Groups of Homotopy Spheres

Daniel A.P. Galvin

May, 2020

Department of Mathematics,
Durham University

Contents

'This piece of work is a result of my own work except where it forms an assessment based on group project work. In the case of a group project, the work has been prepared in collaboration with other members of the group. Material from the work of others not involved in the project has been acknowledged and quotations and paraphrases suitably indicated.'

Chapter 0

Introduction

In 1956, John W. Milnor published a paper entitled 'On manifolds homeomorphic to the 7-sphere', in which he proved the existence of 'exotic spheres', manifolds homeomorphic, but not diffeomorphic, to S^n. This was a very surprising result for topologists, who at the time may have regarded the smooth structure imposed on toplogical manifolds as somewhat of a formality. Milnor, along with Michel A. Kervaire, went on to construct a group of these exotic spheres, called the group of homotopy spheres, denoted Θ_n. The relationship between Θ_n and exotic spheres was provided using the h-cobordism theorem, proved by Stephen Smale in 1962. These discoveries led to a great wealth of work being done in the realm of differential topology in the latter half of the 20th century. To this day, differential topology and the theory of smooth manifolds is still an exciting and diverse area to work in for mathematicians. In the course of this project, we will develop the theory necessary to give an exposition on the group of homotopy spheres and its relationship with exotic spheres, culminating in proving the existence of Milnor's original exotic 7-spheres.

Chapters 1 and 2 are devoted to developing the required tools in differential topology. We start with the basic theory of smooth manifolds, before progressing to study smooth maps and diffeomorphisms. Chapters 3 and 4 discuss the theory of vector bundles, an incredibly important object in topology, beginning with definitions and properties in section 3, before moving on to classifying vector bundles over S^n in section 4. Chapter 5 brings the previous material together to study smooth vector bundles, which will be important for later chapters.

With the required theory developed, chapter 6 introduces homotopy spheres, manifolds homotopy equivalent to spheres, and we explore many of their properties. Particularly important is how homotopy spheres behave under the connected sum operation of smooth manifolds, which we define. Chapter 7 then introduces the h-cobordism theorem, which we use to prove the generalised Poincaré conjecture in dimensions ≥ 6. Chapter 8 then defines Θ_n and we use it to deepen our understanding of smooth structures on S^n.

Finally, chapters 9 and 10 work towards proving the existence of exotic spheres. Chapter 9 uses our work on vector bundles to give a brief introduction to characteristic classes and also introduces the signature invariant for manifolds. The highlight of chapter 9 is the Hirzebruch signature theorem, which relates these two concepts together for closed manifolds. Chapter 10 then puts together all of the work in the previous 9 chapters to prove the existence of Milnor's original exotic 7-spheres.

This project is intended to be fully readable to anyone who has learnt about algebraic topology. Topics that are taken to be understood include homology, cohomology, homotopy and the fundamental group, covering spaces and cup and cap products, along with basic notions from point set topology. The reader is also expected to be familiar and confident with concepts from linear algebra, group and ring theory and in chapter 9 there is a short instance of complex analysis.

Chapter 1
Smooth structures on manifolds

(Unless specified, the material in this chapter is based on [1].)

Our first aim will be to understand some basic differential topology with the aim to understand smooth structures on manifolds. First, we will give the definition of a topological manifold, along with some examples, before moving onto smooth manifolds and smooth structures themselves.

Definition 1.1. A *topological manifold of dimension* n is a topological space that is Hausdorff, has countable basis and has the property that every point in the space has a neighbourhood that is homeomorphic to an n-ball for a fixed $n \in \mathbb{N}$.

Topological manifolds are the generalisation of surfaces in $\mathbb{R}^3$. We may refer to the third condition in the above definition as a space being *locally Euclidean.* It would be useful to see some examples.

Example 1.2. (a) The most trivial manifold is simply Euclidean space $\mathbb{R}^n$. All manifolds are locally Euclidean but this is the only manifold which is globally Euclidean.

(b) The n-sphere given as $S^n = \{\mathbf{x} \in \mathbb{R}^{n+1} \mid \|\mathbf{x}\| = 1\}$. This simple example will actually be one of the main objects in this project.

(c) M_g, the surface of genus g, that will be familiar to readers who have studied surfaces. These are formed by successively glueing torii together.

(d) If M^m, N^n are both manifolds of dimension m and n, respectively, then $M^m \times N^n$ is a manifold of dimension $m+n$. This is because the product of the two Hausdorff spaces is similarly Hausdorff; the basis for the product topology will still be countable as the bases for M and N are both countable; and $D^m \times D^n$ is homeomorphic to D^{n+m}.

Seeing some non-examples may also be useful to see the difference between manifolds and general topological spaces. In particular, the first two conditions are somewhat technical and it is not immediately obvious how a space might fail to satisfy them whilst also being locally Euclidean.

(a) The wedge of circles $S^1 \vee S^1$. This is the space formed by identifying one point on each circle together. This space satisfies the first two conditions, but the wedge point has no neighbourhood homeomorphic to D^1.

(b) The real line with two zeroes. This is defined as $\mathbb{R}^\times \cup \{0_1\} \cup \{0_2\}$ and has a topology given by U open if $U \cap \mathbb{R}^\times$ open in $\mathbb{R}$ and if $0_i \in U$ then there exits $a < 0 < b$ with $(a, 0), (0, b) \subset U$. This satisfies the second and third conditions but fails to be Hausdorff as 0_1 and 0_2 cannot be separated by open sets.

(c) The so-called *long line* which is created by identifying the ends of two closed long rays together. A closed long ray is given by the product $[0, 1) \times \omega_1$ where ω_1 is the first uncountable ordinal. One can compare the long line to the real line by noting that the real line is made of a union of only countably many intervals. The effect of this is that while the long line is still Hausdorff and locally Euclidean, it does not admit a countable basis.

Similarly, we can define a *topological manifold with boundary* where we replace the locally Euclidean condition with the condition that every point has a neighbourhood homeomorphic to an n-ball or the half n-ball which is given by $\{(x_1, ..., x_n) \in \mathbb{R}^n | \, \|\mathbf{x}\| < 1, x_1 \geq 0\}$, for a fixed $n \in \mathbb{N}$.

Definition 1.3. A *smooth atlas* on a topological space M is a collection $\{U_\alpha, V_\alpha, \varphi_\alpha\}$ where $\{U_\alpha\}$ is an open cover of M, V_α are open subsets of $\mathbb{R}^n$ for some $n \in \mathbb{N}_0$ and $\varphi_\alpha : U_\alpha \to V_\alpha$ are homeomorphisms such that for the intersection of any two U_α, U_β the transition map:

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

is smooth (in the usual sense for real functions). The maps φ_α are called *charts*.

Definition 1.4. A *smooth structure* is a maximal smooth atlas. A *smooth manifold* is a topological manifold along with a smooth structure.

We can similarly define a *smooth manifold with boundary* in a completely analogous way where the V_α are also allowed to be open subsets of $\mathbb{R}^n_+ := \{(x_1, ..., x_n) \in \mathbb{R}^n \mid x_1 \geq 0\}$. Then a smooth manifold with boundary is a topological manifold with boundary along with a smooth structure on it. If M is a smooth manifold with boundary we will write $\mathring{M}$ for the *interior* of M which is all points in M which have neighbourhoods homeomorphic to an n-ball. Then we define the *boundary* of M to be $\partial M = M \setminus \mathring{M}$.

Lemma 1.5. *The boundary ∂M of a topological manifold with boundary M is well defined as a topological property. That is, ∂M is preserved by homeomorphisms.*

Proof. Let $f : M \to N$ be some homeomorphism. We first show that elements in the boundary of M are sent to elements of the boundary of N. Let $x \in \partial M$ and assume towards a contradiction that $f(x) \notin \partial N$. Then there exists U an open neighbourhood of $f(x)$ with $\varphi : U \to B \subset \mathbb{R}^n$ a homeomorphism where B is the open $n - ball$. Now, since f is a homeomorphism, $f^{-1}(U)$ is open and we have a homeomorphism $\varphi \circ f \mid_{f^{-1}(U)}: f^{-1}(U) \to B$ which means that x has a neighbourhood homeomorphic to an n-ball, which is a contradiction.

Now assume that $x \notin \partial M$ and assume towards a contradiction that $f(x) \in \partial N$. The argument is now identical to the above argument using that f^{-1} is a homeomorphism. ∎

Lemma 1.6. *If M is a smooth manifold with boundary of dimension n, then ∂M is a smooth manifold of dimension $n-1$.*

Proof. Let $x \in \partial M$ and let $\varphi : U \to V$ be a chart containing x. So V is an open subset of $\mathbb{R}^n_+$ and $\varphi(x) \in \mathbb{R}^{n-1} \subset \mathbb{R}^n_+$, since if it wasn't then it would imply that $X \notin \partial M$. So if $(U_\alpha, \varphi_\alpha)$ is a smooth atlas on M then $U_\alpha \cap \partial M$ is mapped by φ_α to an open subset of $\mathbb{R}^{n-1}$. Furthermore, we have that $\varphi_\alpha \mid_{U_\alpha \cap \partial M}$ is a homeomorphism onto its image. All we have left to check is that the corresponding transition maps are smooth. This follows since if we take a transition map for M $\varphi_\beta \circ \varphi_\alpha^{-1}$, which is smooth, then the corresponding transition map for ∂M is simply the restriction of the original transistion map to $\varphi_\alpha(U_\alpha \cap U_\beta) \cap \varphi(\partial M) \subset \mathbb{R}^{n-1}$ which is therefore smooth as a map to $\mathbb{R}^{n-1}$. □

From now on we will call both smooth manifolds with and without boundary smooth manifolds for simplicity. In the cases where it matters, we shall specify that $\partial M = \emptyset$ by saying that M is *closed*.

If we assume $\mathbb{R}^n$ has the standard orientation, then one can see that every chart φ_α determines a generator of $H^n(M, M \backslash p)$ for each point $p \in U_\alpha$. This is what is known as a *local orientation*. If all of these local orientations are compatible, which is equivalent to saying that the determinant of the Jacobian for all of the transistion maps is positive, then we say that we have an *oriented smooth atlas* or an *oriented smooth structure*. If M is an oriented smooth manifold (i.e. one with an oriented smooth structure) then we will write $-M$ for the manifold with the opposite orientation.

For oriented smooth manifolds we have some nice properties for their homologies. We will state these without proof, but you can find the proofs for these in [2].

Theorem 1.7. *Let M be a closed, connected, oriented n-dimensional smooth manifold. Then $H_n(M;\mathbb{Z}) \cong \mathbb{Z}$ and $H_k(M,\mathbb{Z}) = 0$ for all $k > n$.*

The generator of the highest homology group is called the *fundamental class* of M and is written $[M]$. For manifolds with boundary, an analogous statement holds where M has a relative fundamental class $[M, \partial M]$ that generates $H_n(M, \partial M; \mathbb{Z})$.

Theorem 1.8. *Let M be a closed, connected, orientable smooth n-dimensional smooth manifold. Then the map $PD : H^k(M;\mathbb{Z}) \to H_{n-k}(M;\mathbb{Z})$ that sends a cocycle φ to $[M] \frown \varphi$ is an isomorphism.*

This statement is known as *Poincaré duality*. Once again, an analogous version exists for smooth manifolds with boundary. The difference is that if $\partial M = A \cup B$, then the isomorphism is given as $PD : H^k(M, A; \mathbb{Z}) \to H_{n-k}(M, B; \mathbb{Z})$ and cap product is taken with the relative fundamental class instead.

With topological spaces we have the notion of a homeomorphism to act as an equivalence relation between them. For smooth manifolds this role will be played by the notion of a diffeomorphism, which we now begin the process of defining.

Definition 1.9. Let $f : M^m \to N^n$ be a map between smooth manifolds that have atlases

$$\{U_\alpha, \varphi_\alpha(U_\alpha), \varphi_\alpha\}_{\alpha \in \mathfrak{A}}, \ \{W_\beta, \psi_\beta(W_\beta), \psi_\beta\}_{\beta \in \mathfrak{B}}$$

and define $X_{\alpha\beta} \subset \mathbb{R}^m$ and $Y_{\alpha\beta} \subset \mathbb{R}^n$ as:

$$X_{\alpha\beta} = \varphi_\alpha(U_\alpha \cap f^{-1}(W_\beta \cap f(U_\alpha)), \quad Y_{\alpha\beta} = \psi_\beta(W_\beta \cap f(U_\alpha)).$$

Then we say that f is *smooth* if

$$\psi_\beta \circ f \circ \varphi_\alpha^{-1} : X_{\alpha\beta} \to Y_{\alpha\beta} \text{ is smooth } \forall \alpha \in \mathfrak{A}, \forall \beta \in \mathfrak{B}.$$

This definition is essentially saying that a smooth map between M and N is one that induces smooth maps between the charts of M and N. Also, note that since φ_α and ψ_β are homeomorphisms and the composition $\psi_\beta \circ f \circ \varphi_\alpha^{-1}$ is necessarily continuous, f itself must also be continuous.

Definition 1.10. Let M, N be smooth manifolds. A map $f : M \to N$ is a *diffeomorphism* if it is smooth and has smooth inverse.

This gives us our equivalence relation on smooth manifolds. Note that diffeomorphisms are necessarily homeomorphisms since smooth maps are also continuous maps. This means that our equivalence relation preserves not just the smooth structure but also the underlying topological manifold. It is not immediately obvious that this extra equivalence relation is strictly necessary: one might think that given a topological manifold M with two smooth structures $\mathscr{S}$, $\mathscr{T}$, the two resulting smooth manifolds are diffeomorphic if and only if $\mathscr{S} = \mathscr{T}$. This is not correct, as we demonstrate with an example.

Example 1.11. Consider the manifold $\mathbb{R}$. Since we can cover $\mathbb{R}$ with a single chart, we can specify a smooth atlas by giving a homeomorphism $\varphi : \mathbb{R} \to \mathbb{R}$. If we take another smooth atlas defined by a homeomorphism ψ, this new smooth atlas will only be compatible with the previous one if $\varphi \circ \psi^{-1}$ is smooth. It is easy to see that this is not always true; for example, we could take φ the identity map and ψ any continuous but not smooth bijection. (Once such example would be the map equal to the identity on $x < 0$ but equal to the map x^2 on $x \geq 0$.) This means that the smooth structures defined by these atlases are not equal. However, they are diffeomorphic. If we take $f = \psi^{-1} \circ \varphi$ then the map induced through the charts is simply the identity map which is always smooth and clearly has a smooth inverse.

Since we will be studying spheres throughout the project, it will be useful to determine that they do admit a smooth structure. This smooth structure is what is known as the standard smooth structure on S^n.

Proposition 1.12. *S^n is a smooth manifold of dimension n with smooth structure given by the atlas with charts:*

$$U_i^{\pm} = \{(x_1, \dots, x_{n+1}) \in S^n | \pm x_i > 0\},$$
$$\varphi_i^{\pm} : (x_1, \dots, x_{n+1}) \mapsto (x_1, \dots, \hat{x}_i, \dots, x_{n+1}) \in \mathbb{R}^n.$$

Proof. It suffices to show that the given charts above form a smooth atlas on S^n. U_i^+ is disjoint from U_i^- so it suffices to check the transisition map for $U_i^{\pm} \cap U_j^{\pm}$ when $i \neq j$. Consider $U_i^+ \cap U_j^+, i < j$. Then the map $\varphi_j \circ \varphi_i^{-1}$ is given as

$$(x_1, \dots, \hat{x}_i, \dots, x_{n+1}) \to (x_1, \dots, x_{i-1}, \sqrt{1 - \sum_{\substack{k=1 \\ k \neq j}}^{n} x_k^2}, x_{i+1}, \dots, \hat{x}_j, \dots, x_{n+1}).$$

This is just a combination of square roots, polynomials and projections which means that it is a smooth map. All of the other cases are completely analogous, and hence these charts form a smooth atlas. □

An obvious question is: given a topological manifold M, can we find smooth structures $\mathscr{S}$, $\mathscr{T}$ such that the resulting smooth manifolds are not diffeomorphic? An equivalent way of stating this question is: given a smooth manifold M, can we find another smooth manifold that is homeomorphic to M, but not diffeomorphic to M? A key question that we will answer by the end of this project is whether there exists a manifold that is homeomorphic to S^n, but not diffeomorphic to S^n with the standard smooth structure. Any such smooth manifold will be called an *exotic sphere*, and it will take a great amount of work to prove one exists.

Chapter 2

Tangent space and diffeomorphisms on manifolds

(Unless specified, the material in this chapter is based on [1] and the Riemannian geometry lecture course given by P. Tumarkin at Durham University 2019-2020.)

We now turn to studying the set of all diffeomorphisms of a manifold M, which forms a group under composition. This object is of great importance in its own right but we only study it here because we will be interested in a certain quotient of it later. Before we can do that, we have to improve our understanding of smooth maps by introducing the concept of a *tangent space*. We denote the space of all smooth maps to $\mathbb{R}$ defined in a neighbourhood of $p \in M$ as $C^\infty(M, p)$.

Definition 2.1. Let M be a smooth manifold. We define the *tangent space* of M at $p \in M$, denoted T_pM, to be the set of all linear maps $\delta : C^\infty(M, p) \to \mathbb{R}$ such that

$$\delta(fg) = f(p)\delta(g) + \delta(f)g(p).$$

The linear maps δ are known as *derivations* or *tangent vectors*. Sometimes we will write $D^\infty(M, p)$ for the set of all derivations on a manifold M at a point $p \in M$.

This is a purely theoretical description, so why are they called tangent vectors? First, note that $T_p(M)$ is a vector space. Then we will characterise $T_p(M)$ by showing that it is equivalent to the set of all directional derivatives, which we define now.

Definition 2.2. Let $\gamma : (-1, 1) \to M$ be a curve in a smooth manifold with $\gamma(0) = p$ and let $f \in C^\infty(M, p)$. A *directional derivative* of f along γ at p is

$$\gamma'(0)(f) = \lim_{t \to 0} \frac{f(\gamma(t)) - f(\gamma(0))}{t} = (f \circ \gamma)'(0).$$

Lemma 2.3. *Directional derivatives are derivations.*

Proof. This is simply an algebraic trick. Let γ be as above and consider f, $g \in C^\infty(M, p)$. Then

$$\begin{aligned}
\gamma'(0)(fg) =& \lim_{t \to 0} \frac{fg(\gamma(t)) - fg(\gamma(0))}{t} \\
=& \lim_{t \to 0} \frac{fg(\gamma(t)) - f(\gamma(t))g(\gamma(0)) + f(\gamma(t))g(\gamma(0)) - fg(\gamma(0))}{t} \\
=& \lim_{t \to 0} \frac{f(\gamma(t))[g(\gamma(t)) - g(\gamma(0))] + g(\gamma(t))[f(\gamma(t) - f(\gamma(0))]}{t} \\
=& f(p)\gamma'(0)(g) + \gamma'(0)(f)g(p)
\end{aligned}$$

as required. □

It is clear that directional derivatives are in one to one correspondence with equivalence classes of curves under the relation $\gamma_1 \sim \gamma_2$ if and only if $\gamma_1'(0) = \gamma_2'(0)$ (for curves $\gamma_1(0) = \gamma_2(0) = p$).

Given a point $p \in M$ and a curve γ passing through p, we know that our point p lies inside an open U, with $\varphi : U \to \mathbb{R}^n$ a chart. Since we only care about the behaviour of γ close to p, we may assume that the image of γ lies entirely inside U also. Then we consider *coordinate curves* $\gamma_i(t) = \varphi^{-1}(\varphi(t) + te_i)$ where e_i is an element of the standard basis for $\mathbb{R}^n$ and write $\frac{\partial}{\partial x_i} \mid_p$ for the directional derivative corresponding to $\gamma_i'(0)$.

Lemma 2.4. *Directional derivatives are equivalent to linear combinations of* $\frac{\partial}{\partial x_i} \mid_p$.

Proof. We start by showing that every directional derivative is a linear combination of $\frac{\partial}{\partial x_i} \mid_p$. Let $p \in M$, $f \in C^\infty(M,p)$ and $\gamma : (-1,1) \to M$ a curve with $\gamma(0) = p$ and assume that γ lies entirely in a single chart $\varphi : U \to V$. Write $\varphi(q) = (x_1(q), x_2(q), ...x_n(q))$. Then

$$\gamma'(0)(f) = (f \circ \gamma)'(0)$$
$$= (f \circ \varphi^{-1} \circ \varphi \circ \gamma)'(0) = ((f \circ \varphi^{-1}) \circ (\varphi \circ \gamma))'(0).$$

Using the chain rule, we see that this last expression equals

$$\sum_{i=1}^{n} \frac{\partial(f \circ \varphi^{-1})}{\partial x_i}(\varphi \circ \gamma(0)) \cdot (x_i \circ \gamma)'(0),$$

which is exactly what we want as $\frac{\partial(f\circ\varphi^{-1})}{\partial x_i}(\varphi \circ \gamma(0)) = \frac{\partial}{\partial x_i} \mid_p (f)$.

Now we need to show that any linear combination of $\frac{\partial}{\partial x_i} \mid_p$ is a directional derivative. We can specify a linear combination by giving a set of numbers $v_i \in \mathbb{R}$, where the specified linear combination is then $\sum_{i=1}^n v_i \frac{\partial}{\partial x_i}$. We show that there exists a directional derivative equal to this linear combination. First, define $\tilde{\gamma}(t) = \sum_{i=1}^n (v_i e_i)t$ and then the curve of interest is $\gamma(t) = \varphi^{-1}(\tilde{\gamma}(t))$. Then by an argument very similar to the above we get that $\gamma'(0)(f) = \sum_{i=1}^n v_i \frac{\partial}{\partial x_i}(f)$ as required. ☐

Lemma 2.5. *Derivations are directional derivatives.*

Proof. Pick $\delta \in D^\infty(M,p)$ for some $p \in M$ a smooth manifold. We want to find a curve γ such that $\gamma'(0)(f) = \delta(f)$ for all $f \in C^\infty(M,p)$. Let $\varphi : U \to V$ be the chart containing p. Now let $\pi_i : V \to \mathbb{R}$ be the map projecting onto the ith coordinate. Then $\pi_i \circ \varphi \in C^\infty(M,P)$ and let $v_i = \delta(\pi_i \circ \varphi(p))$. By the previous lemma we can find a curve γ that is contained entirely within U such that $\gamma(0) = p$ and has $\gamma'(0) = \sum_{i=1}^n v_i \frac{\partial}{\partial x_i} \mid_p$. We want to show that $\gamma'(0)(f) = \delta(f)$ for all $f \in C^\infty(M,p)$.

Before beginning the next stage, note that we may easily assume that $\varphi(p) = 0$ by composing φ with some translation in $\mathbb{R}^n$. Recall from calculus Taylor's theorem, which when applied to the function $f \circ \varphi^{-1}$ gives us (after some elementary calculations):

$$(f \circ \varphi^{-1})(\sum_{i=1}^{n} v_i t e_i) = f(p) + \sum_{i=1}^{n} v_i h_i(t)$$

for some smooth functions h_i, provided that t is small enough. Now since we have that $\varphi^{-1}(\sum_{i=1}^n v_i t e_i) = \gamma(t)$, we have that $\gamma'(0)(f)$ is given by

$$\gamma'(0)(f) = (f \circ \gamma)'(f) \mid_{t=0} = (f(p) + \sum_{i=1}^n v_i h_i(t))' \mid_{t=0} = \sum_{i=1}^n v_i h_i(0).$$

Now if we take any $q \in M$ sufficiently close to p, we get that

$$f(q) = (f \circ \varphi^{-1}(\varphi(q)) = f(p) + \sum_{i=1}^n (\pi_i \circ \varphi)(q) h_i(\varphi(q))$$

which means that if we apply our derivation δ to f we get

$$\delta(f) = \sum_{i=1}^n \delta(\pi_i \circ \varphi)(h_i \circ \varphi(p)) + (\pi_i \circ \varphi(p))\delta(h_i \circ \varphi).$$

But this means we are done, since $\delta(\pi_i \circ \varphi) = v_i$, $h_i \circ \varphi(p) = h_i(0)$ and $\pi_i \circ \varphi(p) = 0$ by assumption. This means that $\delta(f) = \sum_{i=1}^n v_i h_i(0)$ which completes the proof. □

Putting the previous three lemmas together gives us our characterisation of the tangent space.

Proposition 2.6. $T_p(M) = \{$*directional derivatives*$\} = \langle \frac{\partial}{\partial x_i} \mid_p \rangle$.

Note that this shows that $T_p M$ is a real vector space of dimension n. We now see that smooth maps induce linear maps on the tangent space.

Definition 2.7. Let $f : M \to N$ be a smooth map of smooth manifolds. Then the *differential* of f at $p \in M$ $Df_p : T_p M \to T_{f(p)} N$ is given by

$$Df_p(\gamma'(0)) = (f \circ \gamma)'(0)$$

with γ a smooth curve such that $\gamma(0) = p$.

We now give some properties of the differential, including that it is well defined (does not depend on choice of curve γ.)

Lemma 2.8. *Df_p is well defined and linear. Furthermore, if f, g are two smooth maps of smooth manifolds, then $D(g \circ f) = Dg \circ Df$.*

Proof. For well-definedness, note that $Df_p(\gamma'(0))$ is well defined if $Df_p(\gamma'(0))(g)$ does not depend on the choice of curve γ for all $g \in C^\infty(N, f(p))$. Now

$$\begin{aligned} Df_p(\gamma'(0))(g) &= ((f \circ \gamma)'(0))(g) = (g \circ f \circ \gamma)'(0) \\ &= ((g \circ f) \circ \gamma)'(0) = \gamma'(0)(g \circ f), \end{aligned}$$

which clearly does not depend on the choice of γ. Hence Df_p is well defined.

Now consider $D(g \circ f)_p(\gamma'(0))$. By definition this is given by

$$Dg_{f(p)}((f \circ \gamma)'(0)) = Dg_{f(p)} \circ Df_p(\gamma'(0)),$$

which proves the relationship $D(g \circ f) = Dg \circ Df$.

Finally, we have to show that this map is linear. To do so, we will show that for any chart $\varphi : U \to \mathbb{R}^n$, the map $D\varphi$ is linear. Although the differential was not defined at the time, our proof of the second part of 2.4 showed that $D\varphi^{-1}(0)$ is linear. Furthermore, it showed that the map $D\varphi^{-1}(0)$ was an isomorphism of vector spaces, since $ker(D\varphi^{-1}) = \{0\}$. Inverses of linear isomorphisms are themselves linear, so $D\varphi(p)$ must also be linear. Now, since the differential commutes with compositions, we have that for charts φ and ψ on M and N and containing p and $f(p)$, respectively

$$Df = D(\psi^{-1}) \circ D(\psi \circ f \circ \varphi^{-1}) \circ D(\varphi).$$

This means that Df is the composition of three maps where the first and the last we just showed are linear, and the middle map is the classical differential for a smooth map $\mathbb{R}^n \supset U \to \mathbb{R}^m$ and so must be linear. Since the composition of linear maps is linear, this proves that Df is a linear map.

□

The differential gives us extra properties that we can impose on smooth maps and we will use it to extend the definition of topological embeddings to smooth manifolds. Recall that a topological embedding is a map that is a homeomorphism onto its image.

Definition 2.9. Let $f : M \to N$ be a smooth map of smooth manifolds. Then we say that f is an *immersion* if Df is everywhere injective. Similarly, we say that f is a *submersion* if Df is everywhere surjective. We say that f is a (smooth) *embedding* if f is a topological embedding and an immersion.

We would now like to define a notion of equivalence for embeddings. This will be analogous to the notion of homotopy, except we will want our maps to be embeddings throught the 'homotopy'. We make this clear now.

Definition 2.10. Let f, $g : M \to N$ be two embeddings. We say that f and g are *isotopic* if there exists a smooth map $I : M \times [0,1] \to N$ that satisfies

(a) $I(x,0) = f(x)$, $I(x,1) = g(x)$;

(b) $I(x,t)$ is an embedding for all $t \in [0,1]$.

This smooth map I is then called an *isotopy*.

Sometimes we will actually want an even stronger notion of equivalence than isotopy that also takes into account the ambient space of our embeddings.

Definition 2.11. Let $f,\ g : M \to N$ be two embeddings. We say that f and g are *ambient isotopic* if there exists an isotopy $I : N \times [0,1] \to N$ that satisfies

$$I(x,0) = x,\ I(f,1) = g.$$

(The second equation here means that I_1 maps the image of f to the image of g.) This isotopy is then called an *ambient isotopy*.

To see that this stronger condition is sometimes useful, we look at an example.

Example 2.12. Consider two embeddings $\mathbb{R} \to \mathbb{R}^2$, one whose image is the unit circle missing the point $(0,1)$ and another whose image is the unit interval. These two embeddings are isotopic (simply unfurl the circle) but not ambiently isotopic. If they were, then their complements would be homeomorphic but this is not the case. If we remove one point from the complement of the unit interval, then the resulting space is always connected. Whereas, if we take the complement of the circle missing a point and remove the point $(0,1)$ we have a disconnected space. Since restrictions of homeomorphisms are homeomorphisms and connected components are preserved by homeomorphisms, this means that the two complements cannot be homeomorphic.

We now state (without proof) an important result called the *disc theorem* due to Richard Palais, which we will be instrumental for the next topic of discussion. For a proof, see [1].

Theorem 2.13. *Let $f,\ g : D^k \to M^n$ be two embeddings . If $k = n$ and $f,\ g$ are equioriented (i.e. both orientation preserving or both orientation reversing) then f is ambient isotopic to g.*

Definition 2.14. Let M be a smooth manifold. Then we define Diff M to be the set of all orientation preserving diffeomorphisms from M to itself.

Lemma 2.15. Diff M *is a group under composition.*

Proof. The only non-trivial condition to check is that Diff M is closed under composition. For this, all we have to show is that if we have an atlas $\{U_\alpha, \varphi_\alpha\}$ then $\varphi_\alpha \circ g \circ f \circ \varphi_\beta^{-1}$ is smooth (and then similarly for f^{-1}, g^{-1}.) But this is not hard to show as we can find a chart φ_γ such that our map becomes

$$\varphi_\alpha \circ g \circ \varphi_\gamma^{-1} \circ \varphi_\gamma \circ f \circ \varphi_\beta^{-1}$$

which is the composition of two classically smooth functions, hence smooth, since f and g are both smooth. The argument is the same for the inverses. □

Consider the subset $\text{Diff}_0\ M \subset \text{Diff}\ M$ defined as the diffeomorphisms that are isotopic to the identity. This is a subgroup since isotopies can be composed together which means that it is closed under composition, it clearly contains the identity, and f^{-1} is isotopic to the identity if and only if f is.

Now recall that the commutator subgroup of a group G, denoted $[G,G]$ is the normal subgroup generated by the elements of the form $ghg^{-1}h^{-1}$ for all g, $h \in G$. We also have the fact that a quotient group G/N is Abelian if and only if $[G,G] \subset N$. We now prove a lemma relating the commutator subgroup of DiffS^{n-1} and Diff$_0 S^{n-1}$.

Lemma 2.16. $[\text{Diff } S^{n-1}, \text{Diff } S^{n-1}] \subset \text{Diff}_0 \ S^{n-1}$.

Proof. Let f, $g \in$ DiffS^{n-1} and let the northern and southern hemispheres of S^{n-1} be D_+ and D_- respectively. We now use the disc theorem 2.13 to construct $\tilde{f}$ and $\tilde{g}$ isotopic to f and g which are equal to the identity on D_+ and D_-, respectively. To use the disc theorem, we need two embeddings of the disc. For the first, note that $f\mid_{D_+}: D_+ \to S^{n-1}$ is an embedding of the disc in S^{n-1}. For the second, take simply the inclusion map $i : D_+ \to S^{n-1}$, which is clearly an embedding. Now the disc theorem tells us that these two embeddings are ambiently isotopic, which means we have an isotopy $I : S^{n-1} \times [0,1] \to S^{n-1}$ with

$$I(x,0) = x, \ I(D^+,1) = f(D^+).$$

Note that this means that $I(D_-,1) = f(D_-)$ as sets, but not necessarily pointwise. Let $h : D_- \to D_-$ be a diffeomorphism that fixes the boundary of D_- such that

$$I(h(x),1) = f(x) \text{ for } x \in D_-.$$

Then, using this, define the following isotopy:

$$\tilde{I}(x,t) = \begin{cases} I(x,t) & \text{if } x \in D_+ \\ I(h(x),t) & \text{if } x \in D_-. \end{cases}$$

Note that this is a well-defined isotopy since on the intersection of the two hemispheres $h(x) = x$. Then we have that $\tilde{I}(x,1) = f(x)$ and we define $\tilde{f}(x) := \tilde{I}(x,0)$. This gives us $\tilde{f}$ isotopic to f which is equal to the identity map on D_+. The argument to show the existence of $\tilde{g}$ is identical.

Now the commutator of f and g, $fgf^{-1}g^{-1}$ is isotopic to the commutator of $\tilde{f}$ and $\tilde{g}$ since we can compose the respective isotopies in turn. Now clearly $\tilde{f}$ and $\tilde{g}$ commute, so their commutator is the identity map. So we have shown that the commutator of any two elements is isotopic to the identity, and hence this completes the proof. □

We now consider the restriction homomorphism $\partial : \text{Diff } D^n \to \text{Diff } S^{n-1}$. The image of this homomorphism is exactly the diffeomorphisms of S^{n-1} which we can extend to diffeomorphisms over D^n.

Lemma 2.17. $\text{Diff}_0 \ S^{n-1} \subset \partial \, \text{Diff} \, D^n$.

Consider $f \in \text{Diff}_0\ S^{n-1}$; we want to show that it extends to a diffeomorphism over D^n. Since $f \in \text{Diff}_0\ S^{n-1}$ we have an isotopy $I : S^{n-1} \times [0,1] \to S^{n-1}$ where $I(x,0) = x$ and $I(x,1) = f(x)$. The idea here is to use the isotopy itself to perform this extension. That is, by writing every point in D^n as rx where $r \in [0,1]$ and $x \in S^{n-1}$, we can define the extension $F(x) : D^n \to D^n$ as

$$F(rx) = rI(x,r).$$

This clearly extends f, as by setting $r = 1$ we get $F(x) = I(x,1) = f(x)$. One might be worried that this will not be a diffeomorphism at $x = 0$, but we can avert this issue by reparameterising our isotopy to perform the old isotopy from $t = 0$ to $t = 0.9$ and then simply be the identity for $t = 0.9$ to $t = 1$.

Note that putting together the previous two lemmas gives us that both $\text{Diff}_0\ S^{n-1}$ and $\partial\text{Diff}\ D^n$ are normal as any subgroup that contains the commutator subgroup is normal. Let Γ_n be defined as the quotient $\text{Diff}\ S^{n-1}/\partial\text{Diff}\ D^n$.

Proposition 2.18. *Γ^n, the group of diffeomorphisms of S^{n-1} modulo those which extend over D^n is an Abelian group.*

Proof. We noted before that any quotient group G/N is Abelian if the commutator subgroup is contained in N. This is shown by the previous two lemmas. □

We will return to this group later in 8.

Chapter 3
Vector bundles

(Unless specified, the material in this chapter is based on [3] and [4].)

In the previous section we saw the tangent space, where we assigned to every point p in a smooth manifold M a vector space T_pM. This is a specific example of a more general concept called a *vector bundle* which we will now define and study. Vector bundles are of great importance in various settings and we will see some of their applications in differential and algebraic topology in this project.

Definition 3.1. Let E and B be topological spaces. An n-dimensional *vector bundle* with *total space* E over a *base space* B is a continuous map $p : E \to B$ that satisfies the following conditions:

(a) For all $x \in B$, $p^{-1}(x)$ has the structure of a real n-dimensional vector space.

(b) There exists $\{U_\alpha\}$ an open cover of B with homeomorphisms

$$\varphi_\alpha : p^{-1}(U_\alpha) \to U_\alpha \times \mathbb{R}^n$$

such that the following diagram commutes:

$$\begin{array}{ccc} p^{-1}(U_\alpha) & \xrightarrow{\varphi_\alpha} & U_\alpha \times \mathbb{R}^n \\ {\scriptstyle p}\downarrow & \swarrow{\scriptstyle \pi_1} & \\ U_\alpha & & \end{array}$$

where π_1 is the map that projects onto the first coordinate.

(c) The restriction of every homeomorphism φ_α to $p^{-1}(x)$ for any $x \in U_\alpha$,

$$\varphi_\alpha \mid_{p^{-1}(x)} : p^{-1}(x) \to \{x\} \times \mathbb{R}^n,$$

is an isomorphism of vector spaces. These $p^{-1}(x)$ are known as the *fibres* of the vector bundle.

(d) $\{U_\alpha\}$ is maximal with respect to the above properties.

Note that the fact that φ_α maps $p^{-1}(x)$ to $\{x\} \times \mathbb{R}^n$ in c) is a consequence of b). Often when we refer to a vector bundle we will refer to just the total space for simplicity, but that does not mean that the projection map p is not of critical importance. Also note that d) is similar to the maximality condition for smooth structures we saw earlier, in that it mostly a technical one. A vector bundle can be specified by giving some open cover $\{U_\alpha\}$ since it must be contained inside some maximal cover. We now give some examples.

Example 3.2. (a) The *trivial bundle* over a base space B is the space $B \times \mathbb{R}^n$ where the projection map is simply given by $p : (x, \mathbf{v}) \mapsto x$.

(b) Take the unit interval $I = [0, 1]$ and form the trivial line bundle $I \times \mathbb{R}$. We can then quotient this space by the relation $(0, v) \sim (1, -v)$ to form the Möbius bundle, which is homeomorphic to a Möbius strip with its boundary removed. A sufficient open cover to define the vector bundle structure here would be $\{(0, 0.2) \cup (0.4, 1), (0, 0.6) \cup (0.8, 1)\}$. The associated φ_α would then be the trivial maps on each of the two elements in the cover.

(c) Let M be a n-dimensional smooth manifold. Then the *tangent bundle* $TM := \bigcup_{p \in M} T_pM$ with the projection map simply sending each element in T_pM to p is a vector bundle of dimension n. We will return to this in 5.

We can similarly define *orientable vector bundles* to be those vector bundles that we can assign an orientation to each fibre $p^{-1}(x)$ such that for every $x \in B$, every φ_α defined at x maps the orientation on $p^{-1}(x)$ to the standard orientation of $\{x\} \times \mathbb{R}^n$. The Möbius bundle given above is an example of a non-orientable vector bundle.

We now need a notion of equivalence for vector bundles. We will see that there are two such notions, *isomorphism* and *stable isomorphism*. As one might presume, stable isomorphism is a weaker notion than isomorphism but is sometimes useful to consider. We start with isomorphisms and will look at stable isomorphisms later.

Definition 3.3. Let $p_1 : E_1 \to B$ and $p_2 : E_2 \to B$ be two vector bundles (over the same base space B). A vector bundle *morphism* is a continuous map $f : E_1 \to E_2$ such that fibres of E_1 are sent to fibres of E_2 and the consequential restriction map

$$f\mid_{p_1^{-1}(x)} : p_1^{-1}(x) \to p_2^{-1}(x)$$

is a linear map. Similarly, an *isomorphism* of vector bundles E_1, E_2 is a morphism $f : E_1 \to E_2$ such that f is a homeomorphism and the restriction maps given above are isomorphisms of vector spaces. We then say that E_1 and E_2 are *isomorphic* vector bundles.

Actually, if f is a continuous map, then the other conditions imply are that f is a homeomorphism. We show this in the following lemma.

Lemma 3.4. *Let $f : E_1 \to E_2$ be a continuous map between vector bundles $p_1 : E_1 \to B$ and $p_2 : E_2 \to B$ such that fibres of E_1 are sent to fibres of E_2 by an isomorphism of vector spaces. Then f is an isomorphism of vector bundles.*

Proof. We need to show that f is a bijection and that its inverse is continuous. Let $e_2 \in E_2$, then there exists an $x \in B$ such that $e_2 \in p_2^{-1}(x)$. Since f maps $p_1^{-1}(x)$ to $p_2^{-1}(x)$ by a linear isomorphism and $f^{-1}(p_2^{-1}(x)) = p_1^{-1}(x)$, there must exist a unique $e_1 \in E_1$ such that $f(e_1) = e_2$. This means that f is a bijection. To show that f^{-1} is continuous, it sufficies to show that it is continuous locally on E_2. Let $x \in B$, and let $\{U_\alpha\}$, $\{U_\beta\}$ be the associated covers for the two vector bundles. Then by maximality, there exists some open $U \ni x$ that is in both of these covers and hence $p_1^{-1}(U)$ and $p_2^{-1}(U)$ are both homeomorphic to $U \times \mathbb{R}^n$. Then it suffices to check that the induced map $\tilde{f} : U \times \mathbb{R}^n \to U \times \mathbb{R}^n$ has a continuous inverse. By hypothesis, this map must send $(x, \mathbf{v})$ to $(x, h_x(\mathbf{v}))$, where $h_x \in GL_n(\mathbb{R})$. So the inverse map $\tilde{f}^{-1}$ sends $(x, \mathbf{v})$ to $(x, h_x^{-1}(\mathbf{v}))$. Now the continuity of f, and hence the continuity of $\tilde{f}$, implies that the entries in the matrix h_x must continuously depend on x. This in turn implies that the entries of the inverse matrix h_x^{-1} must also depend continuously on x. So $\tilde{f}^{-1}$ is continuous on $p_2^{-1}(U)$. Therefore, f^{-1} is continuous on all of E_2. □

Definition 3.5. Let $p : E \to B$ be a vector bundle. Then a *section* of E is a continuous map $s : B \to E$ such that $p \circ s : B \to B$ is the identity map.

Another way of viewing this definition is that s assigns to every element of B an element in the corresponding fibre. One such example of a section is the *zero section* of a vector bundle which is defined by mapping every element $x \in B$ to the zero vector in $p^{-1}(x)$.

Example 3.6. We can use the concept of sections to prove that the trivial bundle $S^1 \times \mathbb{R}$ is not isomorphic to the Möbius bundle. Remove the zero section from both bundles and the trivial bundle is no longer connected but the Möbius bundle still is. The former is not connected as we can split it into two sets $S^1 \times \mathbb{R}^+$ and $S^1 \times \mathbb{R}^-$ which are both open and disjoint. To see that the latter is still connected, note that we can still find paths to the 'other side' of the zero section by doing a full loop of the S^1 factor. Now if these two vector bundles were isomorphic, then the isomorphism would preserve the zero section since the fibres are taken to fibres by isomorphisms. This would mean that the spaces remaining after removing the zero sections would be homeomorphic. But homeomorphisms preserve connectedness and hence these two vector bundles cannot be isomorphic.

Now we see a lemma that we can use to characterise vector bundles which are isomorphic to trivial bundles, which is sometimes useful to have.

Lemma 3.7. *Let $p : E \to B$ be an n-dimensional vector bundle. $p : E \to B$ is isomorphic to the trivial bundle $B \times \mathbb{R}^n$ if and only if there exist n sections $s_1, ..., s_n$ such that for all $x \in B$, $s_1(x), ..., s_n(x)$ is a basis for $p^{-1}(x)$.*

Proof. Begin by assuming that $p : E \to B$ is a vector bundle isomorphic to $B \times \mathbb{R}^n$, the trivial bundle. We know $\mathbb{R}^n$ has a standard basis $\{\mathbf{e_i}\}$, and we use this to define sections $s_i : x \mapsto (x, \mathbf{e_i})$ which give a basis for $\{x\} \times \mathbb{R}^n$ for each $x \in B$. Then we have a vector bundle isomorphism $f : B \times \mathbb{R}^n \to E$ and we define sections on E by $\tilde{s_i} := f(s_i)$. Since f restricted to a fibre $\{x\} \times \mathbb{R}^n$ is a linear isomorphism onto $p^{-1}(x)$, this means that $\{f(s_i(x))\}$ is a basis for $p^{-1}(x)$.

Now assume that $p : E \to B$ has sections $s_1, ..., s_n$ with $\{s_i(x)\}$ a basis for $p^{-1}(x)$ for all $x \in B$. Then take a map $f : B \times \mathbb{R}^n \to E$ defined by

$$f(x, \sum_{i=1}^{n} a_i \mathbf{e_i}) = \sum_{i=1}^{n} a_i s_i(x).$$

When restricted to a single fibre $\{x\} \times \mathbb{R}^n$ this is an isomorphism of vector spaces as the s_i form a basis of $p^{-1}(x)$ by hypothesis. Now there exists some open $U \in B$ containing x such that $p^{-1}(U)$ is homeomorphic to $U \times \mathbb{R}^n$. If we compose f with this homeomorphism it is clear that the resulting map is continuous as a map $U \times \mathbb{R}^n \to U \times \mathbb{R}^n$, and hence f is itself continuous. Then, using 3.4, we have that f is an isomorphism of vector bundles. This completes the proof. □

We are now interested in constructing vector bundles out of other vector bundles. A simple example of this is the *restriction vector bundle.*

Proposition 3.8. *Let $p : E \to B$ be an n-dimensional vector bundle and let $D \subset B$. Then $p\mid_{p^{-1}(D)} p^{-1}(D) \to D$ is a vector bundle of dimension n.*

Proof. This is very easy since if $\{U_\alpha\}$ is the associated open cover of B for the original vector bundle, then $\{U_\alpha \cap D\}$ is an open cover for D and the φ_α restricted to $U_\alpha \cap D$ clearly take $U_\alpha \cap D$ to $U_\alpha \cap D \times \mathbb{R}^n$ homeomorphically. The rest follows directly from E being a vector bundle. □

A more interesting example is the *Whitney sum* of two vector bundles.

Definition 3.9. Let $p_1 : E_1 \to B$ and $p_2 : E_2 \to B$ be two vector bundles over the same base space. Then the *Whitney sum* of E_1 and E_2 is given by

$$E_1 \oplus E_2 := \{(e_1, e_2) \in E_1 \times E_2 \mid p_1(e_1) = p_2(e_2)\}$$

with the projection map $p : E_1 \oplus E_2 \to B$ given by $p : (e_1, e_2) \mapsto p_1(e_1) = p_2(e_2)$.

We should show that this does actually give us a new vector bundle.

Proposition 3.10. *Let $p_1 : E_1 \to B$ and $p_2 : E_2 \to B$ be two vector bundles of dimension n and m, respectively. Then the Whitney sum $p : E_1 \oplus E_2 \to B$ is a vector bundle of dimesion $n + m$.*

Proof. Before we begin, note that $E_1 \oplus E_2$ inherits the topology from $E_1 \times E_2$ and so is a well defined topological space. Also, note that p is clearly a continuous map by its definition. Now we have to verify conditions (a) through (d). Let $x \in B$ be given. Then $p^{-1}(x) = \{(e_1, e_2\} \in E_1 \times E_2 \mid e_1 \in p_1^{-1}(x), e_2 \in p_2^{-1}(x)\} = p_1^{-1}(x) \oplus p_2^{-1}(x)$ and so has the structure of a $n + m$-dimensional vector space. This shows (a). Now, after suitable intersections of the associated open covers for E_1 and E_2, we can find an open cover $\{U_\alpha\}$ such that E_1 and E_2 are both trivial over each U_α. Then we have homeomorphisms from $p_1^{-1}(U_\alpha)$ and $p_2^{-1}(U_\alpha)$ to $U_\alpha \times \mathbb{R}^n$ and $U_\alpha \times \mathbb{R}^m$, respectively. Using those homeomorphisms, we can easily find a homeomorphism from $p^{-1}(U_\alpha) = \bigcup_{x \in U_\alpha} p_1^{-1}(x) \oplus p_2^{-1}(x)$ to $U_\alpha \times \mathbb{R}^n \times \mathbb{R}^m$, which is exactly what we wanted. This shows (b). (c) is easy, using the fact that if V_1, V_2, W_1, W_2 are four vector spaces with isomorphisms $f_1 : V_1 \to W_1$ and $f_2 : V_2 \to W_2$, then $f : V_1 \oplus V_2 \to W_1 \oplus W_2$ given by $f(\mathbf{v_1}, \mathbf{v_2}) = (f_1(\mathbf{v_1}), f_2(\mathbf{v_2}))$ is an isomorphism of vector spaces. As noted before, (d) is a technical condition and we can assume that this holds since our cover is clearly contained inside some maximal one. This completes the proof. □

We mentioned the concept of *stable isomorphism* earlier, and we will now give the definition.

Definition 3.11. Let $p_1 : E_1 \to B$ and $p_2 : E_2 \to B$ be two vector bundles over the same base space. We say that E_1 and E_2 are *stably isomorphic* if there exists $n, m \in \mathbb{N}_0$ such that $E_1 \oplus (B \times \mathbb{R}^n)$ is isomorphic to $E_2 \oplus (B \times \mathbb{R}^m)$. If $p_1 : E_1 \to B$ is stably isomorphic to the trivial bundle, then we say that E_1 is *stably trivial*.

Clearly isomorphic vector bundles are also stably isomorphic, but there do exist vector bundles that are stably isomorphic but not isomorphic. Similarly, it is clear that the Whitney sum of two trivial bundles is itself trivial, but it is also possible for the Whitney sum of two non-trivial bundles to be trivial which we now give an example of.

Example 3.12. Let ξ be the Möbius bundle, $[0, 2\pi] \times \mathbb{R}$ quotiented by the relation $(0, v) \sim (2\pi, -v)$. Then $\xi \oplus \xi$ is then given by $S^1 \times \mathbb{R}^2$, quotiented by the relation $(0, \mathbf{v}) \sim (2\pi, -\mathbf{v})$. We now show this is isomorphic to the trivial bundle $S^1 \times \mathbb{R}^2$. Take $f : S^1 \times \mathbb{R}^2 \to \xi \oplus \xi$, defined by

$$f : (\theta, \mathbf{v}) \mapsto (\theta, \begin{pmatrix} \cos(\theta/2) & -\sin(\theta/2) \\ \sin(\theta/2) & \cos(\theta/2) \end{pmatrix} \mathbf{v}).$$

Note that this map is well-defined since $f(0, (v)) = (0, \mathbf{v}) \sim (2\pi, -\mathbf{v}) = f(2\pi, \mathbf{v})$. It is clearly continuous and fibres map to fibres by an isomorphism of vector spaces since the map

$$\mathbf{v} \mapsto \begin{pmatrix} \cos(\theta/2) & -\sin(\theta/2) \\ \sin(\theta/2) & \cos(\theta/2) \end{pmatrix} \mathbf{v}$$

is an isomorphism of $\mathbb{R}^2$. So by 3.4, f is an isomorphism.

We can also view the above example geometrically by embedding two orthogonal Möbius bands in the solid torus $S^1 \times D^2$, where the two Möbius bands intersect along the central $S^1 \times \{\mathbf{0}\}$, and then removing the boundaries to form the respective vector bundles. This gives a visual representation of the decomposition $S^1 \times \mathbb{R}^2 = \xi \oplus \xi$.

In linear algebra, we can define inner products on vector spaces to turn them into inner product spaces. It is reasonable to think that we can do something similar with vector bundles, and it turns out that we can, provided that our base space is 'nice' enough. The property that we want our base space to have that makes it 'nice' is *paracompactness*, which we will now define.

Definition 3.13. Let X be a topological space and let $\{U_\alpha\}_{\alpha\in\mathfrak{A}}$ be an open cover of X. Then an open cover $\{V_\beta\}_{\beta\in\mathfrak{B}}$ is a *refinement* of $\{U_\alpha\}_{\alpha\in\mathfrak{A}}$ if for all $\beta \in \mathfrak{B}$ there exists $\alpha \in \mathfrak{A}$ such that $V_\beta \subset U_\alpha$.

Definition 3.14. Let X be a topological space and let $\{\mu_i\}_{i\in\mathfrak{I}}$ be a collection of continuous maps $\mu_i : X \to [0,1]$. We say that$\{\mu_i\}_{i\in\mathfrak{I}}$ is a *partition of unity* if is satisfies the following two conditions:

(a) For $x \in X$, there are only finitely many $i \in \mathfrak{I}$ such that $\mu_i(x) \neq 0$.

(b) $\sum_{i\in\mathfrak{I}} \mu_i(x) = 1$ for all $x \in X$.

Definition 3.15. Let X be a space. We say that X is *paracompact* if for every open cover $\{U_\alpha\}_{\alpha\in\mathcal{A}}$ of X, there exists an partition of unity $\{\mu_i\}$ such that $\{\mu_i^{-1}((0,1])\}$ is a refinement of $\{U_\alpha\}$

We now see that vector bundles with paracompact base spaces admit inner products in the following way. Note that this is not much of a restriction in our interests, as all manifolds are paracompact (although we give no proof here.)

Definition 3.16. Let $p : E \to B$ be a vector bundle. An *inner product* on E is a continuous map $\langle\ ,\ \rangle : E \oplus E \to \mathbb{R}$ such that for each fibre $p^{-1}(x)$,

$$\langle\ ,\ \rangle|_{p^{-1}(x)} \colon p^{-1}(x) \oplus p^{-1}(x) \to \mathbb{R}$$

is an inner product of vector spaces.

Proposition 3.17. *Let $p : E \to B$ be a vector bundle over a paracompact base space B. Then E admits an inner product.*

Proof. Let $\{U_\alpha\}_{\alpha\in\mathfrak{A}}$ be the associated open cover to $p : E \to B$. The standard inner product on $\mathbb{R}^n$ allows us to define an inner product on the trivial bundle $U_\alpha \times \mathbb{R}^n$ easily. Then, since $p^{-1}(U_\alpha)$ is homeomorphic to $U_\alpha \times \mathbb{R}^n$ we can use that inner product to define an inner product $\langle\ ,\ \rangle_\alpha$ on $p^{-1}(U_\alpha)$. Now, as B is paracompact, there exists $\{\mu_i\}_{i\in\mathfrak{I}}$ such that such that $\{\mu_i^{-1}((0,1])\}_{i\in\mathfrak{I}}$ is a refinement of $\{U_\alpha\}_{\alpha\in\mathfrak{A}}$. Equivalently, for every $i \in \mathfrak{I}$ there exists α_i such that $\mu_i^{-1}((0,1]) \subset U_{\alpha_i}$. Then we define the continuous map $\langle\ ,\ \rangle : E \oplus E \to \mathbb{R}$ by

$$\langle e_1, e_2\rangle = \sum_{i\in\mathfrak{I}} \mu_i(p(e_1))\langle e_1, e_2\rangle_{\alpha_i}.$$

The claim is that this is an inner product on E. Formally speaking, there is some small deception present in this formula, since $\langle e_1, e_2\rangle_{\alpha_i}$ requires e_1 and e_2 to be elements of $p^{-1}(U_{\alpha_i})$, but there is no such restriction on the inputs to the whole function. To see that this is justified, note that only finitely many of the μ_i are non-zero, and so the sum only consists of a finite number of terms. The terms that are left are precisely the ones where $e_1 \in \mu_i^{-1}((0,1]) \subset U_{\alpha_i}$. Finally, note that since this function was defined on $E \oplus E$, $e_1 \in U_{\alpha_i}$ implies $e_2 \in U_{\alpha_i}$, which means that the formula is valid. Now that we know the formula is valid, $\langle\ ,\ \rangle$ is an inner product simply because of the fact that it inherits symmetry, bilinearity and positive-definiteness directly from the $\langle\ ,\ \rangle_\alpha$. □

Armed with an inner product on vector bundles (over a paracompact base space) we can use the inner product to define new spaces. These are not themselves vector bundles, but instead are examples of a more general object called *fibre bundles*.

Definition 3.18. Let $p : E \to B$ be an n-dimensional vector bundle with an inner product $\langle\ ,\ \rangle$. Then the corresponding *n-sphere bundle* over B is defined as

$$S(E) := \{e \in E \mid \langle e, e\rangle = 1\}.$$

Similarly the corresponding *n-disc bundle* over B is defined as

$$D(E) := \{e \in E \mid \langle e, e\rangle \leq 1\}.$$

Fibre bundles generalise vector bundles in that the fibres no longer have to be vector spaces and instead can be any fixed topological space. We will not need the full concept of fibre bundles in this project so we will not study them in detail, but we will make use of sphere and disc bundles later.

Another use of inner products on vector bundles is that it allows us to establish the existence of orthogonal complements for vector bundles. Since orthogonal complements for vector spaces involve vector subspaces, we will need the notion of a *sub-bundle* first.

Definition 3.19. Let $p : E \to B$ be a vector bundle. Then a *sub-bundle* $p\,|_{E_0}\, E_0 \to B$ is a vector bundle such that $E_0 \cap p^{-1}(x)$ is a vector subspace of $p^{-1}(x)$ for all $x \in B$.

Theorem 3.20. *Let $p : E \to B$ be a vector bundle with B paracompact and let $p_0 : E_0 \to B$ be a sub-bundle. Then there exists another sub-bundle $E_0^{\perp}$ such that $E = E_0 \oplus E_0^{\perp}$.*

Proof. Let $\dim(E) = n$ and $\dim(E_0) = m$. Then, let $E_0^{\perp}$ be the subspace of E defined as the union of the orthogonal complements of all the fibres $p_0^{-1}(x)$ as a subspace of $p^{-1}(x)$. We want to show that $p\,|_{E_0^{\perp}}\colon E_0^{\perp} \to B$ is a sub-bundle of E and that $E = E_0 \oplus E_0^{\perp}$. To show that $E_0^{\perp}$ is a sub-bundle it suffices to show that it is locally trivial. All of the other conditions are clearly inherited from E. Let $x_0 \in B$ be given, and let U be an open neighbourhood of x such that both E and E_0 are trivial over U. Then by 3.7 there exists $\{t_1, ..., t_n\}$ linearly independent sections of E and $\{s_1, ..., s_m\}$ linearly independent sections of E_0 defined on U. With a potential reordering of the t_i, we may assume that $\{s_1, ..., s_m, t_{m+1}, ..., t_n\}$ are linearly independent sections of E on U. We can then use the Gram-Schmidt orthonormalisation procedure to turn these sections into an orthonormal basis in each fibre $\{\tilde{s}_1, ..., \tilde{s}_n\}$ and note that $\{\tilde{s}_1, ..., \tilde{s}_m\}$ will be a set of linearly independent sections for E_0. Now we can find a homeomorphism from $p^{-1}(U) \to U \times \mathbb{R}^n$ in the obvious way, where each $\tilde{s}_i$ maps to the ith basis element of $\mathbb{R}^n$. Restricting this homeomorphism to just the $\tilde{s}_{m+1}, ..., \tilde{s}_n$ gives us the required homeomorphism to show that $E_0^{\perp}$ is locally trivial.

Finally, the map $E_0 \oplus E_0^{\perp} \to E$ sending $(v_1, v_2) \mapsto (v_1 + v_2)$ is clearly continuous and sends fibres to fibres by an isomorphism of vector spaces. This means we can use 3.4 to show that this map is an isomorphism of vector bundles, completing the proof. □

One might ask whether this orthogonal complement is unique up to isomorphism (i.e. does not depend on our choice of inner product on E), and the answer is yes. To see that, we need the following construction. Let E_0 be a sub-bundle of E. Then we can define the *quotient bundle* E/E_0 to be the union of all of the quotient vector spaces formed by quotienting the fibres of E by the fibres of E_0. Since for vector spaces $V \subset W$, the quotient vector space $W/V = V^{\perp}$, it is not hard to see that E/E_0 is a sub-bundle and is isomorphic to $E_0^{\perp}$. Since E/E_0 does not make any mention to an inner product at all, it follows that $E_0^{\perp}$ does not depend on the choice of inner product either.

We will now see the final concept in this section, that of a *pullback bundle*. The general idea is that, given a vector bundle $p : E \to B$ and map $f : \tilde{B} \to B$, we can 'pull back' the vector bundle E to form a new vector bundle $\tilde{p} : \tilde{E} \to \tilde{B}$.

Proposition 3.21. *Let $p : E \to B$ be a vector bundle and let $f : \tilde{B} \to B$ be a continuous map. Then there exists a vector bundle $\tilde{p} : \tilde{E} \to \tilde{B}$ and a continuous map $\tilde{f} : \tilde{E} \to E$ that takes fibres $\tilde{p}^{-1}(\tilde{x})$ to fibres $p^{-1}(f(\tilde{x}))$ by an isomorphism of vector spaces. Furthermore, $\tilde{p} : \tilde{E} \to \tilde{B}$ is unique up to isomorphism.*

Proof. We will define the vector bundle $\tilde{p} : \tilde{E} \to \tilde{B}$ explicitly and then show that it has the required properties. First, define

$$\tilde{E} := \{(\tilde{x}, e) \in \tilde{B} \times E \mid f(\tilde{x}) = p(e)\}$$

with $\tilde{p}$ given easily as projection onto the first coordinate. This also gives us an obvious continuous map $\tilde{f} : \tilde{E} \to E$ as projection onto the second coordinate. Since $f \circ \tilde{p}((\tilde{x}, e)) = f(\tilde{x}) = p(e) = p \circ \tilde{f}((\tilde{x}, e))$, we get the following commutative diagram:

$$\begin{array}{ccc} \tilde{E} & \xrightarrow{\tilde{f}} & E \\ {\scriptstyle\tilde{p}}\downarrow & & \downarrow{\scriptstyle p} \\ \tilde{B} & \xrightarrow[f]{} & B \end{array}$$

which summaries the information above. To see that this is the required vector bundle, define $G(f) = \{(\tilde{x}, f(\tilde{x}) \in \tilde{B} \times B\}$. Then the restriction of the map $\mathrm{Id} \times p : \tilde{B} \times E \to \tilde{B} \times B$ to $\tilde{E}$ is a map $\varphi : \tilde{E} \to G(f)$. If we post-compose φ with the map π_1 which projects onto the first coordinate, we get the map $\tilde{p}$. Equivalently, we have the following commutative diagram:

$$\begin{array}{ccc} \tilde{E} & \xrightarrow{\varphi} & G(f) \\ & \underset{\tilde{p}}{\searrow} & \downarrow{\scriptstyle \pi_1} \\ & & \tilde{B} \end{array}$$

Now since f is continuous we have that $\pi_1 : G(f) \to \tilde{B}$ is a homeomorphism. This means that if $\varphi : \tilde{E} \to G(f)$ was a vector bundle, then $\tilde{p} : \tilde{E} \to \tilde{B}$ would be one also. But φ is a just a restriction of the map $\mathrm{Id} \times p : \tilde{B} \times E \to \tilde{B} \times B$ which itself is a vector bundle, and restricting a vector bundle to a subspace of the base space gives us a well-defined vector bundle, so $\varphi : \tilde{E} \to G(f)$ must be a vector bundle. From here it is now easy to see that fibres of $\tilde{E}$ are taken by linear isomorphism to fibres of E using the first commutative diagram. Now all that has to be shown is that $\tilde{p} : \tilde{E} \to \tilde{B}$ is unique up to isomorphism.

Assume we have another vector bundle $p_1 : E_1 \to \tilde{B}$ that satisfies the required conditions, with a continuous map $f_1 : E_1 \to E$. Define a map $g : E_1 \to \tilde{E}$ by $g : e_1 \mapsto (p_1(e_1), f_1(e_1))$. Note that $(p_1(e_1), f_1(e_1))$ is indeed an element of $\tilde{E}$ since $f(p_1(e_1)) = p(f_1(e_1))$ by commutativity of the corresponding version of the first diagram for E_1. Now g takes fibres of E_1 to fibres of $\tilde{E}$ by an isomorphism of vector spaces since f_1 takes fibres of E_1 to fibres of E by an isomorphism of vector spaces. This completes the proof by using 3.4. □

If we have a vector bundle $p : E \to B$ and $f : \tilde{B} \to B$ a continuous map, then the pullback bundle associated to f will be written $f^*(E)$. We will be using pullback bundles more in the next section, but to finish off this section let us prove some basic properties of them.

Proposition 3.22. *Let E and F be any two vector bundles over an arbitrary base space B and let $f : C \to B$, $g : D \to C$ be two arbitrary maps of spaces. Then the following properties hold:*

(a) $(f \circ g)^*(E) \cong g^*(f^*(E))$;

(b) $\mathrm{Id}^*(E) \cong E$;

(c) $f^*(E \oplus F) \cong f^*(E) \oplus f^*(F)$.

Proof. We will use 3.21 to prove all three. For (a), first fix $x \in D$. Then we have a continuous map $\tilde{g} : g^*(f^*(E)) \to f^*(E)$ that sends the fibre of $g^*(f^*(E))$ corresponding to x to the fibre of $f^*(E)$ corresponding to $g(x)$ by an isomorphism of vector spaces. Similarly, after fixing $y \in C$, we have $\tilde{f} : f^*(E) \to E$ that sends the fibre of $f^*(E)$ corresponding to y to the fibre of E corresponding to $f(y)$. Then, $\tilde{f} \circ \tilde{g} : g^*(f^*(E)) \to E$ sends the fibre of $g^*(f^*(E))$ corresponding to x to the fibre of E corresponding to $f \circ g(x)$ by an isomorphism of vector spaces. So, by 3.21, we have that$(f \circ g)^*(E) \cong g^*(f^*(E))$.

(b) is clear since $\text{Id} : E \to E$ satisfies all of the required properties of the pullback bundle trivially, so we can use 3.21 like for (a).

(c) is much the same as the proof for (a) except here we use that we have continuous maps $\tilde{f}_E : f^*(E) \to E$ and $\tilde{f}_F : f^*(F) \to F$ to give us the map

$$\tilde{f}_E \times \tilde{f}_F \mid_{f^*(E) \oplus f^*(F)} : f^*(E) \oplus f^*(F) \to E \oplus F$$

which has the required properties to satisfy 3.21. □

It should be stated that everything in this section directly applies also to *complex vector bundles*, where every fibre is given the structure of a complex vector space as opposed to a real one.

Chapter 4
Vector bundles over S^n

(Unless specified, the material in this chapter is based on [3].)

The previous section introduced the concept of vector bundles, which we will now explore in more detail. The goal of this section will be towards classifying vector bundles over S^n. Our main tool to begin with will be the pullback bundles that we defined at the end of the last section. To begin with, we will work with a general base space but will quickly restrict to working with spheres First, we introduce some notation, which is that we shall denote the set of all isomorphism classes of n-dimensional real vector bundles over a fixed base space B by $\text{Vect}^n(B)$.

Proposition 4.1. *Let $f : \tilde{B} \to B$ be a continuous map. Then f induces a map $f^* : \text{Vect}^n(B) \to \text{Vect}^n(\tilde{B})$.*

Proof. This is made easy due to the existence and well-definedness of pullback bundles, since we have the map $f^* : \text{Vect}^n(B) \to \text{Vect}^n(\tilde{B})$ which sends a vector bundle $E \in \text{Vect}^n(B)$ to its pullback bundle $f^*(E) \in \text{Vect}^n(\tilde{B})$. □

Clearly this map depends on the original map f in some way, but the extent to which is unclear. We shall now show that homotopic maps give rise to the same induced map on vector bundles, in much the same way that homotopic maps give rise to the same induced map on homology. However, before we can show this, we will need a series of lemmas.

Lemma 4.2. *Let $a < b \in \mathbb{R}$ and let $p : E \to X \times [a, b]$ be a vector bundle. Further suppose that there exists $a \leq c \leq b$ such that the restrictions $E_1 = p^{-1}(X \times [a, c])$ and $E_2 = p^{-1}(X \times [c, b])$ are both trivial. Then E is the trivial bundle.*

Proof. Let φ_1, φ_2 be the isomorphisms from E_1 and E_2 to $X \times [a, c] \times \mathbb{R}^n$ and $X \times [c, b] \times \mathbb{R}^n$, respectively, which exist since both of these bundles are trivial. Now we define a map $\psi : X \times [c, b] \times \mathbb{R}^n \to X \times [c, b] \times \mathbb{R}^n$ by requiring that ψ takes every fibre $\{x\} \times \{t\} \times \mathbb{R}^n$ first by translating in the second coordinate to $\{x\} \times \{c\} \times \mathbb{R}^n$, then to $\{x\} \times \{c\} \times \mathbb{R}^n$ via $\varphi_1 \circ \varphi_2^{-1}$ and finally by translating back to $\{x\} \times \{t\} \times \mathbb{R}^n$. This is clearly a continuous map that takes fibres to fibres by a linear isomorphism and so ψ is an isomorphism of vector spaces. Now, we use ψ to define $\varphi : E \to X \times [a, b] \times \mathbb{R}^n$ on $e = (x, t, v)$ as:

$$\varphi(e) = \begin{cases} \varphi_1(e) & \text{if } t \leq c \\ \psi \circ \varphi_2(e) & \text{if } t \geq c \end{cases}$$

which is a continuous map since φ_1 and $\psi \circ \varphi_2$ match at $t = c$. Since both also take fibres to fibres via an isomorphism of vector spaces, this means that φ is an isomorphism of vector spaces, which completes the proof. □

Lemma 4.3. *Let $p : E \to X \times I$ be a vector bundle. Then there exists an open cover $\{U_\alpha\}_{\alpha \in \mathfrak{A}}$ such that for all $\alpha \in \mathfrak{A}$ the restriction vector bundle $p^{-1}(U_\alpha \times I) \to U_\alpha \times I$ is trivial.*

Fix $x \in X$. Due the definition of vector bundles, we know that every point $(x,t) \in X \times I$ has an open neighbourhood $U_{x,t}$ for which the restriction of the vector bundle E to $U_{x,t}$ is trivial over. Now, we may assume that $U_{x,t}$ is connected, since if it was not we could simply remove the component not containing (x,t), and so we can write $U_{x,t} = \tilde{U}_{x,t} \times (t^-, t^+)$ for some $t^- < t < t^+$. We then have the infinite open cover $\bigcup_{t\in[0,1]}(t^-, t^+)$, which we know by compactness of $[0,1]$ must contain a finite subcover. Let $U_{x,i}$, $i = 1, ..., k$ be the associated open sets in $X \times I$, which we know must take the form $\tilde{U}_{x,i} \times (t_i^-, t_i^+)$. Now each of the intervals (t_i^-, t_i^+) contain a closed interval I_i such that the restriction of E to $\tilde{U}_{x,i} \times I_i$ is still trivial. By suitably restricting these intervals further, we can assume that they form a partition of $[0,1]$, more specifically we can assume that $I_i = [t_i, t_{i+1}]$ with $1 = t_0 < t_2 < ... < t_{k-1} < t_{k+1} = 1$. Then by repeated use of the previous lemma, we have that $p^{-1}(U_\alpha \times I) \to U_\alpha \times I$ is trivial.

Lemma 4.4. *Let X be compact and let $p : E \to X \times I$ be a vector bundle. Then the restrictions of E over $X \times \{0\}$ and $X \times \{1\}$ are isomorphic.*

This lemma holds more generally for X paracompact, but our applications will only ever require that X is compact and Hausdorff and so we restrict to this case to simplify the proof.

Proof. By the previous lemma, there exists an open cover $\{U_\alpha\}_\alpha$ with the restriction $p^{-1}(U_\alpha \times I)$ trivial. Now, since X is compact, this cover has a finite subcover; call it $\{U_i\}_{i=1,...,m}$. Furthermore, since X is also paracompact (as it is Hausdorff), we have a partition of unity $\{\mu_i\}$ with $\mu_i^{-1}((0,1])$ refining this cover. So we may assume without loss of generality that $U_i = \mu_i^{-1}((0,1])$. Now define a new function

$$\psi_k := \begin{cases} 0 & \text{if } k = 0 \\ \sum_{i=1}^{k} \mu_i & \text{if } k \geq 1. \end{cases}$$

Now, let $G_i(X) := \{(x, \psi_i(x)) \in X \times I\}$ be the graph of ψ_i. Note that we have natural map $h_i : G_i(X) \to G_{i-1}(X)$ which sends $(x, \psi_i(x))$ to $(x, \psi_{i-1}(x))$. h_i is continuous with continuous inverse, since we can write it as $h_i(x, \psi_i(x)) = (x, \psi_i(x) - \mu_i(x))$. This means h_i is a homeomorphism that is the identity map outside of $U_i \times I$. Now, if we let $p_i : E_i \to G_i(X)$ be the restriction vector bundle over $G_i(X)$, we can use h_i to define a continuous map $\tilde{h}_i : E_i \to E_{i-1}$ that is the identity map outside of $p_i^{-1}(U_i \cap G_i(X) \times I)$. Using that $p^{-1}(U_i \times I)$ is trivial, we then define $\tilde{h}_i$ on $p_i^{-1}(U_i \cap G_i(X) \times I)$ as the map sending $(x, \psi_i(x), v) \in U_i \cap G_i(X) \times I \times \mathbb{R}^n$ to $(x, \psi_{i-1}(x), v)$. This is clearly continuous and sends fibres to fibres by an isomorphism of vector spaces and hence is an isomorphism of vector bundles. We can then compose these $\tilde{h}_i$ to form an isomorphism $H := \tilde{h}_1 \circ \tilde{h}_2 \circ ... \circ \tilde{h}_m : G_m(X) \to G_0(X)$; since $G_m(X) = X \times \{1\}$ and $G_0(X) = X \times \{0\}$, this completes the proof.

□

Theorem 4.5. *Let $p : E \to B$ be a vector bundle and let $f, g : \tilde{B} \to B$ be homotopic maps with $\tilde{B}$ compact and Hausdorff. Then $f^*(E)$ and $g^*(E)$ are isomorphic.*

Similarly to the previous lemma, this also holds for $\tilde{B}$ paracompact.

Proof. Since f and g are homotopic, we have a homotopy $H : \tilde{B} \times I \to B$ such that $H_0 = f$ and $H_1 = g$. Now we consider the pullback bundle $H^*(E)$ and its restrictions over $\tilde{B} \times \{0\}$ and $\tilde{B} \times \{1\}$. These restrictions are clearly isomorphic to $f^*(E)$ and $g^*(E)$, respectively. Then we can use the previous lemma to conclude that $f^*(E)$ and $g^*(E)$ must be isomorphic. □

This gives us two powerful corollaries that we will make use of later.

Corollary 4.6. *Let $f : \tilde{B} \to B$ be a homotopy equivalence of compact spaces $B, \tilde{B}$. Then the induced map $f^* : \mathrm{Vect}^n(B) \to \mathrm{Vect}^n(\tilde{B})$ is a bijection.*

Proof. Since f is a homotopy equivalence, there must exist some $g : B \to \tilde{B}$ such that $f \circ g$ and $g \circ f$ are homotopic to the respective identity maps. But then by the theorem $(f \circ g)^* = Id_B^*$ and $(g \circ f)^* = Id_{\tilde{B}}^*$. Then it is a simple application of the basic properties of the pullback bundle established in 3.22 that $f^* \circ g^* = Id_{\mathrm{Vect}^n(B)}$ and $g^* \circ f^* = Id_{\mathrm{Vect}^n(\tilde{B})}$ and hence f^* has an inverse g^* and so must be a bijection. □

Corollary 4.7. *Let B be a compact, contractible space and let $p : E \to B$ be a vector bundle. Then E is trivial.*

Proof. Note that there is only one possible vector bundle over the base space $\{\mathrm{pt}\}$ given by $\{\mathrm{pt}\} \times \mathbb{R}^n$. Then if B is compact and contractible, then by the previous corollary there is a bijection between $Vect^n(B)$ and $Vect^n(\{\mathrm{pt}\})$. This means that $Vect^n(B)$ must contain only one possible element, which is the trivial bundle. □

We will now apply these results to studying a specific class of vector bundles, that is where the base space is a sphere. The sphere is a compact manifold and so all of the results so far about vector bundles will apply in full. The aim is to use *clutching functions* to classify all of the possible vector bundles over an n-sphere, up to isomorphism. That is, to determine $\mathrm{Vect}^k(S^n)$. We begin with the initial claim in this topic.

Proposition 4.8. *Write the n-sphere as a union of its upper and lower hemispheres $S^n = D_-^n \cup D_+^n$ where $D_-^n \cap D_+^n = S^{n-1}$ and let $f : S^{n-1} \to GL_k(\mathbb{R})$ be a continuous map. Define E_f as*

$$E_f := \frac{(D_-^n \times \mathbb{R}^k) \sqcup (D_+^n \times \mathbb{R}^k)}{\sim}$$

where $\partial D_-^n \times \mathbb{R}^k \ni (x, v) \sim (x, f(x)(v)) \in \partial D_+^n \times \mathbb{R}^k.$

Then $p : E_f \to S^n$ sending $(x, v) \mapsto (x, 0)$ is a k-dimensional vector bundle.

Proof. Clearly E_f satisfies that for every $x \in S^n$, $p^{-1}(x)$ has the structure of a k-dimensional vector space. Now we need an open cover such that E_f is trivial over each element of the cover. Let U_+ and U_- be the interiors of two larger embedded discs covering D^n_+ D^n_-, respectively. More precisely, set $U_\pm = \{\mathbf{x} \in S^n \mid \pm x_1 > -\varepsilon\}$. Then define a map from $p^{-1}(U_\pm)$ to $U_\pm \times \mathbb{R}^k$ as follows: first define it on $p^{-1}(D^n_\pm)$ as the inclusion since $p^{-1}(D^n_\pm)$ is already trivial; the remainder of $p^{-1}(U_\pm)$ is homeomorphic to $I \times S^{n-1} \times \mathbb{R}^k$, with $\{0\} \times S^{n-1} \times \{0\}$ being the image of $\partial D^n_\pm \times \mathbb{R}^k$, and we define our map to send $(t, x, v) \mapsto (t, x, (f(x))^{-1}(v))$. Note that for any x, the map $f(x))^{-1}$ is defined since $f \in GL_k(\mathbb{R})$. This map is a well defined since our two definitions match on $\partial D^n_\pm \times \mathbb{R}^k$ due to the equivalence relation in our definition of E_f. This map also clearly takes fibres to fibres by an isomorphism of vector spaces since $f \in GL_k(\mathbb{R})$. Thus, $p : E_f \to S^n$ is a k-dimensional vector bundle. □

The map $f : S^{n-1} \to GL_k(\mathbb{R})$ is called the *clutching function* for E_f. We now show that homotopic clutching functions give rise to isomorphic vector bundles.

Proposition 4.9. *Let $f, g : S^{n-1} \to GL_k(\mathbb{R})$ be homotopic clutching functions for the vector bundles E_f and E_g respectively. Then E_f and E_g are isomorphic.*

Proof. Let $H : S^{n-1} \times I \to GL_k(\mathbb{R})$ be the homotopy between f and g. Then define the space E_H as:

$$E_H := \frac{(D^n_- \times I \times \mathbb{R}^k) \sqcup (D^n_+ \times I \times \mathbb{R}^k)}{\sim}$$

where $\partial D^n_- \times I \times \mathbb{R}^k \ni (x, t, v) \sim (x, t, H_t(x)(v)) \in \partial D^n_+ \times I \times \mathbb{R}^k$.

We claim that E_H is a vector bundle. Proving this is exactly analagous to the proof of the previous proposition. Now, note that the restriction of E_H to $S^n \times \{0\}$ is isomorphic to E_f, and the restriction to $S^n \times \{1\}$ is isomorphic to E_g. Then we can use 4.4, since $S^n \times I$ is compact, to show that E_f is isomorphic to E_g. □

If we write the set of a homotopy classes of maps $X \to Y$ as $[X, Y]$, then the above proposition gives us a map $\Phi : [S^{n-1}, GL_k(\mathbb{R})] \to \mathrm{Vect}^k(S^n)$ which sends a homotopy class of clutching functions to the isomorphism class of associated vector bundles. This map is definitively not a bijection, but it becomes a bijection if we restrict to orientable vector bundles. We shall denote the group of orientation preserving invertible maps $\mathbb{R} \to \mathbb{R}$ as $GL^+_k(\mathbb{R}) := \{X \in M_k(\mathbb{R}) \mid \det(M) > 0\}$ and the set of isomorphism classes of orientable k-dimensional vector bundles over a base space B as $\mathrm{Vect}^k_+(B)$.

Theorem 4.10. *The map $\Phi : [S^{n-1}, GL^+_k(\mathbb{R})] \to \mathrm{Vect}^k(S^n)$ which sends a homotopy class of clutching functions represented by f to the isomorphism class of orientable vector bundles represented by E_f is a bijection.*

Proof. We proceed by explicitly constructing Φ^{-1}. Let $p : E \to S^n$ be a k-dimensional orientable vector bundle. The restriction vector bundles $p^{-1}(D^n_\pm)$ are both trivial by 4.7 as $D^n_\pm$ are both contractible. This means that there exists $\varphi_\pm : p^{-1}(D^n_\pm) \to D^n_\pm \times \mathbb{R}^k$ homeomorphisms that take fibres to fibres as an orientation preserving isomorphism of vector spaces. On $\partial D^n_\pm = S^{n-1}$ both $\varphi_\pm$ are defined, and so we have a map

$$\varphi_+ \circ \varphi_- : S^{n-1} \times \mathbb{R}^k \to S^{n-1} \times \mathbb{R}^k$$
$$(x, v) \mapsto (x, f(x)(v))$$

for some $f(x) \in GL^+_k(\mathbb{R})$. Note that $f : S^{n-1} \to GL^+_k(\mathbb{R})$ which sends $x \mapsto f(x)$ is a continuous map. We then define $\Phi^{-1}(E)$ to be the homotopy class of f. We now need to show that $\Phi^{-1}(E)$ is well defined and that $\Phi \circ \Phi^{-1}$ and $\Phi^{-1} \circ \Phi$ are both the identity map. If we assume that $\Phi^{-1}(E)$ is well-defined, the second and third statements are clear by the definitions.

Suppose that we used some other trivialisations $\tilde{\varphi}_\pm : p^{-1}(D^n_\pm) \to D^n_\pm \times \mathbb{R}^k$ resulting in a new clutching function $\tilde{f}$. By composing $\varphi^{-1}_\pm$ with $\tilde{\varphi}_\pm$, or vice versa, we get a continuous map $D^n_\pm \times \mathbb{R}^k \to D^n_\pm \times \mathbb{R}^k$ that sends $(x, v) \mapsto (x, g(x)(v))$ for some continuous map $g : D^n_\pm \to GL^+_k(\mathbb{R})$. Now since $D^n_\pm$ is contractible, we have a homotopy $C : D^n_\pm \times I \to D^n_\pm$ with C_0 the identity map and C_1 the map that sends everything to a single point. Then $g \circ C$ is a homotopy between $(g \circ C)_0 = g$ and $(g \circ C)_1$ which is the constant map. We want to conclude that g is homotopic to the map sending everything to the identity element in $GL^+_k(\mathbb{R})$, which is equivalent to saying that $GL^+_k(\mathbb{R})$ is path-connected, which we will now show.

First, we claim that for any $X \in GL_k(\mathbb{R})$ there exists a path in $GL_k(\mathbb{R})$ between X and a diagonal matrix. Recall from linear algebra that any matrix can be diagonalised by performing a sequence of elementary row operations. We can realise any elementary row operation as a path in $GL_k(\mathbb{R})$ by including a factor of t in front of the row multiple being added; as t varies from $t = 0$ to $t = 1$, a path is traced in $GL_k(\mathbb{R})$ whose end point is the result of the row operation. By concatenating these paths together we get a path from X to a diagonal matrix. Note that none of the entries in the diagonal are zero. Now we can further find a path from this diagonal matrix to a diagonal matrix with entries ± 1 by introducing a factor of $1/(1 + t(\pm\lambda - 1))$ in front of each entry λ, with the sign of $\pm$ depending on whether λ is positive or negative. Then similarly allowing t to vary from $t = 0$ to $t = 1$ gives required path. We can also replace any two -1 entries with two $+1$ entries using a path in $GL_k(\mathbb{R})$ by replacing the 2×2 sub-matrix containing those two entries by the matrix

$$\begin{pmatrix} \cos t & -\sin t \\ \sin t & \cos t \end{pmatrix}.$$

Then as t varies from $t = \pi$ to $t = 0$ we get the required path. Therefore, if our matrix has an even number of -1 entries, then we can find a path to the identity matrix. This is equivalent to saying that our original matrix X has $\det(X) > 0$ which means that $X \in GL_k^+(\mathbb{R})$. This gives us the required result that $GL_k^+(\mathbb{R})$ is path-connected.

This tells us that g is homotopic to the map sending everything to the identity in $GL_k^+(\mathbb{R})$, which means that $\varphi_\pm$ and $\tilde{\varphi}_\pm$ are homotopic. Subsequently, the maps f and $\tilde{f}$ are also homotopic. This means that Φ^{-1} is well-defined and completes the proof. □

Although we restricted to considering only orientable vector bundles in this classification, we can still say a lot about the original object $\mathrm{Vect}^k(S^n)$. Note that we have a natural map $\mathrm{Vect}_+^k(S^n) \to \mathrm{Vect}^k(S^n)$ which forgets the orientation information.

Proposition 4.11. *Let $n \geq 2$. The natural map $\mathrm{Vect}_+^k(S^n) \to \mathrm{Vect}^k(S^n)$ is a surjection and is at most two-to-one.*

Proof. Given any vector bundle over S^n, we can certainly assign an orientation to one fibre over a single point $x_0 \in S^{n-1} \subset S^n$. Then, if we follow the the first section of the proof of 4.10 we get a continuous map $f : S^{n-1} \to GL_k(\mathbb{R})$, $x \mapsto f(x)$ with $f(x_0) \in GL_k^+(\mathbb{R})$. S^{n-1} is path-connected as $n \geq 2$, so this means that $f(x) \in GL_k^+(\mathbb{R})$ for all $x \in S^{n-1}$, not just at x_0. Putting this together with the remainder of 4.10, this gives us that our original vector bundle is orientable. More specifically, if we let $\mathrm{Vect}_0^k(S^n)$ be the set of k-dimensional vector bundles over S^n with an orientation specified for a single fibre corresponding to one point $x_0 \in S^{n-1} \subset S^n$ up to isomorphism preserving the orientation of the fibre corresponding to x_0, then we have a bijection between $\mathrm{Vect}_0^k(S^n)$ and $\mathrm{Vect}_+^k(S^n)$.

The reason for making this observation is that the relationship between $\mathrm{Vect}_0^k(S^n)$ and $\mathrm{Vect}^k(S^n)$ is much easier to understand. The map $\mathrm{Vect}_0^k(S^n) \to \mathrm{Vect}^k(S^n)$ that ignores the orientation of the fibre at x_0 is clearly surjective and, since there are only two orientations we can give to the fibre at x_0, is at most two-to-one. The reason that it may not be two-to-one is that the resulting fibre bundles with opposite orientations chosen at x_0 may still be isomorphic as orientable vector bundles. □

This proof does not work for the case $n = 1$, as S^0 is not path-connected. However, that case is in some ways even simpler as S^0 only consists of two points. If we perform a similar analysis, one can see that there are only two elements in $\mathrm{Vect}^k(S^1)$: the trivial bundle ξ_k and the Whitney sum of the Möbius bundle with ξ_{k-1}. We will not be particularly interested in the case $n = 1$.

We have now reduced the problem of understanding $\mathrm{Vect}^k(S^n)$ to that of understanding $[S^{n-1}, GL_k^+(\mathbb{R})]$. So it is natural for us to try and understand $[S^{n-1}, GL_k^+(\mathbb{R})]$ better. First, we show that we can actually replace $GL_k^+(\mathbb{R})$ with $SO(k)$. This will be a consequence of the following fact.

Lemma 4.12. *$SO(k)$ is a deformation retract of $GL_k^+(\mathbb{R})$.*

Proof. We start by recalling the Gram-Schmidt orthonormalisation process from linear algebra. If we have a basis $\{v_i\}$ for $\mathbb{R}^k$, then we can find an orthonormal basis $\{u_i\}$. First, to simplify notation, let the projection map p be defined as:

$$p(u, v) = \frac{\langle u, v\rangle}{\langle u, u\rangle} u$$

Then the Gram-Schmidt process is given by first setting $\tilde{u}_1 = v_1$ and running the iterative process:

$$\tilde{u}_i = v_k - \sum_{j=1}^{i-1} p(\tilde{u}_j, v_i); \quad u_i = \frac{\tilde{u}_i}{\|\tilde{u}\|}.$$

Since any basis of $\mathbb{R}^k$ defines a matrix in $GL_k(\mathbb{R})$ and any orthonormal basis defines a matrix in $O(k)$, this gives us a map $GL_k(\mathbb{R}) \to O(k)$. If we restrict to only considering bases with the standard orientation, this gives us a map $GL_K^+(\mathbb{R}) \to SO(k)$ that is the identity on $SO(k)$. We now want to turn this process into a homotopy. This can be done easily by introducing appropriate factors of t. For the first of the above equations, simply place a factor of t before the sum and then as t varies from $t = 0$ to $t = 1$, that step of the process is performed continuously. For the second step, this is just a rescaling and so can clearly be performed continuously. Then, by concatenating all of these homotopies together we get the required deformation retract. □

Proposition 4.13. *The map $\mathbb{I} : [S^{n-1}, SO(k)] \to [S^{n-1}, GL_k^+(\mathbb{R})]$ defined by sending a representative $f : S^{n-1} \to SO(k)$ to a continuous map $\tilde{f} : S^{n-1} \to GL_k^+(\mathbb{R})$ by post-composing with the inclusion map is a bijection.*

Proof. Let $f : S^{n-1} \to GL_k^+(\mathbb{R})$ be a continuous map and let $i : SO(k) \to GL_k^+(\mathbb{R})$ be the inclusion map. If we post-compose f with the deformation retract $GL_k^+ \times I \to SO(k)$ we get that f is homotopic to some continuous map $g : S^{n-1} \to SO(k)$, so $\mathbb{I}(g) = f$ and hence $\mathbb{I}$ is surjective. Now let $f, g : S^{n-1} \to SO(k)$ be two continuous maps such that $i \circ f$ is homotopic to $i \circ g$. This means there exists $H : S^{k-1} \times I \to GL_k^+(\mathbb{R})$ with $H_0 = i \circ f$ and $H_1 = i \circ g$, respectively. If we then compose this homotopy with the retraction $R : GL_k^+(\mathbb{R}) \to SO(k)$, we get a homotopy between f and g in $SO(k)$. This means that $\mathbb{I}$ is injective. □

Corollary 4.14. *The map $\Phi \circ \mathbb{I} : [S^{n-1}, SO(k)] \to \mathrm{Vect}_+^k(S^n)$ is a bijection.*

Proof. Both Φ and $\mathbb{I}$ are bijections and compositions of bijections are themselves bijections. □

We can now introduce a group structure on the set $[S^n, SO(k)]$. However, the existence of this group structure is not unique to S^n, so we show it in full generality.

Proposition 4.15. $[X, SO(k)]$ *with multiplication defined on representatives* f, g *as* $(f \cdot g) : x \mapsto f(x)g(x)$ *as multiplication in* $SO(k)$ *and identity the class with representative sending all of* X *to* $\mathrm{Id} \in SO(k)$ *is a group.*

Proof. First, we need to check that this multiplication is well-defined. Let $f_1 \simeq f_2$, $g_1 \simeq g_2 : X \to SO(k)$. Then we have homotopies $H^f, H^g : X \times I \to SO(k)$ which gives us a map $H^f \cdot H^g : X \times I \to SO(k)$ with $(H^f \cdot H^g)_0 = H_0^f \cdot H_0^g = f_1 \cdot g_1$ and $(H^f \cdot H^g)_1 = H_1^f \cdot H_1^g = f_2 \cdot g_2$. Since $SO(k)$ is a topological group, this is a homotopy and thus our multiplication operation is well defined on homotopy classes. As a consequence of this, the identity, as defined above, is well-defined on homotopy classes and has the desired properties. For inverses, let $f : X \to SO(k)$ be a continuous map and then define f^{-1} to be the map that sends $x \mapsto (f(x))^{-1} \in SO()$. This is continuous and well-defined on homotopy classes with a similar argument to the above, using the fact that, since $SO(k)$ is a topological group, the map sending $f(x) \mapsto f(x)^{-1}$ is continuous. Finally, associativity comes directly from the associativity of $SO(k)$. □

In our case, the groups $[S^{n-1}, SO(k)]$ will be isomorphic to the *homotopy groups* $\pi_{n-1}(SO(k)$. These groups are the direct generalisations of the fundamental group π_1. We now take a brief detour to define these groups, before proving the already mentioned isomorphism.

Denote the set of homotopy classes of continuous maps $f : S^n \to X$ that have a fixed $s_0 \in S^n$ and $x_0 \in X$ such that $f(s_0) = x_0$ by $\pi_n(X, x_0)$. The reason that s_0 doesn't make an appearence in the notation is that all of the points in S^n are indistinguishable due to symmetry. We can introduce a group structure on $\pi_n(X, x_0)$ as follows. Let $C : S^n \to S^n \vee S^n$ be the map that collapses the equator S^{n-1} of S^n to a single point with $s_0 \in S^{n-1}$. Then take the map $f \vee g : S^n \vee S^n \to X$ that maps via f on the first wedge factor, and via g on the second, with s_0 the wedge point so that this is a well-defined continuous map. We can then define $f + g$ to be the composition $f \vee g \circ C$. Another way of thinking about this is that $f + g$ maps everything in the north or south hemisphere first by a continuous map to S^n that sends all of the equator to a single point and then by f or g, respectively. The existence of inverses is very similar to the existence of inverses in the fundamental group. $-f$ is defined as $f(\tilde{x})$ where $\tilde{x}$ is obtained from x by reflection in the plane perpendicular to the equator, which swaps just one coordinate to its negative. Then it can be shown that $f + (-f)$ is homotopic to the map sending everything to x_0. The associativity of this operation is clear since $S^n \vee S^n \vee S^n$ is well-defined.

It is clear that these constructions are well-defined up to homotopy and hence the group is defined. Much like with the fundamental group, $\pi_n(X, x_0)$ is isomorphic to $\pi_n(X, x_1)$, provided that X is path-connected. So for path-connected spaces we often simplify the notation and write $\pi_n(X)$. One can find details on the definition of these homotopy groups in [2], along with the analogous definitions for the relative homotopy groups $\pi_n(X, A)$ where $A \subset X$.

Proposition 4.16. *The inclusion map $\pi_n(SO(k)) \to [S^n, SO(k)]$ is a bijection.*

Proof. Since we are free to choose our basepoint, let us assume that x_0 is the identity matrix in $SO(k)$ for simplicity. First we show surjectivity. Let $f : S^n \to SO(k)$ be a representative of some homotopy class. Now f may not have Id in its image, but we can find a homotopy from f to a map that does. Since $SO(k)$ is path-connected, we can find a path $\gamma(t) \in SO(k)$ with $\gamma(0) = \text{Id}$ and $\gamma(1) = f(s_0)^{-1}$. Then $H : S^n \times I \to SO(k)$ defined as $H = f(x)\gamma(t)$ is a homotopy where $H_0 = f$ and $H_1(x) = f(x)f(s_0)^{-1}\gamma$. So $H_1(s_0) = \text{Id}$ and thus our map is surjective. Now we show injectivity. Assume we have two functions $f_0, f_1 : S^n \to SO(k)$ that both map s_0 to the identity matrix and f_0, f_1 both represent the same homotopy class in $[S^n, SO(k)]$. This means we have a homotopy $H : S^n \times I \to SO(k)$ with $H_0 = f_0$ and $H_1 = f_1$, but this homotopy may not preserve the basepoint. Now we have a new homotopy $H_t(s_0)^{-1}H_t$ which clearly maps s_0 to the identity for all t. Since we still have $H_0(s_0)^{-1}H_0(x) = f_0(x)$ and $H_1(s_0)^{-1}H_1(x) = f_1(x)$, this shows injectivity. □

We have now successfully reduced the problem of classifying vector bundles over S^n to that of computing the homotopy groups $\pi_n(SO(k))$. In general, homotopy groups are notoriously difficult to compute, but there are certainly more tools for dealing with them than there are for the sets $\text{Vect}^k(S^n)$. We shall make use of this classification later in 10, where we will use covering spaces to tackle computing one of these homotopy groups.

Chapter 5
Smooth vector bundles

(Unless specified, the material in this chapter is based on [1] and the Riemannian geometry lecture course given by P. Tumarkin at Durham University 2019-2020.)

Previously, we stated that the tangent bundle to a smooth manifold M of dimension n, defined as $\bigcup_{p\in M} T_pM$ was a vector bundle similarly of dimension n. In fact, it is an example of a *smooth vector bundle*. This is essentially a space that can be given both a vector bundle structure and a smooth structure such that the two structures are compatible.

Definition 5.1. Let M be a smooth manifold and let $p : E \to M$ be a vector bundle. Then E is a *smooth vector bundle* if it satisfies the following properties:

(a) There exists a smooth atlas $\{U_\alpha\}$ for M such that E is trivial over each U_α.

(b) If ψ_α are the associated charts and φ_α are the associated trivialisations for $\{U_\alpha\}$ as above, then $\{p^{-1}(U_\alpha)\}$ form a smooth atlas for E with charts φ_α post-composed with the map sending the first coordinate to ψ_α and the identity on the second coordinate.

(c) $p : E \to M$ is a smooth map.

(d) $\{U_\alpha\}$ is maximal with respect to the above two conditions.

There is some slight potential for confusion regarding the dimension of a smooth vector bundle. An n-dimensional smooth vector bundle is one for which the underlying vector bundle is n-dimensional. However, as might be inferred from the fact that we have a smooth atlas on the total space E, it is not hard to show that a smooth vector bundle is a smooth manifold. If the base space M is an m-dimensional smooth manifold and E is an n-dimensional smooth vector bundle, then E is an $n+m$-dimensional smooth manifold.

There is an equivalent and often more useful definition to take for smooth vector bundles. Let $E \to M$ be a vector bundle with $\{U_\alpha, \psi_\alpha\}$ a smooth atlas for M with $\varphi_\alpha : p^{-1}(U\alpha) \to \mathbb{R}^n$ trivialisations. Further let $\tilde{\varphi}_\alpha$ be φ_α post-composed with the projection onto $\mathbb{R}^n$ and let $\phi_{\alpha,x} : \mathbb{R}^n \to p^{-1}(x)$ be the inverse of $\tilde{\varphi}_\alpha$. Now we have that, for $x \in U_\alpha \cap U_\beta$, the map $\tilde{\varphi}_\beta \circ \phi_{\alpha,x} : \mathbb{R}^n \to \mathbb{R}^n$ is an isomorphism since it is linear and clearly has ker $= \{0\}$. Then it is easy to define a map $G_{\alpha,\beta} : U_\alpha \cap U_\beta \to GL_n(\mathbb{R})$ by $G_{\alpha,\beta}(x) = \tilde{\varphi}_\beta \circ \phi_{\alpha,x}$. We now can state the following proposition.

Proposition 5.2. *E, as above, satisfying properites (a), (b) and (d), is a smooth vector bundle if and only if all of the maps $G_{\alpha,\beta} : U_\alpha \cap U_\beta \to GL_n(\mathbb{R})$ are smooth.*

Proof. This statement is actually reasonably straightforward. If E is a smooth vector bundle, then the fact that $(\psi_\alpha, \tilde{\varphi}_\alpha)$ form a smooth atlas on E easily gives us that $G_{\alpha,\beta}$ is smooth. Similarly, if the $G_{\alpha,\beta}$ are smooth, then that together with the fact that (ψ_α) give a smooth atlas on M gives us the required smooth atlas on E, and hence E is a smooth vector bundle. □

Proposition 5.3. *Let $p_1 : E_1 \to M$ be an n_1-dimensional smooth vector and let $p_2 : E_2 \to M$ be an n_2-dimensional smooth vector bundle. Then the Whitney sum $p : E_1 \oplus E_2 \to M$ is an $(n_1 + n_2)$-dimensional smooth vector bundle.*

These $G_{\alpha,\beta}$ are often called the *transistion maps* for the vector bundle.

Proof. Since we can find atlases for M such that E_1 and E_2 are trivial over each chart serparately, we can find some refinement of these altases $\{U_\alpha, \varphi_\alpha\}$ such that E_1 and E_2 are both trivial over each U_α, which means (a) is satisfied. For (c), note that the smoothness of p is inherited directly from the smoothness of p_1 and p_2. Then we assume, as per usual, that (d) is satisfied. Finishing the proof is now a simple application of the previous proposition, since if E_1 and E_2 have transistion maps $G^1_{\alpha,\beta}$ and $G^2_{\alpha,\beta}$, respectively, then the transistion maps $G_{\alpha,\beta}$ for $E_1 \oplus E_2$ are given by the block matrices:

$$\begin{pmatrix} G^1_{\alpha,\beta} & \mathbf{0} \\ \mathbf{0} & G^2_{\alpha,\beta} \end{pmatrix},$$

which is clearly a smooth map as both $G^1_{\alpha,\beta}$ and $G^2_{\alpha,\beta}$ are smooth. □

We will now look at a specific smooth vector bundle that is of primary importance: the *tangent bundle* of a smooth manifold M. This has already been mentioned previously, but we define it again for the sake of completeness.

Definition 5.4. Let M be a smooth manifold and let $TM = \bigcup_{p\in M} T_pM$. Then the *tangent bundle* of M is defined as $\pi : TM \to M$ where π is the natural map sending $v \in T_pM \subset TM$ to p.

Proposition 5.5. *Let M be a smooth manifold of dimension n. The the tangent bundle of M is a smooth vector bundle also of dimension n. Furthermore, it is a smooth manifold of dimension $2n$.*

Proof. Let $\{U_\alpha, V_\alpha, \varphi_\alpha\}$ be a smooth atlas on M. We want to find a smooth atlas $\{\tilde{U}_\alpha, \tilde{V}_\alpha, \tilde{\varphi}_\alpha\}$ for TM. First, we define $\tilde{U}_\alpha := \pi^{-1}(U_\alpha) = \bigcup_{p\in U_\alpha} T_pM$. If $v \in \tilde{U}_\alpha$ then $v \in T_pM$ for some $p \in U$, so $v = \sum_{i=1}^n \lambda_i \frac{\partial}{\partial x_i} |_p$ for some $\lambda_i \in \mathbb{R}$. Then the obvious choice for $\tilde{\varphi}_\alpha$ and $\tilde{V}_\alpha$ is

$$\tilde{\varphi}_\alpha(v) = (\varphi(p), \lambda_1, \lambda_2, ..., \lambda_n) \subset \tilde{V}_\alpha := V_\alpha \times \mathbb{R}^n.$$

It is clear that $\tilde{U}_\alpha$ cover TM and we can topologise TM by enforcing that $\tilde{\varphi}_\alpha$ are all homeomorphisms. Hence TM is a vector bundle of dimension n (π is clearly continuous). We now need to check that the transistion maps are smooth.

If we fix $(\varphi_\alpha(p), \lambda_1, ..., \lambda_n) \in \tilde{V}_\alpha$, then the transistion map $\tilde{\varphi}_\beta \circ \tilde{\varphi}_\alpha^{-1}$ sends $(\varphi_\alpha(p), \lambda_1, ..., \lambda_n)$ to $\tilde{\varphi}_\beta(\sum_{i=1}^n \frac{\partial}{\partial x_i}|_p)$. We then have:

$$\tilde{\varphi}_\beta\left(\sum_{i=1}^n \frac{\partial}{\partial x_i}|_p\right) = \tilde{\varphi}_\beta\left(\sum_{i=1}^n \lambda_i \sum_{j=1}^n \frac{\partial x_j^\beta}{\partial x_i^\alpha}\frac{\partial}{\partial x_j^\beta}|_p\right) =$$

$$= \tilde{\varphi}_\beta\left(\sum_{j=1}^n (\sum_{i=1}^n \lambda_i \frac{\partial x_j^\beta}{\partial x_i^\alpha})\frac{\partial}{\partial x_j^\beta}|_p\right) = \left(\varphi_\beta(p), \sum_{i=1}^n \lambda_i \frac{\partial x_1^\beta}{\partial x_i^\alpha}, \dots, \sum_{i=1}^n \lambda_i \frac{\partial x_n^\beta}{\partial x_i^\alpha}\right) \in \tilde{V}_\beta,$$

which is clearly smooth as the φ_α formed a smooth atlas on M and the other terms are just linear combinations of partial derivatives. This completes the proof of TM being a smooth vector bundle.

To see that TM is also a smooth manifold, we need to conclude that it has a countable basis and is Hausdorff. That TM has a countable basis is true simply by definition since we defined its topology in terms of a countable smooth atlas. Now we show TM is Hausdorff. Let v_1,$v_2 \in TM$ and assume to begin with that $\pi(v_1) = \pi(v_2) = p$. But then there is an open set U containing both v_1 and v_2 homeomorphic to $\tilde{U} \subset \mathbb{R}^{2n}$ which is Hausdorff, so we can use the homeomorphism to easily find disjoint open sets in TM containing v_1 and v_2. Now assume that $\pi(v_1) \neq \pi(v_2)$. Then, since M is Hausdorff, we can find disjoint open sets $U \ni \pi(v_1)$ and $V \ni \pi(v_2)$ in M. Now, $\pi^{-1}(U)$ and $\pi^{-1}(V)$ are clearly disjoint and contain v_1 and v_2, respectively. This completes the proof of TM being a smooth manifold. That it is of dimension $2n$ is clear since the $\tilde{V}_\alpha$ were open subsets of $\mathbb{R}^{2n}$. □

Note that the last section of this proof did not depend on the smooth vector bundle being TM specifically, and the argument generalises without any effort to all smooth vector bundles. That is, if $p : E \to M$ is a smooth k-dimensional vector bundle over an n-dimensional smooth manifold M, then E is a smooth manifold of dimension $n+k$. Also note that for historical reasons, sections of the tangent bundle are referred to as *vector fields.*

Definition 5.6. Let M be a smooth manifold. We say that M is *parallelisable* if TM is the trivial bundle. If TM is only stably trivial, then we say that M is *stably parallelisable.*

Note that S^1 is parallelisable; simply rotate all of the tangent spaces by an angle π. For S^2, this is not true. The fact that it is not parallelisable is a consequence of the famous Hairy-ball theorem, which says that any vector field on S^2 must vanish somewhere. This means that it is impossible for TS^2 to admit two linearly independent vector fields, which means that it cannot be trivial. The Hairy-ball theorem actually extends to all even-dimensional spheres, so S^{2n} is not parallelisable. Actually, the only parallelisable spheres are S^1, S^3 and S^7. This is amazingly linked to the only division algebras being the real numbers $\mathbb{R}$, the complex numbers $\mathbb{C}$, the quarternions $\mathbb{H}$ and the octonions $\mathbb{O}$.

Now we study a certain smooth vector bundle for smooth manifolds called the *normal bundle*. Unlike the tangent bundle, it is not possible to define this without reference to some ambient space. This means that our manifold needs to embedded inside some larger manifold.

Definition 5.7. Let $M \subset M_0$ be smooth manifolds. Then M is called a *submanifold* of M_0 if the inclusion map $M \hookrightarrow M_0$ is an embedding.

Consider M a submanifold of some M_0. Now we can consider $T_M M_0$, the restriction of TM_0 to M. Note that TM is a sub-bundle of this bundle $T_M M_0$ since for any $p \in M$, $T_p M$ must be a vector subspace of $T_p M_0$. Since manifolds are paracompact, this means we can define an inner product on $T_M M_0$ and hence we can find the orthogonal complement of TM using 3.20. This means we can make the following definition.

Definition 5.8. Let M be a submanifold of M_0. Then let the *normal bundle* NM of M be defined as the orthogonal complement $TM^\perp$ where is TM is considered as sub-bundle of $T_M M_0$.

3.20 tells us that NM is a vector bundle, but we do not know that it is a smooth vector bundle. We can turn it into a smooth vector by making a specific choice of inner product on M_0: a *Riemannian metric*. Unfortunately, we do not have the space to define these and prove there existence formally and so we will take it on faith that we can make a choice of inner product that gives NM a smooth structure.

Note that the well-definedness of the normal bundle immediately gives us that S^n is stably parallelisable for all n. First we embed S^n inside $\mathbb{R}^{n+1}$ and then the sum $TS^n \oplus NS^n$ is clearly trivial as $T\mathbb{R}^{n+1}$ is trivial. Now note that NS^n must be a orientable line bundle over S^n, and hence by our classification from 4 is trivial (as $SO(1) = \{\text{pt}\}$ and so $\pi_n(SO(1))$ is trivial).

Very related to the concept of a normal bundle is that of a *tubular neighbourhood*, which we now define.

Definition 5.9. Let M be an k-dimensional submanifold of an n-dimensional smooth manifold M_0. Then a *tubular neighbourhood* of M is a subset $E \subset M_0$ with the structure of a $n-k$-dimensional smooth vector bundle over M with M as the zero section.

If M is closed, then NM will give a tubular neighbourhood for M. This is intuitive, but sadly not immediate. To prove this, we would need to show that we could extend our embedding of M in M_0 to an embedding of NM. Once again, proving this requires referring to a Riemannian structure on our smooth manifolds and so a proof is out of reach of the scope of this project. In fact, the normal bundle is the only such structure that will result in a tubular neighbourhood. These results are collated by the following theorem.

Theorem 5.10. *Let M be a closed submanifold of M_0. Then NM is a tubular neighbourhood of M. Furthermore, any tubular neighbourhood of M is isomorphic to NM.*

For the reader familiar with Riemannian geometry, one can find the proof in [1].

If M is not closed, the situation is more complicated. However, if the boundary of M interacts with the boundary of M_0 in a sufficiently nice manner, then we can say something similar to the above. First, we should state what we mean by this 'nice manner'.

Definition 5.11. Let M be an k-dimensional submanifold of an n-dimensional smooth manifold M_0. Then M is a *neat submanifold* of M_0 if it satisfies the following:

(a) M is a closed subset of M_0;

(b) $\partial M = M \cap \partial M_0$;

(c) For all $x \in \partial M$ we have a chart $\varphi : U \to \mathbb{R}^n_+$ with $x \in U$, such that $\varphi^{-1}(\mathbb{R}^k_+) = U \cap M$ (where $\mathbb{R}^k_+$ is thought to be the subspace of $\mathbb{R}^n_+$ consisting of the last k coordinates.)

This definition is a little opaque so it helps to have the conditions described somewhat. (b) is the requirement that all of ∂M is contained in ∂M_0 and that none of the interior of M is in ∂M_0. (c) is the requirement that M meets ∂M_0 akin to how a k-dimensional hyperplane meets $\partial \mathbb{R}^n_+$.

Definition 5.12. Let M be a neat submanifold of M_0 with tubular neighbourhood E. Then we say that E is a *neat tubular neighbourhood* if $E \cap \partial M_0$ is a tubular neighbourhood of ∂M inside ∂M_0.

We then have an equivalent theorem to 5.10 for neat submanifolds.

Theorem 5.13. *Let M be a neat submanifold of M_0. Then M has a neat tubular neighbourhood. Furthermore, it is unique up to isomorphism.*

Once again, a proof can be found in [1]. We will make use of neat tubular neighbourhoods later in 8.

Chapter 6
Homotopy spheres

(Unless specified, the material in this chapter is based on [1], [2] and [5].)

We now define the most important object of this project, *homotopy spheres*. The rest of this section is dedicated to understanding these homotopy spheres and seeing how they relate to the *connected sum* operation on manifolds, which we define later.

Definition 6.1. A *homotopy n-sphere* is a n-dimensional manifold that is homotopy equivalent to S^n.

This definition makes a great deal of sense, but is generally unwieldy to work with since constructing homotopy equivalences in general may be difficult. To get an equivalent statement we need to use two standard results from algebraic topology.

Theorem 6.2. *[Whitehead] Let X, Y be connected CW-complexes. Let $f : X \to Y$ be a continuous map such that f induces isomorphisms between the homotopy groups:*

$$f_* : \pi_n(X) \xrightarrow{\cong} \pi_n(Y) \text{ for all } n \in \mathbb{N}_0.$$

Then f is a homotopy equivalence.

Theorem 6.3. *[Hurewicz] Let X, A be spaces with $A \neq \emptyset$ simply connected and let the pair (X, A) be $n-1$ connected. Then $H_i(X, A) = 0$ for $i < n$ and $\pi_n(X, A) \cong H_n(X, A)$.*

For proofs of these see [2]. We can use these two results to prove the following lemma.

Lemma 6.4. *Let X, Y be spaces with X simply connected and let $f : X \to Y$ be a continuous map inducing*

$$f_* : H_n(X) \xrightarrow{\cong} H_n(Y)) \text{ for all } n \in \mathbb{N}_0.$$

Then f is a homotopy equivalence.

Proof. First consider the mapping cylinder M_f which is defined as:

$$M_f := \frac{([0,1] \times X) \sqcup Y}{\sim}, \quad (0, x) \sim f(x).$$

This space looks like Y with a cylinder extending from $f(X) \subset Y$ to X. We claim that M_f deformation retracts to Y, which is intuitively obvious. The details of proving this are somewhat more complicated and technical, so we refer to [2] instead of giving the full argument. What this means is that in the context of proving the theorem, we may replace Y with M_f. Furthermore, we may assume that f is the inclusion map into M_f since it is clearly homotopic to the inclusion map. Then we can use the long exact sequence of the pair (M_f, X) to conclude that $H_n(M_f, X) = 0$ for all n, since the inclusion induces isomorphisms $H_n(X) \to H_n(M_f)$ for all n. Then by repeated applications of 6.3 we conclude that $\pi_n(X, M_f) \cong H_n(X, Mf)$ for all n, which implies f induces isomorphisms on all homotopy groups. Then we use 6.2 to conclude that f is a homotopy equivalence. ∎

This gives us our equivalent characterisation of homotopy spheres.

Proposition 6.5. *An n-dimensional manifold M is a homotopy n-sphere if and only if $\pi_1(M) = 0$ and $H_n(M) \cong H_n(S^n)$ $\forall$ $n \in \mathbb{N}_0$.*

Proof. The forwards implication is clear as homotopy equivalence implies that the homotopy and homology groups are equal. For the reverse implication we will use 6.4. First, we need to construct a map from M to S^n. Take a point on M, then there exists a small embedded D^n containing it since M is a manifold. Then identify the boundary of the disc to a single point which gives a map from that disc to S^n. We can then extend the map onto the rest of M by simply mapping everything to the same point we mapped the boundary to. Since by assumption M is simply connected and has isomorphic homology groups to S^n, the map must induce these isomorphisms and hence by 6.4 we are done. □

We now define the notion of *connected sum* which is an operation for composing two smooth manifolds to form a new one. Roughly speaking, the operation is to connect both manifolds via a tube.

Definition 6.6. Let M_1, M_2 be two n-dimensional smooth manifolds with embeddings $h_i : D^n \to M_i$. If our manifolds are oriented we assume h_1 is orientation preserving and h_2 is orientation reversing. Then define the *connected sum* of M_1 and M_2, written $M_1 \# M_2$, as the disjoint union of $M_1 \setminus \{h_1(0)\}$ and $M_2 \setminus \{h_2(0)\}$ quotiented by the identification $h_1(rx) \sim h_2((1-r)x)$ (where $x \in S^{n-1}$ and $r \in [0, 1]$).

We have a choice here in our embeddings of D^n, but the implicit claim here is that $M_1 \# M_2$ does not depend on these. Using the disc theorem (2.13), if we consider the connected sum of M_1 and M_2 instead using $\tilde{h_1} : D^n \to M_1$ then we see that there exists a diffeomorphism $f : M_1 \to M_1$ such that $\tilde{h_1} = f \circ h_1$ and so the resulting manifold will be diffeomorphic to our original $M_1 \# M_2$. (The exact same argument applies for h_2.) We now show that this operation gives us a smooth manifold.

Proposition 6.7. *Let M_1 and M_2 be n-dimensional smooth manifolds. Then $M_1 \# M_2$ is itself a smooth manifold of dimension n.*

Proof. We have three things to show: Firstly, that $M_1 \# M_2$ has a countable basis; secondly, that it is Hausdorff; and thirdly, that it has a smooth structure.

Since M_i has a countable basis, so does $M_i \setminus \{h_i(0)\}$. We now use invariance of domain to transfer these countable bases onto our connected sum. Since the projections $M_i \setminus \{h_i(0)\} \hookrightarrow M_1 \# M_2$ are embeddings, invariance of domain gives us that these are open maps. This means that we have a countable basis of $M_1 \# M_2$ given by the images of the basis elements from our constituent manifolds.

Showing that $M_1\#M_2$ is Hausdorff is simply a routine calculation. Firstly, define (just to simplify notation)

$$\alpha : D^n \setminus \{0\} \to D^n \setminus \{0\}$$
$$\alpha(v) = (1 - |v|)\frac{v}{|v|}.$$

Then note that $g = h_2 \circ \alpha \circ h_1^{-1} : M_1 \setminus \{h_1(0)\} \to M_2 \setminus \{h_2(0)\}$ is injective since all three constituent functions are injective (in fact they are all diffeomorphisms). Then it is a matter of checking cases. Let $x \in M_1 \setminus \{h_1(0)\}$, $y \in M_2 \setminus \{h_2(0)\}$ and assume that $g(x) \neq y$ i.e. x and y really represent different points in $M_1\#M_2$. Then x and $g^{-1}(y)$ have neighbourhoods N_x, $N_y \subset M_1 \setminus \{h_1(0)\}$ disjoint since M_1 is Hausdorff. Furthermore, $g(N_x)$ and $g(N_y)$ must also be disjoint since if they were not, there would be an element in their intersection whose preimage under g would be inside both N_x and N_y. This means that x and y have disjoint neighbourhoods in $M_1\#M_2$ given by the projection maps $M_i \setminus \{h_i(0)\} \hookrightarrow M_1\#M_2$. All other cases proceed very similarly and so are left to the reader.

Finally, we have to see that $M_1\#M_2$ admits a smooth structure. Here we use that g is a diffeomorphism and that the projection maps $M_i \setminus \{h_i(0)\} \hookrightarrow M_1\#M_2$ are open maps. This means that we can take the smooth structures on M_i and map them down onto $M_1\#M_2$ where their union will be compatible since g is a diffeomorphism. □

To see that this can be viewed as joining together M_1 and M_2 via a tube, imagine the case $n = 2$. If you view both M_1 and M_2 embedded in some ambient space, you can imagine puncturing the two embedded discs and then by using the hole open them up. Then the identification corresponds to overlaying those two sections which gives you a tube between the two of them.

Theorem 6.8. *The set of closed, oriented, connected smooth manifolds of dimension n is a commutative monoid under the connected sum operation with identity element S^n.*

Proof. We have three things to show: associativity, commutativity and that S^n acts an the identity.

For the first, note that we can choose any two embeddings h_1,$h_2 : D^n \to M$ such that $h_1(D^n) \cap h_2(D^n) = \emptyset$. Let $h_1 : D^n \to M_1$, $h_2^{\pm} : D^n \to M_2$ and $h_3 : D^n \to M_3$ be four embeddings such that $h_2^{\pm}$ are chosen so that they do not intersect. Since the two embedded discs in M_2 do not intersect, the equivalence relations in the quotient do not interact at all and hence commute. This means that both $M_1\#(M_2\#M_3)$ and $(M_1\#M_2)\#M_3$ are given by the expression:

$$\frac{M_1 \setminus \{h_1(0)\} \sqcup M_2 \setminus \{h_2^{\pm}(0)\} \sqcup M_3 \setminus \{h_3(0)\}}{h_1(rx) \sim h_2^-((1-r)x),\ h_2^+(rx) \sim h_3((1-r)x)}$$

and hence they are both equal.

For commutativity, suppose that to construct $M_1 \# M_2$ we used embeddings h_i. We need to swap which one of these is orientation preserving and which is orientation reversing, and we do this using the map $R : D^n \to D^n$ which reflects one of the coordinates (assume it is the first coordinate without loss of generality). Then construct $M_2 \# M_1$ using $\tilde{h}_i = h_i \circ R$. Now to find a diffeomorphism between $M_1 \# M_2$ and $M_2 \# M_1$, take the map that is the identity on both $M_1 \setminus \{h_1(0)\}$ and $M_2 \setminus \{h_2(0)\}$. To see that this map is well defined, let $x = (x_1, ..., x_n) \in S^{n-1}$ and take $h_1(rx) \in M_1 \# M_2$. This maps to $h_1(rx) \in M_2 \# M_1$ which is equal to $\tilde{h}_1(r(-x_1, x_2, ..., x_n))$. Using the equivalence relation, we see that this is equivalent to $\tilde{h}_2((1-r)(-x_1, x_2, ..., x_n))$, which is simply $h_2((1-r)x)$ as desired. The fact that this map is a diffeomorphism is due to the smooth structures on both connected sums being directly inherited from the smooth structures on M_1 and M_2, and the identity clearly being a diffeomorphism.

Finally, since $S^n \setminus \{h_2(0)\}$ is isotopic to a disc, we can contract the $S^n \setminus \{h_2(0)\}$ back to the original embedded disc $h_1(D^n)$. Note that for the first claim we have crucially used the standard smooth structure on S^n. □

Now, every monoid contains a group made from the set containing all invertible elements in the monoid. We define A^n to be the group of all invertible elements in the monoid of closed, oriented, connected smooth manifolds of dimension n. We now want to understand this group and so we prove the following lemma.

Lemma 6.9. *For closed, oriented, connected smooth manifolds M_1 and M_2, $M_1 \# M_2$ is a homotopy sphere if and only if both M_1 and M_2 are also homotopy spheres.*

Proof. Here we make use of the characterisation of homotopy spheres that we developed earlier. The first step is to compute the homology groups of $M_1 \# M_2$. Note that by using Mayer-Vietoris we can see that for a closed, connected smooth manifold M

$$H_i(M \setminus D^n) \cong \begin{cases} H_i(M) & \text{if } i < n \\ 0 & \text{if } i = n \end{cases}$$

which allows us to apply Mayer-Vietoris a second time directly to the connected sum. Split $M_1 \# M_2$ up into $M_1 \setminus \{h_1(0)\}$ and $M_2 \setminus \{h_2(0)\}$ and notice that both of these are homeomorphic to $M_i \setminus D^n$ and that their intersection is homotopy equivalent to S^{n-1}. Then we can use Mayer-Vietoris, with that decomposition, to show that $H_i(M_1 \# M_2) = H_i(M_1) \oplus H_i(M_2)$ for $0 < i < n$. For $i = 0$ or n, use that $M_1 \# M_2$ has a fundamental class and so $H_n(M_1 \# M_2) \cong \mathbb{Z}$, and that $M_1 \# M_2$ is connected and so $H_0(M_1 \# M_2) \cong \mathbb{Z}$. This gives us that $M_1 \# M_2$ has the homology of the sphere if and only if M_1 and M_2 both have the homology of the sphere.

For the fundamental group we use Seifert-van Kampen. If $n \geq 3$ then this gives us that: $\pi_1(M) \cong \pi_1(M \setminus D^n)$ since S^{n-1} is simply-connected and D^n is contractible. Then, using Seifert-van Kampen again gives that $\pi_1(M_1 \# M_2) \cong \pi_1(M_1) \times \pi_1(M_2)$, the free product of the fundamental groups of the original manifolds. So similarly we get that $M_1 \# M_2$ is simply-connected if and only if M_1 and M_2 are both simply-connected for $n \geq 3$.

For $n < 3$, Seifert-van Kampen does not work as S^{n-1} is no longer simply-connected. Instead, we appeal to the topological classification of surfaces and 1-manifolds. For surfaces, the only homotopy sphere is homeomorphic to S^2 and the connected sum of two surfaces is only homeomorphic to S^2 if both were individually. The argument is similar for 1-manifolds. $\square$

Proposition 6.10. *Elements of A^n are homotopy spheres.*

Proof. Any element in A^n is invertible and so for $M \in A^n$ there exists $-M$ such that $M \# -M = S^n$. Using the previous lemma implies that M is a homotopy sphere. $\square$

It turns out that the elements of A^n are actually topologically spheres (i.e. homeomorphic to S^n). But to prove it, we need a theorem which is known as the generalised Schoenflies theorem., due to B. Mazur. For a proof, see [6].

Theorem 6.11. *Suppose $f : S^{n-1} \times [-1,1] \to S^n$ is a topological embedding. Then $S^n \setminus f(S^{n-1} \times \{0\})$ consists of two components and each component's closure is homeomorphic to D^n.*

Proposition 6.12. *If $M_1 \# M_2$ is homeomorphic to S^n then both M_1 and M_2 are homeomorphic to spheres.*

Proof. Let D be an embedded disc in M_1 and construct $M_1 \# M_2$ using D. There exists $f : M_1 \# M_2 \to S^n$ a homeomorphism by assumption and this gives us a topological embedding of the boundary of D into S^n. Since the boundary of D is an embedding of S^{n-1} this is almost enough to use the generalised Schoenflies theorem, but we need to know that we can extend this embedding to an embedding of $S^{n-1} \times [-1,1]$. We can do this because we can always find another embedded disc D^* such that $D \subset D^*$, and then we extend the embedding of S^{n-1} using the embedding of D^*. The generalised Schoenflies theorem then tells us that the closure of $M_1 \setminus D$ is homeomorphic to D^n which means that M_1 is homeomorphic to the union of two discs identified along their boundaries. Constructing a homeomorphism from M_1 to S^n is then straightforward. The argument is then identical for M_2. $\square$

This means that the group A^n consists of manifolds homemorphic to S^n, but not necessarily diffeomorphic to S^n, and as such any non-trivial elements are exotic spheres. However, it should be noted that this group does not need to contain all possible exotic spheres as it only contains those which are invertible under the connected sum operation. An equivalent way to view A^n is then the group of all invertible smooth structures on S^n, where by an invertible smooth structure we mean one that is invertible using the connected sum as described above.

Chapter 7
The h-cobordism theorem

(Unless specified, the material in this chapter is based on [1] and [7].)

An important result for us is the *h-cobordism theorem* which was proved by S. Smale in 1962. An important corollary of the theorem is the generalised Poincaré conjecture which says that homotopy spheres are homeomorphic to spheres in dimension $n \geq 6$.

Definition 7.1. An n-dimensional *cobordism* between closed manifolds M and N, denoted $(W; M, N)$ is a triple of manifolds such that W is dimension n and $\partial W = M \sqcup N$.

Example 7.2. The *trivial cobordism* is simply the product cobordism $(V \times I; V, V)$. This generalises the idea of a cylinder, which is the trivial cobordism where $V = S^1$. Another example is the so-called *pair of trousers* which is given by removing three discs from a sphere. This can be thought of as a cobordism $(M; S^1, S^1 \sqcup S^1)$.

Definition 7.3. An n-dimensional *h-cobordism* $(W; M; N)$ is an n-dimensional cobordism for which the inclusions $M \hookrightarrow W, N \hookrightarrow W$ are both homotopy equivalences. We then say that M and N are *h-cobordant*.

We can now state the powerful h-cobordism theorem which will form the focus for the rest of this section. For a proof, see [1] or [7].

Theorem 7.4. *Let W, M, N be simply connected, closed smooth manifolds with $dim(W) \geq 6$. Then if $(W; M, N)$ is an h-cobordism, it is necessarily the trivial cobordism. That is, W is diffeomorphic to $M \times I$.*

One might notice that the smooth manifold N has disappeared in the conclusion. This is explained by the following corollary.

Corollary 7.5. *If simply connected, closed smooth manifolds M, N are h-cobordant and $dim(M) = dim(N) \geq 5$ then M is diffeomorphic to N.*

Proof. This is more or less a direct corollary, but it is worth spelling it out since the result is so important. If M, N are h-cobordant then there exists W such that $\partial W = M \sqcup N$ and $\dim(W) \geq 6$. This implies, by 7.4, that W is diffeomorphic to both $M \times I$ and $N \times I$. Then $M \times I$ is diffeomorphic to $N \times I$ and since diffeomorphisms preserve boundaries $M \sqcup M$ is diffeomorphic to $N \sqcup N$, which implies M diffeomorphic to N. □

We now use the h-cobordism theorem to prove some non-trivial results. First we give a characterisation of the n-disc.

Proposition 7.6. *Let W be a simply connected, compact, closed smooth manifold with $dim(W) = n \geq 6$ and a simply connected boundary. Then W is diffeomorphic to D^n if and only if*

$$H_i(W; \mathbb{Z}) \cong \begin{cases} \mathbb{Z} \text{ if } i = 0 \\ 0 \text{ else} \end{cases} .$$

Proof. Since W is a smooth manifold we can find an embedded n-disc D_0. By assumption $H_i(W, D_0) = 0$ for all $i \in \mathbb{N}$ and so by excision we have that $H_i(W \setminus \mathring{D}_0, \partial D_0) = 0$ also. We now show that $(W \setminus \mathring{D}_0; \partial D_0, \partial W)$ is an h-cobordism. This means we need to show that the inclusions of both of these boundaries are homotopy equivalences. Since $H_i(W \setminus \mathring{D}_0, \partial D_0) = 0$ we have that $H_i(W \setminus \mathring{D}_0) \cong H_i(\partial D_0)$ and the inclusion map clearly induces this isomorphism, which means we can use 6.4 to show that this is a homotopy equivalence. For ∂W the argument is much the same but first we need to show that $H_i(W \setminus \mathring{D}_0, \partial W) = 0$. This follows since Poincaré duality for manifolds with boundary gives us that $0 = H_i(W/\mathring{D}_0, \partial D_0) = H^{n-i}(W/\mathring{D}_0, \partial W)$. Then, by the universal coefficients theorem, we have that $H_i(W \setminus \mathring{D}_0, \partial W) = 0$. After that the argument proceeds in exactly the same manner.
Clearly, $W \setminus \mathring{D}_0$, ∂D_0 and ∂W are all simply connected since $n \geq 6$ or due to assumption, which means we can use the h-cobordism theorem. Now this gives us that $W \setminus \mathring{D}_0$ is diffeomorphic to $\partial D_0 \times I$ which itself is diffeomorphic to a disc with a smaller disc removed from inside it. So $W \setminus \mathring{D}_0$ is diffeomorphic to an annulus and hence W is diffeomorphic to an n-disc. □

We now prove an important lemma about isotopies which is useful in a wide variety of situations, but in particular we want it for the final theorem of this section. This is known as the *Alexander trick*.

Lemma 7.7. *Let $f : S^{n-1} \to S^{n-1}$ be a homeomorphism, then it can be extended to a homeomorphism $\tilde{f} : D^n \to D^n$. Furthermore, any two extensions f_1, f_2 are isotopic through extensions.*

Proof. Roughly speaking this is a statement about existence and uniqueness of extensions. For existence it is fairly simple, we can construct one extension explicitly as

$$\tilde{f}(x) = \begin{cases} |x| f(x/|x|) & \text{if } x \in D^n \setminus \{0\} \\ 0 & \text{if } x = 0 \end{cases}$$

which is clearly a homeomorphism that restrics to f on the boundary. One can also see that if f were a diffeomorphism, this extension would still not be a diffeomorphism at $x = 0$.
For uniqueness we show that any extension $\tilde{g}$ is isotopic to the above f through extensions, which suffices to prove the statement. Let $F(x,t) : D^n \times [0,1] \to D^n$ be the isotopy between these defined as

$$F(x,t) = \begin{cases} \tilde{f}(x) & \text{if } |x| \geq t \\ t\tilde{g}(x/t) & \text{if } |x| \leq t. \end{cases}$$

Note that this definition makes sense as for $|x| = t$, $t\tilde{g}(x/t) = |x| f(x/|x|) = \tilde{f}(x)$ since $|x/t| = 1$. In words, what this isotopy does is it replicates $\tilde{g}$ to scale on smaller and smaller discs, while $\tilde{f}$ fills up a larger and larger proportion of the disc. The fact that $\tilde{g}$ completely disappears at $t = 1$ is another way of seeing that this trick does not work if we were to be extending diffeomorphisms. Finally, since any extension is isotopic to a given extension $\tilde{f}$, all extensions are isotopic to one another which completes the proof. □

We can now prove the most important theorem of this section, which is the generalised Poincaré conjecture.

Theorem 7.8. *Let Σ be an n-dimensional homotopy sphere with $n \geq 6$. Then Σ is homeomorphic to S^n.*

This actually holds for all n but the proofs for different values of n are very different. For $n = 0$, 1 or 2 the result follows either trivially or from a classification theorem. For $n = 3$ this is the famous Poincaré conjecture which was proved by Grigori Perelman in 2003. The $n = 4$ case was proved by Michael Freedman following from the disc embedding theorem. The case $n = 5$ will be shown later.

Proof. We first show that we can decompose Σ into a union of two discs. Let D_0 be some embedded n-disc in Σ and consider $\Sigma \setminus \mathring{D}_0$. By considering the (slightly degenerate) cobordism $(\Sigma \setminus \mathring{D}_0; \partial D_0, \emptyset)$, Poincaré duality gives us that $H_i(\Sigma \setminus \mathring{D}_0) = H^{n-i}(\Sigma \setminus \mathring{D}_0, \partial D_0)$. Then excision gives us that this is the cohomology of a point, so $\Sigma \setminus \mathring{D}_0$ has the homology of a point. So, by 7.6 we have that $\Sigma \setminus \mathring{D}_0$ is diffeomorphic to some n-disc.
This means that we can write Σ as the disjoint union of two n-discs where we identify the boundaries using some diffeomorphism f. This means that Σ is what we call a *twisted sphere*. We want to construct a homeomorphism from M to S^n which we can do by first mapping one of the n-discs to one of the hemispheres of S^n via the standard map which is, in fact, a diffeomorphism. We then want to extend this map onto the rest of the sphere, which means extending this diffeomorphism from $S^{n-1} \to S^{n-1}$ onto a map $D^n \to D^n$. This is precisely the Alexander trick (7.7) and so we have found a homeomorphism from our twisted sphere to the standard sphere. □

In our proof of the Alexander trick it was noted that we could not extend the map as a diffeomorphism, which is precisely what gives us the possibility of exotic spheres.

Chapter 8

Groups of homotopy spheres

(Unless specified, the material in this chapter is based on [1] and [5].)

We will now put together the work from previous sections to construct the *group of homotopy spheres* and study it in some detail. This group was first defined and studied by Kervaire and Milnor in 1962, following on from Milnor's work on exotic spheres. We saw in the previous section that for simply-connected smooth manifolds with dimension ≥ 5, h-cobordism was equivalent to diffeomorphism and hence is an equivalent relation. We now show that this is not a special case and that h-cobordism is always an equivalence relation between smooth manifolds. We will actually show the stronger result that it is an equivalence relation between oriented smooth manifolds. Recall that if M and N are oriented smooth manifolds, then M and N are h-cobordant if there exists an oriented smooth manifold W such that $(W; M, -N)$ is an oriented h-cobordism.

Lemma 8.1. *The h-cobordism relation on oriented smooth manifolds is an equivalence relation. That is, it is reflexive, symmetric and transitive.*

Proof. Let M, N, O be oriented smooth manifolds. $(M \times I; M, -M)$ is clearly an oriented h-cobordism and so we have that M is h-cobordant to itself. This means that our relation is reflexive. If $(W; M, -N)$ is an h-cobordism, then $(-W; N, -M)$ must also be an h-cobordism as this is simply flipping the orientations. This gives us that our relation is symmetric. Finally, let $(W_1; M, -N)$ and $(W_2; N, -O)$ be two oriented h-cobordisms. Then $W_1 \cap W_2$ is well defined as an oriented manifold as the intersection over the shared boundary N matches in terms of orientations. This means that $(W_1 \cap W_2; M, -O)$ is an oriented h-cobordism and so our relation is transitive. This completes the proof. □

Now we know that h-cobordism defines an equivalence relation, we can use it do define the *group of homotopy spheres*, the object that gives this section and project its name.

Definition 8.2. The *group of homotopy spheres*, denoted Θ_n, is the set consisting of all h-cobordism classes of homotopy n-spheres.

The name suggests that it can be given a group structure, which we will now show through a series of lemmas.

Lemma 8.3. *Let M be a simply-connected, oriented, n-dimensional smooth manifold. Then M is h-cobordant to S^n if and only if M is the boundary of a contractible smooth manifold.*

Proof. First, assume that M is h-cobordant to S^n. So $(W; M, -S^n)$ is an h-cobordism. Now we can attach an $(n+1)$-disc to the S^n to form a new smooth manifold W', which has $\partial W' = M$. Now because $(W; M, -S^n)$ was an h-cobordism, we get that W' is homotopy equivalent to D^{n+1} and hence is contractible.

Conversely, assume that there exists some contractible W with $\partial W = M$. Compare the situation now to that in 7.6. We have condition that W has the homology of a point and hence, if we follow the proof for **??**, we get that M is h-cobordant to S^n, provided that $n > 1$, since S^1 is not simply-connected. For the $n = 1$ case, by classification of 1-manifolds we have the stronger result that M must be diffeomorphic to S^1. □

Lemma 8.4. *Let Σ be a homotopy sphere. Then $\Sigma \# -\Sigma$ bounds a contractible manifold.*

Proof. We can consider embedding $\mathbb{R}^n \# - \mathbb{R}^n$ in $\mathbb{R}^n \times [-1, 1]$ where our original copies of $\mathbb{R}^n$ sit on $\mathbb{R}^n \times \{-1\}$ and $\mathbb{R}^n \times \{1\}$ and the 'tube' joines them together in the space in between.. The manifold that this bounds has a deformation retract that is diffeomorphic to $\mathbb{R}^n \setminus D^n$.

Since we can embed $\mathbb{R}^n$ inside a disk, we can use this construction to embed $\Sigma \# -\Sigma$ into $\Sigma \times [-1, 1]$. Let $h : \mathbb{R}^n \to \Sigma$ be an embedding and then embed $h(\mathbb{R}^n)$ inside $h(\mathbb{R}^n) \times [-1, 1]$ as above and then extend the embedding to the desired embedding. Then we see that we get the exact same phenomenon, $\Sigma \# -\Sigma$ now bounds a manifold that has a deformation retract diffeomorphic to $\Sigma \setminus D^n$. Since Σ is a homotopy sphere, $\Sigma \setminus D^n$ is homotopy equivalent to $S^n \setminus D^n$ which is contractible, completing the proof. □

Lemma 8.5. *Let M be a simply-connected, oriented smooth manifold. Then M is a homotopy sphere if and only there exists a simply-connected, oriented smooth manifold N such that $M \# N$ bounds a contractible manifold.*

Proof. First assume that M is a homotopy sphere. Then $M \# -M$ bounds a contractible manifold by 8.4.

Now assume that there exists N such that $M \# N$ bounds a contractible manifold. By 8.3, $M \# N$ is h-cobordant to S^n and hence is a homotopy sphere. Then we can use 6.9 to see that M is a homotopy sphere. □

Finally, we would like to show that the connected sum operation respects the h-cobordism relation. That is, if M_1 is h-cobordant to M_2, then $M_1 \# M$ is h-cobordant to $M_2 \# M$. However, to show this, we need to define a generalisation of the connected sum operation.

In the connected sum operation, we glued two smooth manifolds together by identifying two embedded discs (minus a single point in each disc) via a orientation reversing diffeomorphism $\alpha(v)$ (see 6.7). One way to view these discs is as tubular neighbourhoods of the submanifold consisting of a single point. This will be how we generalise the connected sum operation, essentially by replacing single points with a general closed submanifold.

Definition 8.6. Let M_1, M_2 be two n-dimensional smooth manifolds with embeddings $h_i : E \to M_i$ where E is the total space of a smooth vector bundle over a

k-dimensional, compact, closed manifold N such that $h_i(E)$ are tubular neighbourhoods of $h_i(N)$. Then define the *pasting* of M_1 and M_2 along a submanifold, written as $M(h_1, h_2)$ to be the disjoint union of $M_1 \setminus \{h_1(N)\}$ and $M_2 \setminus \{h_2(N)\}$ quotiented by the relation $h_1(v) \sim h_2(\alpha(v))$ where α is defined on each fibre as

$$\alpha(v) : \mathbb{R}^{n-k} \to \mathbb{R}^{n-k}$$
$$\alpha : v \mapsto \frac{v}{|v|^2}.$$

With the connected sum, the embeddings h_i were suppressed in the notation because the result did not depend on them. This is not the case for the pasting operation; there is no analogue of the disc theorem for general embeddings, and certainly it should depend on what the submanifold $h_i(N)$ actually is.

Proposition 8.7. *Let M_1, M_2 be two n-dimensional smooth manifolds with embeddings $h_i : E \to M_i$ where E is the total space of a smooth vector bundle over a k-dimensional, compact, closed manifold N such that $h_i(E)$ are tubular neighbourhoods of $h_i(N)$. Then $M(h_1, h_2)$ is a n-dimensional smooth manifold.*

The proof is entirely analagous to the proof of 6.7 and so is unnecessary to give here.

We can still paste along two non-closed submanifolds, provided that they are neatly embedded and have neat tubular neighbourhoods. By 5.13, the existence of the neat tubular neighbourhood is guaranteed. In particular, the case that we will want for our next proof is that we can paste along two neatly embedded *arcs*. An arc simply being a submanifold diffeomorphic to an interval.

Lemma 8.8. *Let M_1, M_2, M be smooth n-dimensional smooth manifolds with $n \geq 3$ and M_1 h-cobordant to M_2. The $M_1 \# M$ is h-cobordant to $M_2 \# M$.*

Proof. By assumption we have the oriented h-cobordism $(W; M_1, -M_2)$ and we always have the trivial cobordism $(M \times I; M, -M)$. Let A be a neatly embedded arc in W such that one endpoint is on M_1, call this point p, and the other on M_2. Then we paste W and $M \times I$ along the arcs A and $q \times I$ for some $q \in M$. Call the resulting smooth manifold from this pasting W'. Now ∂W is the disjoint union of $M_1 \# M$ and $M_2 \# M$ since it consists of the boundaries of W and $M \times I$ pasted along a single point, which we noted earlier is equivalent to the connected sum operation. So we have the cobordism $(W', M_1 \# M, -M_2 \# M)$; we need to show that the inclusions of the boundaries into W' are homotopy equivalences.

First, consider the long exact sequences of the pairs $(M_1, M_1 \setminus \{p\})$ and $(W, W \setminus A)$. By naturality, the inclusion map $i : (M_1, M_1 \setminus \{p\}) \to (W, W \setminus A)$ induces a map $i*$ on the long exact sequences. Now i induces isomorphisms $H_*(M_1) \to H_*(W)$ and $H_*(M_1, M_1 \setminus \{p\}) \to H_*(W, W \setminus A)$, and so, by the five lemma, we have that it must induce isomorphisms $H_*(M_1 \setminus \{p\}) \to H_*(W \setminus A)$. Now since we assumed $n \geq$, removing a point from M_1 means it is still simply-connected, and similarly removing an arc from W means it is still simply-connected. This means that we can use 6.4 to show that i is a homotopy equivalence.

Then we can use the Mayer-Vietoris sequences for $M_1\#M$ and W', where we write $M_1\#M$ as the union of $M_1\setminus\{p\}$ and $M\setminus\{q\}$, and write W' as the union of $W\setminus A$ and $M\times I\setminus\{q\}\times I$. Then by using that the fact that i was a homotopy equivalence, along with the five lemma again, we get that the inclusion map $j: M_1\#M \to W'$ is a homotopy equivalence. Exactly the same argument then works to show that the other inclusion $M_2\#M \to W'$ is a homotopy equivalence. In conclusion, we now have that $(W', M_1\#M, -M_2\#M)$ is an h-cobordism which completes the proof. □

Theorem 8.9. *h-cobordism classes of oriented n-dimensional smooth manifolds form a monoid under the connected sum operation where the identity element is the h-cobordism class represented by S^n. Furthermore, the group of invertible elements of this monoid is Θ_n.*

Proof. Firstly, 8.8 tells us that connected sum is well defined on h-cobordism classes, which in turn tells that h-cobordism classes of oriented n-dimensional smooth manifolds actually form a monoid under the connected sum operation. Now we know that, 8.5 tells us that the set of invertible elements consists entirely of all n-dimensional homotopy spheres, or Θ_n. □

We now give a proof of the slightly stronger generalised Poincaré conjecture. In some ways this proof makes the proof given in 7.8 irrelevant, but that proof was interesting on its own for introducing twisted spheres and for being more direct.

Theorem 8.10. *Let Σ be an n-dimensional homotopy sphere with $n \geq 5$. Then Σ is homeomorphic to S^n.*

Proof. $\Sigma\# - \Sigma$ is h-cobordant to S^n by 8.9. Then by the h-cobordism theorem (7.4), $\Sigma\# - \Sigma$ is diffeomorphic to S^n. This means that $\Sigma \in A^n$, the group of all invertible smooth structures on S^n. But was already shown in 6.12 that A^n consists of topological spheres, and so Σ is homeomorphic to S^n. □

In the course of that proof, we showed that there was some relationship between Θ_n and A^n. Not only will we now make this relationship clear, but we will also show that there is a relationship between these two groups and Γ^n (see 2.18).

Proposition 8.11. *Let $n \geq 5$. Then the map $f : A^n \to \Theta^n$ sending an element of A^n to its h-cobordism class in Θ^n is an isomorphism of groups.*

Proof. Firstly, the definition of f makes sense since any element in A^n is a homotopy sphere and hence its h-cobordism class is in Θ^n. $\ker(f)$ is the set containing all elements of A^n that are h-cobordant to S^n. Then the h-cobordism theorem 7.4 tells us that these elements are diffeomorphic to S^n and hence $\ker(f)$ is trivial. Now let $\Sigma \in \Theta^n$. Then $\Sigma \# -\Sigma$ is h-cobordant to S^n and hence diffeomorphic to S^n by the h-cobordism theorem. Therefore $\Sigma \in A^n$ and $f(\Sigma) = \Sigma$, so f is surjective. f is clearly a homomorphism which completes the proof. □

When we defined A^n, we proved that it consisted entirely of homotopy spheres, but we did not know whether it contained all homotopy spheres. This proposition essentially means that by using the h-cobordism theorem and changing our equivalence relation to that of h-cobordance rather than diffeomorphism, we have been able to show that A^n does contain all homotopy spheres. This is equivalent to showing that all of the smooth structures on S^n are invertible, at least for $n \geq 5$.

In the proof of 7.8 we introduced the twisted sphere, which is the disjoint union of two n-discs with their boundaries identified by some diffeomorphism $h : S^{n-1} \to S^{n-1}$. If we write $\Sigma(h)$ for this twisted sphere, we can formulate how Γ^n is related to our groups of homotopy spheres.

Proposition 8.12. *The map $F : \Gamma^n \to A^n$ that sends a diffeomorphism $h : S^{n-1} \to S^{n-1}$ to $\Sigma(h)$ is a well-defined injective homomorphism of groups.*

Proof. Notice that it is not actually clear that $\Sigma(h) \in A^n$. However, if we prove that $\Sigma(h)\#\Sigma(g) = \Sigma(hg)$ then this necessarily means that all $\Sigma(h)$ are invertible, since given any $\Sigma(h)$, $\Sigma(h)\#\Sigma(h^{-1})$ would then necessarily be diffeomorphic to S^n and hence $\Sigma(h)$ is invertible. Firstly, we show that this map is injective.

Let h be such that there exists a diffeomorphism $f : \Sigma(h) \to S^n$. We want to show that h is trivial in Γ^n, which means we want to show that it extends to a diffeomorphism $D^n \to D^n$. Since the restriction of f to the southern hemisphere of $\Sigma(h)$ is an embedding of a disc, we can use 2.13 to assume that f sends the southern hemisphere of $\Sigma(h)$ to the southern hemisphere of S^n via the inclusion map. Now f necessarily maps the boundary of the northern hemisphere of $\Sigma(h)$ to the equator of S^n via h, and so f restricted to the northern hemisphere of $\Sigma(h)$ is a diffeomorphism $D^n \to D^n$ that extends h. Hence F is injective.

Consider $\Sigma(h)\#\Sigma(g)$. We can view this as a disjoint union of an n-disc D^n, a cylinder $D^n \times I$ and another disc D^n where the first two boundaries are identified using h and the second two via g. Call this $\Sigma(h, g)$. We now construct a diffeomorphism from $\Sigma(h, g)$ to $\Sigma(gh, \mathrm{Id})$. First, map the disc associated to $\Sigma(g)$ to the disc associated to $\Sigma(\mathrm{Id})$ by the identity. Now we map the cylinder by mapping it via g on each slice $S^{n-1} \times \{t\}$. Now map the final disc, the one associated to $\Sigma(h)$, to the disc associated to $\Sigma(gh)$ via the identity map also. To conclude this is a diffeomorphism we need to check that is well defined on the boundaries where the identifications occur, but this is clear by how the map was defined. Hence $\Sigma(h)\#\Sigma(g)$ is diffeomorphic to $\Sigma(gh, \mathrm{Id})$, which is clearly diffeomorphic to $\Sigma(gh)$. This means that F is a homomorphism.

Finally we should check that F is well defined, that is to say that $\Sigma(he) = \Sigma(h)$ if e is trivial in Γ^n. Since we showed that F is a homomorphism, we only need to check that F maps trivial elements of Γ^n to S^n. Let $h : S^{n-1} \to S^{n-1}$ be a diffeomorphism trivial in Γ^n, which means that h extends to a diffeomorphism $D^n \to D^n$. Then we can construct a diffeomorphism from $\Sigma(h)$ to S^n as follows: first map the southern hemisphere of $\Sigma(h)$ to the southern hemisphere of S^n via the inclusion map. Then necessarily the boundary of the northern hemisphere of $\Sigma(h)$ maps to the equator of S^n via the diffeomorphism h. Then we finish off our diffeomorphism by extending it by using the extension of h onto the rest of the northern hemisphere. This means that F is well defined, which completes the whole proof. □

Theorem 8.13. *Let $n \geq 6$. Then Γ^n, A^n and Θ^n are all isomorphic as groups.*

Proof. We have already shown in 8.11 that we have an isomorphism from A^n to Θ^n. We also just showed that we have an injective homomorphism from Γ^n to A^n, so all we need to do is conclude that the composition of our isomorphisms is surjective and our proof will be complete. Let $\Sigma \in \Theta^n$ be a homotopy sphere. Then in our first proof of the generalised Poincaré conjecture (7.8), we showed that for $n \geq 6$, Σ must be a twisted sphere. This means that there exists some $h \in \Gamma^n$ that is mapped to Σ under our composition, and hence the map is surjective. □

As a consequence of this theorem, we can now say more about exotic spheres. $\Theta^n \cong \Gamma^n$ tells us that every smooth structure on S^n can be given by an atlas containing only two charts, and we already noted that $\Theta^n \cong A^n$ tells us that every smooth structure on S^n is invertible.

One might ask why we could not extend to $n = 5$ in the proof. The answer is that we could not show that 5-dimensional homotopy spheres are twisted spheres, since our proof for $n \geq 6$ rested on the h-cobordism in dimension $n - 1$. Our proof would work if the h-cobordism theorem held in dimension 4, but this is not true. However, for $n \leq 5$, Θ_n is always trivial. For $n = 1$ and 2, this is a consequence of topological manifolds admitting only one smooth structure, up to the diffeomorphism, in dimensions 1 and 2. For Θ_3, the argument is similar but also rests on the Poncaré hypothesis. Finally, Kervaire and Milnor showed that both Θ_4 and Θ_5 are trivial in [5].

It remains to be seen what the structure of Θ^n is for $n \geq 6$. This problem is hard and we will not have the space to deal with it general. Instead, we will now move to showing explicitly that it is possible for these groups to be non-trivial. More specifically, we will show that Θ^7 is non-trivial. Since $7 \geq 5$, this is equivalent to the construction of an exotic 7-sphere.

Chapter 9

Invariants of manifolds

(Unless specified, the material in this chapter is based on [3] and [4].)

If our aim is to construct an exotic sphere, we will need to have some way to determine that it is actually 'exotic'. To do so, we need an invariant of smooth manifolds that is a diffeomorphism invariant. In this section we will describe two invariants that we will use in the final section to construct the desired diffeomorphism invariant. These are the *signature* of a manifold, and *Pontrjagin classes* of vector bundles. We will start with the signature.

Let X be a space and R be a commutative ring. Recall from algebraic topology that the cup product gives us a R-bilinear map $\smile: H^k(X;R) \times H^l(X;R) \to H^{k+l}(X;R)$. Recall also that this map satisfies the graded commutativity relationship $[\varphi \smile \psi] = (-1)^{kl}[\psi \smile \varphi]$. Now consider an n-dimensional manifold M. We would like to be able to use the cup product to define a symmetric R-bilinear map $H^k(M;R) \times H^k(M;R) \to H^n(M;R)$. Clearly we need that $k = n/2$, but for the map to be symmetric we also need that k is even. This means that n must be divisible by four. In conclusion, if M is a $4n$-dimensional smooth manifold, we can define a symmetric R-bilinear map $H^{2n}(M;R) \times H^{2n}(M;R) \to H^{4n}(M;R)$.

Now, if we assume that M is compact, connected and oriented, then M has a fundamental class $[M]$. This gives us a map $H^{4n}(M;R) \to R$ which sends a cocycle $\varphi \mapsto \varphi([M])$. Now, we can form a symmetric bilinear form $H^{2n}(M;R) \times H^{2n}(M;R) \to R$ as the composition of the two maps that we have just constructed. This is called the *intersection form* of M. Any symmetric bilinear form defines a quadratic form by restricting to the diagonal, so we can define a quadratic form $Q : H^{2n}(M;R) \to R$ which sends $[\varphi] \mapsto [\varphi \smile \varphi]([M])$.

Proposition 9.1. *$Q : H^{2n}(M;\mathbb{Q}) \to \mathbb{Q}$ is non-degenerate and the torsion subgroup of $H^{2n}(M;\mathbb{Q}) \subset \ker(Q)$.*

Proof. Non-degeneracy is a simple consequence of Poincaré duality, as this means a non-trivial element of $H^{4n}(M;\mathbb{Q})$ evaluated on the fundamental class must be non-trivial in $H_0(M;\mathbb{Q}) \cong \mathbb{Q}$. The second assertion is similarly simple, since a torsion element must be mapped to $0 \in \mathbb{Q}$ as $\mathbb{Q}$ has no torsion. □

This result means that the intersection form can represented by a $b_{2n} \times b_{2n}$ matrix, where b_{2n} is the $2n$th Betti number. For such a quadratic form, we can define the *signature* to be the number of positive eigenvalues minus the number of negative eigenvalues. Now we can define the signature of a manifold.

Definition 9.2. Let M be an n-dimensional, compact, connected, oriented manifold. The *signature* of M, written $\sigma(M)$ is defined as zero if n is not divisible by 4, and if $n = 4k$ then it is defined as the signature of the quadratic form $Q : H^{2k}(M;\mathbb{Q}) \to \mathbb{Q}$ as defined previously.

Note that if M is not connected, we can define the signature as the sum of the signatures of all the connected components of M. Also, it should be said that we can define the signature of a manifold with boundary in a completely analogous way, where the intersection form is a map on relative cohomology $Q : H^{2n}(M, \partial M; \mathbb{Q}) \to \mathbb{Q}$ but we will leave out the details for expediency.

Since all real symmetric matrices are diagonalisable, an alternative way of defining this would be first tensor with $\mathbb{R}$, and then to choose a basis of $H^{2n}(M; \mathbb{R})$ such that the matrix is diagonal and then simply count the number of positive and negative entries with sign.

To understand the signature more, we will have to now study *characteristic classes*. In essence, a characteristic class is a method for assigning certain cohomology classes to vector bundles. There are four types of characterstic classes: Stefel-Whitney, Chern, Euler and Pontrjagin. It is the last of these, the Pontrjagin classes, that we are interested in, although we will define them using the Chern classes. We begin by giving an axiomatic treatment of the Chern classes. For proofs of their existence, see [3] or [4].

Theorem 9.3. *There exists a unique sequence of functions (up to multiplication by scalars) $c_i : \mathrm{Vect}^*_{\mathbb{C}}(B) \to H^{2i}(B, \mathbb{Z})$, such that the c_i satisfy the following relations for all complex vector bundles:*

(a) $c_0 = 1$*;*

(b) For any continuous $f : \tilde{B} \to B$, $c_i(f^*(E)) = f^*(c_i(E))$;

(c) Let $c = \sum_{i=0}^{\infty} c_i \in H^*(B; \mathbb{Z})$. Then $c(E_1 \oplus E_2) = c(E_1) \smile c(E_2)$;

(d) If $\dim(E) = n$, then $c_i(E) = 0$ for all $i > n$.

Note the slight abuse of notation for (b), where f^* denotes both the induced map on cohomology and the induced pullback map.

For a complex vector bundle $E \to B$, $c_i(E)$ is called the ith *Chern class* of E. We can make these unique by specifying their value on some specific vector bundle and this will also ensure that our Chern classes are non-trivial. To describe this vector bundle, we first need to introduce the *Grassmann manifold.*

Definition 9.4. Let $k \geq n$. Then the *Grassmann manifold* $G_n(\mathbb{R}^k)$ is defined as the set containing all n-dimensional hyperplanes that pass through the origin in $\mathbb{R}^k$.

We can give a topology on $G_n(\mathbb{R}^k)$ in the following manner. Firstly, define an *n-frame* in $\mathbb{R}^k$ to be a set of n linearly independent vectors in $\mathbb{R}^k$. Then let the set of all n-frames in $\mathbb{R}^k$ be denoted $V_n(\mathbb{R}^k)$. This is known as the *Stiefel manifold* and has a natural topology as an open subset of $\prod_{i=1}^{n} \mathbb{R}^k$. Then we have a natural sujection $V_n(\mathbb{R}^k) \to G_n(\mathbb{R}^k)$ that sends an n-frame to the n-dimensional hyperplane spanned by the n-frame. This allows us to view $G_n(\mathbb{R}^k)$ as a quotient space of $V_n(\mathbb{R}^k)$, where we identify all n-frames that span the same hyperplane, and so we can give $G_n(\mathbb{R}^k)$ the quotient topology.

Note that Grassmann manifolds are a direct generalisation of the real projective spaces, in that $\mathbb{RP}^n = G_1(\mathbb{R}^{n+1})$. The name suggests that these are manifolds, but we will not show this.

We can define the infinite Grassman manifold $G_n(\mathbb{R}^\infty)$ by taking the weak limit of $G_n(\mathbb{R}^k)$ as $n \to \infty$. We can do this because we have the sequence of inclusions $G_n(\mathbb{R}^k) \subset G_n(\mathbb{R}^{k+1})$ since clearly every n-dimensional hyperplane in $\mathbb{R}^k$ is also an n-dimensional hyperplane in $\mathbb{R}^{k+1}$.

We can construct a vector bundle over $G_n(\mathbb{R}^k)$, and indeed over $G_n(\mathbb{R}^\infty)$, in the following way.

Definition 9.5. Let $k \in \mathbb{N} \cup \{\infty\}$ and let $E_n(\mathbb{R}^k) = \{(e, v) \in G_n(\mathbb{R}^k) \times \mathbb{R}^k \mid v \in e\}$. Then define the *tautological bundle* as $p : E_n(\mathbb{R}^k) \to G_n(\mathbb{R}^k)$, where p is projection onto the first coordinate.

We claim, without proof, that this is a k-dimensional vector bundle. One can find a proof in [3]. This bundle is tautological in the sense that the fibre associated to every point $e \in G_n(\mathbb{R}^k)$ is e itself. More precisely, $p^{-1}(e) = e$ for all $e \in G_n(\mathbb{R}^k)$.

All of these constructions work similarly well if we replace $\mathbb{R}$ with $\mathbb{C}$ to give the complex Grassman manifolds $G_n(\mathbb{C}^k)$ and the complex tautological bundles $E_n(\mathbb{C}^k)$. In fact, these are the ones that we want.

We can now fix the constant in our Chern classes by specifying their values on $E_1(\mathbb{C}^\infty)$, the tautological line bundle. The base space of this bundle is $G_1(\mathbb{C}^\infty) = \mathbb{CP}^\infty$ and one can compute the cohomology ring of this space as

$$H^*(\mathbb{CP}^\infty; \mathbb{Z}) = \mathbb{Z}[x]$$

where x is the generator of this ring, and in fact of $H^2(\mathbb{CP}^\infty; \mathbb{Z})$. We then demand that $c_1(E_1(\mathbb{C}^\infty) = x$, which now means that our Chern classes are unique.

Lemma 9.6. *For complex vector bundles E_1, E_2, the Chern classes c_k satisfy the following product formula:*

$$c_k(E_1 \oplus E_2) = \sum_{i+j=k} (c_i(E_1) \smile c_j(E_2))$$

.

Proof. This is really just a reformulation of (c) is our definition of the Chern classes. If we take that relation and examine both sides of the equation, the right hand side gives

$$1 + c_1(E_1 \oplus E_2) + c_2(E_1 \oplus E_2) + ...$$

and the left hand side gives

$$(1 + c_1(E_1) + c_2(E_1) + ...)(1 + c_1(E_2) + c_2(E_2) + ...).$$

Expanding these brackets and comparing terms on each side of the equation then gives the required result. □

Another property of the Chern classes is that they are stable isomorphism invariants. To see this, first note that the Chern classes of a trivial bundle are trivial since the trivial bundle is a pullback of $\{\text{pt}\} \times \mathbb{R}^n$ (then we can use properties (a) and (b) from 9.3). Now we are able to use 9.6 to show that the Chern classes are invariant under Whitney sum with a trivial bundle, and hence are stable isomorphism invariants.

We now aim to construct the *Pontrjagin classes* for real vector bundles using these Chern classes. Since Chern classes are for complex vector bundles, it is not immediately clear how we might apply them to real vector bundles. We will do this via the *complexification* of real vector bundles.

Definition 9.7. The *complexification* of a real vector V is the tensor product $V \otimes \mathbb{C} = \{u + iv \mid u, v \in V\}$ where i denotes the imaginary unit. Given a real vector bundle, given by $p : E \to B$, we form the *complexification* of E by taking the complexification of each fibre $p^{-1}(x)$ to form a complex vector bundle over the same base space B. We denote this as $E \otimes \mathbb{C}$.

Definition 9.8. Given a real vector bundle $E \to B$, the ith *Pontrjagin class* $p_i(E) \in H^{4i}(B, \mathbb{Z})$ is defined as

$$p_i(E) = (-1)^i c_{2i}(E \otimes \mathbb{C}).$$

Note that we have ignored the odd Chern classes in this definition. This is because all of the odd Chern classes of a complexification of a real vector bundle are order two, and so can be expressed in terms of other characteristic classes. We will omit the proof of this for expediency, but one can find a proof in [4].

Proposition 9.9. *For real vector bundles E_1, E_2, the Pontrjagin classes p_k satisfy the following product formula, modulo elements of order two:*

$$p_k(E_1 \oplus E_2) = \sum_{i+j=k} (p_i(E_1) \smile p_j(E_2))$$

Proof. Note that for any real vector spaces V, W, we have that $(V \oplus W) \otimes \mathbb{C}$ is isomorphic to $(V \otimes \mathbb{C}) \oplus (W \otimes \mathbb{C})$. This means that if E_1, E_2 are real vector bundles over the same base space, then we have an isomorphism between $(E_1 \oplus E_2) \otimes \mathbb{C}$ and $(E_1 \otimes \mathbb{C}) \oplus (E_2 \otimes \mathbb{C})$. Now we can use 9.6 to get

$$c_k((E_1 \oplus E_2) \otimes \mathbb{C}) = \sum_{i+j=k} (c_i(E_1 \otimes \mathbb{C}) \smile c_j(E_2 \otimes \mathbb{C})).$$

Now if we mod out by elements of order two, this removes all of the odd Chern classes from the formula, giving us

$$c_{2k}((E_1 \oplus E_2) \otimes \mathbb{C}) = \sum_{i+j=k} (c_{2i}(E_1 \otimes \mathbb{C}) \smile c_{2j}(E_2 \otimes \mathbb{C})).$$

Multiplying this equation by $(-1)^k$ then gives us the required result. □

Similarly, one can show that the Pontrjagin classes have properties analogous to 9.3 very easily. That is, taking Pontrjagin classes commutes with taking pullbacks, $p_0 = 1$ and Pontrjagin classes vanish past the dimension of the vector bundle. The proofs follow directly from the results for Chern classes and so we will omit them. Note though that this means that the Pontrjagin classes are also stable isomorphism invariants. Another important property of the Pontrjagin classes is that they do not depend on the orientation of the bundle. The easiest way to prove this uses their relationship with the *Euler classes* and as such we do not have the time to prove this fact, although we will use it later. For more details, see [4].

We now introduce the concept of a *Pontrjagin number* for a $4n$-dimensional smooth manifold M. First, we define a *partition* I of a non-negative integer n to be an unordered sequence $I = i_1, ..., i_k$ such that $\sum_{j=1}^{k} i_j = n$.

Definition 9.10. Let M be a $4n$-dimensional compact, oriented smooth manifold and let I be a partition of n. Then the Ith *Pontrjagin number* $p_I(M)$ of M is defined as the integer

$$p_I(M) = p_{i_1}...p_{i_k}(M) = (p_{i_1}(TM) \smile ... \smile p_{i_k}(TM))([M])$$

where TM is the tangent bundle of M and $[M]$ denotes the fundamental class of M, as before.

Note that, since I is a partition of n, $(p_{i_1}(TM) \smile ... \smile p_{i_k}(TM)) \in H^{4n}(M;\mathbb{Z})$ and so it makes sense to evaluate this on the fundamental class.

We will need the following lemma later which concerns the Pontrjagin numbers of a boundary.

Lemma 9.11. *Let M be the boundary of some smooth $(4n+1)$-dimensional compact, oriented smooth manifold. Then all of the Pontrjagin numbers of M are zero.*

We omit the proof for expediency. More details can be found in [4].

Now our aim is to describe a relationship between the Pontrjagin numbers of a closed manifold and its signature, which is given by the Hirzebruch signature theorem. To understand the statement of this theorem, we need to define *multiplicative sequences* of polynomials. We begin by recalling some terminology from algebra.

Recall that an *algebra* A over a commutative ring R is an R-module with a multiplication operation satisfying distributivity under sums of vectors and multiplication by scalars in R. A *graded ring* R is a ring that can be written as a direct sum of additive groups $R = \bigoplus_{i=0}^{\infty} R_i$ where the multiplication operation satisfies $R_i R_j \subset R_{i+j}$. Every algebra over a ring has an underlying ring structure where we forget about the scalar multiplication operation, and so we define a *graded algebra* over a commutative ring to be an algebra over a ring whose underlying ring structure is graded. We say that a non-zero element a of a graded algebra $A^* = (A_0, A_1, ...)$ is *homogenous of degree* k if $a \in A_k$. Note that the cohomology ring is an example of a graded algebra over a commutative ring.

If $A^* = (A_0, A_1, ...)$ is a graded algebra over a commutative ring R, then we have a subgroup $A = \{1 + a_1 + a_2 + ... \mid a_i \in A_i\}$, where the group operation of multiplication is given by the formula

$$(1 + a_1 + a_2 + ...)(1 + b_1 + b_2 + ...) = 1 + (a_1 + b_1) + (a_2 + a_1 b_1 + b_2) + ...$$

Definition 9.12. Let $x_i \in A_i$. Then a sequence of polynomials $K_1(x_1), K_2(x_1, x_2), ...$ with coefficients in R is called a *multiplicative sequence* if

(a) $K_k(x_1, ..., x_k)$ is homogenous of degree k;

(b) For any $a, b \in A$ and $K(a) := 1 + K_1(a_1) + K_2(a_1, a_2) + ...$, we have that

$$K(ab) = K(a)K(b)$$

where multiplication is defined as $K(a) \in A$;

(c) The above conditions hold for any graded algebra A^*.

Now let $A^* = R[t]$ be the polynomial ring with coefficients in R. Note that this has a natural grading given by the polynomial degree. So A is the group of power series $\{1 + a_1 t + a_2 t^2 + ... \mid a_i \in R\}$. Now we state the following theorem which will allow us to classify multiplicative sequences.

Theorem 9.13. *Let $f(t) = 1 + a_1 t + a_2 t^2 + ...$ be a power series with coefficients in R. Then there exists a unique multiplicative sequence of polynomials $\{K_n\}$ with coefficients in R such that the condition $K(1 + t) = f(t)$ is satisfied.*

We will refer to such a $\{K_n\}$ as being the multiplicative sequence *belonging to* the power series $f(t)$.

Proof. We will only prove the uniqueness statement. A proof of existence can be found in [4]. Assume $\{K_n\}$ is a multiplicative sequence that satisfies $K(1+t) = f(t)$ for some power series f. Now, let $A^* = R[t_1, t_2, ..., t_n]$ be the polynomial ring in n-variables with coefficients in R. Note that this has a similiar grading by total degree, where an element $t_1^{e_1} t_2^{e_2} ... t_n^{e_n}$ has degree $e_1 + e_2 + ...e_n$.

Let $\sigma = (1 + t_1)(1 + t_2)...(1 + t_n) \in A$. Notice that the terms in the expansion $\sigma = 1 + \sigma_1 + \sigma_2 + ...$ are the elementary symmetric polynomials in the variables $t_1, ..., t_n$. We can use that $\{K_n\}$ is a multiplicative sequence to calculate

$$K(\sigma) = K(1 + t_1)K(1 + t_2)...K(1 + t_n) = f(t_1)...f(t_n).$$

Expanding the left hand side we get that $K(\sigma) = K_1(\sigma_1) + ... + K_n(\sigma_1, ..., \sigma_n)$. If we compare the term of degree n on both sides, we see that $K_n(\sigma_1, ..., \sigma_n)$ is given by the nth degree term in $f(t_1)...f(t_n)$, which is entirely determined by the power series f. To finish off the proof, note that the elementary symmetric polynomials σ_i do not have any relations in them, and hence computing $K_n(\sigma_1, ..., \sigma_n)$ for all n determines the multiplicative sequence completely. ☐

Now we will consider the graded subalgebra of $H^*(M;\mathbb{Q})$ given by $\bigoplus_{i=0}^{\infty} H^{4i}(M;\mathbb{Q})$. Note that this is where the Pontrjagin classes of TM lie.

Definition 9.14. Let $\{K_n\}$ be a multiplicative sequence of polynomials with coefficients in $\mathbb{Q}$. Then to a compact, closed, oriented, smooth n-dimensional manifold M we define the K-*genus* $K[M]$ as the evaluation $K_n(p_1, p_2, ..., p_n)([M])$ if $n = 4k$ and 0 otherwise.

Note that since K_n is homogenous of degree n, we have that $K_n(p_1, p_2, ..., p_n) \in H^{4n}(M;\mathbb{Q})$ and so can be evaluated on $[M]$.

Now we can state the Hirzebruch signature theorem, which relates the signature of a closed manifold to its Pontrjagin numbers.

Theorem 9.15. *The signature $\sigma(M)$ of a compact, closed, oriented, smooth $4n$-dimensional manifold M is equal to the L-genus $L[M]$ where $\{L_n\}$ is the multiplicative sequence of polynomials with coefficients in $\mathbb{Q}$ belonging to the power series*

$$\frac{\sqrt{t}}{\tanh\sqrt{t}} = 1 + \frac{1}{3}t - \frac{1}{45}t^2 + ... + \frac{2^{2n}B_n t^n}{(2n)!} + ...$$

where B_n is the nth Bernoulli number*.*

For reference, the Bernoulli numbers are a classical sequence of numbers that were introduced by Jakob Bernoulli and published posthumously in 1713. Their values are well known and can be easily found. The first few values are given below:

$$B_0 = 1,\ B_1 = \frac{1}{6},\ B_2 = -\frac{1}{30},\ B_3 = \frac{1}{42}.$$

A key feature of the Bernoulli numbers is that they alternate in sign after the first two positive terms. With these numbers, we can compute the first few L polynomials, given below:

$$L_1 = \frac{1}{3}p_1,\ L_2 = \frac{1}{45}(7p_2 - p_1^2),\ L_3 = \frac{1}{945}(62p_3 - 13p_2p_1 + 2p_1^3).$$

An intriguing consequence of the signature theorem is that, due the signature being integer valued, these particular linear combinations of Pontrjagin numbers must also be integers for all M, which gives us a highly non-trivial divisibility relationship.

Before we can give a proof of this, we need to introduce a new concept and prove a short lemma concerning the K-genus. In 8, we showed that we can use the h-cobordism relation to form a monoid under the connected sum operation. We can introduce a different algebraic structure of oriented cobordisms (not necessarily h-cobordisms) using the disjoint union operation $\sqcup$ and cartesian product operation $\times$. We denote the set of all oriented cobordism classes as Ω and the set of all n-dimensional oriented cobordism classes as Ω_n.

Proposition 9.16. *The set of all compact, closed, oriented, smooth cobordism classses of manifolds Ω forms a graded ring $\Omega = (\Omega_0, \Omega_1, ...)$ under the disjoint union operation $\sqcup$ and the cartesian product operation $\times$. The zero element is given by $\emptyset$ and the identity element is given by $\{\mathrm{pt}\}$ considered as a 0-manifold with a positive orientation.*

Proof. We need to check that $\sqcup$ and $\times$ are well-defined on oriented cobordism classes. Let M and M' represent the same oriented cobordism class in Ω_m and let N be a representative of a class in Ω_n. Then there exists W such that $\partial W = M \sqcup -M$. $W \sqcup N \times I$ is now clearly an oriented cobordism between $M \sqcup N$ and $M' \sqcup N$. The same argument works on the second term, and hence $\sqcup$ is well defined on Ω. For $\times$, note that $\partial(W \times N) = (\partial(W) \times N) \cup (W \times \partial N) = M \sqcup -M' \times N$ as N is closed. This gives us that $W \times N$ is an oriented cobordism between $M \times N$ and $M' \times N$. The same argument then applies to the second term, and hence $\times$ is well defined on Ω.

To see that additive inverses exist, note that $M \sqcup -M$ is the boundary of $M \times I$, and hence $(M \times I; M \sqcup -M, \emptyset)$ is an oriented cobordism. Clearly $\{\mathrm{pt}\}$ with a positive orientation acts as a multiplicative identity. Distributivity comes from the fact that $(X \sqcup Y) \times Z = (X \times Z) \sqcup (Y \times Z)$ for sets X, Y, Z.

Finally, the grading comes immediately from $M \times N \in \Omega_{n+m}$. □

Since the orientation of a product is only unchanged after reversing the order of the product if both manifolds are of even dimension, this ring satisfies the following graded commutativity relationship: if M is m-dimensional and N is n-dimensional, then $M \times N = (-1)^{nm} N \times M$.

Our proof of 9.15 will now be based on the following proposition.

Proposition 9.17. *Let $\{K_n\}$ be a multiplicative sequence of polynomials with coefficients in $\mathbb{Q}$. The map $\Omega \to \mathbb{Q}$ sending a representative M to its K-genus $K[M]$ is a well-defined ring homomorphism.*

Proof. We can restrict to considering only $4n$-dimensional manifolds since all others are mapped to zero, which trivially satisfies the requirements for being a ring homomorphism. Now, note that the Pontrjagin numbers of a disjoint union satisfy the relation $p_I(M \sqcup N) = p_I(M) + p_I(N)$. Now let M and M' represent the same cobordism class. This means that $M \sqcup -M'$ is the boundary of some $(4n+1)$-dimensional compact, oriented smooth manifold. 9.11 then tells us that $p_I(M \sqcup -M') = 0$, which in turn gives us that $p_I(M) = p_I(M')$. This proves that the Pontrjagin numbers are an oriented cobordism invariant. Since the K-genus is a linear combination of Pontrjagin numbers, this gives us that our map is well defined and that $K(M \sqcup N) = K(M) + K(N)$. All that is left to check is that the map commutes with multiplication.

Let M have total Pontrjagin class $p = 1 + p_1 + p_2 + ...$ and N have total Pontrjagin class $q = 1 + q_1 + q_2 + ...$ Since the tangent bundle of $M \times N$ splits as a Whitney sum $TM \oplus TN$, one can use 9.9 to show that, up to elements of order two, the total Pontrjagin class of $M \times N$ is given by $p \smile q$. We can then use the fact that $\{K_n\}$ is a multiplicative sequence to get that $K(p \smile q) = K(p) \smile K(q)$. Hence, one can show that we have the following relation:

$$K(p \smile q)([M \times N]) = (-1)^{nm} K(p)([M]) K(q)([N])$$

where $m = \dim(M)$ and $n = \dim(N)$. We already restricted to considering the case where both n and m are multiples of 4, and hence our map commutes with multiplication. This completes the proof. □

Note that by tensoring with $\mathbb{Q}$, this gives us a well defined algebra homomorphism $\Omega \otimes \mathbb{Q} \to \mathbb{Q}$.

We now finish this section by returning to proving the Hirzebruch signature theorem.

Proof of 9.15. First, we claim that the map sending a cobordism class represented by M to its signature $\sigma(M)$ is a ring homomorphism $\Omega \to \mathbb{Z}$. To show this we need to prove the same facts as in 9.17 but for the signature. We will omit doing this for brevity, but more details can be found in [4]. After tensoring by $\mathbb{Q}$ this gives us an algebra homomorphism to $\mathbb{Q}$. We already have an algebra homomorphism $\Omega \otimes \mathbb{Q} \to \mathbb{Q}$ sending M to its K-genus $K[M]$, so proving the theorem reduces to showing that these two homomorphisms match on a set of generators for $\Omega \otimes \mathbb{Q}$. Our second claim is that a set of generators is given by the complex projective spaces $\mathbb{PC}^{2n}$, where we ignore their complex structure and consider them as real smooth manifolds of dimension $4n$. Again for details, see [4].

The cohomology ring of $\mathbb{PC}^{2n}$ can be computed as $H^*(\mathbb{PC}^{2n};\mathbb{Q}) = \mathbb{Q}[z]/(z^{2(n+1)})$, where we denote the generator by z for reasons that will be obvious later. Then the signature is given by the signature of $Q : H^{2n}(M;\mathbb{Q}) \to \mathbb{Q}$, which is represented by a 1×1 matrix, with entry equal to $z^n \cdot z^n([M]) = z^{2n}([M]) = 1$. Hence $\sigma(\mathbb{PC}^{2n}) = 1$.

We now seek to compute the L-genus $L[\mathbb{PC}^{2n}]$. To begin, we claim that the total Pontrjagin class of the tangent bundle to $\mathbb{PC}^{2n}$ is given as $p(T\mathbb{PC}^{2n}) = (1+z^2)^{2n+1}$. So,

$$L(p) = L((1+z^2)^{2n+1}) = L(1+z^2)^{2n+1} = \left(\frac{\sqrt{(z^2)}}{\tanh\sqrt{z^2}}\right)^{2n+1} = \left(\frac{z}{\tanh z}\right)^{2n+1},$$

where we have used the definition of multiplicative sequences and that $\{L_n\}$ is the multiplicative sequence belonging to $f(z) = \sqrt{z}/\tanh\sqrt{z}$. Note that the coefficient of the $2n$th term in the expansion $\left(\frac{z}{\tanh z}\right)^{2n+1}$ is exactly the L-genus of $\mathbb{PC}^{2n}$, and so we have reduced the problem to that of finding the coefficient in a power series expansion. This is the reasoning behind the notation, as we will now treat z as a complex variable to compute this coefficient.

Recall from complex analysis that we can use the Cauchy-Taylor theorem to find the coefficient in a power series expansion. If $f(z)$ is a power series expansion about $a \in \mathbb{C}$ with coefficients c_k that converges on some ball of radius r centred at a, then

$$c_k = \frac{1}{2\pi i} \int_{|z-a|=\rho} \frac{f(z)}{(z-a)^{k+1}} dz,$$

where $0 < \rho < r$. In our case, this results in our desired coefficient being given by the integral

$$\frac{1}{2\pi i} \int_{|z|=\rho} \frac{1}{(\tanh z)^{2n+1}} dz,$$

for some $0 < \rho < \pi/2$. We continue by integration by substitution, with the substitution $z = \operatorname{arctanh} u$, which gives $du/(1-u^2) = dz$. Putting this into our integral, we get

$$\frac{1}{2\pi i} \int_{|z|=\rho} \frac{du}{(1-u^2)u^{2n+1}} = \frac{1}{2\pi i} \int_{|z|=\rho} \frac{1+u^2+u^4+...}{u^{2n+1}} du.$$

All of the terms in this integral are holomorphic and so will integrate to zero by Cauchy's theorem, except for the term $u^{2n}/u^{2n+1} = u^{-1}$ which integrates to $2\pi i$, since $u \approx z$ for small enough z. Hence the required coefficient is equal to 1. This means that the L-genus is also equal to 1, which completes the proof. □

Chapter 10
Exotic Spheres

(Unless specified, the material in this chapter is based on [8] and [9].)

We will now put together everything we have learnt so far to show that Θ_7 contains a non-trivial element, and hence exotic spheres exist. We need two things: firstly, a candidate for being an exotic sphere (a homotopy sphere), and secondly, an invariant to distinguish our homotopy sphere from S^n. Our method will be to define an invariant in terms of a manifold that our homotopy sphere bounds, in such a way that, rather miraculously, our invariant does not depend on the bounding manifold in the end. We begin by constructing our candidate.

In 4, we classified all vector bundles over S^n. If we consider 4-dimensional real vector bundles over S^3, we have the bijections:

$$\mathrm{Vect}^4_+(S^4) \to [S^3, SO(4)] \to \pi_3(SO(4)).$$

In fact, all of these vector bundles can be given the structure of smooth vector bundles and hence are smooth manifolds of dimension 8. From the definition of our clutching construction, it is not hard to see that the only obstruction to doing so is that the clutching function $f : S^{n-1} \to SO(4)$ may not be smooth. However, this is not an issue if we use the following result.

Theorem 10.1. *Let $f : M \to N$ be a continuous map of compact, smooth manifolds. Then f is homotopic to a smooth map $M \to N$.*

For a proof, see [1]. The proof uses tubular neighbourhoods, which we defined in 5. In fact, this result holds with much weaker hypotheses, but this version suffices for us.

This means that our vector bundles over S^n can be assumed to be smooth vector bundles. We would like to classify these now, which amounts to calculating $\pi_3(SO(4))$.

Lemma 10.2. $\pi_3(SO(4)) = \mathbb{Z} \oplus \mathbb{Z}$.

Proof. We proceed by constructing a covering map $S^3 \times S^3 \to SO(4)$. First, identify $\mathbb{R}^4$ with the quaternions, denoted by $\mathbb{H}$ and note that this similarly identifies S^3 with the unit quaternions. Now, let $\phi : S^3 \times S^3 \to GL_4(\mathbb{R})$ be the map sending $(u, v) \in S^3 \times S^3$ to the map $\mathbb{R}^4 \to \mathbb{R}^4$ that sends $x \mapsto uxv^{-1}$. Clearly this is an element of $GL_4(\mathbb{R})$ and ϕ is a group homomorphism. Furthermore, note that the entries of $\varphi(u, v)$ when considered as a matrix are polynomials in the coordinates of u and v and hence φ is smooth. We can then conclude, since $S^3 \times S^3$ is connected and $\varphi(1, 1)$ is the identity matrix, $\mathrm{Im}(\phi) \subset GL_4^+(\mathbb{R})$. Note that $\phi(u, v)$ is orthogonal since its transpose is given by $x \mapsto v^{-1}xu$ which is clearly the inverse of $\phi(u, v)$. This means that $\mathrm{Im}(\phi) \subset SO(4)$. We can further show that this inclusion is an equality using a dimension argument, but we will leave out the details for expediency.

We can compute the kernel of this map easily, by noting that if $(u, v) \mapsto \mathrm{Id}$, then $uv^{-1} = 1$, which implies $u = v$. Now we have that $ux = xu$ for all $x \in \mathbb{H}$, and hence u lies in the centre of $\mathbb{H} = \mathbb{R}$. Since $u \in S^3$, we get that $u = \pm 1$ and hence $\ker(\phi) = \{(1,1),(-1,-1)\}$. So we have constructed a surjective smooth map $\phi : S^3 \times S^3 \to SO(4)$ that is a two-fold covering map, as $\ker(\phi) = \{(1,1),(-1,-1)\}$.

We can then use the fact that covering maps induce isomorphisms on homotopy groups higher than the fundamental group (see [2] for details on this). This means that $\pi_3(SO(4)) \cong \pi_3(S^3 \times S^3)$. Now homotopy groups behave well under products of path-connected spaces, since homotopies of maps $S^n \to X \times Y$ correspond uniquely (by the product topology) with homotopies of maps $S^n \to X$ and $S^n \to Y$. So $\pi_3(SO(4)) \cong \pi_3(S^3) \times \pi_3(S^3)$. We can then use the Hurewicz theorem 6.3 to give that $\pi_3(S^3) \cong \mathbb{Z}$, which in turn gives the required result. □

So, we have a one-to-one correspondence between $\mathbb{Z} \oplus \mathbb{Z}$ and $\mathrm{Vect}_+^4(S^4)$ and we can use the covering map to write down this correspondence explicitly. Note that we have the isomorphism $\mathbb{Z} \to \pi_3(S^3)$ given by $i \mapsto (u \mapsto u^i)$. So we can give the correspondence between pairs of integers and clutching functions as

$$(i, j) \mapsto (u \mapsto (x \mapsto u^i x u^j))$$

where we have chosen to introduce a minus sign in the second term to simplify formulas later.

Now, we will write $\xi_{i,j}$ for the sphere bundle associated to the vector bundle given by $(i, j) \in \mathbb{Z} \oplus \mathbb{Z}$. We claim that $\xi_{i,j}$ is a smooth 7-dimensional manifold. This is because it can be given the structure of a submanifold of the smooth vector bundle, provided that the inner product defined on the smooth vector bundle is itself smooth. Much like when we defined the normal bundle in 5.8 we claim that this is always possible.

Proposition 10.3. *Let $i + j = 1$. Then $\xi_{i,j}$, as defined above, is a homotopy sphere.*

Proof. We would like to use 6.5 to show that $\xi_{i,j}$ is a homotopy sphere. That means we need to show $H_k(\xi_{i,j}) \cong H_k(S^7)$ and that $\xi_{i,j}$ is simply connected. We start with the simply connectivity of $\xi_{i,j}$.

We will use that fibre bundles give rise to a long exact sequence of homotopy groups. For a proof of this, see [2]. The long exact sequence for $p : E \to B$ with fibre F is given as (assuming E, B and F are all path-connected):

$$\longrightarrow \pi_n(F) \longrightarrow \pi_n(E) \longrightarrow \pi_n(B) \longrightarrow \pi_{n-1}(F) \longrightarrow \quad .$$

Now since $\pi_1(S^3)$ and $\pi_1(S^4)$ are both trivial, this long exact sequence gives us that $\pi_1(\xi_{i,j})$ must also be trivial. Note that this did not depend require $i + j = 1$ and holds more generally.

Now we must compute the homology groups of $\xi_{i,j}$. Using the fact that $\xi_{i,j}$ is path connected, simply connected and orientable, we get that $H_k(\xi_{i,j}) \cong \mathbb{Z}$ for $k = 0 = 7$ and $H_1(\xi_{i,j})$ is trivial. For the rest of the groups, we can use the Mayer-Vietoris sequence with the decomposition $\xi_{i,j} = (D^4 \times S^3) \cup (D^4 \times S^3)$ where the union is over the identification given by the clutching function $f : S^3 \times SO(4)$. We can easily compute the homologies $H_k(D^4 \times S^3) \cong H_k(S^3)$ and $H_k(S^3 \times S^3)$. The Mayer-Vietoris sequence then immediately gives us that $H_k(\xi_{i,j})$ is trivial for $k = 2, 5$ or 6. The only tricky cases are $k = 3$ and 4. The relevant portion of the Mayer-Vietoris sequence is given below:

$$H_4(\xi_{i,j}) \longrightarrow H_3(S^3 \times S^3) \xrightarrow{h} H_3(S^3) \oplus H_3(S^3) \longrightarrow H_3(\xi_{i,j}) \ .$$

Both of the middle groups are isomorphic to $\mathbb{Z} \oplus \mathbb{Z}$ and so our aim is to show that $h : \mathbb{Z}\oplus\mathbb{Z} \to \mathbb{Z}\oplus\mathbb{Z}$ is an isomorphism, which would in turn give us that the remaining two homology groups are trivial. Let (a, b) be the generators for $H_3(S^3 \times S^3)$, so that a generates the homology class of the equator of the base space S^4, and b generates the homology class of the fibre S^3. If we assume that when a is included into the southern hemisphere it comes in unchanged, then when it is included into the northern hemisphere it undergoes the clutching identification. If we assume that the first factor in $H_3(S^3) \oplus H_3(S^3)$ represents the trivial bundle over the southern hemisphere, then h maps (a, b) to the first factor as just b. We need to understand how h maps the homology classes into the second factor. Our clutching function is a map $f : S^3 \to SO(4)$ that maps $u \mapsto (x \mapsto u^i x u^j)$. We can see what this does to the homology class a by fixing a point in the fibre, say $x = 1$. Then $u \mapsto (1 \mapsto u^{i+j} = u)$. This means the generator for the equatorial S^3 gives us a generator of the fibre S^3 where we have crucially used that $i+j = 1$. Then we can see that $h(a, b) = (b, b-a)$ and hence is an isomorphism, completing the proof. □

Now we need to produce an invariant that we can use to show $\xi_{i,j}$ is not diffeomorphic to S^n. Assume that M is a homotopy sphere and bounds a smooth manifold B. Then, by using the long exact sequence of the pair, we have that the inclusion homomorphism $i : H^4(B, M; \mathbb{Z}) \to H^4(B; \mathbb{Z})$ is an isomorphism as $H^4(M; \mathbb{Z})$ and $H^5(M; \mathbb{Z})$ are both trivial. We can then pull the Pontrjagin class p_1 of TB, the tangent bundle of B, back to $H^4(B, M; \mathbb{Z})$. Then we can construct a pseudo-Pontrjagin number where we evaluate $p_1 \smile p_1$ on the relative fundamental class of B. Write this pseudo-Pontrjagin number as $q(B)$. Then we define our invariant $\lambda(M) = 2q(B) - \sigma(B) \in \mathbb{Z}$, where $\sigma(B)$ denotes the signature of B.

Theorem 10.4. *$\lambda(M)$ modulo 7 does not depend on B.*

Proof. Assume we have two smooth manifolds B_1, B_2 which both bound M. We can form a smooth manifold C as the union of these two manifolds over the boundary M. We can ensure that C is oriented by specifying B_2 to have the opposite orientation to B_1 on M. By definition, C must be a closed manifold and so we can use the Hirzebruch signature theorem (9.15) to calculate its signature. Since $\dim(C) = 8$, we will be using the L_2 polynomial, which gives us that

$$\sigma(C) = \frac{1}{45}(7p_2 - p_1^2)([C]).$$

Consequently, we can multiply both sides by 45 and reduce modulo 7 to get:

$$2p_1^2(C) - \sigma(C) \equiv 0 \bmod 7. \tag{†}$$

We look at $\sigma(C)$. By an easy application of relative Mayer-Vietoris and the long exact sequence of the pair we have a pair of isomorphisms:

$$H^4(B_1, M) \oplus H^4(B_2, M) \to H^4(C, M) \to H^4(C).$$

Let $\varphi \in H^4(C)$ correspond to $(\varphi_1, \varphi_2) \in H^4(B_1, M) \oplus H^4(B_2, M)$ under the above isomorphisms. Note that by our choice of orientations, $[C] = [B_1] - [B_2]$. We now compute $\varphi([C]) = (\varphi_1^2, \varphi_2^2)([B_1] - [B_2]) = \varphi_1^2([B_1]) - \varphi_2^2([B_2])$. Hence $\sigma(C) = \sigma(B_1) - \sigma(B_2)$.

For the Pontrjagin numbers, we can similarly construct another pair of isomorphisms:

$$H^4(B_1, M) \oplus H^4(B_2, M) \to H^4(B_1) \oplus H^4(B_2) \to H^4(C).$$

Then by an argument almost identical to the above argument we see that $p_1(C) = q(B_1) - q(B_2)$. Putting our results for the signature and Pontrjagin numbers together with (†) then gives us that

$$2q(B_1) - \sigma(B_1) \equiv 2q(B_2) - \sigma(B_2) \bmod 7$$

as required, completing the proof. □

It is clear that $\lambda(S^7) = 0$ since S^7 is bounded by D^8 which is contractible and hence both its signature and Pontrjagin classes vanish. The aim is now to compute $\lambda(\xi_{i,j})$. Notice $\xi_{i,j}$ is the boundary of the corresponding disc bundle which we will denote as $B_{i,j}$, which also has a natural structure as an 8-dimensional smooth manifold. We will denote the original vector bundle corresponding to $\xi_{i,j}$ as $E_{i,j}$.

Lemma 10.5. *$p_1(E_{i,j}) = \pm 2(i-j)\mu$ where μ is the generator of $H^4(S^4)$.*

Proof. First we argue that $p_1(E_{i,j})$ must be a linear function of i, j. Note that if f and g are the clutching functions corresponding to vector bundles $E_{i,j}$ and $E_{k,l}$, then $E_{i+k,j+l}$ is given by the clutching function fg. Now we claim that $E_{i+k,j+l}$ is stably isomorphic to $E_{i,j} \oplus E_{k,l}$. To see this, consider f and g as matrices in $SO(4)$, then

$$\begin{pmatrix} fg & \mathbf{0} \\ \mathbf{0} & \mathrm{Id} \end{pmatrix} \begin{pmatrix} g^{-1} & \mathbf{0} \\ \mathbf{0} & g \end{pmatrix} = \begin{pmatrix} f & \mathbf{0} \\ \mathbf{0} & g \end{pmatrix}.$$

So $p_1(E_{i+k,j+l}) = p_1(E_{i,j} \oplus E_{k,l})$ and now we can use 9.9 to get $p_1(E_{i+k,j+l}) = p_1(E_{i,j}) + p_1(E_{k,l})$, which proves linearity.

To finish off, notice that if we reverse the orientation of our fibres in $E_{i,j}$ this is equivalent to conjugating our clutching function with some orientation reversing diffeomorphism $\mathbb{H} \to \mathbb{H}$. If we take, for simplicity, this diffeomorphism to be quaternion conjugation, then a little quarternion algebra shows us that our new vector bundle is isomorphic to $E_{-j,-i}$. But recall that the Pontrjagin classes do not depend on orientation, which means $p_1(E_{i,j}) = p_1(E_{-j,-i})$ and hence $p_1(E_{i,j}) = \pm c(i-j)\mu$ for some constant c. We can calculate this constant by evaluating p_1 on a single $E_{i,j}$, the simplest choice being $E_{1,0}$, to see that $c = 2$. We omit the details of this lengthy calculation for lack of space. □

Proposition 10.6. $\lambda(\xi_{i,j}) = (i-j)^2 - 1 \mod 7$.

Proof. First we want to compute $q(B_{i,j})$, which means we want the Pontrjagin classes of the tangent bundle of $B_{i,j}$. We can split $TB_{i,j}$ up as the Whitney sum of the bundle of tangent vectors to the fibre and the bundle of tangent vectors to the base space. These are given as the pullback bundle of $E_{i,j}$ and the pullback bundle of TS^4, respectively, where both pullbacks are with respect to the projection $\pi_{i,j} : B_{i,j} \to S^4$. Since pullbacks commute with taking Pontrjagin classes this means $p_1(TB_{i,j}) = \pi^*_{i,j}(p_1(E_{i,j} \oplus TS^4))$. But since all spheres are stably parallelisable, we can use 9.9 to conclude $p_1(TB_{i,j}) = \pi^*_{i,j}(p_1(E_{i,j}))$. Now $B_{i,j}$ clearly has S^4 as a deformation retract, and so $H^4(B_{i,j})$ is generated by a single element ν and we know that π^* must map μ to this generator. So $p_1(B_{i,j}) = \pm 2(i-j)\nu$ and hence $q(B_{i,j}) = \pm 4(i-j)^2$.

Now we can calculate $\sigma(B_{i,j})$ easily since we can choose orientations of $B_{i,j}$ such that the quadratic form $Q : H^4(B_{i,j}) \to \mathbb{Q}$ is simply given by squaring and so its signature is $+1$. This means that $\lambda(\xi_{i,j}) = 8(i-j)^2 - 1 \equiv (i-j)^2 - 1$ modulo 7 as required. □

So, by picking i and j such that $i + j = 1$ and $i - j \not\equiv 1$ or -1, $\xi_{i,j}$ is an exotic sphere. For example, $\xi_{3,-2}$ is an exotic sphere. This equivalently means that Θ_7 contains a non-trivial element.

Chapter 11
Conclusion

In the course of this project we have shown that we can construct a group out of homotopy spheres, Θ_n, and for $n \geq 5$ we have shown that it coincides with all possible smooth structures on S^n. The key theorem in all of this is the h-cobordism theorem which allowed us to substitute our equivalence relation of diffeomorphism to h-cobordism and to prove the generalised Poincaré conjecture for $n \geq 5$, allowing us to show that our homotopy spheres were topologically spheres. We then sought to gain a better understanding of Θ_n and we focused on showing that Θ_7 was non-trivial. While this seems like a modest goal, it was a highly non-trivial fact to show and we needed a large amount of theory to tackle the problem.

An obvious continuation of this project would be to learn more about Θ_n. This is what Kervaire and Milnor did in 1962 where they computed that Θ_n was always finite, although the finiteness of Θ_3 depended on the Poincaré hypothesis being true. Once the Poincaré hypothesis was finally proved in 2003, this proved that Θ_3 was trivial and so Θ_n is finite for all n. They also managed to compute these groups in a few cases. In the first twelve groups, Θ_k is trivial for $k = 1$, 2, 3, 4, 5, 6 and 12. $\Theta_7 \cong \mathbb{Z}/28$, $\Theta_8 \cong \mathbb{Z}/2$ and $\Theta_{11} \cong \mathbb{Z}/992$. They also showed that the orders of Θ_9 and Θ_{10} were 8 and 6, respectively.

To perform this computations they used an operation called *surgery*, which is a special case of the pasting operation defined in 8.6, where both submanifolds are spheres (in their original paper, they call surgery *spherical modifications* but this terminology has fallen out of fashion). The aim is to compute Θ_n by finding an exact sequence involving both Θ_n and the subgroup of Θ_n given by homotopy spheres that bound parallelisable manifolds. Then we can compute Θ_n by computing this subgroup and we use surgery to do this. Giving this argument in full would take at least two or three more chapters, and so was not possible to include it, although it builds off of much of the material already developed.

We finish with some discussion on exotic spheres, for which the theory is immensely interesting. Although Θ_4 is trivial, this does not mean that no exotic spheres exist in dimension 4, as the h-cobordism theorem is not applicable to give us the required correspondence. In fact, whether or not exotic spheres exist in dimension 4 is perhaps the largest open problem in differential topology currently. Another very interesting fact about exotic spheres is due to a theorem in geometry, proved in 2009 by S. Brendle and R. Schoen that requires some understanding of Riemannian geometry to parse. It says, suppose M is a geodesically complete, connected, smooth manifold with sectional curvature K 'quarter-pinched', i.e. there exists K_0 such that $K_0/4 < K < K_0$, then M is diffeomorphic to S^n. This tells us that it is impossible to quarter-pinch the curvature of an exotic sphere! This is in direct contrast to the standard sphere, which has constant positive sectional curvature. This theorem was proved via Ricci flow, which was the technique used to prove the Poincaré hypothesis.

Bibliography

[1] Antoni A. Kosinski. *Differential manifolds.* Dover Publications, 2007.

[2] Allen Hatcher. *Algebraic topology.* Cambridge University Press, 2002.

[3] Allen Hatcher. *Vector Bundles & K-Theory.* Nov 2017.

[4] John W. Milnor and James D. Stasheff. *Characteristic classes.* Princeton University Press, 1974.

[5] John W. Milnor and Michel A. Kervaire. Groups of homotopy spheres. *The Annals of Mathematics*, 77(2):504, 1962.

[6] Glen E. Bredon. *Topology and geometry by Glen E. Bredon.* Springer-Verlag, 1993.

[7] John W. Milnor, Laurence C. Siebenmann, and J. Sondow. *Lectures on the h-cobordism theorem.* Princeton University Press, 1965.

[8] John Milnor. On manifolds homeomorphic to the 7-sphere. *The Annals of Mathematics*, 64(2):399, 1956.

[9] Akhil Mathew. On manifolds homeomorphic to the 7-sphere, Mar 2012.