x^I) \wedge (\mathrm{d}\beta_J \wedge \mathrm{d}x^J) \\
&= (\mathrm{d}\alpha) \wedge \beta + (-1)^p \alpha \wedge (\mathrm{d}\beta)
\end{aligned}$$

as desired. □

The exterior derivative turns out to be the unique map $\Omega^\bullet(X) \to \Omega^{\bullet+1}(X)$ making the proposition true.

Corollary 3.2. *Given $F : X \to Y$ and a p-form α on Y, we have* $\mathrm{d}(F^*\alpha) = F^*(\mathrm{d}\alpha)$.

Proof. We have by (iii) and (iv) in the preceding proposition that

$$\begin{aligned}
F^*(\mathrm{d}\alpha) &= F^*(\mathrm{d}\alpha_I \wedge \mathrm{d}y^{i_1} \wedge \cdots \wedge \mathrm{d}y^{i_p}) = F^*(\mathrm{d}\alpha_I) \wedge F^*(\mathrm{d}y^{i_1}) \wedge \cdots \wedge F^*(\mathrm{d}y^{i_p}) \\
&= \mathrm{d}(F^*\alpha_I) \wedge \mathrm{d}(F^* y^{i_1}) \wedge \cdots \wedge \mathrm{d}(F^* y^{i_p}) \\
&= \mathrm{d}((F^*\alpha_I)\, \mathrm{d}(F^* y^{i_1}) \wedge \cdots \wedge \mathrm{d}(F^* y^{i_p})) = \mathrm{d}(F^*\alpha)
\end{aligned}$$

where the sums are over multi-indices $I = \{i_1 < \cdots < i_p\}$. □

3.2 De Rham Cohomology

Definition 3.3. A form α is closed if $\mathrm{d}\alpha = 0$, exact if there is some β with $\alpha = \mathrm{d}\beta$. We write $Z^r(X), B^r(X) \subset \Omega^r(X)$ to denote the spaces of closed and exact r-forms, respectively.

It's clear that $Z^r(X), B^r(X)$ are both real vector spaces. Since $\mathrm{d}^2 = 0$, we have $B^r(X) \subset Z^r(X)$.

Definition 3.4. The r-th de Rham cohomology of X, denoted $H^r_{\mathrm{dR}}(X)$ is the quotient $Z^r(X)/B^r(X)$.

Note that $H^r_{\mathrm{dR}}(X) = 0$ for $r > \dim X$ and for $r < 0$.

Example 3.1. (i) $H^0_{\mathrm{dR}}(X) = Z^0(X)/B^0(X) = \{f : \mathrm{d}f = 0\}/0 \cong \{f : f \text{ locally constant}\} \cong \mathbb{R}^{\{\text{components of } X\}}$.
(ii) $H^r_{\mathrm{dR}}(\{*\})$ is isomorphic to $\mathbb{R}$ for $r = 0$ and vanishes otherwise.
(iii) $H^r_{\mathrm{dR}}(S^1)$ is isomorphic to $\mathbb{R}$ for $r = 0$ and vanishes for $r \neq 0, 1$. As for $H^1_{\mathrm{dR}}(S^1)$, note that a 1-form on S^1 can be written uniquely in the form $f(\theta)\, \mathrm{d}\theta$. All 1-forms are closed due to the vanishing of $\Omega^2(S^1)$. For the exact forms, consider the map $\Omega^1(S^1) \to \mathbb{R}$ via

$$f(\theta)\, \mathrm{d}\theta \mapsto \int_0^{2\pi} f(\theta)\, \mathrm{d}\theta$$

which is $\mathbb{R}$-linear, nonzero, and hence surjective. Its kernel consists of those 1-forms whose integral over $\theta = [0, 2\pi)$ is zero. This is exactly $B^1(S^1)$: If $f\,\mathrm{d}\theta = \mathrm{d}g$, then $f = g'$ and hence the integral vanishes as $g(2\pi) = g(0)$ (notice that g is a smooth function on S^1). Conversely, if the integral vanishes, then

$$g(\theta) = \int_0^\theta f(t)\,\mathrm{d}t$$

must have $\mathrm{d}g = f\,\mathrm{d}\theta$.
Consequently we have $H^1_{\mathrm{dR}}(S^1) \cong \mathbb{R}$.

Proposition 3.3 (Contravariant Functoriality)**.** *If $F : X \to Y$ is a smooth map, then F^* induces a linear map $H^r_{\mathrm{dR}}(Y) \to H^r_{\mathrm{dR}}(X)$ for each r.*

By abuse of notation, we also denote the induced map by F^*.

Proof. If $\alpha \in Z^r(Y)$, then $\mathrm{d}(F^*\alpha) = F^*(\mathrm{d}\alpha) = 0$, so $F^*\alpha$ too is closed. Similarly, if $\alpha \in B^r(Y)$, say $\alpha = \mathrm{d}\beta$, then $F^*\alpha$ is also exact since $F^*\alpha = F^*\mathrm{d}\beta = \mathrm{d}(F^*\beta)$. □

Example 3.2. Recall that $H^1_{\mathrm{dR}}(S^1) \cong \mathbb{R}$, then the map $S^1 \to S^1, z \mapsto z^n$ induces multiplication by n on $H^1_{\mathrm{dR}}(S^1)$.

Proposition 3.4. *The wedge product $\Omega^i(X) \times \Omega^j(X) \to \Omega^{i+j}(X)$ descends to a pairing $H^i_{\mathrm{dR}}(X) \times H^j_{\mathrm{dR}}(X) \to H^{i+j}_{\mathrm{dR}}(X)$ on cohomology, which makes $H^*_{\mathrm{dR}}(X) = \bigoplus_i H^i_{\mathrm{dR}}(X)$ a unital graded-commutative associative $\mathbb{R}$-algebra.*

Proof. Write $[\alpha]$ to denote the cohomology class of the form α. Given $[\alpha], [\beta] \in H^*_{\mathrm{dR}}(X)$, we know that $\mathrm{d}(\alpha \wedge \beta) = (\mathrm{d}\alpha) \wedge \beta + (-1)^{|\alpha|}\alpha \wedge (\mathrm{d}\beta) = 0$. Furthermore, $[\alpha \wedge \beta]$ depends only on $[\alpha]$ and $[\beta]$. Indeed, if $\alpha' = \alpha + \mathrm{d}\gamma, \beta' = \beta + \mathrm{d}\delta$, then

$$\begin{aligned}\alpha' \wedge \beta &= \alpha \wedge \beta + (\mathrm{d}\gamma) \wedge \beta + \alpha \wedge (\mathrm{d}\delta) + (\mathrm{d}\gamma) \wedge (\mathrm{d}\delta) \\ &= \alpha \wedge \beta + \mathrm{d}(\gamma \wedge \beta + (-1)^{|\alpha|}\alpha \wedge \delta + \gamma \wedge (\mathrm{d}\delta))\end{aligned}$$

So $[\alpha' \wedge \beta'] = [\alpha \wedge \beta]$. □

Since F^* commutes with wedge product and $F^*1 = 1$, then the pullback map $F^* : H^*_{\mathrm{dR}}(Y) \to H^*_{\mathrm{dR}}(X)$ is a homomorphism of unital algebras.
Recall

Definition 3.5. Smooth maps $F_0, F_1 : X \to Y$ are smoothly homotopic if there is a smooth map $F : X \times \mathbb{R} \to Y$ such that $F(-, 0) = F_0, F(-, 1) = F_1$.

Proposition 3.5 (Homotopy Invariance)**.** *If $F_0, F_1 : X \to Y$ are smoothly homotopic, then $F_0^*, F_1^* : H^*_{\mathrm{dR}}(Y) \to H^*_{\mathrm{dR}}(X)$ agree.*

Proof. Ehhh let's do it later. Probably. □

Corollary 3.6. *If $F : X \to Y$ is a homotopy equivalence (i.e. there is some $G : Y \to X$ such that $F \circ G$ is smoothly homotopic to id_Y and $G \circ F$ is smoothly homotopic to id_X), then $H^*_{\mathrm{dR}}(Y) \cong H^*_{\mathrm{dR}}(X)$ via F^*.*

Proof. G^* would be an inverse to F^* by the preceding proposition. □

Example 3.3 (Poincaré Lemma)**.** For all n, we have $H^*_{\mathrm{dR}}(\mathbb{R}^n) \cong H^*_{\mathrm{dR}}(\{*\})$.

3.3 Orientations

Definition 3.6. An orientation of an n-dimensional vector space V is a nonzero element of $\bigwedge^n V$ modulo positive rescalings.
An orientation of a vector bundle $E \to X$ of rank k is a nowhere zero section of $\bigwedge^k V$ modulo rescaling by positive smooth functions. The vector bundle is called orientable if an orientation exists. An oriented vector bundle is a vector bundle equipped with a choice of orientation.

Remark. A rank k vector bundle E is orientable if and only if $\bigwedge^k E$ is trivial.

Example 3.4. 1. Any trivial bundle is orientable (thank god).
2. The tautological bundle of $\mathbb{RP}^n$ is not orientable (example sheet).

Definition 3.7. A manifold is orientable if its tangent bundle is orientable. An oriented manifold is a manifold equipped with a choice of orientation for its tangent bundle.

Example 3.5. 1. S^n is orientable for all n.
2. $\mathbb{RP}^n$ is orientable for some but not all n (example sheet).

Definition 3.8. A volume form on an n-manifold X is a nowhere zero n-form on X.

A volume form ω determines an orientation: We can say a basis $e_1, \ldots, e_n$ of T_pX is positively oriented iff $\omega(e_1, \ldots, e_n) > 0$. Conversely, any orientation gives rise to a volume form modulo scaling by positive smooth functions.

3.4 Partition of Unity

Definition 3.9. Given an open cover $\{U_\alpha\}_\alpha$ of a manifold X. A partition of unity subordinate to the cover $\{U_\alpha\}_\alpha$ is a collection of smooth functions $\rho_\alpha : X \to \mathbb{R}_{\geq 0}$ such that:
(i) $\operatorname{Supp} \rho_\alpha = \overline{\{x \in X : \rho_\alpha(x) \neq 0\}} \subset U_\alpha$.
(ii) For every $p \in X$, there is some neighbourhood V of p such that all but finitely many ρ_α are identically zero.
(iii) For every $p \in X$, $\sum_\alpha \rho_\alpha(p) = 1$ (note that this is always a finite sum by (ii)).

It's a fact that there always exists a partition of unity subordinate to any open cover $\{U_\alpha\}_\alpha$ of X.

3.5 Integration

Fix an oriented n-manifold X and a compactly supported n-form ω on X.

Definition 3.10. The integral of ω over X, denoted

$$\int_X \omega$$

is defined as follows:
First, cover X by local coordinate patches $\{U_\alpha\}_\alpha$ with local coordinates x_α^i, positively oriented in the sense that $\bigwedge_i \partial_{x_\alpha^i}$ respects the orientation. Then,

we pick a partition of unity $\{\rho_\alpha\}_\alpha$ subordinate to the cover. Write $\rho_\alpha\omega = f_\alpha \, \mathrm{d}^1_\alpha \wedge \cdots \wedge \mathrm{d}x^n_\alpha$ and we define

$$\int_X \omega = \sum_\alpha \int_{\mathbb{R}^n} f_\alpha \, \mathrm{d}x^1_\alpha \cdots \mathrm{d}x^n_\alpha$$

Remark. The sum is in fact finite, since ω is compactly supported.

Lemma 3.7. *The integral of ω is well-defined, i.e. independent of the choice of chart and the choice of a partition of unity.*

Proof. Cover X alternatively by patches V_β with coordinates $y^1_\beta, \ldots, y^n_\beta$, positively oriented. Take a partition of unity σ_β subject to this cover. Then $\sigma_\beta\omega = g_\beta \, \mathrm{d}y^1_\beta \wedge \cdots \wedge \mathrm{d}y^n_\beta$.
On overlaps $U_\alpha \cap V_\beta$, we have $\sigma_\beta f_\alpha \, \mathrm{d}^1_\alpha \wedge \cdots \wedge \mathrm{d}x^n_\alpha = \sigma_\beta\rho_\alpha\omega = \rho_\alpha\sigma_\beta\omega = \rho_\alpha g_\beta \, \mathrm{d}y^1_\beta \wedge \cdots \wedge y^n_\beta$, so $\sigma_\beta f_\alpha = \rho_\alpha g_\beta J$ where

$$J = \det\left(\frac{\partial y^j_\beta}{\partial x^i_\alpha}\right)_{i,j}$$

So we can just calculate

$$\begin{aligned}
\sum_\alpha \int_{\mathbb{R}^n} f_\alpha \, \mathrm{d}x^1_\alpha \cdots \mathrm{d}x^n_\alpha &= \sum_{\alpha,\beta} \int_{\mathbb{R}^n} \sigma_\beta f_\alpha \, \mathrm{d}x^1_\alpha \cdots \mathrm{d}x^n_\alpha \\
&= \sum_{\alpha,\beta} \int_{\mathbb{R}^n} \rho_\alpha g_\beta J \, \mathrm{d}x^1_\alpha \cdots \mathrm{d}x^n_\alpha \\
&= \sum_{\alpha,\beta} \int_{\mathbb{R}^n} \rho_\alpha g_\beta \, \mathrm{d}y^1_\beta \cdots \mathrm{d}y^n_\beta \\
&= \sum_\beta \int_{\mathbb{R}^n} g_\beta \, \mathrm{d}y^1_\beta \cdots \mathrm{d}y^n_\beta
\end{aligned}$$

Note that $J \geq 0$ since both coordinates are positively oriented. □

3.6 Stokes' Theorem

The fundamental theorem of calculus states that, under reasonable hypotheses, integrating f' over $[a,b]$ is the same as calculating $f(b) - f(a)$. We can write this as

$$\int_{[a,b]} \mathrm{d}f = \int_{\partial[a,b]} f$$

Definition 3.11. A smooth n-manifold with boundary is defined the same way as an ordinary n-manifold, except we allow things to be locally homeomorphic to an open set of $\mathbb{H}^n = \mathbb{R}_{\geq 0} \times \mathbb{R}^{n-1}$.
Given a smooth n-manifold X with boundary and a chart $U \to V$ containing $p \in X$, we say p is in the boundary ∂X if its image in $\mathbb{H}^n$ lands in $\{0\} \times \mathbb{R}^{n-1}$. Otherwise, we say p is in the interior X°.

It's a fact that the notion of being in the boundary/interior is independent of the choice of chart. Smooth functions can be defined analogously as in the usual way (recall that a map $\mathbb{H}^n \times \mathbb{R}^n$ is smooth if it can be extended smoothly to an open set in $\mathbb{R}^n$ containing $\mathbb{H}^n$).

Example 3.6. 1. Any n-manifold is naturally an n-manifold with boundary, with empty boundary.
2. The closed unit ball $X = \{x \in \mathbb{R}^n : |x| \leq 1\}$ is a manifold with boundary whose boundary is S^{n-1} and whose interior is $\{x \in \mathbb{R}^n, |x| < 1\}$.
3. If X is a manifold with boundary and Y an ordinary manifold, then $X \times Y$ is a manifold with boundary with boundary $\partial(X \times Y) = (\partial X) \times Y$. But if Y also has boundary, the $X \times Y$ may not be a manifold with boundary.
4. If X is an n-manifold with boundary, then ∂X is naturally an $(n-1)$-manifold and X° an n-manifold.

Definition 3.12. If X is an oriented n-manifold with boundary, then we orient ∂X as follows: Given $p \in \partial X$, pick a section o_X of $\bigwedge^n T_p X$ representing the orientation of X. Pick a vector $n \in T_p(X)$ transverse to ∂X and pointing outwards (i.e. pointing to the negative direction over a chart). Orient ∂X at p by the unique section $o_{\partial X}$ of $\bigwedge^{n-1} T_p \partial X \subset \Lambda^{n-1} T_p X$ satisfying $o_X = n \wedge o_{\partial X}$.

Theorem 3.8 (Stokes' Theorem)**.** *Given an oriented n-manifold X with boundary and a compactly supported $(n-1)$-form ω on X, we have*

$$\int_X \mathrm{d}\omega = \int_{\partial X} \omega$$

Formally, the right hand side means the integral of $i^*\omega$ (where $i : \partial X \to X$ is the inclusion) over ∂X.

Proof. Cover X by coordinate patches $\{U_\alpha\}_\alpha$ with coordinates (x^i_α) and pick a subordinate partition of unity ρ_α. Then

$$\int_{\partial X} \omega = \int_{\partial X} \sum_\alpha \rho_\alpha \omega = \sum_\alpha \int_{\partial U_\alpha} \rho_\alpha \omega$$

On the other hand,

$$\int_X \mathrm{d}\omega = \int_X \mathrm{d}\left(\sum_\alpha \rho_\alpha \omega\right) = \sum_\alpha \int_{U_\alpha} (\rho_\alpha \omega)$$

So it suffices to prove the result when $X = U_\alpha$ is just a single coordinate patch, i.e. WLOG $X = \mathbb{H}^n$.
Write $\omega = \sum_i \omega_i \, \mathrm{d}x \wedge \cdots \wedge \widehat{\mathrm{d}x^i} \wedge \cdots \wedge \mathrm{d}x^n$. We want to show that

$$\int_{\mathbb{H}^n} \sum_i (-1)^{i-1} \frac{\partial \omega_i}{\partial x^i} \, \mathrm{d}x^1 \cdots \mathrm{d}x^n = \int_{\partial \mathbb{H}^n} \omega$$

By the usual fundamental theorem of calculus, we have

$$\int_{\mathbb{R}^{n-1}} \left(\int_0^\infty \frac{\partial \omega_1}{\partial x^1} \, \mathrm{d}x^1 \right) \mathrm{d}x^2 \cdots \mathrm{d}x^n = \int_{\mathbb{R}^{n-1}} -\omega_1 \, \mathrm{d}x^2 \cdots \mathrm{d}x^n$$

and, for $j > 1$,

$$\int_{\mathbb{R}_{\geq 0} \times \mathbb{R}^{n-2}} \left(\int_{-\infty}^\infty \frac{\partial \omega_j}{\partial x^j} \, \mathrm{d}x^j \right) \mathrm{d}x^1 \cdots \widehat{\mathrm{d}x^j} \cdots \mathrm{d}x^n = 0$$

with the awareness that ω vanishes at infinity since it is compactly supported. But this just gives us what we wanted since $i^*\omega = \omega_1 \, \mathrm{d}x^2 \wedge \cdots \wedge \mathrm{d}x^n$ and ∂X is oriented by $-\partial_{x_2} \wedge \cdots \wedge \partial_{x_n}$. □

Example 3.7. Take $X = \{x \in \mathbb{R}^2 : |x| \leq a\}$. The area of X is the integral of the area form $\mathrm{d}x \wedge \mathrm{d}y$ on X. But $\mathrm{d}x \wedge \mathrm{d}y = (1/2)\mathrm{d}(x\,\mathrm{d}y - y\,\mathrm{d}x)$, so we have

$$\begin{aligned}\int_X \mathrm{d}x \wedge \mathrm{d}y &= \frac{1}{2}\int_X \mathrm{d}(x\,\mathrm{d}y - y\,\mathrm{d}x) = \frac{1}{2}\int_{\partial X} x\,\mathrm{d}y - y\,\mathrm{d}x \\ &= \frac{1}{2}\int_{\partial X} r^2\,\mathrm{d}\theta = \frac{a^2}{2}\int_{\partial X}\mathrm{d}\theta = \frac{a^2}{2}(2\pi) = \pi a^2\end{aligned}$$

Yay (?).

3.7 Applications of Stokes' Theorem

Proposition 3.9 (Integration by Parts)**.** *Suppose X is an oriented n-manifold with boundary, a $(p-1)$-form α on X and an $(n-p)$-form β, at least one of which is compactly supported. Then*

$$\int_X (\mathrm{d}\alpha) \wedge \beta = \int_{\partial X} \alpha \wedge \beta + (-1)^p \int_X \alpha \wedge (\mathrm{d}\beta)$$

Proof. We know that $\mathrm{d}(\alpha \wedge \beta) = (\mathrm{d}\alpha) \wedge \beta + (-1)^p \alpha \wedge (\mathrm{d}\beta)$. Integrating this and applying Theorem 3.8 gives

$$\int_{\partial X} \alpha \wedge \beta = \int_X (\mathrm{d}\alpha) \wedge \beta + (-1)^{p-1} \int_X \alpha \wedge (\mathrm{d}\beta)$$

which is the claim. □

Proposition 3.10. *If X is a compact oriented n-manifold (without boundary), then integration over X defines a linear map $H^n_{\mathrm{dR}}(X) \to \mathbb{R}$.*

Proof. Theorem 3.8 shows that if $[\alpha] = [\alpha']$, then $\alpha' = \alpha + \mathrm{d}\beta$ for some β, hence

$$\int_X \alpha' - \int_X \alpha = \int_X \mathrm{d}\beta = \int_{\partial X} \beta = 0$$

since $\partial X = \varnothing$. □

Corollary 3.11. *If X is a compact orientable n-manifold (without boundary), then $H^n_{\mathrm{dR}}(X) \neq 0$.*

Proof. Fix an orientation on X. Choose a volume form ω representing this orientation. Since $\mathrm{d}\omega$ is an $(n+1)$-form, it must have degree 0. So we get a class $[\omega] \in H^n_{\mathrm{dR}}(X)$. But its integral over X is positive by definition, so $[\alpha] \neq 0$ by the preceding proposition. □

4 Connections on Vector Bundles

Given a vector bundle $E \to B$, an E-valued r-form is a section of $E \otimes \bigwedge^r T^*B$. For a vector space V, a V-valued r-form is a $\underline{V}$-valued r-form. We write $\Omega^r(E)$ to denote the space of E-valued r-forms, and $\Gamma(E) = \Omega^0(E)$ to denote the space of global sections of E.
We also write $\mathfrak{gl}(k, \mathbb{R})$ for the manifold consists of the $k \times k$ real matrices.

4.1 Connections

Fix a vector bundle $E \to B$ of rank k. For a global section s of it, we can write it locally under each trivialisation Φ_α as an $\mathbb{R}^k$-valued function, which we'll denote by v_α.
Do we want to differentiate s by differentiating v_α on each trivialisation? Well, let's see what happens:
The naïve derivative is given by $\mathrm{d}v_\alpha$, which one can view as a local E-valued 1-form over U_α. But if Φ_β is a different trivialisation, we have $v_\beta = g_{\beta\alpha} v_\alpha$. Then if we differentiate this thing and move it back to the α trivialisation, we get $g_{\beta\alpha}^{-1}\,\mathrm{d}(g_{\beta\alpha}v_\alpha) = \mathrm{d}v_\alpha + g_{\beta\alpha}^{-1}(\mathrm{d}g_{\beta\alpha})v_\alpha$. So these $\mathrm{d}v_\alpha$ are trivialisation-dependent due to the extra $g_{\beta\alpha}^{-1}(\mathrm{d}g_{\beta\alpha})v_\alpha$ term.
That is, we have a dependency arising from the $\mathfrak{gl}(k,\mathbb{R})$-valued 1-form $g_{\beta\alpha}^{-1}\,\mathrm{d}g_{\beta\alpha}$. This motivates the definition of a connection.

Definition 4.1. A connection $\mathcal{A}$ on a vector bundle $\pi : E \to B$ of rank k is the data of a $\mathfrak{gl}(k,\mathbb{R})$-valued 1-form $\mathcal{A}_\alpha$ on each trivialisation $\Phi_\alpha : \pi^{-1}U_\alpha \to U_\alpha \times \mathbb{R}^k$ such that $\mathcal{A}_\alpha = g_{\beta\alpha}^{-1}\mathcal{A}_\beta g_{\beta\alpha} + g_{\beta\alpha}^{-1}\,\mathrm{d}g_{\beta\alpha}$.
Given a connection $\mathcal{A}$ of $\pi : E \to B$, the covariant derivative of a section s is the E-valued 1-form $d^{\mathcal{A}}s$ given, under Φ_α, by $\mathrm{d}v_\alpha + \mathcal{A}_\alpha v_\alpha$.

Our earlier calculation then reveals that the covariant derivative is well-defined (i.e. consistent on overlaps). $\mathcal{A}_\alpha$ are known as the local connection 1-forms.

Definition 4.2. The section s is horizontal or covariantly constant for $\mathcal{A}$ if $\mathrm{d}^{\mathcal{A}}s = 0$.

The zero section is always horizontal. Nonzero horizontal sections may not exist, even locally.

Example 4.1 (Trivial connection)**.** Suppose $E \to B$ admits a global trivialisation Φ_α. We define a connection $\mathcal{A}$ by $\mathcal{A}_\alpha = 0$, then define $\mathcal{A}_\beta$ for all other trivialisations by the transformation law of connections. A section is horizontal for $\mathcal{A}$ iff it is locally constant under Φ_α.

Somehow we didn't introduce this notation earlier, but we write $\Gamma(E)$ to denote the space of sections of the vector bundle E.

Lemma 4.1. *Given a connection $\mathcal{A}$ on $E \to B$, the covariant derivative $\mathrm{d}^{\mathcal{A}} : \Gamma(E) \to \Omega^1(E)$ is $\mathbb{R}$-linear and satisfies the Leibniz rule $\mathrm{d}^{\mathcal{A}}(fs) = f\,\mathrm{d}^{\mathcal{A}}s + s \otimes \mathrm{d}f$. Conversely, any $\mathbb{R}$-linear map $\Gamma(E) \to \Omega^1(E)$ satisfying this Leibniz rule arises from a unique connection in this way.*

Proof. $\mathbb{R}$-linearity is obvious and the Leibniz rule can be checked locally under trivialisations:

$$\mathrm{d}^{\mathcal{A}}(fs) = \mathrm{d}(fv_\alpha) + \mathcal{A}_\alpha f v_\alpha = v_\alpha \otimes \mathrm{d}f + f\,\mathrm{d}v_\alpha + fA_\alpha v_\alpha = f\,\mathrm{d}^{\mathcal{A}}s + s \otimes \mathrm{d}f$$

The converse is on example sheet. □

Example 4.2. Given a submanifold $i : X \hookrightarrow \mathbb{R}^N$, $i^*T\mathbb{R}^N$ has a standard trivialisation and hence a trivial connection $\mathcal{A}_0$.

$$\Gamma(TX) \longrightarrow \Gamma(i^*T\mathbb{R}^N) \xrightarrow{\mathrm{d}^{\mathcal{A}_0}} \Omega^1(i^*T\mathbb{R}^N) \longrightarrow \Omega^1(TX)$$

where the last arrow comes from the orthogonal projection $i^*T\mathbb{R}^N \to TX$. This is $\mathbb{R}$-linear and inerits the Leibniz rule from $\mathrm{d}^{\mathcal{A}}$, which gives us a unique connection on TX.

Lemma 4.2. *Any vector bundle admits a connection.*

Proof. Trivialise E over U_α with transition functions $g_{\beta\alpha}$ as usual. Pick a partition of unity subordinate to $\{U_\alpha\}_\alpha$. Now define $\mathcal{A}_\alpha = \sum_\gamma \rho_\gamma g_{\gamma\alpha}^{-1}\,\mathrm{d}g_{\gamma\alpha}$. Let's show that it works. We have

$$\begin{aligned} g_{\beta\alpha}^{-1}A_\beta g_{\beta\alpha} &= \sum_\gamma \rho_\gamma g_{\beta\alpha}^{-1}(g_{\gamma\beta}^{-1}\,\mathrm{d}g_{\gamma\beta})g_{\beta\alpha} = \sum_\gamma \rho_\gamma g_{\gamma\alpha}^{-1}\left(\mathrm{d}(g_{\gamma\beta}g_{\beta\alpha}) - g_{\gamma\beta}\,\mathrm{d}g_{\beta\alpha}\right) \\ &= \sum_\gamma \rho_\gamma g_{\gamma\alpha}^{-1}\,\mathrm{d}g_{\gamma\alpha} - \sum_\gamma \rho_\gamma g_{\beta\alpha}\,\mathrm{d}g_{\beta\alpha} = A_\alpha - g_{\beta\alpha}^{-1}\,\mathrm{d}g_{\beta\alpha} \end{aligned}$$

which is our beloved transformation law. □

4.2 Connections and the Endomorphism Bundle

Fix a rank k vector bundle $E \to B$. Let $\rho : \mathrm{GL}(k,\mathbb{R}) \to \mathrm{GL}(\mathfrak{gl}(k,\mathbb{R}))$ be the representation $\rho(A)(M) = AMA^{-1}$.

Definition 4.3. The endomorphism bundle $\mathrm{End}(E)$ of E is the vector bundle over B of rank k^2 with total space $\coprod_{b\in B}\mathrm{End}(E_b)$, and if E is trivialised over $\{U_\alpha\}_\alpha$ with transition functions $g_{\beta\alpha}$ then $\mathrm{End}(E)$ is trivialised over $\{U_\alpha\}_\alpha$ with transition function $\rho(g_{\beta\alpha})$.

A section M of $\mathrm{End}(E)$ is then locally a matrix-valued function M_α such that $M_\beta = g_{\beta\alpha}M_\alpha g_{\beta\alpha}^{-1}$. Equivalently, $\mathrm{End}(E) = E \otimes E^\vee$ using the fibrewise identificatio $\mathrm{End}(V) = V \otimes V^\vee$.

Lemma 4.3. *Given a connection $\mathcal{A}$ on E and a section Δ of $\Omega^1(\mathrm{End}(E))$, there exists a connection $\mathcal{A}+\Delta$ on E given locally by $\mathcal{A}_\alpha + \Delta_\alpha$. Conversely, every connection $\mathcal{A}'$ on E can be written uniquely as $\mathcal{A}+\Delta$ for some Δ. Consequently, the space of connections on E is an affine space of $\Omega^1(\mathrm{End}(E))$, i.e. it carries a free transitive action of $\Omega^1(\mathrm{End}(E))$.*

Proof. To check that $\mathcal{A}+\Delta$ indeed transforms in the right way, we just calculate

$$\mathcal{A}_\alpha + \Delta_\alpha = g_{\beta\alpha}^{-1}\mathcal{A}_\beta g_{\beta\alpha} + g_{\beta\alpha}^{-1}\,\mathrm{d}g_{\beta\alpha} + g_{\alpha\beta}\Delta_\beta g_{\alpha\beta}^{-1} = g_{\beta\alpha}^{-1}(\mathcal{A}_\beta + \Delta_\beta)g_{\beta\alpha} + g_{\beta\alpha}^{-1}\,\mathrm{d}_{\beta\alpha}$$

Similarly, if $\mathcal{A}'$ is another connection, then $\mathcal{A}' - \mathcal{A}$ transforms correctly for a section Δ of $\Omega^1(\mathrm{End}(E))$. □

4.3 Curvature, Algebraically (Differentially?)

Fix a connection $\mathcal{A}$ on a vector bundle $E \to B$.

Definition 4.4. The exterior covariant derivative is the unique $\mathbb{R}$-linear map $\mathrm{d}^{\mathcal{A}} : \Omega^\bullet(E) \to \Omega^{\bullet+1}(E)$ satisfying the Leibniz rule $\mathrm{d}^{\mathcal{A}}(\sigma\otimes\omega) = (\mathrm{d}^{\mathcal{A}}\sigma)\wedge\omega + s\otimes\mathrm{d}\omega$ for E-valued forms σ and forms ω.

In local trivialisations, an E-valued p-form σ becomes an $\mathbb{R}^k$-valued p-form σ_α and $\mathrm{d}^{\mathcal{A}}\sigma$ is given by $\mathrm{d}\sigma_\alpha + A_\alpha \wedge \sigma_\alpha$. Note that $(\mathrm{d}^{\mathcal{A}})^2$ is not necessarily zero.

Proposition 4.4. *There exists a unique* $\mathrm{End}(E)$*-valued 2-form* F *on* B *such that for any* E*-valued forms* σ *such that* $(\mathrm{d}^{\mathcal{A}})^2\sigma = F \wedge \sigma$.

Proof. In a local trivialisation, $(\mathrm{d}^{\mathcal{A}})^2\sigma$ is given by $\mathrm{d}(\mathrm{d}\sigma_\alpha + \mathcal{A}_\alpha \wedge \sigma_\alpha) + \mathcal{A}_\alpha \wedge (\mathrm{d}\sigma_\alpha + \mathcal{A}_\alpha \wedge \sigma_\alpha) = \mathrm{d}^2\sigma_\alpha + (\mathrm{d}\mathcal{A}_\alpha) \wedge d\sigma_\alpha - \mathcal{A}_\alpha \wedge \mathrm{d}\sigma_\alpha + \mathcal{A}_\alpha \wedge \mathrm{d}\sigma_\alpha + \mathcal{A}_\alpha \wedge \mathcal{A}_\alpha \wedge \sigma_\alpha = (\mathrm{d}\mathcal{A}_\alpha) \wedge \sigma_\alpha + \mathcal{A}_\alpha \wedge \mathcal{A}_\alpha \wedge \sigma_\alpha$. So our only choice is $F_\alpha = \mathrm{d}\mathcal{A}_\alpha + \mathcal{A}_\alpha \wedge \mathcal{A}_\alpha$. You'll check that these transform correctly on example sheet. □

Definition 4.5. F is the curvature of $\mathcal{A}$. We say $\mathcal{A}$ is flat if $F = 0$.

Example 4.3. (i) Trivial connections are flat. Conversely, if $\mathcal{A}$ is flat then it is locally trivial (example sheet).
(ii) Consider $\underline{\mathbb{R}^2}$ on $\mathbb{R} \times S^1$. Suppose we have a connection $\mathcal{A}$ in the form

$$\mathcal{A}_\alpha = f \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \mathrm{d}x + g \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \mathrm{d}\theta$$

So its curvature is

$$\begin{aligned} F_\alpha &= \mathrm{d}\mathcal{A}_\alpha + \mathcal{A}_\alpha \wedge \mathcal{A}_\alpha \\ &= \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \mathrm{d}f \wedge \mathrm{d}x + \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \mathrm{d}g \wedge \mathrm{d}\theta \\ &\quad + fg \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \mathrm{d}x \wedge \mathrm{d}\theta + fg \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \mathrm{d}\theta \wedge \mathrm{d}x \\ &= \left(-\frac{\partial f}{\partial \theta} \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} + \frac{\partial g}{\partial x} \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} - 2fg \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \right) \mathrm{d}x \wedge \mathrm{d}\theta \end{aligned}$$

which is generally nontrivial.

Time for geometric motivations. F measures the degree to which $\mathcal{A}$ fails to be trivial, and being trivial is the same as admitting a fibrewise basis of horizontal sections. It is obstructed precisely when things are "curved", and F measures how deformed an infinitesimal parallelopiped become when passing along $\pi : E \to B$.
We will expand upon this idea using the notion of parallel transport.

4.4 Parallel Transport

Suppose we have a vector bundle $E \to [0, 1]$ on which we have a connection $\mathcal{A}$.

Lemma 4.5. *Given any* $v_0 \in E_0$*, there is a unique horizontal section* s *with* $s(0) = v_0$ *which depends linearly on* v_0.

Proof. In local trivialisations, the condition of s being horizontal is to say $\mathrm{d}v_\alpha + A_\alpha v_\alpha = 0$. As $A_\alpha = M_\alpha \,\mathrm{d}t$ for some $\mathfrak{gl}(k, \mathbb{R})$-valued function M_α, where t is the coordinate on $[0, 1]$. So the horizontality condition is the same as $\mathrm{d}v_\alpha/\mathrm{d}t + M_\alpha v_\alpha = 0$. This is just a linear ODE of order 1, so there is a unique solution locally given the initial condition v_0, and the linear dependence is clear. The local existence implies that there is a fibrewise basis of horizontal sections locally around p for any $p \in [0, 1]$. By the compactness of $[0, 1]$, there are some $0 = a_0 < \cdots < a_N = 1$ such that on a small open interval U_i containing $[a_i, a_{i+1}]$ we have such a basis $s_i^1, \ldots, s_i^k$.

Now write v_0 (uniquely) as $v_0 = \sum_{j=1}^k \lambda_{0,j} s_0^j(0)$. Define s on U_0 by $\sum_j \lambda_{0,j} s_0^j$. Continue this for $a_1, a_2, \ldots$ with initial conditions given by the previous iteration. □

Definition 4.6. The linear map $E_0 \to E_1, v_0 \mapsto s(1)$, where s is as in the preceding lemma, is called the parallel transport of v_0 along $[0,1]$ (with respect to $\mathcal{A}$).

How about over a general base? Suppose $E \to B$ is a vector bundle with a connection $\mathcal{A}$. For a path $\gamma : [0,1] \to B$, it induces a pullback bundle $\gamma^* E$ on $[0,1]$. Furthermore, $\gamma^* A_\alpha$ defines a connection on $\gamma^* E$.

Definition 4.7. Given a vector $v_0 \in E_{\gamma(0)}$, the unique horizontal section of $\gamma^* E$ starting at v_0 is the horizontal lift of γ to E (starting at v_0). The vector $s(1) \in E_{\gamma(1)}$ is the parallel transport of v_0 along γ. Doing this for all v_0 gives a linear map $P_\gamma : E_{\gamma(0)} \to E_{\gamma(1)}$. If γ is a loop, then P_γ is an endomorphism of $E_{\gamma(0)} = E_{\gamma(1)}$ and is known as the monodromy (or holonomy) of $\mathcal{A}$ around γ.

Note that P_γ is always an isomorphism, and its inverse is given by the parallel transport along the curve going backwards.

Example 4.4. 1. Consider TS^2 with orthogonal projection connection. Given a path $\gamma : [0,1] \to S^2$ and $v_0 \in T_{\gamma(0)} S^2$. The horizontal lift is the map $v : [0,1] \to TS^2$ such that $v(t) \in T_{\gamma(t)} S^2 \subset \mathbb{R}^3$ for all t and $\dot{v}(t)$ is orthogonal to $T_{\gamma(t)} S^2$.
2. Let's bring back our lovely connection on the trivial bundle $\underline{\mathbb{R}^2}$ on $\mathbb{R} \times S^1$ of the form

$$\mathcal{A}_\alpha = f \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \mathrm{d}x + g \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \mathrm{d}\theta$$

Consider $\gamma(t) = (t, 0)$. $\gamma^* \mathcal{A}_\alpha$ is given by $f \left(\begin{smallmatrix} 1 & 0 \\ 0 & -1 \end{smallmatrix}\right) \mathrm{d}t$, so

$$\mathrm{d}^{\gamma^* \mathcal{A}} = \frac{\mathrm{d}v}{dt} \mathrm{d}t + f \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} v \, \mathrm{d}t$$

So the horizontal lift v starting at v_0 satisfies

$$\dot{v} + f \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} v = 0$$

which means that

$$v(t) = \begin{pmatrix} e^{-\lambda} & 0 \\ 0 & e^{\lambda} \end{pmatrix} v_0, \lambda = \int_0^t f(x, 0) \, \mathrm{d}x$$

Similarly, the monodromy around $\gamma(t) = (0, 2\pi t)$ is

$$\begin{pmatrix} \cos\phi & \sin\phi \\ -\sin\phi & \cos\phi \end{pmatrix}, \phi = \int_0^{2\pi} g(0, \theta) \, \mathrm{d}\theta$$

4.5 Curvature, Geometrically

Fix a vector bundle $E \to B$ with connection $\mathcal{A}$. Fix also a point $p \in B$ and a trivialisation Φ_α on a chart with local coordinates $(x^i)_i$ around p. Let $\mathcal{A}_\alpha = A_i \, \mathrm{d}x^i$ where the A_i's are $\mathfrak{gl}(k, \mathbb{R})$-valued functions. Similarly let $F_\alpha = F_{ij} \, \mathrm{d}x^i \otimes \mathrm{d}x^j$.
WLOG $p = (0, \dots, 0)$. For $a, b \in \mathbb{R}$ small, let $\gamma_1(t) = ate_i, \gamma_2(t) = ae_i + bte_j$ (in local coordinates). Then let γ_3, γ_4 be the other two sides of the rectangle two of whose sides are given by γ_1, γ_2.

Proposition 4.6. *Write $P_{a,b} = P_{\gamma_4} P_{\gamma_3} P_{\gamma_2} P_{\gamma_1} \in \mathrm{End}(E_p)$. Then*

$$\left.\frac{\partial^2 P_{a,b}}{\partial a \partial b}\right|_{a=b=0} = -F_{ij}(p)$$

A fully rigorous proof can be found in example sheet. We're gonna put a intuitive argument here.

Sketchy stuff. Parallel transport in the x^i direction satisfies $\dot{v} = -A_i v$. So $P_{\gamma_1} = I - aA_i(x) + \cdots$ (leaving out the higher order terms, same below). Similarly, $P_{\gamma_2} = I - bA_j(\gamma_1(1)) + \cdots = I - b(A_j(p) + a(\partial A_j / \partial x^i)(p)) + \cdots$. Consequently,

$$P_{\gamma_2} P_{\gamma_1} = I - aA_i(p) - bA_j(p) + abA_j(p)A_i(p) - ab\frac{\partial A_j}{\partial x^i}(p) + \cdots$$

We can do this for $P_{\gamma_4} P_{\gamma_3}$ in exactly the same way, which gives

$$P_{\gamma_4} P_{\gamma_3} = I + aA_i(p) + bA_j(p) - abA_i(p)A_j(p) + ab\frac{\partial A_i}{\partial x^j}(p) + \cdots$$

Multiply everything out, we get

$$P_{a,b} = I + ab\left(\frac{\partial A_i}{\partial x^j}(p) - \frac{\partial A_j}{\partial x^i}(p) + A_j(p)A_i(p) - A_i(p)A_j(p)\right) + \cdots$$

So we have

$$\left.\frac{\partial^2 P_{a,b}}{\partial a \partial b}\right|_{a=b=0} = \frac{\partial A_i}{\partial x^j}(p) - \frac{\partial A_j}{\partial x^i}(p) + A_j(p)A_i(p) - A_i(p)A_j(p) = -F_{ij}(p)$$

Less exciting than you expected, huh? ☐

Corollary 4.7. *If $v \in E_p$ is such that $F(p)v \neq 0$ (in $E_p \otimes \bigwedge^2 T_p^* B$), then there does not exist a local horizontal section s about p with $s(p) = v$.*

Proof. If such an s exists, then $P_{\gamma_1}(v) = s(\gamma_1(1))$. Similar for other P_{γ_i}'s and we conclude $P_{a,b}(v) = s(p) = v$, which by the preceding proposition gives $-F_{ij}(p)v = 0$, so $F(p)v = 0$. ☐

Example 4.5. Consider the trivial bundle $\underline{\mathbb{R}}$ over $\mathbb{R}^2$ with $A_\alpha = Cx^1 \, \mathrm{d}x^2$. Let's calculate its curvature. We have $P_{\gamma_1} = \mathrm{id} = P_{\gamma_3} = P_{\gamma_2}$. On the other hand, $P_{\gamma_2} = e^{-Cab}$ by integrating $\dot{v} + Cav = 0$ from $t = 0$ to $t = b$. Hence $P_{a,b} = e^{-Cab}$ and therefore $-F_{12} = \partial P_{a,b} / \partial a \partial b|_{a=b=0} = -C$.
Explicitly, if s were a local horizontal section about p given by v_α in a standard trivialisation, then we'd have $\mathrm{d}v_\alpha + A_\alpha v_\alpha = 0$, i.e. $\mathrm{d}v_\alpha + Cx^1 v_\alpha \, \mathrm{d}x^2 = 0$. This means that $\partial v_\alpha / \partial x^2 = 0, \partial v_\alpha / \partial x^2 = -Cx^1 v_\alpha$. Hence $0 = \partial^2 v_\alpha / \partial x^2 \partial x^1 = \partial^2 v_\alpha / \partial x^1 \partial x^2 = -Cv_\alpha$. So if $C \neq 0$ then the only horizontal local section is 0.

Example 4.6. Let's now consider $\underline{\mathbb{R}}$ over S^1 with $A_\alpha = C\,\mathrm{d}\theta$ for some constant C. Local horizontal sections exist and have the form $v_\alpha = Ke^{-C\theta}$ for some constant K. This extends to a global section if $C = 0$. If $C \neq 0$, then we've got ourselves a nontrivial monodromy $e^{-2\pi C}$.
So curvature is the local obstruction to the existence of horizontal sections, and monodromy is the global obstruction.

But sometimes we don't want to choose connections but still want to differentiate (tall order). How would we do that?

5 Flows and Lie Derivatives

5.1 Flows

Fix a manifold X and a vector field v on X. Given $p \in X$, we want to "flow along" v from p. For fun haters, we want to solve the ODE $\dot\gamma(t) = v(\gamma(t))$ subject to $\gamma(0) = p$. By standard ODE theory, solutions exist locally, are unique, and depend smoothly on initial conditions. They are known as integral curves of p. Like how all good things start, we introduce a non-standard definition.

Definition 5.1. A flow domain is an open set U containing $X \times \{0\}$ in $X \times \mathbb{R}$ such that for all $p \in X$, $U \cap (\{p\} \times \mathbb{R})$ is connected.

Definition 5.2. A (local) flow of v comprises a flow domain U and a smooth map $\Phi : U \to X$ such that $\Phi(-, 0) = \mathrm{id}_X$ and $(\mathrm{d}/\mathrm{d}t)\Phi(p, t) = v(\Phi(p, t))$ for all $p \in X$.

So $\gamma(t) = \Phi(p, t)$ is an integral curve starting at p. We sometimes write $\Phi^t(p)$ for $\Phi(p, t)$. The discussion on ODEs means that local flows always exist, and are unique in the sense that if $(U_1, \Phi_1), (U_2, \Phi_2)$ are local flows, then $\Phi_1 = \Phi_2$ on $U_1 \cap U_2$. The uniqueness argument immediately implies

Lemma 5.1. *If Φ is a local flow of v, then $\Phi^{s+t} = \Phi^s \circ \Phi^t$ whenever everything makes sense.*

Proof. Fix some point p and suitable t. Then $\gamma_1(s) = \Phi^{s+t}(p)$ and $\gamma_2(s) = (\Phi^s \circ \Phi^t)(p)$ both satisfies $\mathrm{d}\gamma_i/\mathrm{d}s = v \circ \gamma_i$ and $\gamma_i(0) = q$, so $\gamma_1 = \gamma_2$ which gives what we wanted. ☐

Definition 5.3. A vector field v is complete if it has a global flow, i.e. a flow with flow domain $X \times \mathbb{R}$.

Remark. Unsurprisingly, not all vector fields are complete. For example, the vector field $x^2\partial_x$ on $\mathbb{R}$ isn't complete. However, if v is compactly supported, then it is complete: For each $p \in X$, there is a neighbourhood U_p of p and $\epsilon_p > 0$ such that a flow exists on $U_p \times (-\epsilon_p, \epsilon_p)$. By compactness, we can choose ϵ_p to be uniformly bounded below by some $\epsilon > 0$. Then we get a local flow on $X \times (-\epsilon, \epsilon)$. Again since v is compactly supported, a global flow can be defined by $\Phi^t = (\Phi^{t/N})^{\circ N}$ (note the next lemma) for suitably large N.

5.2 Lie Derivatives

Let's tell some lies (I had to).
Fix a manifold X and a vector field v on X. Let Φ be a local flow of v.

Definition 5.4. For a tensor or form T on X, its Lie derivative $\mathscr{L}_v T$ is defined to be

$$\mathscr{L}_v T = \left.\frac{\mathrm{d}}{\mathrm{d}t}\right|_{t=0} (\Phi^t)^* T$$

This is independent of the choice of Φ by local uniqueness.

Lemma 5.2. $(\mathrm{d}/\mathrm{d}t)(\Phi^t)^* T = (\Phi^t)^* \mathscr{L}_v T$.

Proof. We have

$$\begin{aligned} \frac{\mathrm{d}}{\mathrm{d}t}(\Phi^t)^* T &= \left.\frac{\mathrm{d}}{\mathrm{d}t}\right|_{h=0} (\Phi^{t+h})^* T = \left.\frac{\mathrm{d}}{\mathrm{d}t}\right|_{h=0} (\Phi^t)^*(\Phi^h)^* T \\ &= (\Phi^t)^* \left.\frac{\mathrm{d}}{\mathrm{d}t}\right|_{h=0} (\Phi^h)^* T = (\Phi^t)^* \mathscr{L}_v T \end{aligned}$$

□

Lemma 5.3. *For a function f, we have $\mathscr{L}_v f = \mathrm{d}f(v)$. For a 1-form α, we have*

$$\mathscr{L}_v \alpha = \left(v^j \frac{\partial \alpha_i}{\partial x^j} + \alpha_j \frac{\partial v^j}{\partial x^i} \right) \mathrm{d}x^i$$

Proof. At each p, we have

$$\mathscr{L}_v(f) = \left.\frac{\mathrm{d}}{\mathrm{d}t}\right|_{t=0} f(\Phi^t(p)) = \mathrm{d}f \left(\left.\frac{\mathrm{d}}{\mathrm{d}t}\right|_{t=0} \Phi^t(p) \right) = \mathrm{d}f(v)$$

by chain rule. On the other hand,

$$\begin{aligned} \mathscr{L}_v \alpha &= \left.\frac{\mathrm{d}}{\mathrm{d}t}\right|_{t=0} (\Phi^t)^* \alpha = \left.\frac{\mathrm{d}}{\mathrm{d}t}\right|_{t=0} (\alpha_i \circ \Phi^t)\, \mathrm{d}(x^i \circ \Phi^t) = (\mathscr{L}_v \alpha_i)\, \mathrm{d}x^i + \alpha_i\, \mathrm{d}(\mathscr{L}_v x^i) \\ &= v_j \frac{\partial \alpha_i}{\partial x^j} \mathrm{d}x^i + \alpha_i\, \mathrm{d}v^i = \left(v^j \frac{\partial \alpha_i}{\partial x^j} + \alpha_j \frac{\partial v^j}{\partial x^i} \right) \mathrm{d}x^i \end{aligned}$$

which is what we wanted. □

Lemma 5.4. *For a 1-form α and vector field w, $\mathscr{L}_v(\alpha_i w^i) = (\mathscr{L}_v \alpha)_i w^i + \alpha_i (\mathscr{L}_v w)^i$. For any tensors S, T, we have $\mathscr{L}_v(S \otimes T) = (\mathscr{L}_v S) \otimes T + S \otimes (\mathscr{L}_v T)$.*

Proof. Pullback by Φ^t commutes with contraction and with $\otimes$. The lemma then follows from the ordinary Leibniz rule. □

Corollary 5.5. *For a vector field w,*

$$\mathscr{L}_v w = \left(v^j \frac{\partial w^i}{\partial x^j} - w^j \frac{\partial v^i}{\partial x^j} \right) \partial_{x^i}$$

Proof. By the preceding lemma, we have $\mathscr{L}_v(\alpha_i w^i) = (\mathscr{L}_v \alpha)_i w^i + \alpha_i (\mathscr{L}_v w)^i$ for any 1-form α, so

$$v^j \frac{\partial(\alpha_i w^i)}{\partial x^j} = \left(v^j \frac{\partial \alpha_i}{\partial x^j} + \alpha_j \frac{\partial v^j}{\partial x^i} \right) w^i + \alpha_i (\mathscr{L}_v w)^i$$

Expanding these out,

$$v^j \alpha_i \frac{\partial w^i}{\partial x^j} = \alpha_j w^i \frac{\partial v^j}{\partial x^i} + \alpha_i (\mathscr{L}_v w)^i$$

For this to hold for all α, we must have the desired identity. $\square$

Definition 5.5. The Lie bracket of v, w is $[v, w] = \mathscr{L}_v w = -\mathscr{L}_w v$.

This makes $\Gamma(TX)$ a Lie algebra, i.e. a vector space equipped with a bilinear form $[,]$ satisfying $[v, v] = 0$ and the Jacobi identity

$$[u, [v, w]] + [v, [w, u]] + [w, [u, v]] = 0$$

You'll show on example sheet that if $[v, w] = 0$ then their flows commute.

Lemma 5.6. *If $F : X \to Y$ is a diffeomorphism, then for any vector field v on Y and any tensor T on Y, we have $F^*(\mathscr{L}_v T) = \mathscr{L}_{F^* v}(F^* T)$.*

Proof. We have

$$\begin{aligned} F^*(\mathscr{L}_v T) &= F^* \left.\frac{\mathrm{d}}{\mathrm{d}t}\right|_{t=0} (\Phi^t)^* T = \left.\frac{\mathrm{d}}{\mathrm{d}t}\right|_{t=0} F^*(\Phi^t)^* T \\ &= \left.\frac{\mathrm{d}}{\mathrm{d}t}\right|_{t=0} F^*(\Phi^t)^*(F^*)^{-1} F^* T = \left.\frac{\mathrm{d}}{\mathrm{d}t}\right|_{t=0} (F^{-1} \circ \Phi^t \circ F)^*(F^* T) \end{aligned}$$

But $F^{-1} \circ \Phi^t \circ F$ is a flow of $F^* v$. $\square$

5.3 Homotopy Invariance of de Rham Cohomology

Let's prove Proposition 3.5 (finally!).

Definition 5.6. Given an r-form α and a vector field v, the $(r-1)$-form $\iota_v \alpha$ or $v \lrcorner \alpha$ is defined to be $(\iota_v \alpha)_{i_1 \cdots i_{r-1}} = v^j \alpha_{j i_1 \cdots i_{r-1}}$.

Proposition 5.7 (Cartan's magic formula)**.** *For a vector field v and an r-form α, we have*

$$\mathscr{L}_v \alpha = \mathrm{d}(\iota_v \alpha) + \iota_v(\mathrm{d}\alpha)$$

Proof. Example sheet (where all the magic happen). $\square$

Proof of Proposition 3.5. Let $F : [0,1] \times X \to Y$ be a homotopy between $F_0, F_1 : X \to Y$. We write $F_t = F(t, -)$. Let $i_t : X \to [0,1] \times X$ be the inclusion $x \mapsto (t, x)$. Note that $F_t = F \circ i_t, i_t = \Phi^t \circ i_0$ where Φ^t is the flow of ∂_t on $[0,1] \times X$.
For any form α on Y, we have

$$F_1^* \alpha - F^* \alpha = \int_0^1 \frac{\mathrm{d}}{\mathrm{d}t} F_t^* \alpha \,\mathrm{d}t = \int_0^1 i_0^* \frac{\mathrm{d}}{\mathrm{d}t} (\Phi^t)^* F^* \alpha \,\mathrm{d}t = \int_0^1 i_0^* (\Phi^t)^* \mathscr{L}_{\partial_t}(F^* \alpha) \,\mathrm{d}t$$

Suppose α is closed, then Proposition 5.7 gives $\mathscr{L}_{\partial_t}(F^* \alpha) = \mathrm{d}(\iota_{\partial_t} F^* \alpha)$, therefore

$$\begin{aligned} F_1^* \alpha - F^* \alpha &= \int_0^1 i_0^* (\Phi^t)^* \mathrm{d}(\iota_{\partial_t} F^* \alpha) \,\mathrm{d}t \\ &= \int_0^1 i_t^* \mathrm{d}(\iota_{\partial_t} F^* \alpha) \,\mathrm{d}t = \mathrm{d}\left(\int_0^1 i_t^* \iota_{\partial_t} F^* \alpha \,\mathrm{d}t \right) \end{aligned}$$

which is exact. $\square$

6 Frobenius Integrability

6.1 Foliations

Let X be an n-manifold. If $F : X \to Y$ is a submersion, than X decomposes into "slices" $F^{-1}(q), q \in Y$ which are submanifolds of X with codimension $\dim Y$. A k-foliation on X, loosely speaking, is a local decomposition of X into k-dimensional slices, which need not globally form submanifolds of X.

Example 6.1. Consider $X = T^2 = \mathbb{R}^2/\mathbb{Z}^2$. For $\alpha \in \mathbb{R}$, we can locally slice X into lines of slope α. They form submanifolds globally if α is rational. But this is not the case when α is irrational.

Definition 6.1. At atlas on X is k-foliated if the transition functions $\phi_\beta \circ \phi_\alpha^{-1}$ locally sends $(x, y) \in \mathbb{R}^k \times \mathbb{R}^{n-k}$ to $(\zeta(x, y), \eta(y))$ for some smooth ζ, η.
Two k-foliated atlases are equivalent if their union is k-foliated.

On a k-foliated atlas, we often refer to the local coordinates as $x^1, \ldots, x^k$, $y^1, \ldots, y^{n-k}$. We almost always take foliated atlases to be part of the smooth structure of X, in which case the choice of a foliated atlas is thought of as an extra structure equipped to X.

Definition 6.2. Slices on a manifold with a foliated atlas are locally given by $y = c$ for some constant c.

Example 6.2. Suppose $F : X \to Y$ is a submersion. A k-foliated chart corresponds to local coordinates on X in which F corresponds to projection onto the last $n - k$ components.

6.2 Distributions

Fix an n-manifold X.

Definition 6.3. A k-plane distribution on X is a rank k subbundle D of TX.

Example 6.3. On $\mathbb{R}^3$, $\langle \partial_x, \partial_y \rangle$ and $\langle \partial_x + y\partial_z, \partial_y \rangle$ are 2-plane distribution. We can of course also think of them as $\ker \mathrm{d}z, \ker(\mathrm{d}z - y\,\mathrm{d}x)$, respectively. Notably, $\mathrm{d}z$ is closed but $\mathrm{d}z - y\,\mathrm{d}x$ is not.

In general, a k-plane distribution is locally the kernel of $n - k$ 1-forms.

Example 6.4. If X is equipped with a foliation with foliated coordinates x^i, y^j as usual. Then

$$\langle \partial_{x^1}, \ldots, \partial_{x^k} \rangle = \bigcap_{i=1}^{n-k} \ker \mathrm{d}y^i$$

is a k-plane distribution on X. The fact that this comes from a foliation means that it describes the tangent space to the slices.

Definition 6.4. A k-plane distribution is integrable if it arises from a k-foliation.

If this were the case, then we say the distribution is the tangent distribution to the foliation.

Example 6.5. When $k = 1$, every distribution is integrable. Indeed, 1-plane distributions are just vector fields, which arises from the foliation of X by its (local) integral curves.

6.3 Frobenius Integrability

Theorem 6.1. *A distribution D on X is integrable if and only if D is closed under the Lie bracket, i.e. for all vector field v, w tangent to D (meaning that they lie in D), their Lie bracket $[v, w]$ is also tangent to D.*

Proof. Everything's local, so we can work locally around points $p \in X$.
If D is integrable with local foliated coordinates x, y around p, then $D = \langle \partial_{x^1}, \dots, \partial_{x^k} \langle$ which is closed under the Lie bracket.
Conversely, suppose D is closed under the Lie bracket. We need to find local coordinates x, y such that $D = \langle \partial_{x^1}, \dots, \partial_{x^k} \rangle$. First choose local coordinates, s, t about p such that $D = \langle \partial_{s^1}, \dots, \partial_{s^k} \rangle$ at p (but not necessarily beyond just a point). WLOG p is the origin of these coordinates. Locally, there are unique smooth functions a_{ij} such that $v_i = \partial_{s^i} + \sum_j a_{ij} t^j$ is tangent to D.
Let Φ_i^t be the local flow of v_i and consider the function from an open neighbourhood of $0 \in \mathbb{R}^n$ to X defined via $F(x, y) = \Phi_1^{x^1} \circ \dots \circ \Phi_k^{x^k}(s = 0, t = y)$. F is a parameterisation around p since $F(0) = p$ and $\mathrm{D}_0 F(\partial_{x^i}) = v_i(p) = \partial_{s^i}, \mathrm{D}_0 F(\partial_{y^j}) = \partial_{t^j}$.
It remains to show that $\partial_{x^i} = v_i$. If $\Phi_i^{x_i}$ all commute with each other, then this is clearly true. But this always happen since $[v_i, v_j] = 0$ for all i, j: As D is closed under the Lie bracket, there are some $b_{i,j,l}$ such that $[v_i, v_j] = \sum_l b_{i,j,l} v_l$, but equating coefficients of ∂_{s^l} shows that $b_{i,j,l} = 0$ for all i, j, l. □

Example 6.6. 1. Goin' back to our beloved "weird" distribution $D = \langle \partial_x + y\partial_z, \partial_y \rangle$ on $\mathbb{R}^3$. This is not closed under the Lie bracket, since $[\partial_x + y\partial_z, \partial_y] = [\partial_x, \partial_y] + [y\partial_z, \partial_y] = 0 - \partial_z \notin D$, so D is not integrable. We can also do this by hand. Suppose for the sake of contradiction that D were tangent to a surface $f = 0$, then $\partial f/\partial x + y \partial f/\partial z = \partial f/\partial y = 0$. So

$$0 = \frac{\partial^2 f}{\partial x \partial y} + \frac{\partial f}{\partial z} + y \frac{\partial^2 f}{\partial y \partial z} = \frac{\partial f}{\partial z}$$

Plugging this back in our condition on f shows that in fact $\mathrm{d}f = 0$, but then $f = 0$ cannot be a surface!

7 Connections but the Bundles are on Crack

7.1 Connections on Tangent Bundles

Suppose $\mathcal{A}$ is a connection on $E = TX$ over a manifold X. Given local coordinates $x^1, \dots, x^n$ on X, we get a trivialisation of E by $\partial_{x^1}, \dots, \partial_{x^n}$, which we'll start to refer to as the coordinate trivialisation. We typically write the induced connection 1-form as $\Gamma^i{}_{jk} \, \mathrm{d}x^k$, where i, j are the matrix indices on $\mathfrak{gl}(n, \mathbb{R})$. So for a vector field v, we have $(\mathrm{d}^{\mathcal{A}} v)^i = \mathrm{d}v^i + \Gamma^i{}_{jk} v^j \, \mathrm{d}x^k$.
Note that these $\Gamma^i{}_{jk}$ do not tranform like a tensor of type $(1, 2)$. But we do know that the space of connections on E is an affine space for $\Omega^1(\operatorname{End} E) = \Gamma(E \otimes E^\vee \otimes TX) = \Gamma(TX \otimes T^*X \otimes T^*X)$, the space of tensors of type $(1, 2)$.

Definition 7.1. The solder form Θ is an E-valued 1-form that corresponds to the fibrewise identity map under the identification $E \otimes T^*X = TX \otimes T^*X = \operatorname{End}(TX)$. The torsion of $\mathcal{A}$ is the E-valued 2-form $\mathcal{T} = \mathrm{d}^{\mathcal{A}} \Theta$.

In a coordinate trivialisation, $\Theta = e_i \otimes \mathrm{d}x^i$ where e_i is the i-th standard basis vector, and

$$\mathcal{T} = \mathrm{d}(e_i \otimes \mathrm{d}x^i) + A_\alpha \wedge (e_i \otimes \mathrm{d}x^i) = \Gamma^i{}_{jk} e_j \otimes \mathrm{d}x^k \wedge \mathrm{d}x^i = A_\alpha(e_i) \wedge \mathrm{d}x^i$$

Definition 7.2. $\mathcal{A}$ is torsion-free (or symmetric) if $\mathcal{T} = 0$, i.e. if $\Gamma^i{}_{jk} = \Gamma^i{}_{kj}$.

Proposition 7.1 ((First) Bianchi Identity)**.** $\mathrm{d}^{\mathcal{A}}\mathcal{T} = F \wedge \Theta$.

Proof. $\mathrm{d}^{\mathcal{A}}\mathcal{T} = (\mathrm{d}^{\mathcal{A}})^2\Theta = F \wedge \Theta$. □

Definition 7.3. A curve γ in X is geodesic (with respect to $\mathcal{A}$) if $\dot{\gamma}$ is covariantly constant (as a section of γ^*TX).

This is equivalent to the geodesic equation $(\ddot{\gamma})^i + \Gamma^i{}_{jk}(\dot{\gamma})^j(\dot{\gamma})^k = 0$.

Remark. 1. A connection on TX induced a connection on T^*X and on all bundles of tensors and forms. So we only need one connection on TX in order to differentiate tensors and forms. But if we had taken the covariant derivative of Θ as a tensor of type $(1,1)$ (instead of as a TX-valued 1-form), we would automatically get 0.
2. The curvature of $\mathcal{A}$ is an $\mathrm{End}(E)$-valued 2-form, which we can view as a tensor $F^i{}_{jkl}$ of type $(1,3)$ which is antisymmetric in k, l.
3. Often $\mathrm{d}^{\mathcal{A}}$ is denoted ∇, and the contraction of $\mathrm{d}^{\mathcal{A}}$ with a vector or vector field v is written ∇_v.

7.2 Orthogonal Vector Bundles

Fix a vector bundle $E \to B$ of rank k.

Definition 7.4. An inner product on E is a section of $(E^\vee)^{\oplus 2}$ which is fibrewise symmetric and positive definite as a bilinear form.

Lemma 7.2. *E admits an inner product.*

Proof. Cover E with trivialisations Φ_α over U_α. Define an inner product g_α on $E|_{U_\alpha}$ by taking the standard inner product on $\mathbb{R}^k$ under Φ_α. Pick a partition of unity ρ_α subordinate to U_α and glue 'em: $g = \sum_\alpha \rho_\alpha g_\alpha$. □

Definition 7.5. A vector bundle equipped with an inner product is called an orthogonal vector bundle. A trivialisation of it is orthogonal if the inner product becomes the standard inner product under them.

Remark. Transition functions of orthogonal vector bundles then necessarily have to land in $\mathrm{O}(k)$.

From now on fix an orthogonal vector bundle $E \to B$ with inner product g.

Lemma 7.3. *E can be covered by orthogonal trivialisations.*

Proof. Good ol' Gram-Schmidt. □

Definition 7.6. A connection $\mathcal{A}$ on E is orthogonal if g is covariantly constant with respect to the induced connection on $(E^\vee)^{\oplus 2}$.

Lemma 7.4. *E admits an orthogonal connection. Moreover, the space of orthogonal connections on E is an affine space for $\Omega^1(\mathfrak{o}(E))$-valued 1-forms, where $\mathfrak{o}(E) \leq \mathrm{End}(E)$ is the bundle of skew-adjoint endomorphisms on E.*

Proof. Example sheet. □

Lemma 7.5. *If $\mathcal{A}$ is an orthogonal connection on (E, g), the its curvature is an $\mathfrak{o}(E)$-valued 2-form.*

Example 7.1. Example sheet.

8 Riemannian Geometry

8.1 Riemannian Metric

Definition 8.1. A (Riemannian) metric on X is an inner product on TX. A Riemannian manifold is a pair (X, g) where X is a manifold and g is a Riemannian metric on X.

Since every vector bundle admits an inner product, every manifold admits a Riemannian metric. Given a Riemannian metric g_{ij}, we write g^{ij} for the dual metric on T^*X. This satisfies $g^{ij} = g^{ji}$ and $g^{ij}g_{jk} = \delta^i{}_k$. We denote contractions with g_{ij} or g^{ij} by raising and lowering indices, respectively. For example, $g_{il}T^{ij}{}_k$ is denoted as $T_l{}^j{}_k$ and $g^{ik}S_{ij} = S^k{}_j$ with hopefully won't cause that much confusion. A section $T^i{}_j$ of $\mathrm{End}(TX)$ lies in $\mathfrak{o}(TX)$ iff $T^i{}_j g_{ik} = -T^i{}_k g_{ji}$. When writing coordinate expressions, we write $\mathrm{d}x^i\,\mathrm{d}x^j$ to mean $(1/2)(\mathrm{d}x^i \otimes \mathrm{d}x^j + \mathrm{d}x^j \otimes \mathrm{d}x^i)$.

Example 8.1. On $\mathbb{R}^n$, we have the standard Riemannian metric (the "Euclidean metric") $g_{\mathrm{Euc}} = \sum_i \mathrm{d}x^i\,\mathrm{d}x^i = \sum_i (\mathrm{d}x^i)^2$

8.2 The Levi-Civita Connection

Fix a Riemannian manifold (X, g).

Theorem 8.1 (Fundamental Theorem of Riemannian Geometry)**.** *There is a unique torsion-free orthogonal connection on TX.*

Proof. We'll prove a more general statement: The map from the set of orthogonal connections to $\Omega^2(TX)$ whicih sends a connection to its torsion is a bijection. Fix an orthogonal connection $\mathcal{A}_0$. Any other orthogonal connection $\mathcal{A}$ can be written as $\mathcal{A}_0 + \Delta$ for some $\mathfrak{o}(E)$-valued 1-form Δ. So it suffices to show that the map $\Omega^1(\mathfrak{o}(TX)) \to \Omega^2(TX), \Delta \to \mathcal{T}_{\mathcal{A}_0+\Delta} - \mathcal{T}_{\mathcal{A}_0}$ is a bijection. Now this is a linear map. Also, suppose $\mathcal{A}_0$ is locally given by Γ^i_{jk}, then

$$(\mathcal{T}_{\mathcal{A}_0+\Delta} - \mathcal{T}_{\mathcal{A}_0})^i{}_{jk} = (\Gamma + \Delta)^i{}_{kj} - (\Gamma + \Delta)^i{}_{jk} - (\Gamma^i{}_{kj} - \Gamma^i{}_{jk})$$

So Δ is sent to $\Delta \wedge \Theta$ (which as you'll recall is locally given by $\Delta^i{}_{kj} - \Delta^i{}_{jk}$).
Suppose this linear map is induced by the bundle morphism $F : \mathfrak{o}(TX) \otimes T^*X \to TX \otimes \bigwedge^2 T^*X$ given fibrewise by taking the wedge product with Θ. If this were an isomorphism then we are happy, and we can check this fibrewise.
Note that both bundles have rank $n^2(n-1)/2$ since $\mathfrak{o}(TX) \otimes T^*X = \{\Delta^i{}_{jk} :$

$\Delta_{ijk} = -\Delta_{jik}\}$ and $TX \otimes \bigwedge^2 T^*X = \{T^i{}_{jk} : T^i{}_{jk} = -T^i{}_{kj}\}$. So it suffices to show that $\Delta \mapsto \Delta \wedge \Theta$ is injective, i.e. that if $\Delta^i{}_{jk}$ satisfies $\Delta_{ijk} = -\Delta_{jik}$ and $\Delta^i{}_{jk} = \Delta^i{}_{kj}$, then $\Delta = 0$. But this is clear: $\Delta_{ijk} = -\Delta_{jik} = -\Delta_{jki} = \Delta_{kji} = \Delta_{kij} = -\Delta_{ikj} = -\Delta_{ijk}$. □

Definition 8.2. The unique connection predicted in the preceding theorem is known as the Levi-Civita connection of (X, g). Its components $\Gamma^i{}_{jk}$ are called Christoffel symbols.

Explicitly,

$$\Gamma_{ijk} = \frac{1}{2}\left(\partial_j g_{ik} + \partial_k g_{ji} - \partial_i g_{jk}\right)$$

where $\partial_j = \partial/\partial x^j$.

Example 8.2. The Levi-Civita connection on $(\mathbb{R}^N, g_{\text{Euc}})$ is the trivial connection.

Proposition 8.2. *Suppose $\iota : X \to \mathbb{R}^N$ is an embedding. We know that:*
(a) X inherits a metric $\iota^* g_{\text{Euc}}$ from the Euclidean metric on $\mathbb{R}^N$, hence a Levi-Civita connection.
(b) TX carries a connection given by the orthogonal projection from $\iota^* T\mathbb{R}^N$.
Tthese two connections coincide.

Proof. Example sheet. □

8.3 The Riemann Tensor

Fix a Riemannian manifold (X, g).

Definition 8.3. The curvature of the Levi-Civita connection ∇ is the Riemann tensor $R^i{}_{jkl}$, which is an $\mathfrak{o}(TX)$-valued 2-form, viewed as a tensor of type $(1, 3)$.

It's clear that $R^i{}_{jkl} = -R^i_{jlk}$. We also have $R_{ijkl} = -R_{jikl}$ since it takes values in $\mathfrak{o}(TX)$. It also satisfies the first Bianchi identity $R \wedge \Theta = \mathrm{d}^\nabla \mathcal{T} = 0$, i.e. $R^i{}_{jkl} + R^i{}_{klj} + R^i{}_{ljk} = 0$. Then second Bianchi identity (example sheet) on the other hand gives $\mathrm{d}^\nabla R = 0$.

8.4 Hodge Theory

Let (X, g) be an oriented Riemannian manifold of dimension n. The dual metric g^{ij} gives an inner product on T^*X, hence induces an inner product on all exterior powers of it. Explicitly, if $\alpha^1, \ldots, \alpha^n$ is a local fibrewise basis of 1-forms, then things of the form $\alpha^I = \alpha^{i_1} \wedge \cdots \wedge \alpha^n \in \bigwedge^{i_p}$ are a fibrewise orthonormal basis of p-forms. There is a distinguished unit volume form ω parallel to the orientation. Given a p-form β, there is a unique $(n-p)$-form $*\beta$ such that for any p-forms α we have $\alpha \wedge *\beta = \langle \alpha, \beta \rangle \omega$. Indeed, $*\alpha^I = \pm \alpha^J$ where $J = \{1, \ldots, n\} \setminus I$, and the sign is given by the orientation.

Definition 8.4. The map $* : \Omega^p(X) \to \Omega^{n-p}(X)$ is the Hodge star operator.

By considering its action on the α^I, we see that it's a fibrewise isometry and that $*^2 = (-1)^{p(n-p)} \, \mathrm{id}_{\Omega^p(X)}$.

Example 8.3. Take $\mathbb{R}^3$ with the standard metric and orientation. Then $\omega = \mathrm{d}x^1 \wedge \mathrm{d}x^2 \wedge \mathrm{d}x^3$, so $*\mathrm{d}x^1 = \mathrm{d}x^2 \wedge \mathrm{d}x^3, *(\mathrm{d}x^2 \wedge \mathrm{d}x^3) = \mathrm{d}x^1$ (and similar statements by cyclically permuting the indices).

Now assume X is compact. Then we can define an inner product on $\Omega^p(X)$ by

$$\langle \alpha, \beta \rangle_X = \int_X \langle \alpha, \beta \rangle \omega = \int_X \alpha \wedge *\beta$$

For a $(p-1)$-form α and a p-form β, we have

$$\begin{aligned} \langle \mathrm{d}\alpha, \beta \rangle_X &= \int_X \mathrm{d}\alpha \wedge *\beta = \int_X (\mathrm{d}(\alpha \wedge *\beta) - (-1)^{p-1} \alpha \wedge \mathrm{d} * \beta) \\ &= \int_X (-1)^p \alpha \wedge \mathrm{d} * \beta = \langle \alpha, (-1)^p *^{-1} \mathrm{d} * \beta \rangle \end{aligned}$$

So the operator $\delta = (-1)^p *^{-1} \mathrm{d}*$ is adjoint to d.

Definition 8.5. $\delta : \Omega^p(X) \to \Omega^{p-1}(X)$ is called the codifferential.

Remark. 1. $\delta = (-1)^{np+n+1} * \mathrm{d}*$ is so preferred.
2. $\delta^2 = - *^{-1} \mathrm{d} * *^{-1} \mathrm{d}* = - *^{-1} \mathrm{d}^2 * = 0$.

Definition 8.6. A form α is coclosed is $\delta\alpha = 0$, and coexact if there is some β with $\alpha = \delta\beta$.

Definition 8.7. The Laplace-Beltrami operator $\Delta : \Omega^p(X) \to \Omega^p(X)$ is the operator $\Delta = \mathrm{d}\delta + \delta\mathrm{d}$.

If we view d, δ as $\Omega^\bullet(X) \to \Omega^\bullet(X)$, one might just write $\Delta = (\mathrm{d} + \delta)^2$.

Definition 8.8. A form α is harmonic if $\Delta\alpha = 0$. Denote by $\mathcal{H}^p(X)$ the space of harmonic p-forms.

It turns out that α is harmonic iff it is both closed and coclosed.

Theorem 8.3. *The map $\mathcal{H}^p(X) \to H^p_{\mathrm{dR}}(X), \alpha \mapsto [\alpha]$ is an isomorphism.*

So every cohomology class has a unique harmonic representative.
The intuitive reason why this is true is the following: $\mathcal{H}^p(X) = \ker \Delta = \ker(\mathrm{d}) \cap (\ker \delta) = (\ker d) \cap (\operatorname{Im} d)^\perp \cong \ker d / \operatorname{Im} d = H^p_{\mathrm{dR}}(X)$. Alas, we don't know if our spaces are finite-dimensional and this course doesn't assume functional analysis as a prerequisite. Nonetheless, we're still gonna blackbox the following fact from analysis to assist the proof.

Theorem 8.4 (Hodge Decomposition)**.** *The space $\mathcal{H}^p(X)$ is finite-dimensional and we have orthogonal decompositions*

$$\Omega^p(X) = \mathcal{H}^p(X) \oplus \mathrm{d}\delta\Omega^p(X) \oplus \delta\mathrm{d}\Omega^p(X) = \mathcal{H}^p(X) \oplus \mathrm{d}\Omega^p(X) \oplus \delta\Omega^p(X)$$

Proof. Analysis. □

Proof of Theorem 8.3. It suffices to show that $\ker \mathrm{d} = \mathcal{H}^p(X) \oplus \mathrm{d}\Omega^{p-1}(X)$. Since harmonic forms are exact hence closed, we know $\mathcal{H}^p(X) \oplus \mathrm{d}\Omega^{p-1}(X) \leq \ker \mathrm{d}$. On the other hand, the preceding theorem gives the expression $\mathcal{H}^p(X) \oplus \mathrm{d}\Omega^{p-1}(X) = (\operatorname{Im} \delta)^\perp$, so it suffices to show that $\langle \ker \delta, \operatorname{Im} \delta \rangle = 0$. But for any $\alpha \in \ker \mathrm{d}$ and β we have $\langle \alpha, \delta\beta \rangle = \langle \mathrm{d}\alpha, \beta \rangle = 0$. □

9 Lie Groups and Principal Bundles

9.1 Lie Groups and Lie Algebras

With $1+\epsilon$ hours of lectures left, let's tell some more lies (yes, I did not have to).

Definition 9.1. A Lie group is a manifold G equipped with a group structure such that multiplication $m : G \times G \to G$ and inversion $i : G \to G$ are smooth. An embedded Lie subgroup of a Lie group G is a submanifold of G which is also a subgroup of G. It's clear that embedded Lie subgroups are Lie groups.

Example 9.1. $\mathrm{GL}(n, \mathbb{R})$ is a Lie group and $\mathrm{SL}(n, \mathbb{R}), \mathrm{O}(n), \mathrm{SO}(n)$ are embedded Lie subgroups. Similarly, $\mathrm{SL}(n, \mathbb{C}), \mathrm{U}(n), \mathrm{SU}(n)$ are embedded Lie subgroups of $\mathrm{GL}(n, \mathbb{C})$.

For each $g \in G$, we get diffeomorphisms $L_g, R_g, C_g : G \to G$ given by $x \mapsto gx, x \mapsto xg, x \mapsto gxg^{-1}$, respectively. Guess what, we call them left-translation, right-translation and conjugation by g, respectively. This goes without saying but their inverses are given by $L_{g^{-1}}, R_{g^{-1}}, C_{g^{-1}}$.

Definition 9.2. A tensor T on G is left (resp. right, conjugation) invariant iff $(L_g)_*T = T$ (resp. $(R_g)_*T = T$, $(C_g)_*T = T$) for all $g \in G$. It is bivariant if it's left and right invariant.

Lemma 9.1. *For any $h \in G$, the map from the space of left (resp. right) invariant tensor fields of type (p, q) to the space of tensors of type (p, q) at h given by $T \mapsto T_h$ is an isomorphism.*

Proof. For any left-invariant T, we have $T_g = (L_{gh^{-1}})T_h$ for any $g \in G$. So the map is injective. Conversely, given a tensor T_h at h, $T_g = (L_{gh^{-1}})T_h$ defines a tensor field whose value at h is T_h, hence the map is surjective. □

Definition 9.3. The Lie algebra $\mathfrak{g}$ of a Lie group G is T_eG.

Example 9.2. $\mathfrak{gl}(n, \mathbb{R})$ is the Lie algebra of $\mathrm{GL}(n, \mathbb{R})$. The Lie algbera of $\mathrm{SL}(n, \mathbb{R})$ is $\mathfrak{sl}(n, \mathbb{R}) = \{A \in \mathfrak{gl}(n, \mathbb{R}) : \operatorname{Tr} A = 0\}$ (noting $D_I \det = (1/2) \operatorname{Tr}$). Similarly, the Lie algebra of $\mathrm{O}(n)$ is $\mathfrak{o}(n) = \{A \in \mathfrak{gl}(n, \mathbb{R}) : A^\top + A = 0\}$.

We had previously met Lie algebras when we defined a Lie bracket for vector fields. This new definition doesn't stray far from that. For any $\xi \in \mathfrak{g}$, let ℓ_ξ denote the corresponding left-invariant vector field.

Lemma 9.2. *The Lie bracket of left invariant vector fields is left invariant.*

Proof. $(L_g)_*[v, w] = [(L_g)_*v, (L_g)_*w] = [v, w]$. □

Definition 9.4. The Lie bracket on $\mathfrak{g}$ is defined by $[\xi, \eta] = \zeta$ where ζ is the unique element of $\mathfrak{g}$ such that $[\ell_\xi, \ell_\eta] = \ell_\zeta$.

This Lie bracket is alternating, is bilinear and satisfies the Jacobi identity since the Lie bracket for vector fields does. This makes $\mathfrak{g}$ a Lie algebra in the abstract sense.

9.2 Lie Group Actions

Definition 9.5. A left action of a Lie group G on a manifold X is a group action $G \times X \to X$ which happens to be smooth. Similarly for right actions.

Example 9.3. 1. The action of $\mathrm{GL}(n, \mathbb{R})$ on $\mathbb{R}^n$ is a Lie group action. So is the action of $\mathrm{O}(n)$ on $S^{n-1} \subset \mathbb{R}^n$.
2. The left regular, right regular, and conjugation action of G on itself are also Lie group actions.
3. The adjoint action/representation of G on $\mathfrak{g}$ is $\mathrm{Ad}_g(\xi) = (C_g)_*\xi$.

Definition 9.6. Given a left Lie group action $\sigma : G \times X \to X$, the infinitesimal action of $\xi \in \mathfrak{g}$ on $x \in X$ is $\xi \cdot x = D_{(e,x)}\sigma(\xi, 0) = [\gamma(t)x]$ where γ is any curve representing ξ. Similarly for right actions.

9.3 Principal Bundles

Fix a Lie group G and a base manifold B.

Definition 9.7. A principal G-bundle P over B is a vector bundle on B except trivialisations are $\Phi_\alpha : \pi^{-1}U_\alpha \to U_\alpha \times G$, and on overlaps $\Phi_\beta \circ \Phi_\alpha^{-1}(b, g) = (b, g_{\beta\alpha}(b)g)$ for (necessarily smooth) maps $g_{\beta\alpha} : U_\alpha \cap U_\beta \to G$.

Example 9.4. Given a rank k vector bundle $E \to B$, its frame bundle $F(E) \to B$ is the principal $\mathrm{GL}(k, \mathbb{R})$-bundle with $F(E)_b$ being the set of ordered bases for E_b. Similarly, if E has an inner product, we can construct the orthonormal frame bundle $F_{\mathrm{o}}(E)$ which is a principal $\mathrm{O}(k)$-bundle with $F_{\mathrm{o}}(E)_b$ the set of ordered orthonormal basis of E_b.

Remark. 1. Many definitions about vector bundles get carried over to principal G-bundles (e.g. sections, gluing constructions, etc.). Although now we don't get a zero section anymore.
2. A principal G-bundle P naturally carries a right G-action defined in trivialisations, i.e. $\Phi_\alpha^{-1}(b, x)g = \Phi_\alpha^{-1}(b, xg)$.
3. Sections s over $U \subset B$ correspond to trivialisations Φ over U: Given Φ, we get $s(b) = \Phi^{-1}(b, e)$; given s, we get $\Phi(s(b)g) = (b, g)$.

9.4 Connections

Fix a principal G-bundle $P \to B$. By abuse of notation, we write $R_g : P \to P$ to denote the right action of $g \in G$ on P.

Definition 9.8. A connection $\mathcal{A}$ on P is a $\mathfrak{g}$-valued 1-form on P such that:
(i) $\mathcal{A}(p \cdot \xi) = \xi$ for any $p \in P, \xi \in \mathfrak{g}$.
(ii) $R_g^*\mathcal{A} = \mathrm{Ad}_{g^{-1}} \mathcal{A}$.

Given a local section s_α (or equivalently a trivialisation), the local connection 1-form $\mathcal{A}_\alpha$ is $s_\alpha^*\mathcal{A}$.

Lemma 9.3. *On overlaps we have* $\mathcal{A}_\alpha = \mathrm{Ad}_{g_{\beta\alpha}^{-1}} \mathcal{A}_\beta + (L_{g_{\beta\alpha}^{-1}})_* \mathrm{d}g_{\beta\alpha}$. *Conversely, given a system* $\mathcal{A}_\alpha$ *transforming as such, we get a unique connection out of them.*

Proof. Example sheet. □

Remark. If $P = F(E)$, a connection on P is equivalent to a connection on E.

Definition 9.9. The curvature of $\mathcal{A}$ is a $\mathfrak{g}$-valued 2-form $\mathcal{F}$ on P given by $\mathcal{F} = \mathrm{d}\mathcal{A} + (1/2)[\mathcal{A} \wedge \mathcal{A}]$ where $[(\sum_i \xi_i \otimes \alpha_i) \wedge (\sum_j \eta_j \otimes \beta_j)] = \sum_{i,j} [\xi_i, \eta_j] \otimes \alpha_i \wedge \beta_j$. $\mathcal{A}$ is flat if $\mathcal{F} = 0$.