A very mild introduction to surface group representations and Higgs bundles

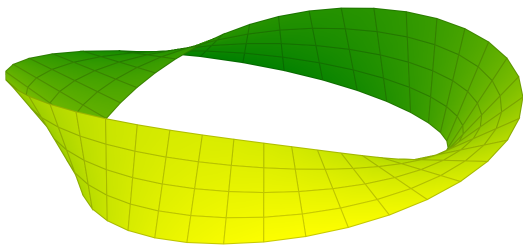

A geometrical and topological introductory promenade

Roberto Rubio
Weizmann Institute of Science

February 19, 2018

Disclaimer

These are lecture notes about geometry and topology leading to the initial study of surface group representations and the definition of Higgs bundle. They reflect the contents of the course Introduction to surface group representations and Higgs bundles, taught at the Weizmann Institute of Science in the first semester of 2017/2018. Note that these notes are better understood with the images from the lectures.

They are meant to be the first iteration of a future set of improved and more complete lecture notes, so the text has not been carefully proofread. If you find any typos, mistakes, or if there is something not sufficiently clear, I will appreciate if you can let me know by email: roberto . rubio @ uab . es

Conventions: we omit sets of indices, which may be countable or uncountable, like in $\{(U_i, c_i)\}$. When taking two arbitrary elements, we use i and j. Some definitions are given inside a paragraph by using **bold type**, or in the case of formulas by the symbol $:=$.

Acknowledgements

The first version of these notes greatly benefited from the interaction with the students and interested listeners of the course Introduction to surface group representations and Higgs bundles at the Weizmann Institute of Science in the first semester of 2017/18. These were Jonathan Bauch, Guy Kapon, Anna Rohova, Yaniv Shahar, Aviv Shalit, Michal Shavit, Ary Shaviv, Raz Slutsky, Stephan Snigerov, Tomer Solberg, Talya Vaknin, Itamar Vigdorovich, Leonid Vishnevsky and Benny Zack. I take this opportunity to thank all of them, and also the Weizmann Institute of Science, for making this course possible.

Contents

1 Preliminaires **1**
- 1.1 Manifolds . . . 1
- 1.2 Bundles . . . 6
- 1.3 Surfaces . . . 15
- 1.4 More on bundles and their sections . . . 23
- 1.5 Connections . . . 29

2 Surface group representations **37**
- 2.1 The set of representations . . . 37
- 2.2 The action of the group . . . 39
- 2.3 A flat bundle for a representation . . . 42
- 2.4 Flat bundles and flat connections . . . 44
- 2.5 A representation for a flat connection . . . 45

3 The definition of Higgs bundle **47**
- 3.1 Some results on complex geometry . . . 47
- 3.2 A funny way to recover a flat connection . . . 49
- 3.3 What if I have a Higgs bundle? . . . 51
- 3.4 The Hitchin–Kobayashi correspondence . . . 52
- 3.5 Where to look next . . . 53

Chapter 1

Preliminaires

1.1 Manifolds

Each mathematical theory has a starting concept. In set theory, it is sets. In group theory, we look at groups. In linear algebra, we focus on vector spaces. In topology, we go nowhere without a topological space. In calculus, we study $\mathbb{R}^n$ in and out. And in differential geometry, we work with differentiable manifolds. This is the first concept we want to talk about.

An important part of our early mathematical life was spent in $\mathbb{R}^n$. We started measuring distances, angles and drawing shapes in $\mathbb{R}^2$. We studied one or multi-variable real functions: continuity, differentiability, integrability... That took years and A differentiable manifold is just a way to assemble different copies of $\mathbb{R}^n$ in possibly an unusual way. For instance, a sphere. Some of the things we knew of $\mathbb{R}^n$ will still work, but others are very different: the angles of a triangle do not add up to 180 degrees.

In order to define differentiable manifolds, it is useful to talk about charts first. Let $n \in \mathbb{N}$. An **n-dimensional chart** on a set M consists of a pair (U, c) where U is a subset of M and c is a bijection between U and an open subset A of $\mathbb{R}^n$, $c : U \to A \subset \mathbb{R}^n$. This chart allows us to already do some things on $U \subset M$. For instance, we can say that a function $f : U \to \mathbb{R}$ is smooth (i.e., of class $\mathcal{C}^\infty$) when $f \circ c^{-1} : A \to \mathbb{R}$ is.

If we have two n-dimensional charts (U_i, c_i), (U_j, c_j) on a set and the intersection $U_{ij} := U_i \cap U_j$ is not empty, the map

$$c_{ij} := c_i \circ c_j^{-1} : c_j(U_{ij}) \to c_i(U_{ij})$$

is called a **change of chart**. This is actually the concept that allows us to do things consistently. As we said, a function $f : U_i \cap U_j \to \mathbb{R}$ is smooth if and only if $f \circ c_i^{-1}$ is, but also if and only if $f \circ c_j^{-1}$ is. Since this must happen to every f, the changes of chart $c_{ij} = c_i \circ c_j^{-1}$ and $c_{ji} = c_j \circ c_i^{-1}$ must be smooth. In this case, or if $U_i \cap U_j = \emptyset$, we say that the charts (U_i, c_i), (U_j, c_j) are **compatible**.

A chart allows us to do calculus on $U \subset M$. With two compatible charts (U_i, c_i), (U_j, c_j) we can do it on $U_i \cup U_j$. In order to do calculus on the whole of M, we need to cover M with compatible chart of the same dimension. This is called an **atlas**: a collection of charts $\{(U_i, c_i)\}$ compatible among them such that $\bigcup U_i = M$. With an atlas on a set M, we have seen that we have a well-defined concept of differentiable function on M.

A set with an atlas is almost a differentiable manifold, and is actually our working definition. As for the formal definition, the thing is that we can give different atlases to the same set, say $\{(U_i, c_i)\}$, $\{(U'_j, c'_j)\}$, which could lead to exactly the same concept of smooth functions. By the argument above, this will be the case exactly when all the charts are compatible among them. Since being compatible is an equivalence relation[1], we can, given an atlas, take all the compatible charts and form a **maximal atlas**. It should not come to a surprise that we refer to the integer n as the **dimension** of M. One can define a differentiable manifold as a set with a maximal atlas.

Despite the maximal atlas, this is a very minimal definition of a manifold, and perhaps not the one you may have seen before. Many definitions start with a topological space. This does not mean that we are not interested in having a topology[2] for the set M. The point is that we automatically get one by taking the usual topology in $\mathbb{R}^n$, coming from the balls of the (usual) Euclidean metric, as this is the topology used to do calculus in $\mathbb{R}^n$. The open sets in M are generated by the preimage by our charts of open sets (i.e., $c_i^{-1}(A_i)$ with $A_i \subset c_i(U_i)$ an open set).

Fine print 1.1. We have defined our manifolds to be finite-dimensional. This simple choice of topology gets much more complicated when dealing with infinite-dimensional manifolds... is there a "usual" topology in $\mathbb{R}^\infty$?

A couple of properties that are usually included in the definition are Haussdorf and have a countable basis of open sets (known as second-countable). We will actually come across an example of a non-Hausdorff manifold in this course, but as it will be an exception, you may probably feel homier with the following definition.

[1] This means reflexive, a chart is compatible with itself; symmetric, if a chart U is compatible with a chart V, then V is compatible with U; and transitive, if U is compatible with V and V is compatible with W, then W is compatible with U.

[2] A topology is a collection of subsets of M including the empty set and the total set that is closed under union or finite intersection of sets. This collection consists of the so-called open sets.

Definition 1.1. A **differentiable manifold** is a set with a maximal atlas such that the induced topology is Haussdorf and second-countable.

Fine print 1.2. There are some nice theorems that only work for Haussdorf and second-countable, like the fact that manifolds are metrizable spaces or that any n-dimensional manifold can be embedded, or seen as a "differentiable subspace" (formally, embedded submanifold) of $\mathbb{R}^{2n+1}$.

We have talked so far about smoothness, as we were thinking about calculus. If we were only concerned about continuos maps and topology, we do not need the changes of chart to be smooth, but just continuous. A set with a maximal atlas with continuous changes of chart is a **topological manifold**. On the other hand, we may want to talk about holomorphicity, and in this case our charts need to map onto open sets of the standard complex Euclidean space $\mathbb{C}^n$. Let us call them complex charts, and use complex atlas for an atlas made of complex charts. A set with a maximal complex atlas with holomorphic changes of chart is a **complex manifold**. We are focusing on differentiable manifolds, but we will talk about topological manifolds and definitely work with complex manifolds.

Everything so far has been very theoretical, so let us show several examples.

Example 1.2. The following are examples of differentiable manifolds:

1. The very $\mathbb{R}^n$, or generally any finite-dimensional vector space, is a differentiable manifold. We can take only one chart, $\mathrm{id} : \mathbb{R}^n \to \mathbb{R}^n$, or we could take many charts. For any collection of open sets $\{A_i\}$ covering $\mathbb{R}^n$, consider the collection of charts $\{(A_i, \mathrm{id}_{|A_i})\}$. All the transtion functions are identity maps and hence differentiable.

2. The circle $\mathrm{S}^1 = \{e^{i\theta} \mid \theta \in [0, 2\pi]\}$ is a 1-dimensional manifold. Let us consider the two subsets

$$\mathrm{S}^1_w = \{e^{i\theta} \mid \theta \in (0, 2\pi)\}, \qquad \mathrm{S}^1_e = \{e^{i\theta} \mid \theta \in (-\pi, \pi)\}.$$

We have charts $c_w : \mathrm{S}^1_w \to (0, 2\pi)$, given by $c_w : e^{i\theta} \mapsto \theta$, and $c_e : \mathrm{S}^1_e \to (-\pi, \pi_1)$, given by $c_e : e^{i\theta} \mapsto \theta$. We have $c_w(\mathrm{S}^1_w \cap \mathrm{S}^1_e) = (0, \pi) \cup (\pi, 2\pi)$, and $c_e(\mathrm{S}^1_w \cap \mathrm{S}^1_e) = (-\pi, \pi) \cup (0, \pi)$, and the change of chart $c_e \circ c_w^{-1}$ is given by the identity on $(0, \pi)$ and $\mathrm{Id} - 2\pi$ on $(\pi, 2\pi)$.

Alternatively, one can give charts by considering the half circles

$$\mathrm{S}^1_t = \{e^{i\theta} \mid \theta \in (0, \pi)\}, \qquad \mathrm{S}^1_l = \{e^{i\theta} \mid \theta \in (\frac{\pi}{2}, \frac{3\pi}{2})\},$$
$$\mathrm{S}^1_b = \{e^{i\theta} \mid \theta \in (\pi, 2\pi)\}, \qquad \mathrm{S}^1_r = \{e^{i\theta} \mid \theta \in (-\frac{\pi}{2}, \frac{3\pi}{2})\},$$

which project to the x-axis (S^1_t and S^1_b) or y-axis (S^1_l and S^1_r). The changes of chart over non-empty intersections are maps defined in quarters of circle sending $\cos\theta \mapsto \sin\theta$ or $\sin\theta \mapsto \cos\theta$. These two maps are plus or minus the map $x \mapsto \sqrt{1-x^2}$, which is smooth because x is not equal to 1 (those are exactly the points missing in the interesection).

3. The sphere is also a manifold as can be seen with spherical angular coordinates or projections on six hemispheres.

4. Any open subset of a manifold is again a manifold. For instance, the interval $I = (-1, 1)$ is a manifold.

5. When we act by a group on a manifold and we look at the quotient, or space of orbits, things are much more subtle. This is not an easy subject at all, but it is worth give some intuition about it. For instance, $\mathbb{Z}_2 = \{\pm 1\}$ acts on both $\mathbb{R}^2$ and S^2 by $(x, y) \mapsto (x, -y)$ and $(x, y, z) \mapsto (-x, -y, -z)$. We then have that $\mathbb{R}^2/\mathbb{Z}^2$ is essentially the upper-half plane
$$\{(x, y) \in \mathbb{R}^2 \mid y \geq 0\},$$
which is not a manifold, as it has a border: the points where $y = 0$. On the other hand, for $\mathrm{S}^2/\mathbb{Z}^2$, we can take an atlas of S^2 such that the open sets satisfy $U_i \cap (-1) \cdot U_i = \emptyset$. This atlas descends to an atlas of $\mathrm{S}^2/\mathbb{Z}^2$ and we get a manifold structures. This manifold describes the possible directions of lines passing through the origin of $\mathbb{R}^3$ and is called the **projective plane**.

6. Given any two manifolds M, N, their Cartesian product $M \times N$ is again a manifold.

7. The cylinder $\mathrm{S}^1 \times (-1, 1)$ can be given charts $(c_w, \mathrm{Id}) : \mathrm{S}^1_w \times (-1, 1) \to (0, 2\pi) \times (-1, 1)$ and $(c_e, \mathrm{Id}) : \mathrm{S}^1_e \times (-1, 1) \to (0, 2\pi) \times (-1, 1)$. The differentiability of the changes of chart follows from the one in S^1 for the charts c_w, c_e, as we have
$$(c_e, \mathrm{Id}) \circ (c_w, \mathrm{Id})^{-1} = (c_e \circ c_w^{-1}, \mathrm{Id}),$$
and analogously for $(c_w, \mathrm{Id}) \circ (c_e, \mathrm{Id})^{-1}$.

8. The Möbius band can be understood as a segment doing half a turn along a circle. If we describe the circle by $\{(\cos\theta, \sin\theta, 0)\}$, the turning segment is on the plane generated by $\{(\cos\theta, \sin\theta, 0), (0, 0, 1)\}$. As

we want it to do half a turn, we combine these two vectors with the coefficients $\cos\frac{\theta}{2}, \sin\frac{\theta}{2}$ and add it to the point in the circle:

$$m(\theta, r) := (\cos\theta, \sin\theta, 0) + r\cos\frac{\theta}{2}(\cos\theta, \sin\theta, 0) + r\sin\frac{\theta}{2}(0, 0, 1).$$

The Möbius band is then given by

$$\{m(\theta, r) : \theta \in [0, 2\pi], r \in (-1, 1)\}.$$

Finding charts for the Möbius band is easy if we consider a projection to S^1 given by

$$\pi : m(\theta, r) \mapsto e^{i\theta} \in S^1.$$

Define the charts $\pi^{-1}(S^1_w) \to (0, 2\pi)$, given by $m(\theta, r) \mapsto (\theta, r)$ and $\pi^{-1}(S^1_e) \to (-\pi, \pi)$ given by $m(\theta, r) \mapsto (\theta, r)$. When looking at the change of chart, one has that

$$(c_e, \mathrm{Id}) \circ (c_w, \mathrm{Id})^{-1} = \begin{cases} (c_e \circ c_w^{-1}, \mathrm{Id}) & \text{on } (0, \pi) \times (-1, 1), \\ (c_e \circ c_w^{-1}, -\mathrm{Id}) & \text{on } (-\pi, 0) \times (-1, 1), \end{cases}$$

as $m(\theta, r) = m(\theta - 2\pi, -r)$.

9. Any discrete collection of points is naturally a 0-dimensional manifold.

Fine print 1.3. The line $\mathbb{R}$ has the trivial chart $\mathrm{id} : \mathbb{R} \to \mathbb{R}$, but it could also have the chart $x \mapsto x^3$. These two are not compatible as the map $x^{\frac{1}{3}}$ is not smooth at 0, so we get two differentiable structures, the usual real line and the, say, cubic real line.

We are used to look at manifolds inside Euclidean space, mostly sitting into $\mathbb{R}^3$. However, we can define the cylinder and the Möbius band very easily in $S^1 \times D \subset \mathbb{R}^4$, where D is the unit disk:

$$\begin{aligned} &\{(e^{i\theta}, r) \mid \theta \in [0, 2\pi], r \in (-1, 1)\}, \\ &\{(e^{i\theta}, re^{i\frac{\theta}{2}}) \mid \theta \in [0, 2\pi], r \in (-1, 1)\}. \end{aligned} \tag{1.1}$$

The torus is also easily described in $\mathbb{R}^4$.

$$S^1 \times S^1 = \{(e^{i\theta}, e^{i\sigma}) \mid \theta, \sigma \in [0, 2\pi]\}.$$

Given a manifold, we know when to say that a map $M \to \mathbb{R}$ is smooth. A map $f : M \to N$ between a manifold M with atlas $\{(U_i, c_i)\}$ and another manifold N with atlas $\{(V_i, b_i)\}$ is said to be a **smooth map** if $c_i^{-1} \circ f \circ b_j$ is smooth for any i, j. In order to check the smoothness we can, and practically must, take any atlas but not a maximal one. If f is smooth with respect to a particular choice of atlases, by the compatibility, it will be for a maximal one.

If a smooth map $f : M \to N$ has a smooth inverse, we call it a **diffeomorphism**. In this case we say that M and N are diffeomorphic and we write $M \cong N$.

Example 1.3. The manifolds $M \times N$ and $N \times M$ are diffeomorphic. The projection map $M \times N \to M$ is smooth but not a diffeomorphism unless N is just a point.

Fine print 1.4. Diffeomorphism between the two real lines. The map $t \mapsto t^3$ from the cubic line to the usual line is a diffeomorphisms, as it is the identity when seen through the charts.

There are many things one usually studies on manifolds before getting to the concept we shall introduce next. But this is not a course on differentiable manifolds and we want to get to the point.

1.2 Bundles

Definition 1.4. A **fibre bundle** over a manifold M with typical fibre a manifold F is a manifold E together with a differentiable and surjective map $\pi : E \to M$ such that there exists an open cover $\{U_i\}$ of M and a diffeomorphism $\varphi_i : \pi^{-1}(U_i) \to U_i \times F$ for each U_i satisfying $\mathrm{pr}_{U_i} \circ \varphi_i = \pi$.

Note that the open cover $\{U_i\}$ and the diffeomorphisms $\{\varphi_i\}$ are not fixed. The definition says that M is locally trivial for some choice of open cover $\{U_i\}$. A particular choice of open sets and diffeomorphisms $\{(U_i, \varphi_i)\}$ is called a **trivialization**. We will refer to a single pair (U_i, φ_i) as a **local trivialization**. For any point $x \in M$, we define $E_x := \pi^{-1}(x)$, the **fibre at x**, which is diffeomorphic to F by using any local trivialization, but not in a canonical way, as we can change the trivialization.

A trivial example of a fibre bundle is a product manifold $M \times F$, where we cover M with the total open set M. In particular, the cylinder is a fibre bundle, say, with fibre I over S^1. The Möbius band is also a fibre bundle, similarly to the charts of Example 1.2, we have now local trivializations

$$\begin{aligned} \pi^{-1}(\mathrm{S}^1_w) &\to \mathrm{S}^1_w \times I & \qquad \pi^{-1}(\mathrm{S}^1_e) &\to \mathrm{S}^1_e \times I \\ m(\theta, r) &\mapsto (e^{i\theta}, r) & m(\theta, r) &\mapsto (e^{i\theta}, r). \end{aligned}$$

Note that we are not using any charts for S^1_w or S^1_e.

If we cut the cylinder or the Möbius band into two pieces, transversally to S^1, we get exactly the same, two rectangular pieces. They way these

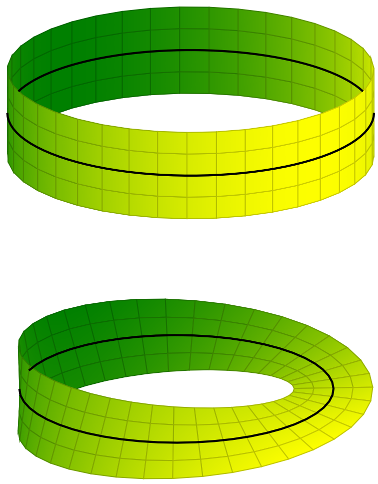

Figure 1.1: The cylinder and the Möbius band

pieces are glued to recover the original manifolds is what determines their global bundle structure. This simple idea is actually the way we think about bundles.

Given two local trivializations $(U_i, \varphi_i), (U_j, \varphi_j)$, consider the map

$$U_{ij} \times F \xrightarrow{\varphi_j^{-1}} \pi^{-1}(U_{ij}) \xrightarrow{\varphi_i} U_{ij} \times F.$$

By the compatibility of the charts with the projection π, we have

$$\mathrm{pr}_{U_i} \circ (\varphi_i \circ \varphi_j^{-1}) = \pi \circ \varphi_j^{-1} = \mathrm{pr}_{U_j}.$$

Since we are in the intersection U_{ij}, this means that, for $x \in U_{ij}$ and $r \in F$,

$$(\varphi_i \circ \varphi_j^{-1})(x, r) = (x, g_{ij}(x, r))$$

for some smooth map $g_{ij} : U_{ij} \times F \to F$, which we call the **transition maps** from U_j to U_i. We will sometimes write $g_{ij}(x, -) \in \mathrm{Diff}(F)$, and eventually omit the point x.

Example 1.5. Since we did a nice choice of charts, we have already computed the transition functions for the cylinder and the Möbius band. For the cylinder we have

$$g_{ew}(x, -) = \begin{cases} \mathrm{Id} & \text{for } x \in \mathrm{S}^1_t, \\ \mathrm{Id} & \text{for } x \in \mathrm{S}^1_b. \end{cases}$$

whereas for the Möbius band,

$$g_{ew}(x,-) = \begin{cases} \mathrm{Id} & \text{for } x \in \mathrm{S}^1_t, \\ -\mathrm{Id} & \text{for } x \in \mathrm{S}^1_b. \end{cases}$$

As a general principle, any definition we do for bundles will be the same as for manifolds, but respecting the fibres. For instance, we say that two bundles E, E' over the same base are **diffeomorphic (as bundles)** if there exists a diffeomorphism $f : E \to E'$, as manifolds, such that $\pi' \circ f = \pi$.

$$\begin{tikzcd} E \arrow[rr, "f"] \arrow[dr, "\pi"'] & & E' \arrow[dl, "\pi'"] \\ & M & \end{tikzcd}$$

Note that the transition maps trivially satisfy, for any point $x \in M$,

i) $g_{ii} = \mathrm{Id}$,

ii) $g_{ji} = g_{ij}^{-1}$.

Moreover, if there are triple intersections $U_{ijk} := U_i \cap U_j \cap U_k$, we must have, for any $x \in M$,

iii) $g_{ij} \circ g_{jk} \circ g_{ki} = \mathrm{Id}$.

We are going to see that all the information about the bundle, up to diffeomorphism, is contained in $\{(U_{ij}, g_{ij})\}$. This is why we make this definition abstract: for an open cover $\{U_i\}$ of a manifold M, a collection

$$\{(U_{ij}, g_{ij} : U_{ij} \times F \to F)\}$$

satisfying the the properties i), ii), iii) above is called a **cocycle of transition maps**.

We do prove now, at the same time, that any cocycle determines a bundle, and that the cocycle coming from a bundle recovers a diffeomorphic bundle to the original one. Given a cocycle associated to an open cover $\{U_i\}$, consider all the sets $U_i \times F$ and glue them following g_{ij}: so $(x, v) \sim (x, w)$ when $(x, v) \in U_i \times F$, $(x, w) \in U_j \times F$ and $v = g_{ij}(x, w)$. This means doing a quotient

$$E' = \coprod_i U_i \times F / \sim .$$

The set E' is given charts around $[(x, v)]$ by taking charts around $x \in U_i$ and $v \in F$. A different choice gives compatible charts by the compatibility of the charts of M and F and the fact that the maps g_{ij} are smooth. One has to check that the resulting topology is Hausdorff and second countable. In order to separate two points $[(x, v)]$ and $[(y, w)]$, with $x \neq y$, we take two separating neighbourhoods of $x, y \in M$, say, N_x and N_y, and consider $\pi^{-1}(N_x)$ and $\pi^{-1}(N_y)$. If $x = y$, we choose a trivialization over an open set U_i, take two separating neighbourhoods of $v, w \in M$, say, N_v and N_w, and consider $U_i \times N_v$ and $U_i \times N_w$. For second countable, as M is second countable, we can take the open cover $\{U_i\}$ to be second countable (by considering as open sets the basic open sets contained in some U_i). Since for each U_i we are taking products of one of the countable charts of U_i with one of the countable charts of F, we have a countable number of charts, which bring the second countable topology of $\mathbb{R}^n$ to a second countable topology in M.

The projection $\pi : E' \to M$, given by $\pi([(x, v)]) = x$ is well defined, differentiable and surjective. For the local triviality it is immediate that $\pi^{-1}(U_i) \cong U_i \times F$. Moreover, if the cocycle $\{(U_{ij}, g_{ij})\}$ came from a fibre bundle E, there is a diffeomorphism $E \to E'$ given by $e \mapsto [\varphi_i(e)]$ for any i such that $\pi(e) \in U_i$.

For arbitrary bundles E, E', a diffeomorphism of bundles $f : E \to E'$ can be see through the trivializations $\{(W_i, \varphi_i)\}$, $\{(V_j, \psi_j)\}$. To start with, the open covers may be different, but taking the intersections $W_i \cap V_j$ and restricting the maps φ_i, ψ_j, we can assume that the open cover is the same, say $\{U_i\}$. Over an open set U_i, we have

$$\begin{tikzcd} U_i \times F \arrow[rrr, bend right, "(\mathrm{pr}_{U_i}, f_i)"'] & \pi^{-1}(U_i) \arrow[l, "\varphi_i"'] \arrow[r, "f"] & \pi'^{-1}(U_i) \arrow[r, "\psi_i"] & U_i \times F, \end{tikzcd}$$

so the diffeomorphism is determined by a map $f_i : U_i \times F \to F$, in such a way that

$$f_{|\pi^{-1}(U_i)} = \psi_i^{-1} \circ (\mathrm{pr}_{U_i}, f_i) \circ \varphi_i.$$

On an intersection U_{ij} we must have that the identities for $f_{|\pi^{-1}(U_{ij})}$ agree:

$$\psi_i^{-1} \circ (\mathrm{pr}_{U_i}, f_i) \circ \varphi_i = \psi_j^{-1} \circ (\mathrm{pr}_{U_j}, f_j) \circ \varphi_j,$$

or in terms of changes of chart, over a point $x \in U_{ij}$,

$$f_i = h_{ij} \circ f_j \circ g_{ji}, \tag{1.2}$$

where g_{ij}, h_{ij} are the transition maps for E and E' respectively.

Just as giving a cocycle $\{(U_{ij}, g_{ij})\}$ determines a bundle, giving $\{(U_i, f_i)\}$, with smooth maps $U_i \times F \to F$, such that (1.2) is satisfied, determines a bundle diffeomorphism. This is important, so let us draw it:

$$\begin{tikzcd}
U_j \times F \arrow[rrr, "(\mathrm{pr}_{U_j}, f_j)"] & & & U_j \times F \arrow[dd, dashed, bend left, "(\mathrm{pr}_{U_{ij}}, h_{ji})"] \\
& \pi^{-1}(U_{ij}) \arrow[ul, "\varphi_j"'] \arrow[dl, "\varphi_i"] \arrow[r, "f"] & \pi'^{-1}(U_{ij}) \arrow[ur, "\psi_j"] \arrow[dr, "\psi_i"'] & \\
U_i \times F \arrow[uu, dashed, bend left, "(\mathrm{pr}_{U_{ij}}, g_{ji})"] \arrow[rrr, "(\mathrm{pr}_{U_i}, f_i)"'] & & & U_i \times F,
\end{tikzcd}$$

where we replace U_i, U_j by U_{ij} when we look at $(\mathrm{pr}_{U_{ij}}, g_{ji})$ and $(\mathrm{pr}_{U_{ij}}, h_{ji})$.

With this generality, fibre bundles are as complicated as manifolds, if not more, but the idea is that the fibres will have some extra structure we know well. We will then be concerned only about how the fibres glue together globally and take for granted the manifold structure of the base. Let us start with vector spaces.

Definition 1.6. A **vector bundle** is a fibre bundle $V \to M$ such that the generic fibre F is a vector space and the transition maps between any two local trivializations are linear, i.e., $g_{ij}(x, -) \in \mathrm{GL}(F)$.

The dimension of the vector space F is called the **rank of the vector bundle** V. Note that the dimension of V as a manifold satisfies

$$\dim V = \dim M + \operatorname{rk} V.$$

Example 1.7. The cylinder and the Möbius band can easily be upgraded to vector bundles if we replace the interval $(-1, 1)$ with $\mathbb{R}$, that is, the cylinder becomes $\mathrm{S}^1 \times \mathbb{R}$. In the case of the Möbius band, to avoid self-intersections in $\mathbb{R}^3$, it is better to look at it as

$$\{(e^{i\theta}, re^{i\frac{\theta}{2}}) \mid \theta \in [0, 2\pi], r \in \mathbb{R}\}.$$

In both cases, we can find linear transition functions, as in Example 1.5. Analogously one can define the n-twisted cylinder as

$$\{(e^{i\theta}, re^{in\frac{\theta}{2}}) \mid \theta \in [0, 2\pi], r \in \mathbb{R}\},$$

whose transition functions are Id on S^1_t and $(-1)^n \,\mathrm{Id}$ on S^1_b. You can see a twisted cylinder, embedded in $\mathbb{R}^3$, not in $\mathbb{R}^4$ in Figure 1.2.

Note that the fibres of a vector bundle have canonically the structure of a vector space, as the transition maps are linear, that is, the vector space

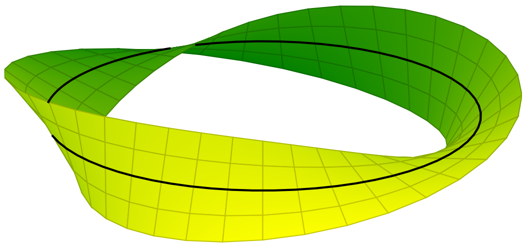

Figure 1.2: The 2-twisted cylinder

structure given by two different trivializations will coincide. The zero element, sum of vectors and scalar products are thus well defined. For example, for $v, w \in V_x$ we use a trivialization restricted to x, $\varphi_{i|x}$, to define the sum

$$v + w := \varphi_{i|x}^{-1}(\varphi_{i|x}(v) + \varphi_{i|x}(w)).$$

When using a different trivialization $\varphi_{j|x}$ we have $\varphi_{j|x} = g_{ji}(x, -) \circ \varphi_{i|x}$, with $g_{ji}(x, -)$ linear. This implies

$$\varphi_{j|x}^{-1}(\varphi_{j|x}(v) + \varphi_{j|x}(w)) = \varphi_{i|x}^{-1}(\varphi_{i|x}(v) + \varphi_{j|x}(w)).$$

Notice that this does not mean that V_x is canonically isomorphic to F.

Fine print 1.5. Actually, an alternative way to define a vector bundle is to ask that the fibres are vector spaces and that the local trivializations preserve this structure. The transition maps will automatically be linear.

When we look at a **diffeomorphism of vector bundles** $f : V \to V'$, we must ask the restriction $f_x : V_x \to V'_x$ to be a linear map. Or alternatively, we ask that the maps $f_i : U_i \times F \to F$ are actually given by $f_i : U_i \to \mathrm{GL}(F)$. We can prove now, just by knowing $\mathrm{GL}(\mathbb{R}) \cong \mathbb{R}^*$, that the cylinder and the Möbius band are not diffeomorphic. If there was a diffeomorphism f sending the cylinder to the Möbius band, there would exist maps $f_e : \mathrm{S}^1_e \to \mathrm{GL}(\mathbb{R})$, $f_w : \mathrm{S}^1 \to \mathbb{R}$ such that, by (1.2) and Example 1.5,

$$f_e = f_w \text{ on } \mathrm{S}^1_t, \qquad\qquad f_w = -f_e \text{ on } \mathrm{S}^1_b.$$

This is not possible as f_e and f_w must have a constant sign on S^1_e and S^1_w respectively. However, witt the notation of Example 1.7, it is possible to write a diffeomorphism between the cylinder and the 2-twisted cylinder just by taking $f_e(x, -) = \mathrm{Id}$, $f_w(x, -) = \mathrm{Id}$. Actually, the same argument easily applies to see that the $2k$-cylinders are diffeomorphic to the cylinder, whereas the $2k + 1$-cylinders are diffeomorphic to the Möbius band.

Fine print 1.6. We are implicitly calculating the first Stiefel-Whitney class, which is something lying, in the case of bundles over S^1, into $H^1(\mathrm{S}^1, \mathbb{Z}_2) \cong \mathbb{Z}_2$. The cylinder has trivial first Stiefel-Whitney class. This class, in the case of vector bundles of rank 1 completely determines the bundle up to diffeomorphism.

Fine print 1.7. Why have we waited till now to show that the cylinder and the Möbius band are not diffeomorphic? It is much easier to see that their vector bundle versions are not diffeomorphic as vector bundles. To see that the are not diffeomorphic as fibre bundles, we would need to use, for instance, the concept of orientation. Positive and negative maps would be replaced by orientation-preserving and orientation-reversing.

Fine print 1.8. If we want to describe mathematically the fact that we are not able to deform the 2-twisted cylinder into the cylinder, we will talk about isotopy in $\mathbb{R}^3$. We will mention isotopy in Section 1.3. Here you can ready how we mathematically prove that they are not isotopic. `https://math.stackexchange.com/questions/2383751/how-is-a-doubly-twisted-cylinder-different-from-a-cylinder`

However, there may be a way to deform one into the other if we embed them in a larger space. Indeed, it is possible to do that already in $\mathbb{R}^4$. You can see the formular for an isotopy in the following link: `https://math.stackexchange.com/questions/2271970/twists-and-half-twists-on-ribbons-in-mathbbr4`

Another possibility is to take as the fibre a Lie group[3] G and ask the transition maps to be elements of G.

Definition 1.8. A **principal G-bundle** for a Lie group G is a fibre bundle $E \to M$ with fibre G, such that the transition functions are maps $U_{ij} \to G$, where G acts on the fibre G by left multiplication.

The fibres of a vector bundle are vector spaces, the fibres of a principal G-bundle are... not groups! This may be shocking, but there was a way of defining the 0 element in the fibres of a vector bundle as 0 is preserved by $\mathrm{GL}(F)$. However, e is not preserved by the left action of G, quite the opposite. We do have, though, a right G-action on each fibre. If we see a point, through a trivialization, as (m, g), we act with $h \in G$ by $(m, g) \cdot h = (m, gh)$. This action is well defined, as if the same point is given by $(m, g_{ij}g)$ in a different trivialization, we also have that $(m, g_{ij}g) \cdot h = (m, g_{ij}gh)$ corresponds to (m, gh) in the first trivialization. This action is free (non-identity elements have no fixed points), and transitive (any two points of the fibre are connected by the action of some g). A space with such an action is called a **G-torsor**.

Fine print 1.9. Choosing an element of a G-torsor as the identity defines a group structure isomorphic to G.

The fibres of a principal bundle are G-torsors. Actually, an alternative way of defining a principal G-bundle is as a fibre bundle $E \to M$ with a fibre-preserving action of G that is free and transitive on the fibres. Or you

[3] A group with a differentiable structure so that the product and the inverse are smooth maps.

may also find the definition of a manifold E with a free right action of a Lie group G, such that the natural projection $E \to E/G$ satisfies the local triviality condition.

Given a vector bundle $E \to M$, we have transition maps $g_{ij} : U_{ij} \to \mathrm{GL}(F) \simeq \mathrm{GL}_n$. This cocycle determines a principal bundle called the **frame bundle** of E. To give a more concrete description, recall that the fibre $E_m := \pi^{-1}(m)$ is a vector space. Consider the set of all possible bases of E_m. A choice of basis can be seen as an isomorphism $\varphi : \mathbb{R}^n \to E_m$, since a basis is determined by the image of the orthonormal basis of $\mathbb{R}^n$. The frame bundle of E has the set of all possible bases as the fibre and the group GL_n acts on the bases, seen as isomorphisms, by $\varphi \cdot g = \varphi \circ g$.

Conversely, if we have a principal G-bundle $E \to M$ and G acts on a vector space F, i.e., $\rho : G \subset \mathrm{GL}(F)$, we can define a vector bundle by composing the transition maps $g_{ij} : U_{ij} \to G$ of E with ρ, so that we get a set of transition maps $\rho \circ g_{ij} : U_{ij} \to \mathrm{GL}(F)$ that determine a vector bundle, called the **associated vector bundle**. More concretely, this bundle is given by $E \times F$ quotiented by the equivalence relation $(e, f) \sim (eg, g^{-1}f)$ for all $e \in E$, $f \in F$ and $g \in G$.

So far we have given simple examples and constructions of bundles. Let us move to a more involved example, the tangent bundle of a manifold. In the mental image we have of a surface, sitting on $\mathbb{R}^3$, we have a clear notion of a tangent space at a point. This strong intuition relies on the actual embedding on $\mathbb{R}^3$ and is not available if we work with abstract manifolds. It is, say, extrinsic, and we want something intrinsic. An intrinsic way to think about a tangent vector on p on a surface would be looking at a curve $\gamma : (-\epsilon, \epsilon) \to M$ passing through p at time 0. The tangent vector, in the extrinsic way, is $\gamma'(0)$. There are many curves giving the same vector. We want to say that two curves γ, σ are equivalent if $\gamma'(0) = \sigma'(0)$... but we cannot do this, it does not make sense for an abstract manifold. Instead, we can take a chart c around p and move this to $\mathbb{R}^n$, where we can take derivatives. Two curves are hence equivalent if $(c \circ \gamma)'(0) = (c \circ \sigma)'(0)$. This clearly defines an equivalence relation and a tangent vector at p is an equivalence class of curves passing through p (for any $(-\epsilon, \epsilon)$. The tangent space at a point $p \in M$ is

$$T_pM = \{\gamma : (-\epsilon, \epsilon) \to M \mid \gamma(0) = p\}/(\gamma \sim \sigma \leftrightarrow (c \circ \gamma)'(0) = (c \circ \sigma)'(0)).$$

This definition does not depend on the choice of chart, as for c_i, c_j,

$$(c_j \circ \gamma)'(0) = D(c_j \circ c_i^{-1})_p(c_i \circ \gamma)'(0). \tag{1.3}$$

What kind of structure does T_pM have, if any? In our intuitive image of a tangent space, we get a plane, a vector space. Do we have this structure for the equivalence classes of curves? We do, thanks to the chart:

$$[\gamma]+[\sigma]=[c^{-1}(c\circ\gamma+c\circ\sigma)],\qquad\qquad \lambda[\gamma]=[c^{-1}(\lambda(c\circ\sigma))],$$

and the zero element is the constant curve $\gamma(t) = p$. We define the following n tangent vectors coming from the coordinates of the chart $c(q) = (x_1(q), \dots, x_n(q))$

$$\frac{\partial}{\partial x_{i\,|p}} := [c^{-1}(t \mapsto (0,\dots,0,\underbrace{t}_{i},0,\dots,0))].$$

At this point, the $\frac{\partial}{\partial x_{i\,|p}}$ are just notation, which will be more meaningful later. We see that $\{\frac{\partial}{\partial x_{i\,|p}}\}$ generate T_pM, as

$$[\gamma] = \sum (x_i \circ \gamma)'(0) \frac{\partial}{\partial x_{i\,|p}}. \tag{1.4}$$

Indeed

$$(c\circ\gamma)'(0) = ((x_1\circ\gamma)'(0), \dots, (x_n\circ\gamma)'(0)).$$

And they are actually linearly independent as $\sum a_i \frac{\partial}{\partial x_{i\,|p}} = 0$ implies

$$(a_1, \dots, a_n) = (0, \dots, 0).$$

This is a very geometrical way of understanding the tangent space to a point, but we also want to see vectors in action. This is done by generalizing the concept of directional derivative. Given $[\gamma] \in T_pM$ and a smooth function $f : U \to \mathbb{R}$ on an open neighbourhood U of p, set

$$[\gamma](f) := (f\circ\gamma)'(0).$$

This action is linear and satisfies, for f, g defined on U,

$$\begin{aligned}[\gamma](fg) = (fg\circ\gamma)'(0) &= ((f\circ\gamma)(g\circ\gamma))'(0)\\ &= [\gamma](f)g(p) + f(p)[\gamma](g).\end{aligned}$$

This means that $[\gamma]$ acts as a derivation. We will see at other moment what this means.

We may have already forgotten why we were doing all this. We wanted a meaningful example of a vector bundle. We just associated to each point p of any manifold M a vector space T_pM. Can they all together be given the structure of a vector bundle? As a set, we have $TM = \cup_{p\in M} T_pM$, and the projection π is quite clear $v \in T_pM$ is sent to p. But we need to give charts, or even better trivializations. We use (1.4), to give a trivialization on $\pi^{-1}(U_\alpha)$

$$\begin{aligned} \cup_{p\in U_\alpha} T_pM &\to U_\alpha \times \mathbb{R}^n \\ \sum a_i \frac{\partial}{\partial x_i}_{|p} &\mapsto (p, (a_i)). \end{aligned}$$

These trivializations, when composed with a chart on U_α will give charts for TM. It is very easy to show that for these charts, π is smooth and the trivializations are diffeomorphisms. Now, by (1.3) and again (1.4), the transition functions on $U_{\alpha\beta}$ are given by

$$D(c_\alpha \circ c_\beta^{-1}),$$

which are smooth, as $c_\alpha \circ c_\beta^{-1}$ is smooth, and most importantly linear, as they are just a matrix!

Fine print 1.10. By giving first the local trivializations on a set, which we still do not know it is a manifold, we make the calculations much simpler, but we have to be sure about all the things we must formally check.

The latest examples have been quite abstract. We were actually doing better with the cylinder, the Möbius band... let us look at surfaces so that we get more tangible examples.

1.3 Surfaces

Surfaces are 2-dimensional manifolds, that is it. We have already seen some surfaces: the sphere, the projective plane, the torus, the cylinder and the Möbius band. The two latter are optimal examples to understand the concepts of fibre bundle, transition maps, diffeomorphisms... but they are not compact, and we are going to be concerned with compact surfaces. We also assume connectedness, as a non-connected surface would just be a disjoint union of connected ones, actually a countable union in order to preserve the second countability. Note that compactness and connectedness are both topological conditions. Actually, let us start working with topological surfaces.

A way to construct more surfaces is the **connected sum**, of which we give the idea. Given two surfaces, we remove a disk in each and we glue along the boundary circle. It is then possible to give charts for this new object so that it is a topological manifold. We just have to take care of the boundary we are using to glue. A good example is the connected sum of g

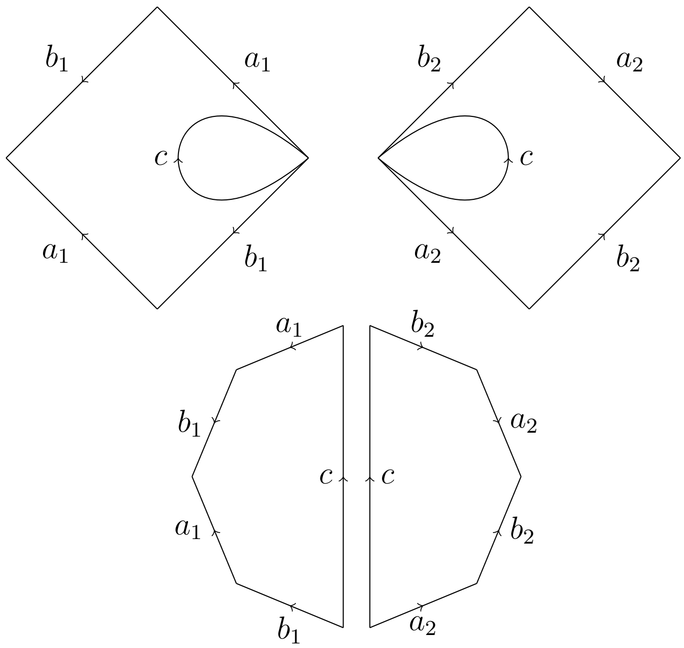

Figure 1.3: The connected sum of two tori.

tori, which intuitively is a surface with g holes. This can be easily seen in terms of polygons.

Connected compact topological surfaces are classified up to homeomorphism by the following theorem.

Theorem 1.9. *Any connected compact surface is homeomorphic to either the sphere, a connected sum of g torus or a connected sum of g projective planes.*

This is a big theorem, but let us show the path of the proof. One has first to see that compact surfaces admit a triangulation, another big theorem. By cutting this triangulation, one obtains a polygon where the edges are identified, as in a square with identified edges for the Möbius band. The next step is to find a normal form for these polygons with identified edges. This is done by getting rid of some repetitions, transforming it into a polygon whose vertices are all identified at one point, and finally cutting this polygon into two pieces and using the identified edges to glue them again until getting to a normal form. The resulting polygon corresponds to one of the surfaces of the theorem. However, this is not the end, as one needs to prove that all these surfaces are different. To do this one uses orientability and the Euler characteristic (the integer #vertices – #edges + #faces of any triangulation, which is well defined). They are indeed complete invariants of topological surfaces, and they are easily computable from the polygon model. The Euler class of the sphere is 2, of the connected sum of g torus is $2-2g$, and of the connected sum of g projective planes is $2-g$. The nonorientable compact surfaces are the connected sum of g projective planes, with $g \geq 1$.

The following will apply for a general topological space, which may or may not be a manifold. We introduce now the notion of **homotopy**. Let X, Y be two topological spaces. We say that two continuous maps $f, g : X \to Y$ are homotopic when there exists a continuous map $H : X \times [0,1] \to Y$ such that $H(x,0) = f(x)$, $H(x,1) = g(x)$ for $x \in X$. We will first deal with homotopic loops on X. Let I be the interval $[0,1]$. A **path** is a continuous map $\gamma : I \to X$. If moreover $\gamma(0) = \gamma(1)$ it is called a **loop**. In the case of paths γ, γ' such that $\gamma(0) = \gamma'(0)$ and $\gamma(1) = \gamma'(1)$, we can ask for **fixed end point homotopy**, a homotopy H such that $H(0,s) = \gamma(0)$ and $H(1,s) = \gamma(1)$ for any $s \in [0,1]$. Fixep end point homotopy defines an equivalence relaction.

On the other hand, given any two paths γ, γ' such that $\gamma(1) = \gamma'(0)$, we define the juxtaposition of paths by

$$(\gamma\gamma')(t) = \begin{cases} \gamma(2t) & \text{for } t \in [0, \frac{1}{2}], \\ \gamma'(2t-1) & \text{for } t \in (\frac{1}{2}, 1]. \end{cases}$$

The **fundamental group of X based on a point** $x \in X$ is the set

$$\pi_1(X,x) = \{ \text{ loops } \gamma : [0,1] \to X \text{ such that } \gamma(0) = x\} \Big/ \begin{array}{c}\text{fixed end point}\\ \text{homotopy}\end{array}$$

together with the product given by the juxtaposition of paths,

$$[\gamma][\gamma'] = [(\gamma\gamma')].$$

The first thing to check is that this operation is well defined (it does not depend on the representatives chosen) and is indeed a group operation. The identity element is given by the class of the constant path $\gamma(t) = x$ for $t \in I$.

The next thing to check is that the definition of the fundamental group up to isomorphism does not depend on the base point. If the manifold is path connected (that is, for any two points $x, x' \in M$ there is a path σ connecting them, $\sigma(0) = x$, $\sigma(1) = x'$), the isomorphism is given by

$$\begin{aligned} \pi_1(X,x) &\to \pi_1(X,x') \\ [\gamma] &\mapsto [(\sigma(\gamma\sigma^{-1}))]. \end{aligned}$$

So if X is path connected, we will then talk about the **fundamental group** of X and denote it by $\pi_1 X$, without caring about the base point.

When a space has trivial fundamental group, we say that it is **simply connected**. This is the case for instance, of a point, a line, a plane, or in general $\mathbb{R}^n$.

The first and most important example is the fundamental group of the circle: $(\mathbb{Z}, +)$, generated by $t \mapsto e^{2\pi i t}$. It should not surprise us, and is not difficult to prove, that the fundamental group of a product of a topological spaces is the direct product of their fundamental groups.

$$\pi_1(X \times Y) \cong \pi_1 X \times \pi_1 Y.$$

Thus, the fundamental group of the torus is $(\mathbb{Z} \times \mathbb{Z}, +)$.

The fundamental group is clearly the same for homeomorphic spaces, but we can go much further. We define a **retraction** of X to a subset $A \subset X$ as a continuous map $r : X \to A$ such that $r_{|A} = \mathrm{Id}_A$. We say that A is a **deformation retract** of X when the identity map $X \to X$ is homotopic to $r : X \to A$ via a homotopy H satisfying $H(a, s) = a$ for any $a \in A$. and $s \in [0, 1]$ In this case, the retraction commutes with the fixed end point homotopy and the fundamental groups of X and A are isomorphic. For example, a point is a deformation retract of a disk, and they both have trivial fundamental group. A torus without a point can be retracted to a figure eight, of two circles joined at a point, and their fundamental group is isomorphic, but we do not know either yet.

When we said the most important fundamental group was the one of the circle, we were really serious, as long as we have the Seifert-Van Kampen theorem. Given a path-connected topological space, consider two open subspaces $U, V \subset X$ such that $U \cup V = X$. The inclusions $U \cap V \to U$ and $U \cap V \to V$ give group homomorphisms $u : \pi_1(U \cap V) \to \pi_1 U$ and $v : \pi_1(U \cap V) \to \pi_1 V$. Seifert-Van Kampen theorem states that we can compute the fundamental group by taking the so-called amalgamated free product

$$\pi_1 U *_{\pi_1(U \cap V)} \pi_1 V,$$

which is the free product[4] of $\pi_1 U$ and $\pi_1 V$ quotiented by the normal subgroup generated by $u(c)v(c)^{-1}$ for $c \in \pi_1(U \cap V)$. In terms of (possibly infinite) presentations

$$\begin{aligned} \pi_1 U &= \langle a_1, \ldots \mid r_1, \ldots \rangle, \\ \pi_1 V &= \langle b_1, \ldots \mid s_1, \ldots \rangle, \\ \pi_1(U \cap V) &= \langle c_1, \ldots \mid t_1, \ldots \rangle, \end{aligned}$$

[4]The free product is just, when taking presentations, the union of generators and relations. For instance, the free product of $\mathbb{Z}$ with $\mathbb{Z}$ is not $\mathbb{Z} \times \mathbb{Z}$, but the free group of two generators.

this means that

$$\pi(X) = \langle a_1, \dots, b_1, \dots \mid r_1, \dots, s_1, \dots, u(c_1)v(c_1)^{-1}, u(c_2)v(c_2)^{-1}, \dots \rangle,$$

where the relations $t_1, \dots$ do not play any role.

For instance, in the figure eight one can take as open sets the total minus the highest point and the total minus the lowest point. These open sets both retract to a circle and their intersection retracts to a point, all as deformation retracts. Consequently, the fundamental group of the figure eight, and hence of a torus without a point, is the free group of two generators.

Even though we know what the fundamental group of a torus is, it is a good exercise to compute it again using Seifert-Van Kampen's theorem. Let y be a point in the torus and V a small disk around it. Define $U = T \setminus \{y\}$. We have in mind the model of the square with identified edges and take a point y in the middle of the square. We base all the fundamental groups on a point x_1 in $U \cap V$. We know that $\pi_1 U$ is the free group of two generators represented by $d^{-1}ad, d^{-1}bd$, whereas $\pi_1 V$ is just trivial. The intersection $U \cap V$ retracts to a circle. Take a loop c containing y in its interior as the generator of $\pi_1(U \cap V)$. Its image in $\pi_1(V)$ is trivial and its image in $\pi_1(U)$ can be computed by looking at the model:

$$u([c]) = [d^{-1}aba^{-1}b^{-1}d] = [d^{-1}ad][d^{-1}bd][d^{-1}a^{-1}d][d^{-1}b^{-1}d].$$

So we deduce, by getting rid of the conjugation by d, that

$$\pi_1 T = \langle a, b \mid aba^{-1}b^{-1} \rangle.$$

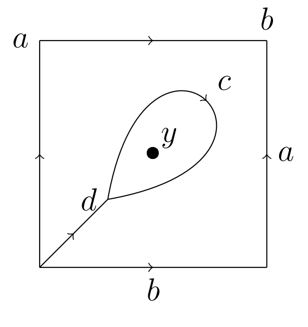

Figure 1.4: Computation of π_1 of a torus using Seifert-van Kampen.

Similarly, we compute the fundamental group of the connected sum of g torus. This is what we are really going to use.

Proposition 1.10. *The fundamental group of a compact connected orientable surface of genus g is*

$$\pi_1\Sigma_g \cong \{a_1, b_1, \dots, a_g, b_g \mid [a_1, b_1] \dots [a_g, b_g] = 1\}.$$

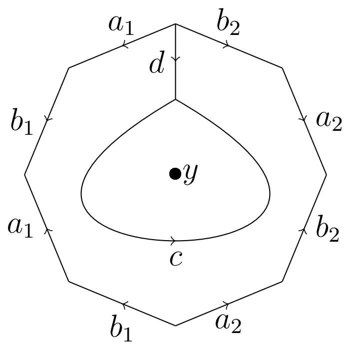

Figure 1.5: Computation of $\pi_1\Sigma_2$ using Seifert-van Kampen.

The fundamental group can endowed with a topology. Its source is a topology on the space of maps $\mathrm{Hom}(\mathrm{S}^1, X)$. In order to motivate this topology, consider three topological spaces T, X and Y. We know what it means for a map $\Psi : T \times X \to Y$ to be continuous, but if we write it as $\hat{\Psi} : T \to \mathrm{Hom}(X, Y)$, where we have not put any topology yet on $\mathrm{Hom}(X, Y)$, we cannot say anything. The idea behind the so-called compact-open topology is to define a topology in $\mathrm{Hom}(X, Y)$ such that any map Ψ is continuous if and only if $\hat{\Psi}$ is continuous. R.H. Fox defined this topology (On topologies for function spaces) and proved that this is the case under some mild assumptions (for instance, X and T being first countable is sufficient). The **compact-open topology** is defined for the space of continuous maps between any two topological spaces X, Y as the topology generated by the sets

$$C(K, U) = \{f \in \mathrm{Hom}(X, Y) \mid f(K) \subset U\},$$

where K is any compact set of X and U is any open set of Y. Give $\mathrm{Hom}(\mathrm{S}^1, X)$ the compact-open topology, consider the subspace $\mathrm{Hom}_x(\mathrm{S}^1, X)$ of loops that start and end at a fixed base point x, and finally consider the quotient topology when we quotient by the equivalence relation of fixed end point homotopy.

Fine print 1.11. One can define other topologies in $\mathrm{Hom}(X, Y)$, usually by defining what it means to converge. If we talk about pointwise convergence, that would be equivalent to the point-open topology, which is generated by $C(K, U)$ where K is just a point[5]. So the topology in $\mathrm{Hom}(X, Y)$ given by pointwise convergence is coarser (has the same or less open subsets) than the compact-open topology. On the other hand, if we consider the topology of compact convergence, we need Y to be a metric space and we say that $\{f_n\}$ converges to f when for $\epsilon > 0$ we can find N such that $\sup_{x\in K} d(f_n(x), f(x)) < \epsilon$ for any compact set $K \subset X$. In the case of Y being a metric space, the compact-open topology is exactly the topology of compact convergence. Note that the topologies above defined using convergence can be defined for the space of all maps from X to Y, whereas for the compact-open topology we do need to consider the continuous maps $\mathrm{Hom}(X, Y)$.

If we have a group with a topology, the first question is whether it is a topological group. And the answer in general is no. The fundamental group is a quasi-topological group, where only the product by the left or the product by the right, but not the product map, are continuous. In our case, as we are dealing with manifolds, we will be safe, as we can say much more about the fundamental group.

To start with, we defined the fundamental group for path-connected topological spaces, but we like to consider connected manifolds, which is in general a weaker condition. However, in the case of manifolds they are the same, as the set of points connected by a path to a base point x is both open and close, by using a local chart.

We chose our manifolds to be second countable and, as a consequence of that, the fundamental group will be countable. This is done by covering the manifold with a countable number of coordinate balls (preimages of balls by a chart), and defining a set I by choosing a point for every intersection of any two balls, including the intersection of a ball with itself. This set is countable. For every ball B and any two points $x, x' \in B$, choose a path $p^B_{x,x'}$ connecting them and lying in B. We get a countable number of paths. Choose a base point x_0 in I and consider the set P of loops starting at x_0 that they are a finite product of some $p^B_{x,x'}$. Again, there is a countable number of loops. The last step is to see that any loop γ is homotopic to one in P, so the fundamental group is a subset of P and must be countable. For this we use the compactness of S^1, the set of the preimages by γ of the coordinate balls is an open cover of S^1. We can choose a finite subcover, so that the image of the loop lies in this finite number of balls. In each ball the path is homotopic to some $p^B_{x,x'}$, preserving the end point x_0 if necessary, so we can make any path homotopic to a path in P. More details and an image can be found in Prop. 1.9 of [Lee13].

Assignment 6: prove that the topology of the fundamental group in the case of manifolds is the discrete topology.

[5]Points are always compact as, given any cover, one can always find a finite subcover, consisting actually of any open set of the original cover.

Hence, in the case of topological manifolds, the topology is always discrete and hence we do have a topological group.

We have come all the way here to define a new manifold. Fix any $x_0 \in M$ and define the **universal cover** $\tilde{M}$ of a manifold M as the set

$$\tilde{M} = \{\text{paths } \gamma : [0,1] \to M \text{ starting at } x_0\} \Big/ \begin{array}{c}\text{fixed end point}\\ \text{homotopy.}\end{array}$$

The universal cover of a manifold is again a manifold. We define a chart for $[\gamma]$ by using a chart (U, φ) around $\gamma(1) \in M$ whose image is a ball in M.

$$U_{[\gamma]} = \{[\sigma\gamma] \mid \sigma \text{ is a path from } \gamma(1) \text{ to any } x \in U\}.$$

Since U is simply connected, two paths $\sigma\gamma$, $\tau\gamma$ are (fixed end point) homotopic if and only if the end points of σ and τ are the same, so $U_{[\gamma]}$ is bijective to U and we can use φ to map it onto an open subset of $\mathbb{R}^n$. If we have $U_{[\gamma]} \cap U'_{[\gamma']} \neq \emptyset$ we also have that $U \cap U' \neq \emptyset$ and the change of chart in $\tilde{M}$ is given exactly by the change of chart in M. Thus, $\tilde{M}$ is a manifold. The best example to have in mind is $M = \mathrm{S}^1$ and $\tilde{M} = \mathbb{R}$.

We interrupted our review on bundle theory to introduce surfaces and the fundamental group because the universal cover of a manifold is perhaps the most important example of a principal bundle in this course. Note that we have a canoncial projection $p : \tilde{M} \to M$ given by $p([\gamma]) = \gamma(1)$. We define an action of $\pi_1 X$, more concretely $\pi_1(M, x_0)$, on $p^{-1}(x)$ for any $x \in M$. Given $\alpha \in \pi_1(M, x_0)$ and $[\gamma] \in p^{-1}(x) \subset \tilde{X}$ we define a right action by

$$[\gamma] \cdot [\alpha] = [\gamma\alpha].$$

Given any other $[\gamma'] \in p^{-1}(x)$, the path $\gamma^{-1}\gamma$ is a loop based on x_0, so

$$[\gamma^{-1}\gamma'] = [\alpha],$$

for some $[\alpha] \in \pi_1 M$ i.e., $[\gamma'] = [\gamma\alpha]$, so the action is transitive. Also, if $[\gamma][\alpha] = [\gamma]$, then $[\gamma^{-1}\gamma\alpha]] = 1_{\pi_1 M}$, i.e., $[\alpha] = 1_{\pi_1 M}$, so the action is free. Thus, the fibres are $\pi_1 M$-torsors and $\tilde{M}$ is actually a principal $\pi_1 M$-bundle. In order to give a trivialization, from a simply connected chart U_α in M and $x \in \pi^{-1}(U_\alpha)$, we define a chart $p{-}1(U_\alpha) \cong U_\alpha \times \pi_1(M, x_0)$ by identifying the path connected component of x in $\pi^{-1}(U_\alpha)$ with $U_\alpha \times \{1\}$.

Fine print 1.12. We have taken as the definition of the universal cover of a manifold the proof of the existence of a universal cover in the general sense. In the usual exposition of this theory, one first defines covering maps as surjections $p : C \to X$ such that any point $x : X$ has an open neighbourhood U with $p^{-1}(U)$ a disjoint union of open sets mapping, each of them, homeomorphically to U by the map p. One then defines a universal cover to be a covering space covering any covering space of the initial space, or, equivalently, to be a simply connected covering space. In categorical terms, one can say that the universal cover is an initial object for the category of covering maps of a given pointed space as objects and covering maps between the domains of the covering maps as morphisms. This universal cover may or may not be exist. A universal cover of a topological space X exists if and only if the space is connected, any point x in an open set U has an open neighbourhood $U' \subset U$ that is pathwise connected (locally path-connected), and any point x has an open neighbourhood where loops based at x are contractible in X (semilocally simply connected). Actually, when the two latter conditions are satisfied then topology of the fundamental group is discrete. Hatcher's Algebraic Topology, Sect. 1.3, from p64.

Fine print 1.13. An example of a space without a universal cover and such that the fundamental group is not a topological group is given by the Hawaiian earring. You can check `https://wildtopology.wordpress.com/2013/11/23/the-hawaiian-earring/` to know more about the Hawaiian earring. For the proof that the multiplication is not continuous, see `https://arxiv.org/pdf/0909.3086.pdf`.

Fine print 1.14. After talking so much about surfaces, you may also wonder what happens for 1-dimensional manifolds: a connected 1-dimensional manifold is homeomorphic to either the circle or the open interval $(0, 1)$.

1.4 More on bundles and their sections

We define a **section** s, sometimes called cross section, of a fibre bundle $\pi : E \to M$ on an open set $U \subset M$ as a smooth map $s : U \to E$ such that $\pi \circ s = \mathrm{Id}_U$. When $U = M$ we say that we have a global section. It is very easy to find global sections in trivial bundles $M \times F$, as for any $f \in F$, the map $x \to (x, f)$ defines one.

In the case of vector bundles, since we have a canonically defined vector space structure, we can define the zero section, mapping any $x \in M$ to $\varphi_i^{-1}(x, 0)$ whenever $x \in U_i$ for a trivialization (U_i, φ_i). This does not depend on the choice trivialization as the transition maps are linear. For a vector bundle, we say that a section vanishes or is **vanishing** at $x \in U$ if $\varphi_i(s(x)) = (x, 0)$ whenever $x \in U_i$.

For a principal bundle P, we cannot do the same with the identity element, as it is not preserved by left multiplication. In this case we get something more involved: over an open set U, trivializations and local sections of a principal G-bundle are in correspondence. On the one hand, for any trivialization φ we define $x \mapsto \varphi(x, e)$. On the other hand, given a section s, define the trivialization $p \in P \mapsto (x, g)$ where $x = \pi(p)$ and g is the unique element of G such that $p \cdot g = s(x)$.

What do sections do for us in vector bundles? They may tell us information about the bundles. For instance, you will never find a nowhere vanishing section on the Möbius band, whereas it is easy to do that for the cylinder, as it is a trivial bundle. This is no coincidence.

* * *

Let us look again at cocycles of transition maps, as we took some advantage from them when passing from vector to principal bundles and back. We can do much more.

Let us consider two vector bundles E, E' over the same manifold M with cocycles of transition maps defined over the same open cover:

$$\{U_{ij}, g_{ij} : U_{ij} \to \mathrm{GL}(V)\}, \qquad \{U_{ij}, h_{ij} : U_{ij} \to \mathrm{GL}(W)\}.$$

We can define a new cocycle by

$$\{U_{ij}, g_{ij} : U_{ij} \to \mathrm{GL}(V) \times \mathrm{GL}(W) \subset \mathrm{GL}(V \oplus W)\},$$

which defines a vector bundle with typical fibre $V \oplus W$. This vector bundle is denoted by $E \oplus E'$ and is called the **direct sum** of E and E'. It is possible to define this by considering the set $\cup_{x\in M} E_x \oplus E'_x$ and defining trivializations by $(\pi, pr_2\varphi, pr_2\varphi')$ using the trivializations φ of E and φ' of E' on the same open set.

Fine print 1.15. The cocycle takes values in $\mathrm{GL}(V) \times \mathrm{GL}(W)$, so it is more exact to say that we get a principal $\mathrm{GL}(V) \times \mathrm{GL}(W)$-bundle. The fact that the principal bundle has structure group $\mathrm{GL}(V) \times \mathrm{GL}(W)$ means that the vector bundle bundle with typical fibre $V \oplus W$ is not any vector bundle, but one coming from this direct sum operation.

Another construction is the **tensor product** of the vector bundles E and E', which we denote by $E \otimes E'$. The transition maps are given in this case by the tensor product of matrices $g_{ij} \otimes h_{ij}$. In this case, it may be clearer to consider the set

$$\cup_{x\in M} E_x \otimes E'_x$$

and consider trivializations

$$e \otimes e' \mapsto (\pi(e), pr_2\varphi(e) \otimes pr_2\varphi'(e')).$$

For a single vector bundle we can consider its **dual bundle** by considering the inverse of the dual of the transition maps, that is, $t_{ij}(x) = g_{ji}(x)^*$ for $x \in U_{ij}$. They define a cocycle of transition maps, as both the dual and the inverse are contravariant -they invert the order of the product- and hence their composition is covariant -it preserves the order of the product. Again, one can give trivializations for the set $\cup_{x\in M} E^*_x$.

A very important example of a dual bundle is the **cotangent bundle**, the dual space of the tangent bundle.

We also have the **determinant bundle** of E, which has $\det g_{ij}(x)$ as transition maps. The result is a line bundle, which can be obtained also by considering the set $\cup_{x\in M} \wedge^{\mathrm{rk}\, E} E_x$ and using $\det \varphi(x)$ as trivializations.

In the case of principal G-bundles we may be tempted to take the product of the cocycles, but this will not be in general satisfy the cocycle condition. A sufficient condition for this to happen is that the group G is abelian. In this case we can talk about the product of two principal G-bundles. An important example of this are principal $\mathbb{C}^*$-bundles, which correspond to complex line bundles. The product of two principal $\mathbb{C}^*$-bundles corresponds to the tensor product of the corresponding complex line bundles.

* * *

We defined tangent vectors at a point $p \in M$ as equivalence classes of curves on M. We saw that they define derivations at p, i.e., $\mathbb{R}$-linear maps from functions defined around p to $\mathbb{R}$,

$$D(fg) = D(f)g(p) + f(p)D(g).$$

We argue now that any derivation at p is indeed given by a curve. It is enough to do that in a chart.

For a function defined around a point p, it is a theorem that there exists a neighbourhood and a smooth function $g_i(x)$ defined on it such that

$$f(x) = f(p) + \sum_i (x_i - p_i) g_i(x),$$

where the functions g_i satisfy $g_i(p) = \frac{\partial f}{\partial x_i}(p)$.

Fine print 1.16. I will add a sketch of the proof.

By applying the property of a derivation, we first see that a derivation of a constant function is zero. This is clear for the constant function 1,

$$D(1) = D(1 \cdot 1) = D(1)1(p) + 1(p)D(1) = 2D(1),$$

and follows by $\mathbb{R}$-linearity for any other $D(c) = cD(1) = 0$.

We then have

$$Df = \sum_i D(x_i) g_i(p) + \sum_i (p_i - p_i) g_i(x) = \sum_i D(x_i) \frac{\partial f}{\partial x_i}(p),$$

so D is a linear combination of the derivations $\frac{\partial}{\partial x_i}_{|p}$ with coefficients $D(x_i)$. In general, derivations come equipped with the commutator operation,

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

In the case of T_pM, this is just zero, as

$$\Big[a\frac{\partial}{\partial x_i}_{|p}, b\frac{\partial}{\partial x_j}_{|p}\Big](f) = ab\Big(\frac{\partial^2 f}{\partial x_j \partial x_i} - \frac{\partial^2 f}{\partial x_i \partial x_j}\Big)(p) = 0,$$

by the symmetry of second derivatives.

When one upgrades this picture from vectors in T_pM to vector fields around p (sections of the tangent bundle) one can define the Lie bracket in the same way

$$[X,Y](f) = X(Y(f)) - Y(X(f)),$$

but we do not get zero anymore, as in the expression above a and b would become functions around p and the bracket would be

$$\Big[a\frac{\partial}{\partial x_i}, b\frac{\partial}{\partial x_j}\Big](f) = a\frac{\partial b}{\partial x_i}\frac{\partial f}{\partial x_j} - b\frac{\partial a}{\partial x_j}\frac{\partial f}{\partial x_i}.$$

Note that for two vector fields X,Y, the value of $[X,Y](f)(p)$ does not necessarily coincide with $[X_p, Y_p](f)$, since the latter is always zero as the commutator of derivations at p.

Fine print 1.17. When we looked at the vector space structure of T_pM we did everything by using curves. However, if we had to write the bracket in terms of curves, we would not get a nice expression, just a linear combination of the expressions above.

The Lie bracket is clearly linear and skew-symmetric, and it is a derivation of itself (Jacobi identity)

$$[X,[Y,Z] = [X,[Y,Z]] + [Y,[X,Z]].$$

A bracket with these three properties on a vector space gives the vector space the structure of a Lie algebra.

In the case of vector fields, note that the vector space is the infinite-dimensional vector space of functions defined around p.

Fine print 1.18. In this case of vector fields there is an action of $f \in \mathcal{C}^\infty(M)$ on a vector field X by $(fX)(p) = f(p)X(p) \in T_pM$. The Lie bracket of vector fields also satisfies the so-called Leibniz rule

$$[X, fY] = X(f)Y + f[X,Y].$$

* * *

We defined the tangent bundle of a manifold. We can also define the tangent of smooth maps $f : M \to N$. At a point m we define

$$\begin{aligned} Tf_p : T_m M &\to T_{f(m)} N \\ [\gamma] &\mapsto [f \circ \gamma], \end{aligned}$$

which extends to a map $Tf : TM \to TN$, which we will denote also by f_* and call push-forward. For the composition of two smooth maps $f : M \to N$, $g : N \to L$ we have

$$T_p(f \circ g) = T_g(p)f \circ T_p g.$$

You may remember the definition of exponential map in riemannian geometry. The idea behind it is just that the riemannian metric allows you to pick up a unique representative curve (the unique local geodesic) of the vector, seen as an equivalence class of curves. Once we associate γ to $X \in T_p M$, the geodesic exponential of X is the point $\gamma(1) \in M$ if defined. The geodesic exponential is a local homeomorphism from the tangent space at a point onto a neighbourhood of the point.

In the case of a Lie group G, we have a different way to pick up a unique representative curve. Any left translation L_g is a smooth map and its differential defines a map $TL_g : T_e G \to T_g G$. For $X \in T_e G$, it is a theorem, by existence and uniqueness of ordinary differential equations (the same one used for the existence of geodesics), that there is an only curve γ such that the tangent vector of γ at $\gamma(t)$ is precisely $TL_{\gamma(t)} X$. We could write this like

$$\gamma'(t) = TL_{\gamma(t)} X.$$

The exponential is then defined by $X \mapsto \gamma(1) \in G$. In the case of a Lie group, since the action of $\{L_g\}$ is transitive, we have that these curves are defined for all time, so the exponential is defined as a map $\exp : T_e G \to G$. This does not mean that it is surjective, but that is a different story.

A last and very important remark is that $T_e G$ is endowed with the structure of a Lie algebra, but not just with the zero bracket. We can identify $T_e G$ with left-invariant vector fields, which means that for $X_e \in T_e G$ we have $X_g = TL_g X_e \in T_g G$. We then use the Lie bracket of vector fields to get another left-invariant vector field (this is something to be checked explicitly), which is identified with another element of $T_e G$. We will use the notation

$$\mathfrak{g} := T_e G$$

for the Lie algebra $T_e G$ with the bracket of left-invariant vector fields.

* * *

We finish this section by giving some ideas about Čech cohomology. We have been working a lot on a suitable open cover and functions defined on its open sets or on two-fold intersections. An open cover consisting of simply connected open sets whose intersections (of any number of them) are simply connected is called a **good cover**. For instance, the cover of the sphere with two hemispheres is not good, as the intersection is not simply connected. However, since we are assuming second countability, which implies paracompactness, a smooth manifold always admits a good cover.

Fine print 1.19. A proof of that fact can be found in `https://ncatlab.org/nlab/show/good+open+cover#ExistenceOnParacompactManifolds`

Fix a good cover $\{U_i\}$ and define the following sets of so-called 0, 1 and 2-Čech cochains.

$$\begin{aligned} C^0 &:= \{ \{f_i\} \mid f_i : U_i \to G\} \\ C^1 &:= \{ \{g_{ij}\} \mid g_{ij} : U_{ij} \to G\} \\ C^2 &:= \{ \{h_{ijk}\} \mid h_i : U_{ijk} \to G\}. \end{aligned}$$

Define a differential operator δ by

$$\begin{aligned} \delta : C^0 &\to C^1 \\ \{f_i\} &\mapsto \{f_i f_j^{-1}\}, \end{aligned}$$

where $f_j^{-1} : U_j \to G$ denotes the map $x \mapsto (f_j(x))^{-1}$ for $x \in U_j$, and the juxtaposition denotes the product (not the composition) of the two functions restricted to the intersection of their domains (that is, $U_i \cap U_j$). Following the same notation, define

$$\begin{aligned} \delta : C^1 &\to C^2 \\ \{g_{ij}\} &\mapsto \{g_{ij} g_{jk} g_{kj}\}. \end{aligned}$$

It is easy to check that $\delta^2 = 0$, and some of these formulas they are actually familiar, as we will see. Let us start with the 0-cocycles, that is, the 0-cochains in the kernel of δ. Being in the kernel means that $f_i f_j^{-1}$ is the identity in the intersection U_{ij}, or alternatively that $f_i = f_j$ on U_{ij}. Hence, $\{f_i\}$ actually defines a global function $M \to G$, and this is precisely the zeroth Čech cohomology group $\check{H}^0(M, G)$.

Next, 1-cocycles: here we have our cocycles of transition maps, which give principal G-bundles. And what are the 1-coboundaries, elements in the image of δ? They are principal G-bundles with very special transition maps given by $\{f_i f_j^{-1}\}$. These are precisely trivial bundles, as we saw in (1.2). And if we take the first cohomology group, 1-cocycles modulo 1-coboundaries, we obtain, also by (1.2), principal G-bundles up to diffeomorphism. Thus, if we choose a nice open cover, the first Čech cohomology group, denoted by

$$\check{H}^1(M, G),$$

gives a classification of the principal G-bundles.

Fine print 1.20. Čech cohomology can be defined in general, but the differential gets a bit more complicated. We just wanted to show what the first two cohomology groups are. On the other hand, in general, instead of an open cover, the cohomology is defined as a direct limit of open covers ordered by refinement.

1.5 Connections

When you studied differential geometry you talked a lot about curvature. This notion did not have to do so much with the metric, but with the connection that was canonically attached to it, the Levi-Civita connection.

For a fibre bundle $\pi : E \to M$, consider the tangent map of the projection, $\pi_* : TE \to TM$. Its kernel $\ker \pi_*$ defines a subset $VE \subset TE$, which has a clear projection $\pi : VE \to E$. This set is indeed a bundle itself, a subbundle. By looking at the trivializations of TE, as a bundle over E, all the fibres of VE (preimages of a point) have the same rank, and then the restriction of the trivializations of TE will give a set of trivializations for VE.

The inclusion $E_{\pi(x)} \subset E$ comes together with a map $T_x(E_{\pi(x)}) \to T_xE$, which happens to be an inclusion. Indeed, for an open trivializing neighbourhood U around $\pi(x)$, we have $E_{|U} \cong U \times F$, which restricts to $E_{\pi(x)} \cong \{\pi(x)\} \times F$. By choosing a chart around x coming from a product of charts of $U \times F$, say $c = (c_1, c_2)$ with image in $\mathbb{R}^n \times \mathbb{R}^k$ with $k = \dim F$, we see that

$$T_x(E_{|U}) \cong T_{c_1(x)}(\mathbb{R}^n) \oplus T_{c_2(x)}(\mathbb{R}^k) \subset \{0\} \times T_{c_2(x)}(\mathbb{R}^k) \cong T_x(E_{\pi(x)})$$

and the inclusion follows. One can similarly prove that the map $T_xE \to T_{\pi(X)}M$ is surjective.

The image of the inclusion $T_x(E_{\pi(x)}) \to T_xE$ lies in $V_x(E)$ as $\pi(E_{\pi(x)}) = \{x\}$, and the map

$$\begin{aligned} T_x(E_{\pi(x)}) &\to V_xE \\ [\gamma] &\mapsto [\gamma] \end{aligned} \tag{1.5}$$

is actually an isomorphism, as the dimension of T_xE is $\dim M + \dim F$, whereas, from the surjectivity of $T_xE \to T_{\pi(X)}M$, the dimension of V_xE is

$$\dim V_xE = \dim T_xE - \dim M = \dim F = \dim T_x(E_{\pi(x)}).$$

Fine print 1.21. It is not true in general that if $i : N \to M$ is an inclusion, the map $i_* : T_xN \to T_xM$ is an inclusion. This property, the differential being injective, is called being an immersion. And, by the way, an immersion is not necessarily injective! Like the immersion of the Klein bottle in $\mathbb{R}^3$, which self-intersects.

To sum up, the vertical bundle is the vector bundle whose fibres are the tangent spaces to the fibres.

A **connection at** $x \in E$ is then a complement H_xE to V_xE inside T_xE, that is, a vector subspace H_xE such that $T_xE = V_xE \oplus H_xE$. Note that V_xE is canonically defined while there are many choices for H_xE. If we had a metric, we could take H_xE to be the orthogonal complement, but that is only one possible choice, there are infinitely many more.

A **connection** of E is the smooth choice of a connection for every $x \in E$, that is $\{H_xE\}_{x\in E}$. In order to make sense of the word smooth, we give a few definitions of theory of distributions (as in differential geometry, nothing to do with the distributions related to generalized functions). A **distribution** is a choice of a subspace $D_x \subset T_xM$ for every point $x \in M$. A distribution is smooth when around any $x \in M$, say in an open neighbourhood V_x, there exist smooth (locally defined) vector fields $X^1, \ldots, X^r$ such that D_y is generated by $X^1_y, \ldots, X^r_y$ for $y \in V_x$. Our distribution $\{H_xE\}_{x\in E}$ has the additional property of having always the same rank. We call this to be a **regular distribution**, although many authors include this hypothesis in the definition of distribution.

A natural way of obtaining distributions is through regular foliations of manifolds. A **regular foliation** is a decomposition of a manifold into disjoint submanifolds of the same dimension, called leaves, a disjoint collection $\{N_\alpha\}$ such that $M = \cup_\alpha N_\alpha$ with the property that locally there are charts

$$U \to \mathbb{R}^n \cong \mathbb{R}^k \times \mathbb{R}^{n-k}$$

where the leaves N_α are locally given by the preimage of $\mathbb{R}^k \times \{a\}$. The tangent spaces of a regular foliation define a regular distribution, which satisfies an additional property, their sections are closed under the Lie bracket of vector fields. This is a very important property and has its own name. We say that a distribution $\{D_x\}$ is **involutive** when the smooth vector fields with image in D_x, which we denote by $\Gamma(D_x)$, satisfy

$$[\Gamma(D_x), \Gamma(D_x)] \subset \Gamma(D_x).$$

When a distribution is given by the tangent vector fields of a foliation, we say that the foliation is **integrable**.

Frobenius' theorem states that an involutive regular distribution is integrable, that is, must come from a foliation. In other words, involutivity is not only a necessary condition for integrability but it is actually sufficient in the case of regular distributions. A good example for this is the vertical bundle, which comes from the regular foliation $\{E_{\pi(x)}\}_{x\in M}$, and is hence involutive.

Fine print 1.22. Distributions may be not regular, sometimes referred to as generalized distributions. This is the case, for $M=\mathbb{R}^2$ with polar coordinates (r,θ), of

$$\langle r(-\sin\theta,\cos\theta)\rangle \subset T_{(r,\theta)}\mathbb{R}^2,$$

which gives a one-dimensional subspace for every point apart from the origin, where it gives the zero vector space. It is clear by the definition that this is a smooth distribution as it is globally generated by a smooth vector field. There is a foliation associated to, integrating if you want, this singular distribution. It is given by the origin and concentric circles centered at the origin (this picture was the inspiration to define the foliation). This is a decomposition of the plane into leaves that do not have the same dimension. There are some points where the dimension is lower. In these points, and only in these points, the leaves will not satisfy the chart property above either. If one wants to make sense of singular foliations, one usually says that they are the ones coming from involutive singular distributions.

Fine print 1.23. Distributions may not be integrable. A good example for this is $\mathbb{R}^3$ and the kernel of the differential form $dz + xdy$. A distribution given as the kernel of an ideal of differential forms (where the algebra operation is the wedge product) is integrable if and only if this ideal is closed under the exterior derivative. When we do this for only one form α, being a differential ideal is the same as satisfying $\alpha\wedge d\alpha = 0$. This is not the case for $\alpha = dz + xdy$, as $\alpha\wedge d\alpha = dz\wedge dx\wedge dy$. This non-integrable distribution is also the simplest example of a contact form, leading to contact geometry, the geometry described my maximallly non-integrable distributions.

A connection is equivalently defined as a surjective bundle map

$$\Phi : TE \to VE$$

such that $\Phi^2 = \mathrm{Id}$, giving for each vector its vertical part. This can be though as a 1-form with values in the bundle VE, i.e., $\Phi\in\Omega^1(E,VE)$. The horizontal complements to VE are then defined by $HE = \ker\Phi$. If Φ is giving the vertical part of a vector, the map $\mathrm{Id}-\Phi : TE\to TE$ gives the horizontal part. This point of view is suitable to introduce the curvature as an operator.

We mentioned above that the bracket of vertical vector fields is again vertical. What about the bracket of horizontal vector fields? The measurement of how far is the bracket of horizontal vector fields from being horizontal is measured by the curvature. The **curvature** of the connection Φ is given by

$$\Omega(X,Y) = \Phi[(\mathrm{Id}-\Phi)(X),(\mathrm{Id}-\Phi)(Y)],$$

so that $\Omega\in\Omega^2(E,VE)$. This means taking the horizontal parts of two vector fields, doing their bracket and looking at the vertical part. If the curvature is zero, the bracket of horizontal fields is again horizontal, and we say that the connection is **flat**.

Being flat means that the distribution of horizontal subspaces is integrable, i.e., there exists a foliation of E by submanifolds such that the tangent spaces to these submanifolds are exactly the horizontal subspaces.

Example 1.11. The following, although seemingly trivial, is a fundamental example. For a trivial bundle $U \times G$, the vertical subspace at each (u, g) is given by $\{0\} \times T_g G$. A complement, and hence a distribution, is given by $T_u U \times \{0\}$. This is clearly integrable, as there are integral submanifolds $U \times \{g\}$ whose tangent space at (u, g) is precisely $T_u U \times \{0\}$.

One very basic question is whether connections exist for any manifold. The answer is yes and uses the connections of the example above together with partitions of unity.

Fine print 1.24. One can actually describe the space of all connections. The differences of any two connections lies in a (infinite dimensional) vector space, so the space of connection is an affine space.

When the fibre bundle has extra structure, we will ask the connection to keep this extra structure. Let E have the structure of a principal G-bundle. For each $g \in G$ and $x \in E$, we have a bundle map $R_g : E \to E$ whose differential gives $(R_g)_* : T_x E \to T_{xg} E$. A **principal connection** on E is a connection such that

$$H_{xg} E = (R_g)_* H_x E. \tag{1.6}$$

On principal G-bundles we will only talk about principal connections and just say connection.

On a principal bundle we can understand the vertical subbundle via the group action. For each $x \in E$ we have a map $L^x : G \to E$. Its differential at the identity, by recalling the notation $\mathfrak{g} = T_e G$, gives $(L^x)_* : \mathfrak{g} \to T_x E$. Exactly as in 1.5, this map is an isomorphism

$$(L^x)_* : \mathfrak{g} \to V_x E. \tag{1.7}$$

The inverse of this map can be combined with $\Phi \in \Omega^1(E, VE)$ to define

$$\theta := -(L^x)_*^{-1} \circ \Phi \in \Omega^1(E, \mathfrak{g}).$$

This map θ is zero in vertical vectors and gives the corresponding element of $\mathfrak{g}$ via 1.7 for horizontal vectors.

The equivalent of condition (1.6) in terms of θ is proved to be

$$(R_g)^* \theta = \mathrm{Ad}(g^{-1}) \theta.$$

Fine print 1.25. Similarly the curvature can be regarded as $\Omega \in \Omega^2(E, \mathfrak{g})$ and the Maurer-Cartan formula is satisfied,

$$\Omega = d\theta + \frac{1}{2}[\theta, \theta].$$

Remark 1.12. A good reference for connections on principal G-bundles is Figueroa O'Farrill's lecture notes on Gauge Theory, `https://empg.maths.ed.ac.uk/Activities/GT/` or on Spin Geometry, Lecture 5, `https://empg.maths.ed.ac.uk/Activities/Spin/`.

You may have been wondering for a while where the connections you studied in differential geometry are, even whether that covariant derivative has something to do with all we have done so far. The connections above, usually when expressed as differential forms taking values elsewhere, are sometimes known as Ehresmann connections. We will mention its relation with the sometimes called Koszul connections, or covariant derivatives. Let us look at that.

Given a connection on a principal bundle E, we shall define a covariant derivative in any associated vector bundle $E(F)$, or $E \times_G F$ with typical fibre a vector space F where G acts. We mean

$$E(F) := \{(x, f) : x \in E, f \in F\}/\langle (x, f) \sim (xg, g^{-1}f) \text{ for } g \in G \rangle.$$

To do that we first understand sections of $E \times_G F$ as certain maps from E to F. A section is a map $s : M \to E(F)$ such that $\pi \circ s = \mathrm{Id}$, so for $m \in M$, we have $s(m) = [(x, f)]$ and we could start by defining a map by $x \mapsto f$. If we have a different representative $s(m) = [(x', f')]$, we should also have $x' \mapsto f'$. As $x' = xg$ for some g, and $f' = g^{-1}f$, the map we are defining is G-equivariant, that is, $\varphi(xg) = g^{-1}f$. Conversely, any G-equivariant map $E \to F$ gives rise to a section of $E(F)$. This can be written as

$$\Omega^0(M, E(F)) \cong \Omega^0_G(E, F).$$

We explain this notation: $\Omega^0(M, E(F))$ denotes functions on M taking values on the bundle $E(F)$. This means that we map $m \mapsto f(m) \otimes v(m)$, where $f(m) \in \mathbb{R}$ and $v(m) \in E(F)$. As $f(m)$ is a function, we can multiply $v(m)$ by $f(m)$ and get $f(m)v(m) \in E(F)$, so we have a section. Note that v already defines a section. Conversely, any section s corresponds to the function $1 \otimes F$. On the other hand, $\Omega^0_G(E, F)$ denotes the G-equivariant maps from E to F.

Thus, sections of $E(F)$ are in correspondence with G-equivariant maps $E \to F$. Something similar happens with one-forms, but in this case we have to ask for an extra condition, horizontality. We say that a differential form is horizontal when its value only depends on the horizontal part. If we denote by h^* the dual of the horizontal projection, we can write horizontality as $h^*\omega = \omega$. Of course $h^*\omega$ is horizontal for any ω. A horizontal and G-invariant differential form is said to be **basic**. It is proved that one-forms taking values in $E(F)$ correspond to basic one-forms on E taking values in F (see Section 2.2.1 or Proposition 5.7 of Figueroa-O'Farrill's notes if you want more details). Similarly, as before we would have

$$\Omega^1(M, E(F)) \cong \Omega^1_{basic}(E, F).$$

With this reinterpretation we can now give a covariant derivative on $E(F)$. Given a section s of $E(F)$, we regard it as a G-equivariant map $E \to F$. This could be thought as a function α on E taking values in F. We take the exterior derivative $d\alpha$ of this function and get a 1-form on E, again with values in F. This is again G-invariant, but it is not necessarily horizontal. This is the point where we use the connection. We take the horizontal projection of this form, $h^*(d\alpha)$, and get a basic form. By the correspondence above, this is a one-form with values in $E(F)$. Let us make it clear:

$$\begin{tikzcd}
\Omega^0(M, E(F)) \arrow[rr, "\nabla"] \arrow[d] & & \Omega^1(M, E(F)) \\
\Omega^0_G(E, F) \arrow[r, "d"] & \Omega^1_G(E, F) \arrow[r, "h^*"] & \Omega^1_{basic}(E, F). \arrow[u]
\end{tikzcd}$$

We have thus obtained a map

$$\nabla : \Omega^0(M, E(F)) \to \Omega^1(M, E(F)) \tag{1.8}$$
$$s \mapsto \nabla s. \tag{1.9}$$

which for each vector field $X \in \Gamma(TM)$ gives a section $\nabla_X s \in \Gamma(E(F))$.

We recall the properties of a covariant derivative. A covariant derivative on a fibre bundle E is an operator $\nabla : \Omega^0(M, E) \to \Omega^1(M, E)$ satisfying

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

From this formula it is clear that the difference of two connections satisfies

$$(\nabla - \nabla')(fs) = f(\nabla - \nabla')(s),$$

that is, it is given by a 1-form with values on the bundle of endomorphisms,

$$\nabla - \nabla' \in \Omega^1(\mathrm{End}\, E).$$

Fine print 1.26. A connection on a principal G-bundle E defines a connection on any associated fibre bundle, for a fibre with an action of G that is not necessarily a vector space. In this generality we will not have a covariant derivative as we know it, since the properties of a covariant derivative, like for $\nabla_X fs$, make sense for vector bundles but not for a general fibre bundle. For more details on the induced connection see Section 19.8 of [Mic08].

Chapter 2

Surface group representations

You may have seen before the definition of a representation of a group H in a vector space V as a map $H \times V \to V$ with some properties. When V is finite dimensional, a simple way to encompass all the properties is the definition or a representation of H in V as a homomorphism of groups

$$\rho : H \to \mathrm{GL}(V).$$

The vector space may come with extra structure (euclidean or hermitian metric, symplectic structure, etc.) and we may want to talk about a representation that preserves this structure. We would have to replace $\mathrm{GL}(V)$ by $O(V)$, $U(V)$, $\mathrm{Sp}(V)$, etc. So in order to have a general setting, we just say that a group H represents into a group G if we have a homomorphism $\rho : H \to G$. In this chapter we look at the case where H is a surface group (the fundamental group of a compact connected orientable surface) and G is a Lie group.

2.1 The set of representations

Let Σ be a compact connected orientable surface of genus g. Denote its fundamental group by $\pi_1\Sigma$. Let G be a Lie group. The set of all representations of $\pi_1\Sigma$ into G is the set of all homomorphisms

$$\mathrm{Hom}(\pi_1\Sigma, G).$$

As the topology of $\pi_1 G$ is discrete, all these homomorphisms are moreover continuous, and we can endow this set with the compact-open topology.

As the group $\pi_1\Sigma$ is finitely presented by

$$\pi_1\Sigma \cong \{a_1, b_1, \ldots, a_g, b_g \mid \prod_{j=1}^{g} [a_j, b_j] = 1\},$$

giving a representation $\rho \in \mathrm{Hom}(\pi_1\Sigma, G)$ is the same as giving $2g$ elements of G, say,

$$A_1, B_1, \dots A_g, B_g,$$

which correspond to the images of the $\rho(a_j)$ and the $\rho(b_j)$, such that

$$\prod_{j=1}^{g}[A_j, B_j] - \mathrm{Id} = 0. \tag{2.1}$$

In other words, by the choice of the generators, any $\rho \in \mathrm{Hom}(\pi_1\Sigma, G)$ can be seen inside G^{2g} as a tuple $(A_1, \dots, A_g, B_1, \dots, B_g)$ satisfying (2.1), i.e., the vanishing set of a function.

Assume $G \subset \mathrm{GL}(n, \mathbb{R})$. By looking at the matrix coefficients, one usually sees $\mathrm{GL}(n, \mathbb{R})$ as an open set $\mathbb{R}^{n^2}$, but we can see it as a vanishing set of $\mathbb{R}^{n^2+1}$ if we consider

$$\{(A, x) \mid A \in \mathbb{R}^{n^2} \simeq \mathrm{Mat}(n, \mathbb{R}), x \in \mathbb{R}, \det A \cdot x = 1\}.$$

The set G^{2g} corresponds then to the vanishing of some polynomials equations in $(\mathbb{R}^{n^2+1})^{2g}$, and $\mathrm{Hom}(\pi_1\Sigma, G)$ corresponds to the vanishing of some extra polynomial equations[1] coming from (2.1). The vanishing set of polynomials in an affine space (like $(\mathbb{R}^{n^2+1})^{2g}$) is an object called **affine algebraic set** and is the very starting point of algebraic geometry.

We spent a lot of time talking about manifolds in Chapter 1, so someone who sees these affine algebric sets for the first time may think they are manifolds, as $\mathrm{GL}(n, \mathbb{R})$ is. This is very far from being true, as some examples show immediately. The vanishing set of $xy = 0$ consists of the two axis, which is not a manifold, as no neighbourhood of the origin (a cross) is homeomorphic to an open set of $\mathbb{R}$. We could say that this was not a good example as it can be seen as the union of two vanishing sets $x = 0$ and $y = 0$, it is reducible. For an irreducible example, take the vanishing set of $y^2 - x(x^2 + 1) = 0$.

---

[1] Note that the components of (2.1) are not necessarily polynomials, as we have some determinants as denominators, but we can either multiply by them, or use the fact that we have added $\det A^{-1}$ as a variable in our description of $\mathrm{GL}(n, \mathbb{R})$ inside $\mathbb{R}^{n^2+1}$.

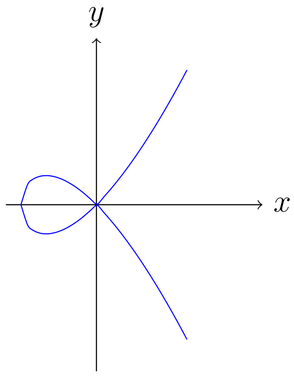

However, an algebraic set is a manifold apart from the singular points, where intuitively the dimension of the tangent space is not the dimension of the manifold. The point is that having the vanishing of polynomials means that the vanishing is not too bad, it has codimension at least one, like the points in curves that we just saw, so it is very mild bad behaviour. If we were to do this with smooth functions, the vanishing set could go really wild.

In general, if G is not inside $\mathrm{GL}(n, \mathbb{R})$, it is known that Lie groups are analytic manifolds, and one can give $\mathrm{Hom}(\pi_1\Sigma, G)$ has the structure of a real analytic variety, and by variety we mean that it looks locally as a real analytic set, that is, as the vanishing set of analytic functions.

2.2 The action of the group

We have a natural left action of $(\varphi, \gamma) \in \mathrm{Aut}\,\pi_1\Sigma \times \mathrm{Aut}\,G$ on $\rho \in \mathrm{Hom}(\pi_1\Sigma, G)$ given, for $a \in \pi_1\Sigma$ by

$$((\varphi, \gamma) \cdot \rho)(a) = \gamma(\rho(\varphi^{-1}(a)).$$

Note that the action of $\mathrm{Inn}\,\pi_1\Sigma$ is contained in the action of $\mathrm{Inn}\,G$. Given $\beta \in \pi_1\Sigma$ acting by conjugation c_β, we have

$$((c_\beta, 1) \cdot \rho)(a) = \rho(\beta a \beta^{-1}) = \rho(\beta)\rho(a)\rho(\beta)^{-1} = ((1, c_{\rho(\beta)}) \cdot \rho)(a),$$

where $c_{\rho(\beta)}$ is the conjugation in G by $\rho(\beta)$. Thus, one can just look at the action of $\mathrm{Out}\,\pi_1\Sigma := \mathrm{Aut}\,\pi_1\Sigma / \mathrm{Inn}\,\pi_1\Sigma$, and $\mathrm{Aut}\,G$, whose action can be decomposed in the action of $\mathrm{Inn}\,G$ and the action of $\mathrm{Out}\,G := \mathrm{Aut}\,G/\mathrm{Inn}\,G$. Although we will only look at $\mathrm{Inn}\,G$, it is good to be aware of the other actions.

Fine print 2.1. In our case, the group $\mathrm{Out}\,\pi_1\Sigma$ coincides with the mapping class group of Σ, which, by definition is the group of diffeomorphisms up to isotopy, or, in the smooth category, $\mathrm{Diff}\,M/\mathrm{Diff}_0 M$, where $\mathrm{Diff}_0 M$ denotes the identity component of the group of diffeomorphisms. For instance, the mapping class group of the sphere is $\mathbb{Z}_2$, and a generator is given by any mirror image of the sphere $(x, y, z) \mapsto (-x, y, z)$. The mapping class group may depend on the category (topological, smooth, etc.) we are working with.

From now on we will only consider the action of G on $\mathrm{Hom}(\pi_1\Sigma, G)$ by conjugation, that is, for $a \in \pi_1\Sigma$ and $g \in G$,

$$(g \cdot \rho)(a) = g\rho(a)g^{-1}.$$

Let us look at some example. When G is abelian, the relation of the commutators is trivially satisfied and we have

$$\mathrm{Hom}(\pi_1\Sigma, G) \simeq G^{2g}.$$

Moreover

$$\mathrm{Hom}(\pi_1\Sigma, G) \equiv \mathrm{Hom}(\pi_1/[\pi_1\Sigma, \pi_1\Sigma], G) \equiv \mathrm{Hom}(H_1(\Sigma), G)$$

where the last equivalence follows from the fact that the abelianization of the fundamental group is the first homology group (Hurewicz theorem).

When $G = \mathbb{R}$, we get $\mathrm{Hom}(H_1(\Sigma), \mathbb{R}) \simeq H^1(\Sigma, \mathbb{R})$, by the universal coefficient theorem. As Σ is a genus g surface we have $H^1(\Sigma, \mathbb{R}) \simeq \mathbb{R}^{2g}$. A worthwhile remark is that the cup product[2] defines a map

$$H^1(\Sigma, \mathbb{R}) \times H^1(\Sigma, \mathbb{R}) \to H^2(\Sigma, \mathbb{R}).$$

By choosing a volume form in Σ, which we can as Σ is orientable, we get a symplectic form (non-degenerate closed 2-form) in $H^1(\Sigma, \mathbb{R})$. This symplectic form is indeed a general feature.

We want to show next with an example how the orbit space for the action of G by inner automorphisms can result in a non-Hausdorff space.

Let Σ be a genus 2 surface and $G = \mathrm{SL}(2, \mathbb{R})$. Note that if we choose images of the generators of $\pi_1\Sigma$ in such a way that $\rho(a_1) = \rho(b_2)$ and $\rho(a_2) = \rho(b_1)$ the relation is automatically satisfied.

Fix a real number $a > 1$ and define the matrix

$$A := \left(\begin{smallmatrix} a & 0 \\ 0 & a^{-1} \end{smallmatrix}\right).$$

Give two representations ρ_1, ρ_2 by

$$\rho_1(a_2) = \rho_1(b_1) = \rho_2(a_2) = \rho_2(b_1) = g$$

[2]Cohomology class of the wedge product of two representatives of cohomology classes $[\alpha] \cup [\beta] = [\alpha \wedge \beta]$.

and

$$\rho_1(a_1) = \rho_1(b_2) = \left(\begin{smallmatrix} 1 & 1 \\ 0 & 1 \end{smallmatrix}\right)$$

$$\rho_2(a_1) = \rho_2(b_2) = \left(\begin{smallmatrix} 1 & 0 \\ 1 & 1 \end{smallmatrix}\right)$$

We check that ρ_1 is not in the same $\mathrm{SL}(2,\mathbb{R})$-orbit of ρ_2. If there were, there would exist $S = \left(\begin{smallmatrix} x & y \\ z & t \end{smallmatrix}\right) \in \mathrm{SL}(2,\mathbb{R})$ such that $S\rho_1(h)S^{-1} = \rho_2(h)$ for $h \in \pi_1\Sigma$. This would mean that S would commute with $A = \rho_1(a_2) = \rho_2(a_2)$,

$$\left(\begin{smallmatrix} x & y \\ z & t \end{smallmatrix}\right)\left(\begin{smallmatrix} a & 0 \\ 0 & a^{-1} \end{smallmatrix}\right) = \left(\begin{smallmatrix} ax & a^{-1}y \\ az & a^{-1}t \end{smallmatrix}\right) = \left(\begin{smallmatrix} ax & ay \\ a^{-1}z & a^{-1}t \end{smallmatrix}\right) = \left(\begin{smallmatrix} a & 0 \\ 0 & a^{-1} \end{smallmatrix}\right)\left(\begin{smallmatrix} x & y \\ z & t \end{smallmatrix}\right),$$

i.e., $y = z = 0$, and hence $t = x^{-1}$. But then, for $x = a_1$ we would have

$$\left(\begin{smallmatrix} x & 0 \\ 0 & x^{-1} \end{smallmatrix}\right)\left(\begin{smallmatrix} 1 & 1 \\ 0 & 1 \end{smallmatrix}\right)\left(\begin{smallmatrix} x^{-1} & 0 \\ 0 & x \end{smallmatrix}\right) = \left(\begin{smallmatrix} 1 & x^2 \\ 0 & 1 \end{smallmatrix}\right) \neq \left(\begin{smallmatrix} 1 & 0 \\ 1 & 1 \end{smallmatrix}\right).$$

We shall see that the orbits of ρ_1 and ρ_2 cannot be separated by open sets by finding a sequence that converges to both of them at the same time. To claim this we are using the characterization of Fine print 1.11: the compact-open topology when the target space is a metrizable (any manifold is) is the uniform convergence over compact subsets. As moreover the source ($\pi_1\Sigma$) is discrete, it coincides with the point convergence. We can think about the topology by looking at converging sequences, which converge if and only if they do for every point.

We first define a sequence φ_n of representations converging to ρ_1. Define φ_n by $\varphi_n(a_2) = \varphi_n(b_1) = A$,

$$\varphi_n(a_1) = \varphi_n(b_2) = \left(\begin{smallmatrix} (1+a^{-2n})^{1/2} & 1 \\ a^{-2n} & (1+a^{-2n})^{1/2} \end{smallmatrix}\right).$$

We have that φ_n converges to ρ_1, whereas

$$A^n\varphi_n A^{-n} = \left(\begin{smallmatrix} (1+a^{-2n})^{1/2} & a^{-2n} \\ 1 & (1+a^{-2n})^{1/2} \end{smallmatrix}\right)$$

converges to ρ_2, so $[\rho_1]$ and $[\rho_2]$ cannot be separated by two open sets and the orbit space is not Hausdorff.

In this situation, one can produce a Hausdorff space by identifying points that do not have disjoint open neighbourhood. This identification generates, but is not necessarily, an equivalence relation. This process gives the largest quotient that is Hausdorff. This has two problems. In general, everything could collapse to a point, like in $\mathbb{R}^2$ with the $\mathbb{R}^+$ action $\lambda \cdot (x, y) = (\lambda x, \lambda y)$. And second, it may not be so easy to understand if one thinks about collapsing orbits. A solution to both problems is the GIT (Geometric Invariant Theory) quotient: we forget about some points from the beginning in such a way that the quotient will be Hausdorff and will have "as many orbits as possible" (this is not a precise statement). Naively, we get rid of the orbits that are too bad, and from the remaining we identify the ones with non-disjoint closure. In the example of $\mathbb{R}^2$, we just have to forget about the origin in order to get a circle as a quotient. If we want to do this in our case, we should start with the completely reducible, or semisimple, representations (those that can be decomposed as a sum of irreducible ones).

Fine print 2.2. This is not the best example of a non-Hausdorff quotient to take in your pocket. Instead, consider the action of $\mathbb{R}^*$ on $\mathbb{R}^2 \setminus \{(0,0)\}$ by $(x,y) \sim (ax, a^{-1}y)$. The orbits are then hyperbolas, the x-axis without the origin, and the y-axis without the origin. The resulting topology is like the line with double origin (where the neighbourhoods of one origin do not include the other origin). A visual way to understand this is by drawing the line $y = 1$. The point $(0,1)$ represents the y-axis, and each other point represents a hyperbola. The x-axis is the second origin, which cannot be separated from y-axis by two disjoint open sets. Thus, the quotient space is not T2, but it is still T1, as all points are closed.

Fine print 2.3. It is interesting to have in mind the relation between Hausdorff and metrizable, that is, a topological space such that its topology comes from some metric. Any metrizable space is Hausdorff, as we can take balls. However, the converse is not necessarily true. Nagata-Smirnov metrization theorem states that a topological space is metrizable if and only if it is regular (a non-empty closed set and a point outside it can be separated by open neighbourhoods, axiom T_3), Hausdorff (this is not necessarily a consequence of regular) and has a base that is a union of countably many locally finite collections of open sets (this is called a σ-locally finite base).

2.3 A flat bundle for a representation

We have discussed the space of all representations and the action of the group on it. We go back now to the study of a single representation and associate a fibre bundle to it.

Recall that the universal cover $\tilde{\Sigma}$ is a principal $\pi_1\Sigma$-bundle over Σ. Given a representation $\rho : \pi_1\Sigma \to G$, we consider the product manifold $\tilde{\Sigma} \times G$ and quotient it by the right action of the group $\pi_1\Sigma$ defined by

$$(x, g) \cdot \alpha = (x \cdot \alpha, \rho(\alpha^{-1})g).$$

The resulting space

$$E_\rho = \tilde{\Sigma} \times G / \langle (x, g) \sim (x \cdot \alpha, \rho(\alpha)^{-1}g) \text{ for } \alpha \in \pi_1\Sigma \rangle$$

has a projection $\pi : E_\rho \to \Sigma$, given by $\pi : [(x, g)] \mapsto p(x)$ where $p : \tilde{\Sigma} \to \Sigma$ is the projection of the universal cover. It moreover has an action of G on the right, for $h \in G$,

$$[(x, g)] \cdot h = [(x, gh)],$$

which is well defined by the associativity of the group. The action of $\pi_1\Sigma$ is free and transitive on the fibres of the projection π, so this looks a lot like a principal G-bundle...

In order to prove this claim formally, we use trivializations, as we are going to need them anyway. Take an open set $U_i \subset \Sigma$ such that $p^{-1}(U_i)$ is a disjoint union of open sets U_i^α, each of them homeomorphic to U_i. We define the trivialization φ_i^α by

$$\begin{aligned} \varphi_i^\alpha : E_{\rho|U_i} &\to U_i \times G \\ [(x,g)] &\mapsto (p(x), g) \end{aligned}$$

where (x,g) is taken such that $x \in U_i^\alpha$. We look now at the transition maps between U_i^α and U_j^β. We will have a transition map whenever $U_{ij} := U_i \cap U_j \neq \emptyset$. Consider the subsets $U_{ij}^\alpha \subset U_i^\alpha$ and $U_{ij}^\beta \subset U_j^\beta$ that project onto U_{ij}. There exists an element $\Gamma_{ij}^{\beta\alpha} \in \pi_1\Sigma$ such that $U_{ij}^\beta \cdot \Gamma_{ij}^{\beta\alpha} = U_{ij}^\alpha$. We then have

$$\begin{array}{ccccccc} U_i \times G & \to & E_{\rho|U_i} & & E_{\rho|U_j} & \to & U_j \times G \\ (p(x), g) & \mapsto & [(x,g)] & = & [(x \cdot \Gamma_{ij}^{\beta\alpha}, \rho(\Gamma_{ij}^{\beta\alpha})^{-1} g)] & \mapsto & (p(x), \rho(\Gamma_{ij}^{\beta\alpha})^{-1} g) \end{array}$$

The transition maps are given by constant maps $\rho(\Gamma_{ij}^{\beta\alpha})^{-1} = \rho(\Gamma_{ij}^{\alpha\beta}) \in G$.

Fine print 2.4. Recall that the property that for each point in Σ there exist an open neighbourhood U whose preimage in $\tilde{\Sigma}$ consists of a disjoint union of open sets homeomorphic to U (through the projection) follows from the fact that $\tilde{\Sigma}$ is a principal $\pi_1\Sigma$-bundle, the discreteness of $\pi_1\Sigma$ and the way we defined charts for every point (just by concatenation of paths).

Fine print 2.5. It may seem we are taking too many trivializations, as we have countably infinitely many for each open set U_i, whereas one would certainly be enough. This is not an issue but an advantage, as we thus see that the fibre bundle structure is canonical, that is, does not depend on any choice.

These trivializations with smooth transition maps, together with an atlas of Σ, define an atlas for E_ρ. If we want to properly check that we get a manifold, it remains to check that the induced topology is second countable and Hausdorff. The second countability just follows from the fact that $\pi_1\Sigma$ is discrete and both Σ and G are second countable.

Proving Hausdorffness requires a bit more of work. Consider $[(x,g)]$ and $[(y,h)]$. If $p(x) \neq p(y)$, consider two separating neighbourhoods $U_x, U_y \subset \Sigma$. The open sets $\pi^{-1}(U_x), \pi^{-1}(U_y) \subset E_\rho$ separate $[(x,g)]$ and $[(y,h)]$. When $p(x) = p(y)$, we have $[(y,h)] = [(x,g')]$ for some $g' \in G$. In order to separate $[(x,g)]$ and $[(x,g')]$, we use a trivialization φ_i^α such that $x, x' \in U_i$. We separate g and g' in G by U_g and $U_{g'}$, and consider a small open neighbourhood

V_x of x. We then have that $(\varphi_i^\alpha)^{-1}(V_x \times U_g), (\varphi_i^\alpha)^{-1}(V_x \times U_{g'}) \subset E_\rho$ separate $[(x, g)]$ and $[(x, g')]$.

We have thus associated to any representation $\rho : \pi_1\Sigma \to G$ a principal G-bundle E_ρ, which has a very special property, it admits locally constant transition maps (hence constant if the domain is connected). Actually, we are getting these transition maps also from the representation, as they are the image of the map ρ.

A choice of trivializations with locally constant transition maps on a fibre bundle E is sometimes called a **flat structure** on E, and the bundle E is said to be a **flat bundle**. We next see why.

2.4 Flat bundles and flat connections

Having constant transition maps can be interpreted geometrically in terms of a connection. For each trivialization take the trivial connection of Example 1.11 and pull it back to the bundle. This will define a global smooth connection since the transition maps are constant. We can spell it out if we wish. For a trivialization $\varphi_i : E|U_i \to U_i \times G$, we have the differential of its inverse $(\varphi_i^{-1})_* : T(U \times G) \to TE$ and we define

$$HE := (\varphi_i^{-1})_*(TU \times \{0\}).$$

This is well defined globally as

$$(\varphi_j^{-1})_*(TU \times \{0\}) = (\varphi_j^{-1} \circ \varphi_i)_*(\varphi_i^{-1})_*(TU \times \{0\}) = (\varphi_i^{-1})_*(TU \times \{0\}),$$

since $\varphi_j^{-1} \circ \varphi_i$ is the product of the identity (on the open set U) times a constant map (on G, the transition map).

Conversely, a flat connection determines trivializations with locally constant transition maps. A flat connection determines a foliation integrating the horizontal distribution. One then proves that in a sufficiently small open set of U, the foliation is such that any leaf intersected with $E_{|U}$ is an open set that maps diffeomorphically to U via the projection π (this is easy for principal bundles, as we can do it around a point of E and then use the G-action). By choosing, for some point $x_0 \in U$, an identification with the typical fibre $\varphi : E_{x_0} \simeq F$. We define a trivialization, for $x \in E_{|U}$ belonging to the leave $L_x \subset E$ passing through x, by

$$x \mapsto (\pi(x), \varphi(L_x \cap E_{x_0})).$$

The trivializations thus defined have transition maps given by locally constants maps, as the group G acts freely and transitively on the leaves of the foliation. They are locally constant as it is an element g that moves a leaf to another leaf, independently over which point of the open set we are.

As we made the choice of φ, the real equivalence is between flat structures up to equivalence (by maps that must be constant on any open set) and flat connections.

Fine print 2.6. This equivalence is known in general, for any fibre bundle, as the Riemann-Hilbert correspondence. If you are interested in this for complex-analytic spaces, take a look at Conrad's notes: `math.stanford.edu/~conrad/papers/rhtalk.pdf`

Fine print 2.7. If the curvature is not flat, we can still relate the holonomy of the connection with the curvature. This is Ambrose-Singer's theorem.

2.5 A representation for a flat connection

Let E be a principal G-bundle with a flat connection. For any $x \in E$ and a path γ in M starting at $\pi(x)$, the horizontal lift of γ is the unique path $\tilde{\gamma}$ in E starting at x such that $\pi \circ \tilde{\gamma} = \gamma$ and its tangent vectors at each point lie in the horizontal distribution. If we think of the flat connection as a distribution that determines a flat structure, that is, trivializations $\{\varphi_i\}$ to $U_i \times G$ with constant transition maps, we can actually see the lift. Locally, given a point $x \in E$ with $\varphi_i(x) = (u, g)$, a path $\gamma(t)$ in U_i is lifted to $\varphi_i^{-1}(\gamma(t), g)$. These local paths glue well as the transition maps are constant.

If the path we are lifting is a loop, we will go back to the same fibre, to some point $x \cdot g$, for a uniquely defined $g \in G$, as the action of g is free and transitive. As the distribution is G-equivariant, the element g only depends on the path γ (and on $\pi(x)$), not on the starting point of the fibre.

Now the slogan, and what you will find written in many places without proof, is that the flatness of the curvature makes that the horizontal lift depend only on the homotopy class, so we get a map

$$\rho : \pi_1 \Sigma \to G.$$

This is true and in most of the references they will think of connections as covariant derivatives. This is the case of Chapter 13 of [Tau11], where a detailed proof is given.

In our approach, the flat connection is given by an integrable horizontal distribution. It is possible to then take a set of trivializations with constant transition maps. In this case, the definition above of the horizontal lift gives an algorithm to compute the holonomy representation. Cover your curves with a finite sequence of good open sets $\{U_i\}_{i=1}^r$, the holonomy is given by $g_{r\,r-1} \dots g_{32} g_{21}$. We just have to prove that this gives the identity for paths homotopic to the identity. We then cover not only the curve but the bounded component it determines with good open sets. If we do this with one open set, the monodromy is clearly the identity by the definition of the horizontal lift. For two open sets, it follows from the condition $g_{ij}g_{ji} = \mathrm{Id}$. For three open sets, it follows from the cocycle condition $g_{ij}g_{jk}g_{ki} = \mathrm{Id}$. For an arbitrary number of open sets, we do it by induction. By modifying the path in one open set, we get a homotopic path with the same holonomy that now can be covered with one less open set, allowing us to apply induction.

(The arguments above are of course best, and perhaps only, understood with the images from the lectures, which at some point in the future will appear in these lecture notes.)

Remark 2.1. Recall that by good we mean connected and simply connected with connected and simply connected intersections (you can think of balls, for instance).

In order to move from here to Higgs bundles, it will be convenient to look at vector bundle counterpart of all this. For $G = \mathrm{GL}(n, \mathbb{C})$ as a real group, we can restate the correspondence in terms of complex vector bundles (these are bundles whose fibres are complex vector spaces, but whose transition functions are not necessarily holomorphic). To each representation $\rho : \pi_1\Sigma \to \mathrm{GL}(n, \mathbb{C})$ we can associate a flat complex vector bundle $E_\rho(\mathbb{C}^n)$ or equivalentely $\tilde{\Sigma} \times_G \mathbb{C}^n$, which carries a flat connection seen as covariant derivative.

Chapter 3

The definition of Higgs bundle

Higgs bundles were not introduced from the study of surface group representations, but by doing it this way, we give a motivation and we will be easily acquainted with the Hitchin–Kobayashi correspondence.

In this chapter, we will just refer to the main ideas, very sketchily, and give some references. There are several inaccuracies, but we hope it is for the best of a first aproach to this theory. Purposely, the desired rigor of the previous chapters will turn at the end of this chapter into some sort of mild carelessness for the sake of a better understanding.

3.1 Some results on complex geometry

Higgs bundles will be defined as holomorphic objects. Apart from a mention when we talked about changes of charts or transition maps, we have not really dealt with holomorphicity. We need some results on complex geometry.

We start with complex structures on manifolds.

- Let M be a complex manifold. By pulling back to the manifold the linear complex structure on $T_z\mathbb{C}^n \cong \mathbb{C}^n$, for any $z \in \mathbb{C}^n$, we get an almost complex structure, that is, a bundle map $J \in \operatorname{End}(TM)$ such that
$$J^2 = -\operatorname{Id}.$$
The complexification $T_{\mathbb{C}}M$ can be decomposed into the $+i$ and $-i$ eigenspaces of J, namely,
$$T_{\mathbb{C}}M = T^{1,0}M \oplus T^{0,1}M.$$
We have a dual complex structure $J^* \in \operatorname{End}(T^*M)$, which analogously decomposes the differential forms into the $(1,0)$-forms, which vanish on

$T^{0,1}M$, and the $(0,1)$-forms, which vanish on $T^{1,0}M$.

$$\Omega^1_{\mathbb{C}}(M) = \Gamma(T^*_{\mathbb{C}}M) = \Omega^{1,0}(M) \oplus \Omega^{0,1}(M).$$

Thus, the operator $d : \Omega^0_{\mathbb{C}}M \to \Omega^1_{\mathbb{C}}M$ decomposes into

$$d = \partial + \bar{\partial},$$

where $\partial : \Omega^0_{\mathbb{C}}(M) \to \Omega^{1,0}(M)$ and $\bar{\partial} : \Omega^0_{\mathbb{C}}(M) \to \Omega^{0,1}(M)$. Note that we are just making this claim for $\Omega^0_{\mathbb{C}}(M)$.

- The decomposition into $(1,0)$ and $(0,1)$-forms allows us to define (p,q)-forms, $\Omega^{p,q}(M)$, the space of forms generated by wedge product of p $(1,0)$-forms and q $(0,1)$-forms. In our mind, they look like sums of expressions like

$$dz_1 \wedge \ldots \wedge dz_p \wedge d\bar{z}_1 \wedge \ldots \wedge d\bar{z}_q.$$

- If we start with an almost complex structure, $J \in \mathrm{End}(TM)$ such that $J^2 = -\mathrm{Id}$, it may not necessarily come from a holomorphic structure on the manifold. That is the case when an integrability condition is satisfied. This is stated in several ways. One is that the sections of $T^{(1,0)}$ are involutive with respect to the Lie bracket, $[\Gamma(T^{(1,0)}), \Gamma(T^{(1,0)})] = \Gamma(T^{(1,0)})$. This can be restated in terms of the so-called Nijenhuis tensor. Lastly, in terms of d, the almost complex structure given by J is integrable if we have $d = \partial + \bar{\partial}$ for any $\Omega^{p,q}(M)$.

- This type decomposition applies also to forms with values in a vector bundle. In particular, a covariant derivative

$$\nabla : \Omega^0(M,E) \to \Omega^1(M,E),$$

decomposes into $(1,0)$ and $(0,1)$ parts, $\nabla = \nabla' + \nabla''$.

For vector bundles, the notation is a bit clumsy. A “complex manifold” is a complex manifold, no doubt about it. A complex vector bundle is not a complex manifold, but a vector bundle whose fibres are complex vector spaces. The analogous concept to complex manifold for vector bundles will be holomorphic vector bundle.

- A holomorphic vector bundle $E \to M$ is a vector bundle $\pi : E \to M$ such that E, as a manifold, is a complex manifold, the base M is a complex manifold, and the map π is holomorphic. Of course, a holomorphic vector bundle is in particular a complex vector bundle.

- We will focus on complex vector bundles $E \to M$ with M a complex manifold. In this case, a holomorphic structure on E is equivalently given by a $\bar{\partial}$-operator, a linear map

$$\bar{\partial}_E : \Omega^0(M, E) \to \Omega^{0,1}(M, E),$$

satisfying the Leibniz rule and squaring to zero. This map is sometimes called a partial connection. This is a meaningful name, as it is telling us that for a complex vector bundle, in the presence of a complex structure on M, a connection such that $\nabla''^2 = 0$ is carrying in particular a holomorphic structure for E.

3.2 A funny way to recover a flat connection

We saw how surface group representations correspond to bundles with a flat connection, passing via flat structures. The equivalence between surface group representations and flat structures (trivializations with locally constant transition maps) is quite direct. The equivalence between flat structures and flat connections is a bit trickier in some more generality and is referred to as the Riemann-Hilbert correspondence.

Now we are going to relate all these representations and connections with an object called Higgs bundle, and the way to do it will not be so clear. If a holomorphic structure on E, when M is complex, is a $\bar{\partial}$-operator, something weaker than a connection on E (such that $(\nabla'')^2 = 0$), one could expect to recover a connection from a holomorphic structure, the $\bar{\partial}$-operator, plus something else. This something else would be the $(1, 0)$-part of the connection, but this is not quite interesting. There is a seemingly intricate way to do it, and precisely because it is not so clear it is actually interesting. Just keep your faith.

- Recall that a hermitian metric on a complex vector space W is a linear pairing $h(\cdot, \cdot)$ such that, for $v, w \in W$, we have $h(v, w) = \overline{h(w, v)}$.

- The geometric version of this is a hermitian metric h on a manifold: a smoothly varying hermitian metric h_p on the vector spaces $(T_pM)_\mathbb{C}$ for $p \in M$.

- It is a fact, which requires a proof, that a hermitian metric h on a holomorphic vector bundle over a complex manifold uniquely determines a compatible connection, called the Chern connection. It is the unique connection $\nabla_{h,\bar{\partial}_E}$ such that, for $X, Y \in \Gamma(E)$, we have

$$dh(X, Y) = h(\nabla_{h,\bar{\partial}_E} X, Y) + h(X, \nabla_{h,\bar{\partial}_E} Y)$$

and $\nabla'' = \bar{\partial} \otimes \mathrm{Id}_E$. Details about this can be found in Section 1.4 of [Kob14].

But we were interested in flat connections.

- The difference between any two connections is an element of $\Phi \in \Omega^1_{\mathbb{C}}(M, \operatorname{End} E)$, so a way of recovering a flat connection from the Chern connection is by choosing wisely Φ such that $\nabla_{h,\bar{\partial}_E} + \Phi$ is flat.

- The space $\Omega^1_{\mathbb{C}}(M, \operatorname{End} E)$ decomposes into $\Omega^{1,0}(M, \operatorname{End} E) \oplus \Omega^{0,1}(\operatorname{End} E)$.

- The hermitian metric allows us to produce an element of $\Omega^{0,1}(M, \operatorname{End} E)$ from an element of $\varphi \in \Omega^{1,0}(M, \operatorname{End} E)$. If we write φ as a combination of $\alpha \otimes T$ where $\alpha \in \Omega^{1,0}(M)$ and $T \in \Gamma(\operatorname{End} E)$, we define

  $$(\alpha \otimes T)^* := \bar{\alpha} \otimes T^{*_h},$$

  where $h(TX, Y) = h(X, T^{*_h}Y)$ is the adjoint endomorphism with respect to h.

- Since we want to get moreover to a holomorphic object, we will consider holomorphic $\varphi \in \Omega^{1,0}(M, \operatorname{End} E)$, which is usually denoted by $\bar{\partial}_E \varphi = 0$, and define

  $$\Phi = \varphi + \varphi^*.$$

- From a holomorphic structure on E, we would like to recover a flat connection by choosing a hermitian metric h and a holomorphic $\varphi \in \Omega^{1,0}(M, \operatorname{End} E)$ such that

  $$\nabla_{h,\bar{\partial}_E} + \varphi + \varphi^{*h}$$

  is flat.

When we say we would like we are completely honest: such a $\varphi \in \Omega^{1,0}(\operatorname{End} E)$ may simply not exist. And perhaps this method will not recover all the possible flat connections.

Fine print 3.1. If we denote by F_∇ the curvature of the connection ∇ we have

$$F_{\nabla_{h,\bar{\partial}_E} + \varphi + \varphi^{*_h}} = F_{\nabla_{h,\bar{\partial}_E}} + [\varphi, \varphi^{*_h}].$$

This together with the holomorphicity condition $\bar{\partial}_E \varphi = 0$ are known as Hitchin's equations.

We were concerned with surface group representations and then with flat connections on bundles over an orientable surface. The fact of the base being a manifold means that we can always choose a complex structure. In general, the hypothesis of having a complex structure on M is very strong. For instance, it is not known whether the six-sphere S^6 admits or not a complex structure, but we are safe for surfaces. Thus we can try to do the above process. Before we do that, we forget about the hermitian metric and put together the holomorphic objects we have got.

Definition 3.1. A **Higgs bundle** (E, φ) over a Riemann surface Σ is a holomorphic vector bundle $E \to \Sigma$ together with a Higgs field φ, a holomorphic $\varphi \in \Omega^{(1,0)}(\operatorname{End} E)$.

Remark 3.2. The bundle of holomorphic top $(n, 0)$-forms, where n is the complex dimension of the manifold, is called the canonical bundle and denoted by K. That is why the Higgs field is usually denoted by $\varphi \in H^0(\operatorname{End}(V) \otimes K)$, where H^0 stands for holomorphic sections.

3.3 What if I have a Higgs bundle?

The previous Section has described how to get a holomorphic object, which we defined as Higgs bundle, from a flat connection. What if we start with the Higgs bundle? If we trace back the process, it all boils down to the crucial choice of a hermitian metric h. It is crucial because the resulting connection will be

$$\nabla_{h,\bar{\partial}_E} + \varphi + \varphi^*,$$

and we want this connection to have zero curvature. In order words, we are solving the equation

$$F_{\nabla_{h,\bar{\partial}_E}} + [\varphi, \varphi^{*_h}], \tag{3.1}$$

for the metric h.

Fine print 3.2. When we talked about Hitchin's equations, h was fixed and φ is the unknown.

Solving equation (3.1) may and may not be possible. The surprising fact is that it depends on a topological condition on the Higgs bundle. Let us sketch what the ingredients of this definition are.

We first introduce the **degree**. We saw in Section 1.3 that an orientable surface minus a point retracts to a bunch of circles. It can be proved that the bundle retracts to the retracted base and a complex line bundle over a bunch of circles must be trivial, and hence the original bundle restricted to the base minus a point is also trivial. Thus, a complex line bundle can be described with only one set of transition maps, as the ones seen in Problem 2 of Assignment 5. The transition map can be described by a homomorphism $\mathrm{S}^1 \to \mathrm{S}^1$. These homomorphisms are always of the form $z \mapsto z^n$ for $n \in \mathbb{Z}$ (we are somehow defining the degree of a map here) and the number n is the degree of the complex line bundle. When we have a general bundle, its degree is defined as the degree of its determinant bundle, as defined in Section 1.4.

We define the **slope** of a bundle **F** as the quotient of the degree by the rank,

$$\mu(F) := \frac{\deg F}{\det F}.$$

A subbundle $F \subset E$ of a Higgs bundle (E, φ) is called φ-invariant if $\varphi(F) \subset F \otimes \Omega^{1,0}_{hol}$. A Higgs bundle is called **stable** when for each φ-invariant strict subbundle F, one has

$$\mu(F) < \mu(E).$$

Finally, a Higgs bundle is called **polystable** if it is a sum of stable bundles of the same slope.

We have defined all this to finish this section the following remarkable fact: a Higgs bundle admits solutions of (3.1) if and only if it is polystable.

3.4 The Hitchin–Kobayashi correspondence

The somehow involved process of passing from flat connections to Higgs bundles and back has a very neat version in terms of spaces of parameters, or moduli spaces. These correspond to equivalences classes, and are expressed as quotients by the group of symmetries, as the inner action of G on the representations in Section 2.2.

$$\frac{\mathrm{Hom}^{red}(\pi_1\Sigma, \mathrm{GL}(n,\mathbb{C}))}{\mathrm{GL}(n,\mathbb{C})} \cong \frac{\{\nabla \text{ red. conn.} \mid F_\nabla = 0\}}{\mathcal{G}} \simeq_{J_\Sigma} \frac{\{\text{Polystable } (E,\varphi)\}}{\mathcal{G}^{\mathbb{C}}}.$$

There actions of the real and complex gauge groups $\mathcal{G}$ and $\mathcal{G}^c$, groups of automorphism of bundles, such that there is an isomorphism at the level of moduli spaces. The Hitchin–Kobayashi correspondence refers to the second isomorphism, which depends on the choice of a complex structure J_Σ on the surface, but we include both to see the whole picture at once.

Somehow, if we started looking at the moduli space of surface group representations —which can be seen as a moduli space of flat connections—, Higgs bundles give another presentation of this moduli space, consisting of holomorphic objects.

By the way, each of these moduli spaces, although they are isomorphic, has a different name. From left to right: Betti, de Rham and Dolbeault moduli spaces.

3.5 Where to look next

We point at two directions where one uses the properties of the moduli space of Higgs bundle to say something about any of these moduli spaces.

Morse theory

Morse theory is a way to obtain topological information about a smooth manifold by means of a differentiable function satisfying certain properties, called Morse function. The typical example, which can be found almost everywhere, including Wikipedia (`https://en.wikipedia.org/wiki/Morse_theory`), is that of the torus. The height function is a Morse function and it is possible to give a cell decomposition of the manifold, by starting with a disk and adding cells everytime the Morse function has a critical point.

It seems unproportionate to get all the way to holomorphic objects if we just want something smooth, especially since the moduli space of flat connections is The point is that the moduli space of Higgs bundles has a Morse function that is defined in a very simple way: it is just the norm of the Higgs field. You can see this in page 96 of [Hit87], where the indices of the critical points of the moduli space for $n = 2$ and the dimensions of the homology groups (or Betti numbers) is explicitly computed.

HyperKähler structure

Even though we keep claiming that by passing to Higgs bundles one gets holomorphic objects, this is not the first time complex geometry appears. Actually, as we are looking at $\mathrm{GL}(n, \mathbb{C})$, one can see from the realization as a moduli space of flat connections that there is a complex structure, say I. By fixing a connection, it is essentially multiplying by i the difference of any connection with the initial connection.

Higgs bundles give also a complex structure, this time coming from the complex structure we put on the surface, which is not the same as the one we already had. Call it J. Let us say that corresponds to multiply by i the Higgs field. A very good place to see how these structures are properly defined in a simple case is Section 5 of [GX08].

Moreover, the two complex structures I and J satisfy

$$IJ = -JI$$

and hence the composition $K = IJ$ defines another complex structure (formally, an almost complex structure that happens to be integrable). But we do not have one, two or three complex structures, as any expression of the form

$$aI + bK = cK$$

for real numbers a, b, c such that $a^2 + b^2 + c^2 = 1$ defines again a complex structure. We have a 3-sphere of complex structure. The complex structure does not come on its own, but with a symplectic structure (formally, on the smooth part), which makes all these structure simultaneously complex and symplectic in a very special way, they are so-called Kähler structures. The structure that we obtain on the moduli space is a **hyperKähler** structure, which is not such a common thing to encounter in mathematics.

But that is already already a different story for a new trip...

Background material

[GX08] William M. Goldman and Eugene Z. Xia. Rank one Higgs bundles and representations of fundamental groups of Riemann surfaces. *Mem. Amer. Math. Soc.*, 193(904):viii+69, 2008.

[Kob14] Shoshichi Kobayashi. *Differential geometry of complex vector bundles*. Princeton Legacy Library. Princeton University Press, Princeton, NJ, [2014]. Freely available at `mathsoc.jp/publication/PublMSJ/PDF/Vol15.pdf`.

[Lee13] John M. Lee. *Introduction to smooth manifolds*, volume 218 of *Graduate Texts in Mathematics*. Springer, New York, second edition, 2013. Available at `link.springer.com/book/10.1007%2F978-0-387-21752-9`.

[Mic08] Peter W. Michor. *Topics in differential geometry*, volume 93 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2008. Freely available at `www.mat.univie.ac.at/~michor/dgbook.pdf`.

[Tau11] Clifford Henry Taubes. *Differential geometry*, volume 23 of *Oxford Graduate Texts in Mathematics*. Oxford University Press, Oxford, 2011. Bundles, connections, metrics and curvature.

Further reading

[Bra12] S. B. Bradlow. Introduction to Higgs bundles. Slides for the GEAR Retreat 2012 mini-course. University of Illinois at Urbana-Champaign. Available online at http://www.math.illinois.edu/~dowdall/talk_resources--GEAR2012/Bradlow/ Videos of the lectures available at http://gear.math.illinois.edu/media/retreat12/archive/?destinationID=Td-eUCb05kivUIqLpZpXDQ&categoryID=FyV2yvoTu0C-mtdsjXVcSQ&pageIndex=1&pageSize=10

[Got08] P. Gothen. Higgs bundles. Slides for the International School on Geometry and Physics: moduli spaces in geometry, topology and physics, CIEM, 2008. Available online at http://cmup.fc.up.pt/cmup/pbgothen/Higgs_Bundles_CIEM_200802.pdf

[BGGW] S.B. Bradlow, O. García-Prada, W. Goldman, A. Wienhard. Representations of surface groups: background material for AIM workshop Available online at http://www.math.illinois.edu/~dowdall/talk_resources--GEAR2012/Bradlow/some_resources/III.2-aim_intro.pdf

[BGG] S. Bradlow., O. García-Prada, P. B. Gothen. WHAT IS... a Higgs Bundle? Notices Amer. Math. Soc. 54 (2007), no. 8, 980–981. Available online at http://www.ams.org/notices/200708/tx070800980p.pdf

[Wen12] R. Wentworth. Higgs Bundles and Local Systems on Riemann Surfaces. Lecture notes for the 3rd International School on Geometry and Physics at the CRM, 2012. Geometry and Quantization of Moduli Spaces, 165–219, Springer 2016. Available online at http://www.math.umd.edu/~raw/papers/barcelona.pdf

The original papers:

[Hit87] N. J. Hitchin. The self-duality equations on a Riemann surface. *Proc. London Math. Soc. (3)*, 55(1):59–126, 1987. Available online at `https://people.maths.ox.ac.uk/hitchin/hitchinlist/Hitchin%20THE%20SELF-DUALITY%20EQUATIONS%20ON%20A%20RIEMANN%20SURFACE%20(PLMS%201987).pdf`

[Sim92] C. T. Simpson. Higgs bundles and local systems. *Inst. Hautes Études Sci. Publ. Math.*, (75):5–95, 1992. Available online at `http://www.numdam.org/article/PMIHES_1992__75__5_0.pdf`