Morse Theory and Flow Categories

A Thesis
Presented to
The Division of Mathematical and Natural Sciences
Reed College

In Partial Fulfillment
of the Requirements for the Degree
Bachelor of Arts

Maxine Elena Calle

May 2020

Approved for the Division
(Mathematics)

Kyle M. Ormsby

Acknowledgements

First and foremost, I must thank Kyle Ormsby all his help. He got stuck as my advisor and I imagine he's gotten rather sick of me at this point, but he has been so helpful throughout this process. I'm very grateful to him for a number of things, including saying lots of things that make sense, putting up with me saying lots of things that make no sense, and preventing me from naming this thesis "Go with the Flow." It is thanks to Kyle that I know more about homotopy theory than I did before.

I would also like to thank Angélica Osorno, who I have learned a lot from both in and out of mathematics. There are many things in this thesis that I understand better because of conversations with her. She has been a wonderful mentor and inspiration throughout my time at Reed.

There are many other people without whom this thesis would not have been. I am indebted to Corey Dunn for telling me that I might like Morse theory, and to Frederico Marques's lecture which reminded me what Corey had said. I am profoundly grateful to Ralph Cohen for sharing insights on the material in Chapter 4. Without his generosity, I would have had far less to say. Many thanks also goes to David Ayala, whose helpful observations were crucial to both the successful completion of this thesis and the preservation of my sleep schedule within the last week.

To the faculty and staff that have impacted my time at Reed, I extend heartfelt gratitude. I would especially like to thank Paul Hovda, for enlightening me as to the importance of giving things names; Jerry Shurman, for teaching my most beloved math class at Reed; and Irena Swanson, for encouraging me to stick with math in the first place. I'm grateful to the entire Mathematics Department for their support, and in particular Lisa Mackola has always been a great help with the logistical side of things.

On a more personal note, I blame my mom for getting me hooked on math in the first place, and Nannie for encouraging me with multiplication problems over the phone. I greatly appreciate Laura and Dennis for their hospitality; they did not have to host me here in Portland, but they did, and that is very sweet of them. Maia Houck has known me through all my phases in life, including the middle school years, and still wants to live with me; I love her for that (among so many other things).

Lucas Williams has been a great friend to me over the years, even if he refused to speak to me when we took Physics together. To be fair, I also refused to speak to him. I'm very glad we got over that. I'm also very glad that he is willing to go to Tom Yum with me so often. There are many things that I want to thank him for, but I'm planning to spread them out over the years, so I'm going to stop here.

I could not have made it through my four years here at Reed without my wonderful friends. Big love goes to Laurel Schuster, for being my day-one; to Zia Pollis, for her warmth and incredible mochas; to Lauren Chacon, for her understanding and excellent sense of fashion; to Sadie Baker-Wacks, for understanding when I don't text back; to Ema Chomsky, for flipping me upside down; to Elena Turner, for all our strange connections; to Robert Irving and Tobias Rubel Janssen, for carving out a slice of thesis heaven on the second floor of the library; to Alex Richter and Sophia Carson, for their sweetness; to Ryan Kobler and Marisa Agger, for the late nights in Eliot Hall; and to all the folks in the aerial acrobatics troupe, STEMGeMs, the 2019 CMRG, and the 2018 CSUSB REU, for the community we made. I thank you all for caring enough to ask about my thesis, and for being polite enough to listen to me talk about it.

Finally, I do not gratefully thank COVID-19 for helping me to finish out my senior year from my bedroom. As much as I like my own room, there's nothing like basking in the sunlight on the great lawn. Reed College, I will remember you fondly.

Index of Notation

A list of some symbols used throughout this thesis, ordered more or less alphabetically:

TM	the tangent bundle of M. 82
Top	the category of topological spaces and continuous maps. 33, 90
TopCat	the category of categories internal to **Top** and continuous functors. 34
$\pitchfork$	transverse. 16, 17
γ_q^p	$= \gamma_q\|_{\mathrm{dom}\gamma_p}$, a truncated flow line. 65
tw	$: \mathbf{Cat} \to \mathbf{Cat}$, the twisted arrow functor. 47
$W^u(a)$	the unstable manifold of a critical point a. 14

Table of Contents

List of Figures

Abstract

The aim of this thesis is to cover the basics of Morse theory with an eye towards understanding the flow category of a Morse function. Morse theory is classically a branch of differential topology, and seeks to draw topological conclusions about a manifold by studying the differentiable functions on them. In particular, the main theorem in the unpublished preprint of Cohen, Jones, and Segal [CJS95b] states that the topological structure of a manifold can be completely recovered (in some cases, only up to homotopy) from the classifying space of the flow category of a Morse function on that manifold. We begin with an overview of classical Morse theory, including the Morse-Smale transversality condition and Morse homology, then proceed to introduce the concepts needed to understand the Cohen-Jones-Segal theorem. We conclude with an overview of the proof of this theorem and a few illustrative examples.

The first chapter of this thesis is based upon the first part of the book by Audin and Damian [AD14]. The following chapters focus upon the work in [CJS95b], but we have expanded upon and reformulated the content in some places. In particular, our proof of the general case of the Cohen-Jones-Segal theorem in Section 4.2.1 fixes an error in the original proof, using methods suggested to the author by Cohen.

Dedication

This thesis is dedicated to my mother, Michelle,
who inspires me to make the most of every day.

Introduction

Morse theory[1] revolves around the idea that we can understand the topological structure of smooth manifolds by studying differentiable functions on them. In particular, we can gain insight by examining the critical points of such a function. When the critical points are non-degenerate, the function is said to be *Morse*. We can pair a Morse function f with a *pseudo-gradient flow* X adapted to it, and then investigate the structure of the manifold in terms of the flow lines between critical points. In certain cases, when the pair (f, X) satisfies the so-called *Morse-Smale* condition, we get a well-formed chain complex called the *Morse complex*. As it turns out, the *Morse homology* of this complex is isomorphic to the singular homology of the manifold, and consequently information about the topology of a manifold can be completely unlocked by understanding the structure of critical points of a Morse function and flows between them.

We can also store information about critical flows and gradient flow lines in the *flow category* $\mathscr{C}_f$. The objects in this category are the critical points of f, and the morphisms are the "broken" flow lines between critical points. We can turn the flow category into a topological space via its *classifying space* (the geometric realization of the nerve of $\mathscr{C}_f$). The main focus of this thesis is the following theorem of Cohen-Jones-Segal in the unpublished preprint [CJS95b] that establishes a connection between this classifying space and the original manifold:

Theorem. *Let $f\colon M \to \mathbb{R}$ be a Morse function on a manifold M and let $\mathscr{C}_f$ denote the flow category of f. Then the classifying space of $\mathscr{C}_f$ is homotopy equivalent to M. Moreover, in the Morse-Smale case, we have a homeomorphism.*

The goal of this thesis is to give an overview of Morse theory, develop the simplicial homotopy theory needed to understand the work of Cohen-Jones-Segal, and then present of proof of this theorem. The original Cohen-Jones-Segal paper was never published, due in part to errors in the proofs (discussed in Chapter 4), although the result is widely cited. To the author's knowledge, there is nothing in the literature that address the error in the non-Morse-Smale case — although excellent progress has been made in the Morse-Smale case — and we hope that our contribution will fill this gap. To get a feel for the theory, we will present a simple example in the following section, although this may be skipped without consequence.

[1] No known connection to Morse code.

A Heuristic Example

Imagine you are a small bug on the surface of a mountainous landscape. As a bug, you are very interested in reaching the lowest possible altitude (to get away from those dastardly, hungry birds), but unfortunately you have little-to-no concept of the global topography of the landscape. All you can do is look around your surrounding area, figure out which direction will help you descend, and head that way.

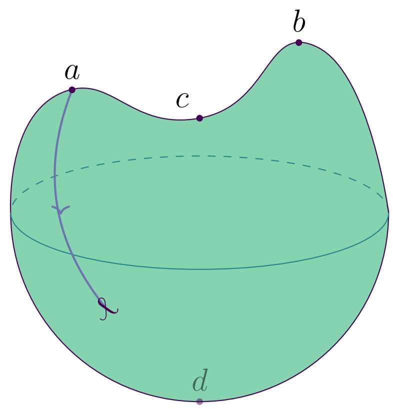

Figure 1: A bug's path. A bug's path of steepest descent traces out a gradient flow line on the alternate sphere. There are four critical points (places where the surface appears flat nearby): a, b, c, and d.

Your path of steepest descent traces out a *gradient flow line* on the surface of the landscape. It is only later that you learn that in fact you were crawling around on a "dented sphere" (which the author is fond of calling the alternate sphere), illustrated in Fig. 1. Since you are a bug who likes to dabble in pure mathematics, you are interested in understanding the topological structure of your roaming grounds. There are four critical points on this surface: the two peaks, the saddle point in the middle, and the minimum on the bottom. The alternate sphere can be covered in gradient flow lines, each of which connects two of the critical points. Morse theory provides ways for us to deduce information about the topology of the surface by just knowing about the structure of the critical points and gradient flow lines.

One way to store this information is in the *flow category*, the category whose objects are the critical points and whose morphisms are flow lines between critical points, including those "broken" flow that visit other critical points along the way. This splits our space up into different pieces (formally, the homspaces), each piece associated with a different pair of critical points, which is illustrated for the alternate sphere in Fig. 2.

If we manipulate and glue these pieces together in a certain way, we get the *classifying space* of the flow category. The main result that this thesis is concerned with says that this construction recovers much of the interesting topological information about the surface you started out with. To get an intuition for this result, we shall go through the construction for the alternate sphere (being a bit dishonest at times, despite our good intentions).

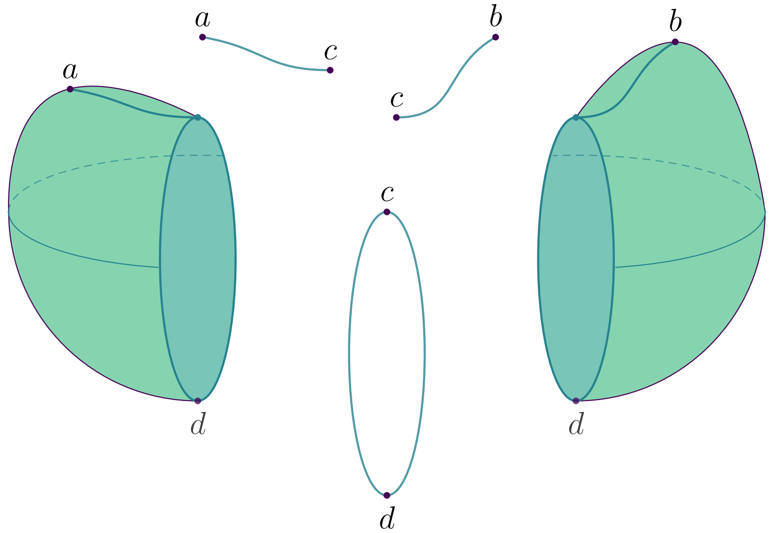

Figure 2: Alternate sphere partitioned by flows. Every pair of critical points can be associated with some portion of the alternate sphere, namely the collection of points that lie on some gradient flow between the two critical points.

We begin with the our four critical points and draw a line segment between two critical points for every flow that connects them. Thus there are two lines connecting c, d, one apiece for a, c and b, c, and a one-parameter family of lines for each of a, d and b, d. Now, for every pair of flows that can be composed, we complete and fill in the triangle for their associated line segments, as shown in Fig. 3. In this case, there is no interesting way for us to compose three or more flow lines, and so we stop at this step. Hopefully the illustration makes it seem plausible that this space we have arrived at is topologically "the same" as the alternate sphere.

A Brief History

Morse theory is the namesake of the American mathematician H. C. Marston Morse (1892–1977) whose 1934 publication *Calculus of Variations in the Large* is widely credited with introducing the techniques in differential topology that developed into Morse theory. Working in the 1920s, Morse was well-aware of the work of Poincaré, Brouwer, Birkhoff, and others in the fledgling field of "analysis situs," a field which is now known as topology. It was the insight of Morse to connect the new notions of topology with analysis; his work was met with great acclaim, earning him numerous accolades (an appointment at the Institute for Advanced Study, the National Science Medal, over twenty honorary awards. . .), but he was "in a sense a solitary figure, battling the *algebraic topology*. . . [he] always saw topology from the side of Analysis, Mechanics, and Differential Geometry" [Bot80, p.1, emphasis in original]. Nonetheless, the ideas of Morse theory — although essentially simple — have remained resilient, often guiding the development of many areas of geometry, topology, and even physics.

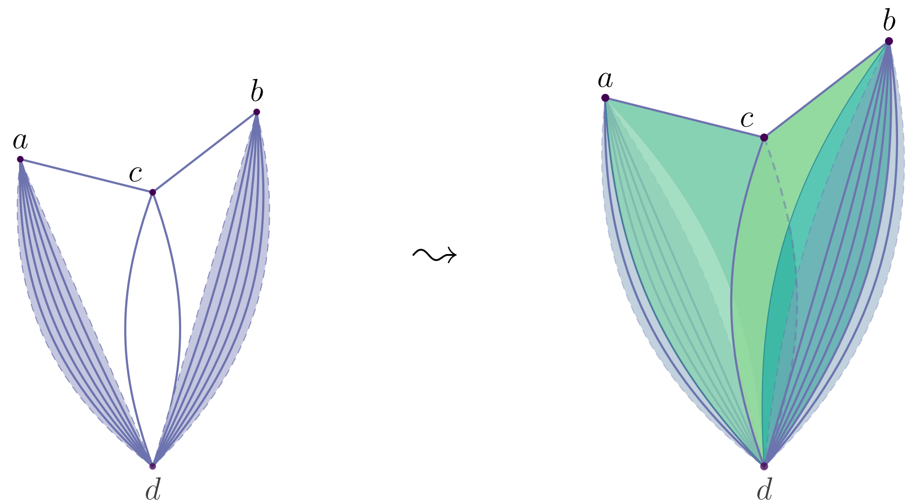

Figure 3: The classifying space of the alternate sphere. Composable flows are "glued in" to the classifying space. On the left, we show the single gradient flow lines glued to their appropriate end points. On the right, we have glued in four triangles (two for each triple a, c, d and b, c, d), corresponding to the "broken" flows that result from composing two genuine gradient flows.

One of the early results of Morse theory was the work of the Raoul Bott (1923–2005). Using Morse-theoretic techniques, Bott was able to make progress on understanding the homotopy theory of Lie groups, ultimately leading to a proof the Bott Periodicity Theorem in 1957, which in turn played an important role in the development of K-theory (cf. [Gue01]). Later, joint work of Bott and Michael Atiyah (1929–2019) made use of Morse theory and the Yang-Mills equations to study the cohomology of bundles over Riemann surfaces, and this collaboration laid out some of the foundational ideas for mathematical gauge theory (cf. [Fre11]). The initial work of Bott heralded the beginning of a new era of Morse theory, which was primarily concerned with investigating the topology of manifolds. The book [Mil63] by John Milnor (1931–) is often cited as the most important exposition on the subject; since the publication of this book, Morse theory has become a standard topic in geometry and topology curricula.

The "modern" approach to Morse theory involves the gradient flow lines of a Morse function, and Bott also had a hand in influencing this next stage of Morse theory. Bott's Ph.D. student, the renowned Stephen Smale (1930–), had a particularly strong impact on the theory via the development of the Palais-Smale compactness condition and the Morse-Smale complex. This work led him to a proof of the Poincaré conjecture for $n \geq 4$ as well as a formulation of the h-cobordism theorem. Additionally, Bott's 1979 lectures on his joint project with Atiyah inspired the theoretical physicist Edward Witten (1951–) to engage with Morse theory, and subsequently find a new approach in the 1980s, using the de Rham complex (cf. [Bot88]). These ideas, influenced and motivated in part by physics, were then extended by Andreas Floer (1956–1991) and involved Morse theory with new developments in symplectic geometry, mathematical gauge theory, topological quantum field theory.

These connections with advanced and deep mathematics exemplify the tendency of Morse-theoretic ideas to be useful in unexpected places. Much of the current research involves this "field-theoretic" approach to Morse theory, such as the work of Cohen-Jones-Segal [CJS95a], Fukaya [FOOO09], and others.

Overview

Classical Morse theory has its roots in differential topology, and consequently assumes familiarity with topological manifolds, calculus on manifolds, and basic algebraic topology. We have provided a rather terse summary of some of the most important ideas in Appendix A.1, however the conscientious reader may be better served by picking up a differential topology textbook (for instance, the author found [Hir76] particularly useful, in addition to the classical references [Mil56, Mun61]). In the appendix, we also briefly review essential concepts from algebraic topology (Appendix A.2) and category theory (Appendix A.3), although more in-depth treatments may be found in other sources, such as [Hat02] and [Rie16], respectively.

Chapter 1 follows the first part of the book by Audin and Damian [AD14] fairly closely, covering the basics properties and techniques of classical Morse theory. After presenting some of the most important definitions in Section 1.1, such as Morse functions, the Morse index, and the flow of a Morse function, we then discuss the stable and unstable manifolds of the flow and the subsequent decomposition of M. The Morse-Smale condition, defined in Section 1.2, ensures that this decomposition is reasonably well-behaved. The remainder of the chapter is dedicated to exploring a few of the developments of Morse theory, including the topology of sublevel sets in Section 1.3.1, Morse homology in Section 1.3.2, and applications of Morse homology in Section 1.3.3.

Chapter 2 introduces the flow category of a Morse function f. The objects in this category are the critical points of f, and the morphisms are the compactified moduli spaces of flow lines discussed in Section 2.1.2. This category can be equipped with a topological structure, and so is an example of a topological category (that is, a category internal to **Top**); Section 2.2 discusses topological categories in general before examining the flow category specifically.

In Chapter 3, we discuss some concepts from simplicial homotopy theory, building up to the definition of the classifying space of a (small) category in Section 3.1.2. In addition to the discussion of simplicial sets and their geometric realization, we also examine the edgewise subdivision of a simplicial set and its relation to the twisted arrow functor. In the latter half of the chapter, we seek to establish conditions between topological categories that induce a homotopy equivalence on the classifying spaces. Section 3.2.1 and Section 3.2.2 recall the basics of (co)fibrations and homotopy pullbacks, respectively. Using the fact that geometric realization preserves homotopy equivalence for Reedy cofibrant simplicial spaces, we show in Section 3.2.3 that a continuous functor on topological categories that is an levelwise homotopy equivalence (meaning we have homotopy equivalences at both the object and morphism level) induces a homotopy equivalence on the classifying spaces.

Finally, Chapter 4 presents the main theorem relating the classifying space of the flow category to the original manifold. We detail some illustrative examples in Section 4.1 before delving into the proofs in Section 4.2. While the proof of the Morse-Smale case (Section 4.2.2) closely follows the original [CJS95b], the proof of the general case given in Section 4.2.1 is original work of the author, following suggestions shared by Cohen.

Notation and Conventions

Unless explicitly stated otherwise, we assume M is a smooth, closed, finite-dimensional Riemannian manifold and that f is a smooth function on M. A list of notation for specific concepts and definitions can be found in the index preceding the table of contents. As is the case with many mathematical papers, it is difficult to not quickly run out of new letters to use. In any case, we typically follow these notational conventions:

$a, b, c, \ldots$	critical points of a Morse function
$\mathscr{C}, \mathscr{D}, \ldots$	categories
$f, g, \ldots$	functions
$H, h, \ldots$	homotopies
$i, j, k, \ldots$	indices
$M, N, \ldots$	manifolds
$n, m, k, \ldots$	dimensions
$p, q, r, \ldots$	arbitrary points on a manifold
$t, s, \ldots$	times
$v, w, \ldots$	tangent vectors
$V, W, \ldots$	arbitrary vector fields
$X, Y, Z, \ldots$	pseudo-gradient vector fields, topological spaces

Luckily, we rarely need to talk about more than two or three of the same type of thing at once. Although it is quite likely that the author has messed up her notational consistency in some places, she hopes you will forgive her and read on.

Chapter 1

Morse Theory

A (smooth) function $f\colon M \to \mathbb{R}$ on a (smooth, closed, finite dimensional Riemannian) manifold M can be thought of as assigning a topography on M, where the values of f define regions of constant elevation (known in mathematical language as level sets). Requiring that such a function is *Morse* (Definition 1.1.4) imposes a sense of stability on the critical points of f, places where the manifold is locally "flat." Specifically, a Morse function must have isolated critical points (Corollary 1.1.8), thus outlawing any kind of "plateau" behavior. We can classify the critical points of f by their *Morse index* (Definition 1.1.6), which roughly describes the number of directions we could descend from the critical point. For instance, a local minimum will have Morse index 0 and a local maximum will have Morse index equal to $\dim(M)$.

To gain more geometric insight on M, we examine the *(pseudo-)gradient flow* in Section 1.1.2. Again, if we imagine the Morse function f as prescribing some notion of height or altitude on our manifold, the gradient flow describes the path of steepest descent. This allows us to define the *stable* and *unstable manifolds* of a critical point (Definition 1.2.1), which yield a "cell decomposition" of M. In order for guarantee that this decomposition is well-behaved, we need to impose a further notion of stability on f, known as the *Morse-Smale* condition (Definition 1.2.9).

Having developed these fundamental concepts, we devote the latter half of the chapter to exploring the power of Morse theory. One of the original motivations of Morse theory in [Mil63] is to understand how the structure of

$$M^\alpha = \{p \in M \mid f(p) \leq \alpha\}$$

changes as α changes. As long as α does not pass a critical value of f, there is no significant topological change (Theorem 1.3.1, Theorem 1.3.4). We can gain even further insight into the topological structure of M using the *Morse complex* and its resulting homology, defined in Section 1.3.2. Remarkably, the Morse homology depends neither on the choice of Morse function nor pseudo-gradient (Theorem 1.3.7), and is ultimately isomorphic to the cellular homology of M (Theorem 1.3.9). In this way, information about the topology of M can be completely unlocked by understanding the structure of critical points of a Morse function and flows between them. Using Morse homology, we recover many familiar topological invariants, such as the Betti numbers, Euler characteristic, and Poincaré polynomial. Finally, we state the infamous *Morse inequalities*, which bound the number of critical points a Morse function on M in terms of the Betti numbers.

1.1 Definitions and Basic Properties

We begin, crucially, by building up to the definition of a Morse function. This allows us to define the Morse index and Morse Lemma. We then cover the basics of gradient flow and generalized pseudo-gradient flow, which are crucial to the more modern perspective. Most of this preliminary material is drawn from the beginnings of [AD14, Mil63], although the reader is also invited to see [CIN06, Mat02] for additional coverage.

1.1.1 Morse Functions

Given a (smooth) function $f\colon M \to \mathbb{R}$ on a (smooth) manifold M, we can study the critical points of this function to reveal structural information about the underlying space. If we think of f as describing the topography of the manifold, critical points are the places where manifold is locally "flat."

Definition 1.1.1. A *critical point* of $f\colon M \to \mathbb{R}$ is a point $a \in M$ such that $(df)_a = 0$. We denote the collection of critical points of f by $\mathrm{Crit}(f)$.

At a critical point of f, we can define the second-order derivative, usually called the Hessian. In general, a function on a manifold may not have a second derivative that is independent of a choice of chart; however the Hessian is well-defined (as a bilinear form) on the vector subspace $\ker(df)_a \subseteq T_aM$.

Definition 1.1.2. For $a \in \mathrm{Crit}(f)$, the *Hessian* of f is a symmetric, bilinear form on T_aM,

$$(d^2f)_a : T_aM \times T_aM \to \mathbb{R}.$$

To define $(d^2f)_a$ on $v, w \in T_aM$, we extend both v, w to local vector fields V, W in some neighborhood of a. Then

$$(d^2f)_a(v, w) = (V(Wf))(a) = v(Wf).$$

We can think of $v(Wf)$ as something akin to the directional derivative of the smooth function $Wf\colon M \to \mathbb{R}$ along the tangent vector v. To see that the form is symmetric, note that

$$0 = (df)_a([V, W]_a) = V(Wf)(a) - W(Vf)(a) = v(Wf) - w(Vf)$$

for $a \in \mathrm{Crit}(f)$.

The same computation shows that the form is well-defined on critical points: Suppose that we had instead chosen different extensions $\tilde{V}$ of v and $\tilde{W}$ of w. Then $\tilde{V}_a = v = V_a$, and moreover

$$v(\tilde{W}f) = w(Vf) = v(Wf),$$

since $(df)_a([V, \tilde{W}]_a) = 0 = (df)_a([V, W]_a)$. Hence the definition of the Hessian is independent of the local extensions of v, w.

Remark 1.1.3. If we choose a local coordinate system $\phi = (x_1, \ldots, x_n)$ in a neighborhood U of a, the matrix

$$H_{ij} = \frac{\partial^2 f}{\partial x_i \partial x_j}(a)$$

represents $(d^2 f)_a$ with respect to the basis $\frac{\partial}{\partial x_1}\Big|_a, \ldots, \frac{\partial}{\partial x_n}\Big|_a$.

We can use the language of bilinear forms to describe corresponding properties of the Hessian. The most crucial notion for Morse theory is the non-degeneracy of critical points.

Definition 1.1.4. A critical point is *non-degenerate* if its Hessian is non-singular. A function is said to be a *Morse function* if all its critical points are non-degenerate.

In a sense, non-degeneracy imposes a notion of stability onto our critical points. This idea is made more precise by the alternate perspective offered in [AD14, Exercise 2]. Recall (or see Appendix A.1) that the cotangent bundle T^*M has a natural manifold structure and the differential map $df\colon M \to T^*M$ is a section of this bundle. In this terminology, the collection of critical points of f is the pre-image of the zero section under df. It follows that a point is a non-degenerate critical point of f exactly when the submanifold $df(M)$ is transverse (see Definition 1.2.5) to the zero section at the point in question.

It is certainly easy to construct examples of functions with degenerate critical points, so we might worry that Morse functions are few and far between. On the contrary, Morse functions are both abundant and generic (cf. [AD14, §1.2] and [Mil63]). More specifically, every compact manifold M admits many Morse functions and every smooth function on M can be approximated by Morse functions. Some common examples of Morse functions include the square distance to a point and the height function (cf. [AD14, §1.4]), the latter of which we will explore on the sphere.

Example 1.1.5 (The height function on S^n)**.** Consider the n-sphere

$$S^n = \left\{ x = (x_1, \ldots, x_{n+1}) \mid \sum_{i=1}^{n+1} x_i^2 = 1 \right\} \subseteq \mathbb{R}^{n+1}$$

under the height function $f(x_1, \ldots, x_{n+1}) = x_{n+1}$. Restricting to the sphere, we can write f as a function of n variables, since $x_{n+1} = (1 - \sum_{i=1}^n x_i^2)^{1/2}$ (where the square root is taken to be positive in northern hemisphere and negative in southern hemisphere).

We naturally expect this function to have two critical points, namely the north and south poles. To verify our intuition, we check the derivative, whose i^{th} component is

$$\frac{\partial f}{\partial x_i} = \frac{-x_i}{x_{n+1}}$$

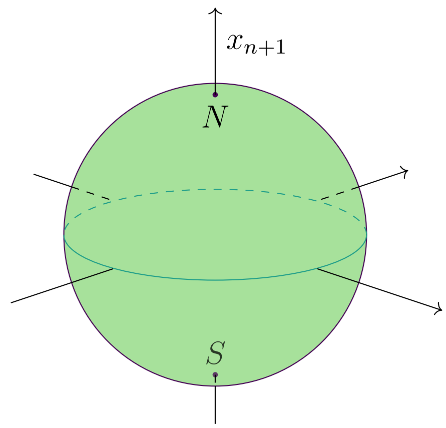

Figure 1.1: The height function on the n-sphere (when $n = 2$).

for $1 \leq i \leq n$, where we think of x_{n+1} as a function of n variables. Thus $(df)_x$ vanishes precisely when $x_1 = \cdots = x_n = 0$. This implies that $x_{n+1}^2 = 1$, which is to say that x must be $(0, \ldots, 0, \pm 1)$. Moreover, these critical points are nondegenerate, as is clear by examining the Hessian, whose components are

$$\frac{\partial^2 f}{\partial x_i \partial x_j} = \begin{cases} \frac{-1}{x_{n+1}} + \frac{x_i^2}{x_{n+1}^3} & i = j \\ \frac{x_i x_j}{x_{n+1}^3} & i \neq j \end{cases}$$

for $0 \leq i, j \leq n$. At either pole, the off-diagonal entries will vanish and we will be left with $\frac{-1}{x_{n+1}}$ down the diagonal. Therefore the second derivative matrix is either the identity (at the south pole) or its negative (at the north pole), both of which are certainly non-singular.

We are used to classifying different types of critical points based on the nearby behavior of the space, and this information is typically gained by inspecting the second derivative. For instance, in the example above, we know the north pole is a maximum and the south pole is a minimum, which can be verified in lower dimensions by checking the sign of the determinant of the second derivative matrix. We can generalize this idea using the Morse index, which describes the number of linearly independent directions in which f is decreasing.

Definition 1.1.6. The *Morse index* $\operatorname{ind}(a)$ of a critical point a is the index of $(d^2 f)_a$, that is, the maximum dimension of a subspace upon which the Hessian at a is negative-definite.

The Morse index can be counted as the number of negative entries in the diagonalization of $(d^2 f)_a$. A local minimum is thus a critical point with Morse index 0, while a local maximum has index equal to the dimension of the manifold. When f is a function of two variables, the critical points of index 1 are commonly called saddle points. In addition to providing information about the manifold around a critical point, the Morse index also completely determines the behavior of f at this point.

Theorem 1.1.7 (The Morse Lemma)**.** *Given a critical point $a \in \mathrm{Crit}(f)$, there exists a neighborhood $\Omega(a)$ of a and a diffeomorphism $\phi\colon (\Omega(a), a) \to (\mathbb{R}^n, 0)$ such that*

$$(f \circ \phi^{-1})(x_1, \ldots, x_n) = f(a) - \sum_{i=1}^{\mathrm{ind}(a)} x_i^2 + \sum_{i=\mathrm{ind}(a)}^{n} x_i^2.$$

The student of multivariable calculus will recall that a function is closely approximated in a small neighborhood by a quadratic function associated to its second order derivative. The Morse Lemma yields a stronger result, declaring the two to be equal after a possible change of chart. The neighborhood $\Omega(a)$ that appears in the lemma statement is called a *Morse chart*; these charts are discussed in far more detail in [AD14, §2.1]. The following corollary is immediate:

Corollary 1.1.8. *Critical points of a Morse function are isolated.*

Proof. We offer two proofs of this corollary, one with the Morse Lemma and one without. Let $a \in \mathrm{Crit}(f)$ and suppose that $f(a) = c$. Then, by the Morse Lemma, there is a diffeomorphism ϕ such that $(f \circ \phi^{-1})(x_1, \ldots, x_n) = c - \sum_i x_i^2 + \sum_i x_i^2$. In these coordinates, we have $\frac{\partial(f \circ \phi^{-1})}{\partial x_i} = \pm 2x_i$. So the Jacobian is 0 just when $x_i = 0$ for all i, which occurs only at $0 = \phi(a)$.

We can also prove that a nondegenerate critical point of a function is isolated without using the Morse lemma. As discussed previously, the zero section Z is transverse to $df(M)$ at any critical point of a Morse function. This implies (via Proposition 1.2.8) that $df^{-1}(Z) = \mathrm{Crit}(f)$ is a submanifold of M of dimension $\dim M - \dim(T^*M) + \dim(Z) = 0$, which is to say $\mathrm{Crit}(f)$ is discrete. □

In particular, this corollary implies that Morse functions on compact manifolds have finitely many critical points.

1.1.2 Gradient Flow

In order to obtain information about the geometry of M from f, we can examine the gradient vector field. On a Riemannian manifold (M, g), we can define the *gradient* of $f\colon M \to \mathbb{R}$ at $p \in M$ to be the unique vector field ∇f determined by

$$g(\nabla_p f, v) = (df)_p(v)$$

for every $v \in T_pM$. When $M = \mathbb{R}^n$, and g is the flat metric on $\mathbb{R}^n$, the gradient is precisely the first derivative matrix

$$\nabla f = \left[\frac{\partial f}{\partial x_1}, \ldots, \frac{\partial f}{\partial x_n}\right],$$

where the x_i give the standard coordinate system on $\mathbb{R}^n$.

Definition 1.1.9. A *gradient flow line* is a curve

$$\varphi\colon J \to M$$

for an open interval $J \subseteq \mathbb{R}$ that satisfies the differential equation

$$\frac{d\varphi}{dt} + \nabla_\varphi f = 0.$$

The gradient flow lines are the integral curves of the vector field $-\nabla f$. As we then expect, the flow lines tell us how to "descend" along f.

Proposition 1.1.10. *The function f is non-increasing along gradient flow lines, and is strictly decreasing along a gradient flow line that does not contain a critical point.*

Proof. Let $\varphi\colon J \to M$ be a gradient flow line, and consider $f \circ \varphi\colon J \to \mathbb{R}$. The derivative of this composition is

$$\begin{aligned}\frac{d}{dt} f(\varphi(t)) &= (df)_{\varphi(t)} \left(\frac{d}{dt}\varphi(t)\right) \\ &= \left\langle \nabla_{\varphi(t)} f, \frac{d\varphi}{dt}(t) \right\rangle \\ &= \left\langle \nabla_{\varphi(t)} f, -\nabla_{\varphi(t)} f \right\rangle \\ &= -\left|\nabla_{\varphi(t)} f\right|^2 \leq 0\end{aligned}$$

with equality precisely when $\varphi(t)$ is a critical point of f. This computation shows that f is non-increasing along φ, and is strictly decreasing whenever the image of φ does not contain any critical points. □

Note that if the image of φ does contain a critical point a, then in fact it must be the constant curve $\varphi(t) = a$. One can quickly verify that the constant curve satisfies the necessary ordinary differential equation, and it follows that this curve is the unique solution. Thus the above proposition shows that there are two kinds of flow lines, the constant (or *steady state*) flow lines at critical points and the flow lines that stay away from critical points (but may get arbitrarily close) along which f is strictly decreasing.

Theorem 1.1.11. *Given any $p \in M$ there is a unique gradient flow line*

$$\varphi_p\colon \mathbb{R} \to M$$

such that $\varphi_p(0) = p$. The flow map $\Phi\colon M \times \mathbb{R} \to M$ given by $\Phi(p,t) = \varphi_p(t)$ is smooth.

We call this φ_p the *minimal* flow for p. The proof of this theorem involves some standard technology from the theory of ordinary differential equations, as detailed in the proof of [CIN06, Theorem 4.6].

Theorem 1.1.12. *For any gradient flow line $\varphi\colon \mathbb{R} \to M$, there exist critical points $a, b \in \operatorname{Crit}(f)$ such that*

$$\lim_{t\to-\infty} \varphi(t) = a \quad \text{and} \quad \lim_{t\to\infty} \varphi(t) = b.$$

We say that $a =: s(\varphi)$ is the starting point and $b =: e(\varphi)$ is the ending point of φ.

Proof. We will show that φ has a limit as $t \to \infty$, and that this limit is a critical point of f; the proof for $t \to -\infty$ is similar. We proceed by contradiction, supposing that there is some time t_0 at which φ leaves the (finite) union

$$\Omega = \bigcup_{a \in \mathrm{Crit}(f)} \Omega(a),$$

never to return. Here $\Omega(a)$ is the Morse chart of a (given in Theorem 1.1.7). Since $M \setminus \Omega$ contains no critical points, the derivative $\frac{d}{dt} f(\varphi(t)) < 0$ on $M \setminus \Omega$ by Proposition 1.1.10, so there is some $\varepsilon > 0$ with $\frac{d}{dt} f(\varphi(t)) \leq -\varepsilon$. Then for all $t \geq t_0$,

$$f(\varphi(t)) - f(\varphi(t_0)) = \int_{t_0}^{t} \frac{d}{ds}(f \circ \varphi)(s) ds \leq -\varepsilon(t - t_0).$$

However, this implies that $f(\varphi(t)) \leq -\varepsilon(t - t_0) + f(\varphi(t_0))$ and so $\lim_{t\to\infty} f(\varphi(t)) = -\infty$, which is absurd since we are working on a compact manifold. Thus $\lim_{t\to\infty} \varphi(t) \in \Omega(a)$ for some $a \in \mathrm{Crit}(f)$. We wish to show that the limit is in fact a, which follows from the same argument as above, taking arbitrarily small neighborhoods of a. That is, for any neighborhood $U \subseteq \Omega(a)$ of a, we have $\lim_{t\to\infty} \varphi(t) \in U$, which is to say $\lim_{t\to\infty} \varphi(t) = a$. □

It can sometimes be useful to work within the context of a generalized gradient vector field, often called a *gradient-like* or *pseudo-gradient* vector field. This type of vector field has two essential properties: first, it points "in the same direction" as $-\nabla f$, vanishing precisely on the critical points of f; second, it is equal to $-\nabla f$ close to critical points.

Definition 1.1.13. A *pseudo-gradient (field)* adapted to f is a vector field X on M such that

(i) $(df)_p(X_p) \leq 0$ with equality if and only if $p \in \mathrm{Crit}(f)$,

(ii) in a Morse chart (given by Theorem 1.1.7) in the neighborhood of a critical point, X coincides with the negative gradient for the canonical metric on $\mathbb{R}$.

Pseudo-gradients are good substitutes for gradients in that they exhibit the same behavior without relying on an inner product. Moreover, any Morse function on any manifold admits a pseudo-gradient field, and the two theorems above will also hold for pseudo-gradients (see [AD14, §2.1]). The flow along the (pseudo-)gradient of a Morse function allows us develop a very rich and fruitful theory, as we will see with Morse homology in Section 1.3.2 and the flow category in Chapter 2.

1.2 The Morse-Smale Condition

The following subsection introduces the stable and unstable manifolds of a Morse function, with respect to some pseudo-gradient flow. After a brief digression into transversality, we can formulate the Morse-Smale condition. This material is pretty standard for most discussions of Morse theory, and our primary references are [AD14] and [CIN06].

1.2.1 Stable and Unstable Manifolds

The stable and unstable manifolds help us formalize the idea of how the flow is "attracted to" and "repelled from" critical points. These ideas, found more generally in the study of dynamical systems, are of particular importance in the context of Morse theory, providing a basis for formulating and proving essential subject material surrounding Morse homology and the flow category.

Definition 1.2.1. Let $a \in \mathrm{Crit}(f)$ and φ be the flow of a pseduo-gradient adapted to f. Define the *stable manifold* to be

$$W^s(a) = \{x \in M \mid e(\varphi_x) = a\},$$

and the *unstable manifold* to be

$$W^u(a) = \{x \in M \mid s(\varphi_x) = a\}.$$

Less formally, $W^s(a)$ is the collection of all points whose flow-lines "end up at" a, and $W^u(a)$ is the collection of points whose flow-lines "emanate from" a. These manifolds are sometime referred to as the *ascending* and *descending* manifolds of a, respectively.

Example 1.2.2. Picking up where Example 1.1.5 left off, we will find the stable an unstable manifolds of S^n under the height function $f(x_1, \dots, x_n, x_{n+1}) = x_{n+1}$. The flow lines are the lines of longitude (flowing from the north to the south) and the constant flow at the north pole N and the south pole S. The stable and unstable manifolds are

$$\begin{aligned} W^s(N) &= \{N\} & W^u(N) &= S^n \setminus \{N\} \\ W^s(S) &= S^n \setminus \{S\} & W^u(S) &= \{S\}. \end{aligned}$$

We can make these computations more explicit in the $n = 2$ case, using spherical coordinates. The spherical coordinate mapping (for a sphere of radius 1) is

$$(\theta, \phi) \mapsto (x, y, z) = (\cos\theta \sin\phi, \sin\theta \sin\phi, \cos\phi),$$

where $\theta \in [0, 2\pi]$ is angle from the positive x-axis to the xy-projection of the point and $\phi \in [0, \pi]$ is the angle from the positive z-axis to the point. Recall that (for a sphere of fixed radius 1) the gradient in spherical coordinates is given by $\nabla f = (\frac{\partial f}{\partial \theta}, \frac{1}{\sin\phi}\frac{\partial f}{\partial \phi})$ (assuming we are away from a critical point). The height function is thus $f(\theta, \phi) = \cos\phi$, with $-\nabla f = (0, 1)$. Therefore the unique solution to the initial value problem

$$\frac{d\varphi}{dt} + \nabla_\varphi f = 0, \text{ with } \varphi(0) = (\theta_0, \phi_0)$$

is just $\varphi = (\theta_0, t + \phi_0)$. That is, we keep the longitude constant and increase the polar angle ϕ steadily as time increases. Translating back into (x, y, z)-coordinates, we see φ precisely describes a meridian line at the longitude θ_0, flowing towards the south

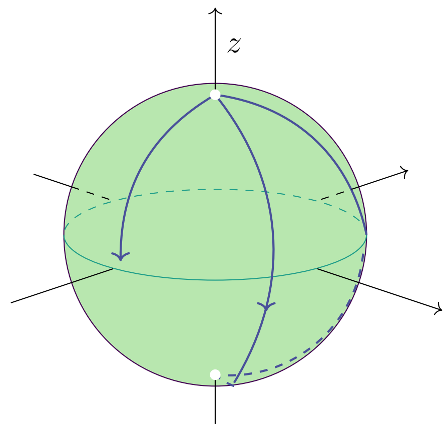

Figure 1.2: Gradient flow lines on the sphere. Every point besides the north a south pole lies on meridian line, which flows towards the south pole.

pole. It is not hard to believe that such a meridian line has $s(\varphi) = N$ and $e(\varphi) = S$, and this fact can be quickly verified via limit computations. We have shown that the gradient flow consists of a one-parameter family of integral curves in addition to the constant solutions id_N and id_S. Moreover, for every $p \in S^n$ besides N or S, the minimal flow φ_p belongs to this one-parameter family. This observation implies that the stable and unstable manifolds are exactly as stated above.

In this example, the stable and unstable manifolds are nice submanifolds of M, topologically equivalent to open disks. This behavior holds in general, via the following important theorem from dynamical systems.

Theorem 1.2.3 (Stable Manifold Theorem)**.** *The stable and unstable manifolds of $a \in \mathrm{Crit}(f)$ are submanifolds of M that are diffeomorphic to open disks, with*

$$\dim(W^u(a)) = \operatorname{codim}(W^s(a)) = \operatorname{ind}(a).$$

Moreover, every $p \in M$ lies on some flow line (which has its limit as $t \to -\infty$ at some critical point) and unstable manifolds of distinct critical points will be disjoint, so we can decompose M in terms of these submanifolds.

Theorem 1.2.4. *The unstable manifolds partition M into disjoint sets*

$$M = \bigcup_{a \in \mathrm{Crit} f} W^u(a).$$

This decomposition in terms of unstable manifolds roughly resembles a CW complex, with one (open) k-cell for each critical point of index k. We would hope that this decomposition describes M as a bonafide CW complex, but unfortunately this is not always the case. The next subsection develops a condition we can impose on f to ensure that we get an appropriately nice decomposition.

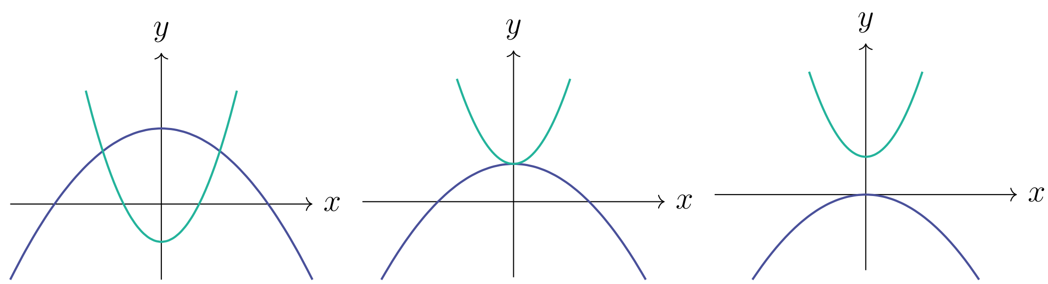

Figure 1.3: Examples of intersecting curves. The curves are transverse in the left-most plot, are tangent (not transverse) in the center plot, and are vacuously transverse in the right-most plot.

1.2.2 The Morse-Smale Condition

Before formulating the Morse-Smale condition, we take a brief interlude to discuss transversality. Transversality formalizes the notion of "general position" in differential topology, and can be seen as the complementary notion to tangency. For instance, as illustrated in Fig. 1.3, two curves on a surface intersect at a point transversely if and only if the point is not a tangent point, meaning the tangent lines (in the tangent plane to the surface) are distinct. Two submanifolds that do not intersect are vacuously transverse.

Definition 1.2.5. Let N and N' be submanifolds of M. We say that N and N' are *transverse at* $p \in M$ if either $p \notin N \cap N'$ or

$$p \in N \cap N' \text{ and } T_pN + T_pN' = T_pM,$$

Two submanifolds N and N' are *transverse*, denoted $N \pitchfork N'$, when this property holds for all $p \in M$.

Remark 1.2.6. Transversality is both a *generic* and a *stable* notion, which is to say that it can be attained through a small deformation and, once attained, is preserved by small deformations. We can see this in the illustrations above, where two submanifolds of $\mathbb{R}^2$ that are not transverse become so after a tiny deformation.

Proposition 1.2.7 ([AD14, Theorem A.3.1])**.** *If $N \pitchfork N'$, then $N \cap N'$ is a submanifold of M whose codimension is equal to* $\operatorname{codim}(N) + \operatorname{codim}(N')$.

If $\dim(N) + \dim(N') < \dim(M)$, then transversality means the absence of intersection. For example, two curves in $\mathbb{R}^3$ intersect transversely precisely when they do not intersect at all.

We can extend transversality to smooth maps. If $f : M \to N$ is a differentiable map and $N' \subseteq N$ is a submanifold, then f is *transverse* to N' at $p \in N$ if either $p \notin f(M) \cap N'$ or

$$p = f(q) \in f(M) \cap N' \text{ for some } q \in M, \text{ and } (T_qf)(T_qM) + T_pN' = T_pN.$$

As before, if f and N' are transverse at all $p \in N$, we say they are *transverse* and denote the relation by $f \pitchfork N'$. We say two maps f and f' are transverse, written

$f \pitchfork f'$, if their images (seen as smooth submanifolds of a common ambient manifold) are transverse.

Proposition 1.2.8 ([AD14, Proposition A.3.5])**.** *If $f \pitchfork N'$ and $f^{-1}(N')$ is non-empty, then $f^{-1}(N') \subseteq M$ is a submanifold of dimension* $\dim(M) - \dim(N) + \dim(N')$.

With this understanding of trasversality, we can now define the Morse-Smale condition. Essentially, this additional stipulation imposes a degree of stability on the stable and unstable manifolds. While the unstable manifolds always partition M, this decomposition may not be a CW complex, as we shall see in Example 1.2.13; requiring that the Morse-Smale condition is satisfied fixes this issue.

Definition 1.2.9. A pseudo-gradient field X adapted to a Morse function $f : M \to \mathbb{R}$ satisfies the *Smale-condition* (sometimes called the *transversality condition*) if

$$W^u(a) \pitchfork W^s(b) \quad \text{for all} \quad a, b \in \mathrm{Crit}(f).$$

We say the pair (f, X) is *Morse-Smale*. If $X = -\nabla f$ with respect to some metric g, we might instead say (f, g) is Morse-Smale.

Remark 1.2.10. Smale [Sma61] showed that Morse-Smale pairs not only exist, but are dense in the sense that every Morse function f with psuedo-gradient X can be replaced by a Morse-Smale pair (f', X') that is 'close' to (f, X) in a C^1 sense, as explained in [AD14, §2.2.c].

Certain stable and unstable manifolds always intersect transversely; for example, $W^u(a) \pitchfork W^s(b)$ when $a = b$ or when $f(a) \leq f(b)$ for $a \neq b$. In the first scenario, the claim follows from the Stable Manifold Theorem and the fact that $W^u(a) \cap W^s(b) = \{a\}$; in the second, the manifolds are vacuously transverse since f decreases along flow lines. By Proposition 1.2.7, if the pseudo-gradient satisfies the Smale condition, then for all $a, b \in \mathrm{Crit}(f)$, we have

$$\mathrm{codim}(W^u(a) \cap W^s(b)) = \mathrm{codim}(W^u(a)) + \mathrm{codim}(W^s(b)),$$

and so $\dim(W^u(a) \cap W^s(b)) = \mathrm{ind}(a) - \mathrm{ind}(b)$. Thus a Morse-Smale pair will not flow from a to b if $\mathrm{ind}(a) < \mathrm{ind}(b)$.

In any case, we denote the intersection $W^u(a) \cap W^s(b)$ by $W(a, b)$, which consists of all points on the trajectories connecting a to b,

$$W(a, b) = \{x \in M \mid s(\varphi_x) = a \text{ and } e(\varphi_x) = b\}.$$

In the Morse-Smale case, $W(a, b)$ is a submanifold of M. The group $\mathbb{R}$ acts on $W(a, b)$ by translations in time, and this action is free (meaning all the stabilizers are trivial) when $a \neq b$. Consequently, the quotient

$$\mathcal{M}(a, b) := W(a, b)/\mathbb{R}.$$

is a manifold of dimension $\mathrm{ind}(a) - \mathrm{ind}(b) - 1$. Note then that the Morse-Smale condition also prohibits flow lines between points with the same Morse index. We think of $\mathcal{M}(a, b)$ as the collection of flow lines from a to b, although we can also view $\mathcal{M}(a, b)$ as a submanifold of M by identifying it with the following space.

Definition 1.2.11. Let $\alpha \in (f(b), f(a))$ and define $W(a,b)^\alpha = W(a,b) \cap f^{-1}(\alpha)$.

If $a \neq b$, then $W(a,b)^\alpha$ is a submanifold of dimension $\dim(M) - 1$. To see this is the case, note that α is a regular (non-critical) value of $f|_{W(a,b)}$ and so $f|_{W(a,b)}^{-1}(\alpha) = W(a,b)^\alpha$ is a submanifold of codimension 1 by the regular value theorem (see [AD14, §A.2.c]). Furthermore, if $\alpha, \beta \in (f(b), f(a))$, then $W(a,b)^\alpha$ is diffeomorphic to $W(a,b)^\beta$. This is just to say that the diffeomorphism type of $W(a,b)^\alpha$ does not depend on α.

Theorem 1.2.12. *For critical points $a \neq b$, the map $W(a,b)^\alpha \times \mathbb{R} \to W(a,b)$ given by $(p,t) \mapsto \varphi_p(t)$ is a diffeomorphism. Moreover, we have $W(a,b)^\alpha \cong \mathcal{M}(a,b)$.*

Proof. We prove the second part of the theorem, pointing the reader to [CIN06, Proposition 7.7] for proof of the first part. Note that $\mathbb{R}$ acts on both the domain (on the right factor by addition) and the codomain (by flow) and the map $(p,t) \mapsto \varphi_p(t)$ is $\mathbb{R}$-equivariant— meaning that $s \cdot (p,t) = s \cdot \varphi_p(t)$— we may quotient out by $\mathbb{R}$ on both sides to get the desired result. □

Although it can be helpful to view $\mathcal{M}(a,b)$ as a submanifold of M, with an induced topology, we can also topologize this space as a subspace of continuous maps; this approach will prove more fruitful in developing the flow category in Chapter 2.

To conclude this section, we examine a prototypical example in topology: the torus.

Example 1.2.13 (Morse functions on T^2)**.** We present three different "height" functions on the 2-dimensional torus, illustrated in Fig. 1.4. The first of these height

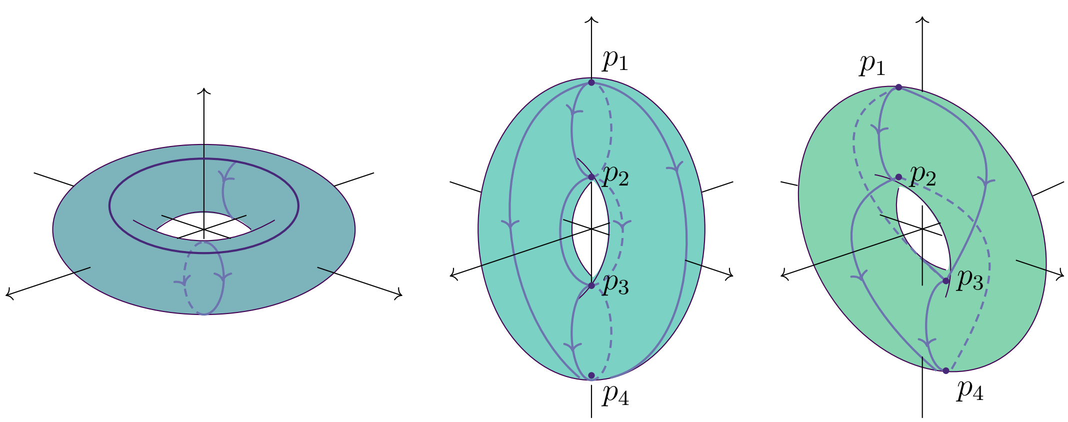

Figure 1.4: Height functions on the 2-dimensional torus. The first height function is not a Morse function, the second is Morse but not Morse-Smale, and the third is Morse-Smale.

functions is not a Morse function, as is evident by the circles of critical points on the top and bottom of the torus. However, if we stand the torus on its end, then the height function is a Morse function.[1] This embedding is given by

$$(\theta, \phi) \mapsto (x, y, z) = (b\cos(\phi),\ (a + b\sin(\phi))\cos(\theta),\ (a + b\sin(\phi))\sin(\theta))$$

with $0 < b < a$. Thus the height function is $f(\theta, \phi) = (a + b\sin(\phi))\sin(\theta)$, with gradient

$$\nabla f = [(a + b\sin(\phi))\cos(\theta),\ b\cot(\phi)\sin(\theta)]\,.$$

We can see we have four critical points at $(\theta, \phi) = (\pm\frac{\pi}{2}, \pm\frac{\pi}{2})$. We will say $p_1 = (\frac{\pi}{2}, \frac{\pi}{2})$, $p_2 = (\frac{\pi}{2}, -\frac{\pi}{2})$, $p_3 = (-\frac{\pi}{2}, \frac{\pi}{2})$, and $p_4 = (-\frac{\pi}{2}, -\frac{\pi}{2})$, in descending order of height. Examining the picture (or the Hessian), we see that $\mathrm{ind}(p_1) = 2$, $\mathrm{ind}(p_2) = \mathrm{ind}(p_3) = 1$, and $\mathrm{ind}(p_4) = 0$. There are two flows from p_1 to p_2, two from p_2 to p_3, and two flows from p_3 to p_4. There are also two one-parameter families of flows from p_1 to p_4. However, we note that this set-up is not Morse-Smale, since there are flows between critical points of the same index.

We can see how the absence of the Morse-Smale condition affects the decomposition of M into unstable manifolds. Since we have a flow between two critical points both of the same index, we are in the awkward situation of having to attach an edge in the middle of another edge.

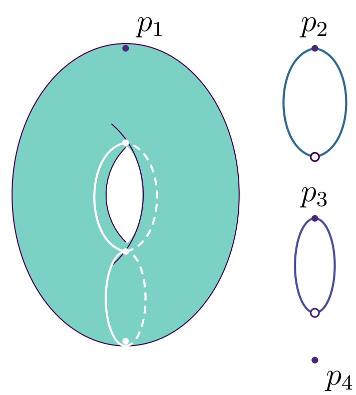

Figure 1.5: Decomposition of the vertical torus into unstable manifolds. Since the vertical torus is not Morse-Smale, the decomposition of T^2 into unstable manifolds is not a CW complex.

To fix the issue, we tilt the torus slightly, so that it is not standing on its edge, as illustrated in the rightmost torus of Fig. 1.4. Although the unstable manifolds do not change their diffeomorphism type (as the index of the critical points has not changed), there are no longer any flow lines between p_2 and p_3. Instead, the flows from both p_2 and p_3 end at p_4, and so we attach both edges to this singular vertex, yielding a true CW complex.

[1]For the sake of visualization, we talk about using the "same" height function and "moving" the torus. The more accurate way to think about this shift is to keep the underlying space the same— the torus— and change the function. The value that the points take under this new function is then the "height" they are prescribed in the picture.

1.3 A Survey of the Theory

The first part of this section covers the classical material from Milnor [Mil63], specifically the cellular decomposition of M understood in terms of sublevel sets. We then turn to the Morse complex and its resulting homology, as in [AD14, Chapters 3–4] and [Hut02, Chapters 2–4]. Given this homology theory, we have access to a plethora of applications, including the Euler characteristic, the Poincaré polynomial, and the Morse inequalities. We end by exploring these concepts via the sphere, torus, and projective space.

1.3.1 Sublevel Sets

Recall that the *level set* of a function $f\colon M \to \mathbb{R}$ for some value $\alpha \in \mathbb{R}$ is the set of all $p \in M$ such that $f(p) = \alpha$. We define the *sublevel set* of f at α to be

$$M^\alpha = \{p \in M \mid f(p) \leq \alpha\}.$$

One of the original results of classical Morse theory is that the topology of the (sub)level sets does not change, so long as we do not cross a critical value of f.

Theorem 1.3.1. *Let $\alpha, \beta \in \mathbb{R}$ and suppose that f does not attain a critical value in $[\alpha, \beta]$. If $f^{-1}([\alpha, \beta])$ is compact, then M^β is diffeomorphic to M^α.*

Proof. The idea of the proof is to push M^β down to M^α along the trajectories perpendicular to the level sets $f^{-1}(c)$ for $\alpha \leq c \leq \beta$. These trajectories are described by the flow of a pseudo-gradient vector field X. Let $\rho\colon M \to \mathbb{R}$ map

$$p \mapsto -\frac{1}{(df)_p(X)} \quad \text{for} \quad p \in f^{-1}([\alpha, \beta]),$$

and vanish outside of a compact neighborhood of $f^{-1}([\alpha, \beta])$. Then $Y = \rho X$ is a vector field with compact support, so its flow φ is complete. For $p \in M$, consider the function $\psi_p\colon \mathbb{R} \to \mathbb{R}$ which sends $t \mapsto f \circ \varphi_p(t)$. This function, given a moment in time, checks where a buoyant grain dropped at point p (at time 0) has flowed to under φ, and evaluates f at that point. We can think of ψ_p as describing the "altitude" of the flow line φ_p over time. If $\varphi_p(t) \in M$ is in $f^{-1}([\alpha, \beta])$, then the derivative of ψ_p is

$$\begin{aligned}
\frac{d}{dt}\left(f \circ \varphi_p(t)\right) &= (df)_{\varphi_p(t)}\left(\frac{d}{dt}\varphi_p(t)\right) && \text{chain rule,}\\
&= (df)_{\varphi_p(t)}\left(Y_{\varphi_p(t)}\right) && \text{definition of flow,}\\
&= Y_{\varphi_p(t)}(f) && \text{definition of differential,}\\
&= \rho(\varphi_p(t))X_{\varphi_p(t)}(f) && \text{definition of } Y,\\
&= -1 && \text{since } \varphi_p(t) \in f^{-1}([\alpha, \beta]).
\end{aligned}$$

Thus $f \circ \varphi_p(t) = f(p) - t$. Now for any $p \in f^{-1}([\alpha, \beta])$, we have $f \circ \varphi_p(\beta - \alpha) \leq \alpha$, thus the diffeomorphism $\varphi_Y^{\beta-\alpha}\colon p \mapsto \varphi_p(\beta - \alpha)$ sends M^β to M^α. $\square$

A similar diffeomorphism shows that M^α is a deformation retract of M^β. Specifically, we define the retraction $r : M^\beta \times [0,1] \to M^\beta$ which maps

$$(p,t) \mapsto \begin{cases} p & \text{if } f(p) \leq \alpha; \\ \varphi_p(t(f(p)-\alpha)) & \text{if } \alpha \leq f(p) \leq \beta. \end{cases}$$

To see that r indeed defines a deformation retraction, observe that $r(p,0) = \varphi_p(0) = p$, $r(M^\beta, 1) = M^\alpha$, and $r|_{M^\alpha} = id_{M^\alpha}$. The previous theorem (in conjunction with the Morse Lemma) yields Reeb's theorem.

Theorem 1.3.2 (Reeb's theorem)**.** *Let M be a compact manifold of dimension n. If there is a Morse function on M with only two critical points, then M is homeomorphic to S^n.*

Remark 1.3.3. Reeb's theorem does *not* entail that M is diffeomorphic to a sphere. The classic reference for this remark is [Mil56], where Milnor constructs a manifold that is homeomorphic but not diffeomorphic to S^7.

On the other hand, when $[\alpha, \beta]$ *does* contain a critical point of f, the sublevel sets M^α and M^β differ in their topology.

Theorem 1.3.4. *Suppose that $f(a) = \alpha$ and $\mathrm{ind}(a) = k$ for $a \in \mathrm{Crit}(f)$. For $\varepsilon > 0$ sufficiently small, suppose that $f^{-1}([\alpha - \varepsilon, \alpha + \varepsilon])$ is compact and contains no critical points of f other than a.*[2] *Then (for ε sufficiently small) the space $M^{\alpha+\varepsilon}$ has the homotopy type of $M^{\alpha-\varepsilon}$ with a k-cell attached (namely, $W^u(a)$).*

In combination with Theorem 1.3.1, this theorem tells us that the topology of M^α does not change until α passes the value of a critical point, and when α passes the value of a critical point with index k, we attach a k-cell to M^α. We thus get a handlebody of M in terms of unstable manifolds, with a k-cell for each critical point of index k, as discussed in Theorem 1.2.4.

1.3.2 Morse Homology

One important development of classical Morse theory is Morse homology, which uses information about the critical points and flows on a compact manifold equipped with a Morse-Smale pair (f, X) to develop a homology theory. Given the discussion in the previous subsection, it is perhaps unsurprising that Morse homology of M turns out to be isomorphic to its cellular homology, thus yielding information about the topology of the underlying manifold. Morse homology begins as all homology theories do, by constructing a chain complex.

Definition 1.3.5 (Morse complex)**.** Let $\mathrm{Crit}_k(f)$ denote the set of critical points of f with index k, and define the vector space over $\mathbb{Z}/2\mathbb{Z}$

$$C_k(f) = \left\{ \sum_{a \in \mathrm{Crit}_k(f)} m_a a \mid m_a \in \mathbb{Z}/2\mathbb{Z} \right\}.$$

[2]It may be the case that f has two critical points with the same critical value. However, we can often resolve this situation by a slight perturbation of f (see [AD14, §2.2.c]).

We define the boundary operator $\partial_k^X : C_k(f) \to C_{k-1}(f)$ by specifying its behavior on the basis elements. Given a critical point $a \in \mathrm{Crit}_k(f)$, the operator ∂_k^X sends a to a linear combination of points in $\mathrm{Crit}_{k-1}(f)$,

$$\partial_k^X(a) = \sum_{b \in \mathrm{Crit}_{k-1}(f)} m_X(a,b)b,$$

where $m_X(a,b) \in \mathbb{Z}/2\mathbb{Z}$ is the number (mod 2) of trajectories of X going from a to b. In other words, $m_X(a,b)$ is the modulo 2 cardinality of $W(a,b)$.

Of course, to verify that we indeed have a complex, we must show that $\partial_X^2 = 0$. For $a \in \mathrm{Crit}_k(f)$, we calculate

$$\partial_{k-1}^X(\partial_k^X(a)) = \sum_{c \in \mathrm{Crit}_{k-2}(f)} \left(\sum_{b \in \mathrm{Crit}_{k-1}(f)} m_X(a,b) m_X(b,c) \right) c.$$

Therefore it suffices to show that the inner sum is zero. The heart of the proof is to think of this number as the cardinality of the disjoint union

$$\coprod_{b \in \mathrm{Crit}_{k-1}(f)} W(a,b) \times W(b,c),$$

and show that this set of points is the boundary of a manifold of dimension 1. Since the boundary of a 1-manifold consists of an even number of points, and we are computing modulo 2, the desired result follows. We point the reader to [AD14, §3.1–3.2] for more details.

Remark 1.3.6. We have defined the chain complex over $\mathbb{Z}/2\mathbb{Z}$ since it is sometimes useful to work over a field (see, for instance, the proof of Proposition 1.3.13). Moreover, taking coefficients in $\mathbb{Z}/2\mathbb{Z}$ simplifies our considerations since we can ignore the signs, which are often algebraic representations of orientation. (To paraphrase Hatcher [Hat02], $\mathbb{Z}/2\mathbb{Z}$ homology is a natural tool in the absence of orientability.)

We can also define a Morse complex over $\mathbb{Z}$, yielding *integral homology*, so long as we are careful about orientation. Roughly, if we fix an orientation on the spaces of trajectories (by choosing orientations of the stable manifolds), then $\mathcal{M}(a,b)$ is an oriented compact manifold of dimension 0 whenever $\mathrm{ind}(a) - \mathrm{ind}(b) = 1$, and so is a finite number of points each with some $\pm$ orientation. We now define $m_X(a,b)$ to be the sum of these signs, noting that this sum modulo 2 is the coefficient in the $\mathbb{Z}/2\mathbb{Z}$ definition, and define the chain complex and boundary operator just as above. Essentially the same proofs work to show that we get a well-defined homology. Although we will not discuss integral Morse homology much further in this thesis, we will denote it by $HM_*(M;\mathbb{Z})$ to distinguish it from the modulo 2 homology $HM_*(M;\mathbb{Z}/2\mathbb{Z})$.

The chain complex clearly depends on our choice of a Morse-Smale pair (f,X), but remarkably the resulting homology groups are independent of this choice.

Theorem 1.3.7. *The homology of the Morse complex depends neither on the function nor the vector field.*

More precisely, this theorem says that for any two Morse-Smale pairs (f, X) and (f', X') on a (compact) manifold M, there is a morphism of complexes $(C_*(f), \partial_X) \to (C_*(f'), \partial_{X'})$ that induces an isomorphism on the homology. The proof relies on choosing a suitable deformation of f into f'. Henceforth we denote the boundary operator for the Morse homology by merely ∂, unless the vector field provides additional important context.

Remark 1.3.8. This result can be extended to manifolds with boundary provided that the critical points and flows between them stay sufficiently far from the boundary. For instance, given a cobordism M whose boundary decomposes as $\partial M = \partial_+ M \cup \partial_- M$, we can define the Morse complex $(C_*(f), \partial)$ in the same way as in the case without boundary, given appropriate conditions on X (although ultimately the homology is not dependent on X). This allows us to define the *relative Morse homology* $HM_*(M, \partial_+ M; \mathbb{Z}/2)$ using the complex $(C_*(f), \partial)$ (see [AD14, §3.2.d, §3.5, and §4.1]).

Given a homology theory, the natural question to ask is whether it is isomorphic to any homology theories we are familiar with. In any case, the answer will be interesting: either we have discovered a new way of thinking about something we already knew, or Morse homology contains some new information that other homology theories do not. As it turns out, the Morse homology of M is isomorphic to the cellular homology of M (see Appendix A.2). This isomorphism result is perhaps less surprising given the decomposition we saw in Theorem 1.2.4. In fact, a Morse-Smale pair (f, X) yields a cellular composition in terms of the unstable manifolds, from which we get a cellular complex $(K_*(f), \partial_X)$. The proof in [AD14, §4.9] takes the cells to be the compactification of the unstable manifolds $\overline{W}^u(a)$ for each $a \in \mathrm{Crit}(f)$.

Theorem 1.3.9. *There is an isomorphism $F\colon C_*(f) \to K_*(f)$ with $\partial_X \circ F = F \circ \partial$. Consequently, the Morse homology of M is isomorphic to its cellular homology.*

More specifically, F is the map that sends $a \mapsto \overline{W^u}(a)$. We may take coefficients in $\mathbb{Z}/2\mathbb{Z}$ or in $\mathbb{Z}$ (so long as we are mindful of orientations). It is well-known that cellular homology is isomorphic to singular homology, and so Morse homology agrees with the latter as well.[3] Thus manifolds of the same homotopy type will have the same Morse homology groups.

Morse homology can also yield information about the homotopy groups of a manifold. For instance, if M admits a Morse function with no critical points of index 1, then M is simply connected, since the fundamental group will be trivial. We can see this phenomenon in our example of the height function on S^n.

Example 1.3.10 (The n-sphere)**.** As we saw in Example 1.1.5, the height function on S^n has two critical points: the north pole N with index n and the south pole S with index 0. One can verify that indeed $(f, -\nabla)$ is Morse-Smale, and so we have a well-defined Morse complex, with $C_n(f) = \mathbb{Z}/2\mathbb{Z}[N]$, $C_0(f) = \mathbb{Z}/2\mathbb{Z}[S]$, and $C_k = 0$

[3] One can also show directly that Morse homology is isomorphic to singular homology, without going through cellular homology, as in [Hut02, §3]. The idea here is to flow a given simplex along X, then send it to the sum of the critical points that it "hangs on."

for $k \neq 0, n$, and all ∂_k are trivial.[4] Thus the homology groups are

$$HM_k(S^n; \mathbb{Z}/2) = \begin{cases} \mathbb{Z}/2 & k = 0, n \\ 0 & \text{otherwise.} \end{cases}$$

Of course, there are many other Morse functions on the sphere that will have more critical points of differing indices. For instance, on S^2, we can imagine "denting" the sphere so that it has two peaks with a valley in the middle, so that we get a new "height" function which looks something like the illustration in Fig. 1.6. Although this

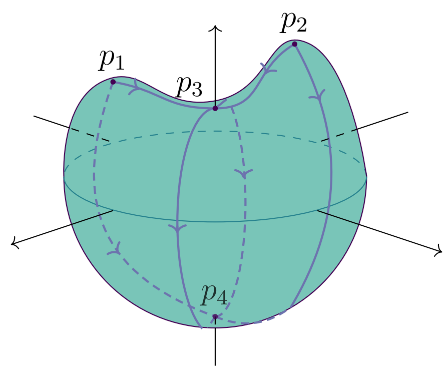

Figure 1.6: The alternate sphere.

manifold is still diffeomorphic to the sphere, the function now has two local maxima p_1, p_2 (of index 2), a saddle point p_3 (of index 1), and a minimum p_4 of (index 0). Now, our chain complex is comprised of $C_2 = \mathbb{Z}/2\mathbb{Z}[p_1] \oplus \mathbb{Z}/2\mathbb{Z}[p_2]$, $C_1 = \mathbb{Z}/2\mathbb{Z}[p_3]$, $C_0 = \mathbb{Z}/2\mathbb{Z}[p_4]$, and $C_k = 0$ for all $k \geq 3$. To understand the behaviors of the boundary operators, we can count the flow lines by hand to see that

$$\partial_2(p_1) = \partial_2(p_2) = p_3 \text{ and } \partial_1(p_3) = 2p_4 = 0.$$

Thus we have $\ker(\partial_1) = \operatorname{im}(\partial_2) = \mathbb{Z}/2\mathbb{Z}[p_3]$, $\ker(\partial_2) = \mathbb{Z}/2\mathbb{Z}[p_1 + p_2]$, $\ker(\partial_0) = \mathbb{Z}/2\mathbb{Z}[p_4]$ and all other images and kernels are trivial, so that

$$HM_k = \begin{cases} \mathbb{Z}/2\mathbb{Z} & k = 0, 2 \\ 0 & \text{otherwise.} \end{cases}$$

As we expect, the Morse homology of the alternative sphere matches the homology of S^2, even though the complexes are quite different.

Example 1.3.11 (The tilted torus)**.** Recall from Example 1.2.13 that if we tilt the torus slightly, the gradient flow of the height function f becomes Morse-Smale. There are four critical points of this function: p_1 of index 2, p_2 and p_3 of index 1, and p_4 of index 0. We can compute the Morse homology of the torus using the following chain complex:

$$\cdots \to 0 \xrightarrow{\partial_3} \mathbb{Z}/2\mathbb{Z}[p_1] \xrightarrow{\partial_2} \mathbb{Z}/2\mathbb{Z}[p_2] \oplus \mathbb{Z}/2\mathbb{Z}[p_3] \xrightarrow{\partial_1} \mathbb{Z}/2\mathbb{Z}[p_4] \xrightarrow{\partial_0} 0.$$

[4] Triviality is immediate except when $k = 1$, but one can quickly verify that in this case there are two flow lines from N to S, and so indeed $\partial_1(N) = 0$.

By counting the trajectories connecting the critical points naively, with reference to Fig. 1.4, we find

$$\partial_2(p_1) = 2p_2 + 2p_3 = 0, \quad \text{and} \quad \partial_1(p_2) = \partial_1(p_3) = 2p_4 = 0,$$

so the resulting homology groups are

$$HM_k = \begin{cases} \mathbb{Z}/2\mathbb{Z} & k = 0, 2; \\ \mathbb{Z}/2\mathbb{Z} \oplus \mathbb{Z}/2\mathbb{Z} & k = 1; \\ 0 & k \geq 3. \end{cases}$$

This is indeed the homology of the torus T^2 (cf. [Hat02, §2.1]).

In the next and final part of this chapter, we will see how truly powerful Morse homology is in uncovering structural information about the underlying manifold.

1.3.3 The Morse Inequalities and Other Applications

In this final section of the chapter, we present some general results using Morse homology. Since this theory is well known, we do not provide proofs but instead direct the interested reader to [AD14, §I.4].

We saw in the previous subsection that Morse homology is an invariant not only under homotopy type, but diffeomorphism type as well. This idea can be restated in the language of category theory.

Theorem 1.3.12. *Morse homology is a covariant functor from* **Diff***, the category of differentiable manifolds and smooth maps, into* $\mathbf{Ab}_*$*, the category of graded Abelian groups and graded homomorphisms.*

More precisely, a smooth map $u\colon M \to N$ induces a morphism

$$u_*\colon HM_*(M) \to HM_*(N),$$

that satisfies the functorality axioms. Furthermore, Morse homology is a so-called *homotopy functor*, meaning that if we have a smooth map

$$u\colon I \times M \to N$$

with $u(t, x) = u_t(x)$, then $(u_0)_* = (u_1)_*$. Thus Morse homology sends homotopy equivalences to isomorphisms.

Homology theory is motivated in part by the topologist's desire to develop topological invariants. One particularly well-known invariant is the *Euler characteristic*, defined for a finite CW complex C_* as the alternating sum

$$\chi(C_*) = \sum_k (-1)^k c_k,$$

where $c_k = \dim(C_k)$, that is, the number of k-cells in C_k. There are plenty of well-known results relating the Euler characteristic to homology groups, and now we can state some such results in the Morse context.

Proposition 1.3.13. *The number of critical points of a Morse function (modulo 2) depends only on the manifold M, not the function. Moreover,*

$$\sum_k (-1)^k c_k = \sum_k (-1)^k \dim HM_k(M;\mathbb{Z}/2).$$

Proof. Let M be a manifold of dimension n with a Morse-Smale pair (f, X). Considering our Morse complex (whose homology is independent of f and X), we have

$$\begin{aligned}
\#\mathrm{Crit}(f) &= \sum_{k=0}^{n} c_k \\
&= \sum_{k=0}^{n} \dim(C_k) \\
&= \sum_{k=0}^{n+1} \dim\ker\partial_k + \dim\operatorname{im}\partial_k && \text{by rank-nullity,} \\
&= \sum_{k=0}^{n} \dim\ker\partial_k + \dim\operatorname{im}\partial_{k+1} && \text{since } \dim\ker\partial_{n+1} = \dim\operatorname{im}\partial_0 = 0, \\
&\equiv \sum_{k=0}^{n} \dim\ker\partial_k - \dim\operatorname{im}\partial_{k+1} && \text{computing modulo 2,} \\
&= \sum_{k=0}^{n} \dim HM_k(M;\mathbb{Z}/2).
\end{aligned}$$

by rank-nullity again. Our applications of rank-nullity give $\dim C_k = \dim\ker\partial_k + \dim\operatorname{im}\partial_k$ and $\dim\ker\partial_k = \dim\operatorname{im}\partial_{k+1} + \dim HM_k(M;\mathbb{Z}/2)$, and so by a substitution and multiplication by $(-1)^k$,

$$(-1)^k \dim(C_k) = (-1)^k \dim HM_k(M;\mathbb{Z}/2) + (-1)^k(\dim\operatorname{im}\partial_k + \dim\operatorname{im}\partial_{k+1}).$$

Summing over $k = 0, \ldots, n$ gives the desired equality. □

This result allows us to define the Euler characteristic of a manifold.

Definition 1.3.14. The k^{th} *Betti number* of M is the integer

$$\beta_k(M) = \dim HM_k(M;\mathbb{Z}/2\mathbb{Z}),$$

and the *Euler characteristic* of M is

$$\chi(M) = \sum_k (-1)^k \beta_k(M).$$

We can also define the more general *Poincaré polynomial* as $P_M(t) = \sum_k \beta_k(M)t^k$. The Künneth formula in [AD14, §4.2] tells us that $P_{M_1\times M_2}(t) = P_{M_1}(t)P_{M_2}(t)$. Note that taking $t = -1$ recovers the Euler characteristic, and hence $\chi(M_1 \times M_2) = \chi(M_1)\chi(M_2)$ as well. We can consider these polynomials for some examples we saw at the end of Section 1.3.2:

- $P_{S^n}(t) = 1 + t^n$ and $\chi(S^n) = 1 + (-1)^n$,
- Writing $T^n = S^1 \times \cdots \times S^1$ (n times), we have $P_{T^n}(t) = \sum_{k=0}^{n} \binom{n}{k} t^k$ and $\chi(T^n) = 0$,

More examples are given in [AD14, Chapter 4]. We can make Proposition 1.3.13 more precise by comparing the number or critical points to the Betti numbers directly, rather than the dimension of the homology groups. Restating the proposition in this way, we have $\#\mathrm{Crit}(f) \geq \sum_{k=0}^{n} \beta_k$. This result is known as the Morse inequalities.

Theorem 1.3.15 (Morse Inequalities)**.** *The number of critical points of a Morse function on a manifold M is greater than or equal to the sum of the dimensions of the Morse homology groups $HM_*(M; \mathbb{Z}/2\mathbb{Z})$ of this manifold. More precisely, we have*

$$\beta_k \leq c_k$$

for all $k \geq 0$.

When the Morse inequalities are true equalities, the function is said to be a *perfect Morse function.*[5] The Morse inequalities have some interesting consequences; for instance, every Morse function on S^n must have at least two critical points (with one of index n and one of index 0).

There are many more beautiful applications of Morse homology, including the Poincaré duality, connections with the fundamental group, and proofs of the Brouwer Fixed Point Theorem and Borsuk-Ulam Theorem. There is undoubtedly more to be said about the remarkable tools that Morse theory bestows upon us, but this thesis is too small to contain it all.

[5]The archetypal example of a perfect Morse function (on complex projective space $\mathbb{CP}^n$) appears in Milnor [Mil63, §4].

Chapter 2

The Flow Category

The flow category $\mathscr{C}_f$ of a Morse function $f\colon M \to \mathbb{R}$ is the category whose objects are the critical points of f and whose morphisms are pieced-together "broken" flow lines between them. The idea is that $\mathscr{C}_f$ encodes information about the underlying manifold, which we can recover via various procedures that we will explore in following chapters. The goal of this chapter is to just define the flow category, and most of the work will be to precisely describe the morphisms of $\mathscr{C}_f$.

In Section 1.2.2, we defined the moduli space of flow lines as the quotient space $\mathcal{M}(a, b) = W(a, b)/\mathbb{R}$, where $W(a, b)$ consists of the points of M residing upon a flow line φ "starting" at a and "ending" at b. Of course, what we really mean is that $\varphi(t)$ has its limit at a (respectively b) as $t \to -\infty$ (respectively $t \to \infty$), which places us in a rather unfortunate situation of being unable to concatenate flow lines in a way that makes sense. To remedy this situation, we reparametrize φ to obtain the *height-parametrized gradient flow* γ (Definition 2.1.1), which we can then glue together as desired. For instance, if γ_1 connects critical points a to c, and γ_2 connects critical points c to b, then we can form the *broken flow* $\gamma_1 \circ \gamma_2$ from a to b that is the natural concatenation of the two. Assembling all such broken flow lines, we get the *moduli space of broken flow* $\overline{\mathcal{M}}(a, b)$ (Definition 2.1.4), the compactification of $\mathcal{M}(a, b)$. These $\overline{\mathcal{M}}(a, b)$ give the morphisms between the objects a and b in $\mathscr{C}_f$.

To record information about the topology of M in the flow category, we equip $\mathscr{C}_f$ with a bit more structure, namely, by viewing it as a *topological category* (a category internal to **Top**, see Definition 2.2.2). This additional structure amounts to endowing the objects and morphisms with a topology and ensuring certain maps (composition, source, target, and identity) are continuous. By viewing the objects as a (discrete) subspace of M, and the morphisms as subspaces of $\mathrm{Map}([f(b), f(a)], M)$ under the compact open topology, we can define the flow category as a genuine topological category (Definition 2.2.7).

2.1 Broken Flow

Following [CJS95b], we reparameterize the flow so that we may glue flow lines together and so form the moduli space of broken flow lines. The space of broken flow is also discussed in [AD14] (albeit without the reparameterization, under a different topology, and with other motivations), and the details of the height-parameterized flow perspective are outlined in [CIN06, §4.5, §8.4–8.5, §9.1].

2.1.1 Reparameterizing the Flow

Let φ be a flow from points a to b (for some critical points $a \neq b$). Then $h(t) := f(\varphi(t))$ is strictly decreasing, so we can think of h like the height. Note that

$$\frac{dh}{dt} = \frac{d}{dt}(f \circ \varphi) = \nabla_\varphi f \frac{d\varphi}{dt} = - \left|\nabla_\varphi f\right|^2 .$$

Moreover, h is a diffeomorphism $\mathbb{R} \to (f(b), f(a))$ and so we may talk about the smooth curve h^{-1}. The reparameterized flow lines we will consider are those

$$\gamma(t) = \varphi(h^{-1}(t))\colon (f(b), f(a)) \to M.$$

Note that $f(\gamma(t)) = t$ since $f \circ \gamma = f \circ (\varphi \circ h^{-1}) = h \circ h^{-1}$. Thus γ has the same image as φ, but now the parameter represents the value of f, which we can think of as the height. Furthermore, we can extend γ to a continuous map on $[f(b), f(a)]$ by setting $\gamma(f(b)) = b$ and $\gamma(f(a)) = a$. To see that this map is continuous, note that

$$\begin{aligned}
\lim_{t \to f(b)^+} \gamma(t) &= \lim_{t \to f(b)^+} \varphi(h^{-1}(t)) \\
&= \lim_{t \to f(b)^+} \varphi(\varphi^{-1} f|_{\mathrm{im}(\varphi)}^{-1})(t) \\
&= \lim_{t \to f(b)^+} f|_{\mathrm{im}(\varphi)}^{-1}(t) \\
&= f|_{\mathrm{im}(\varphi)}^{-1}(f(b)) \\
&= b,
\end{aligned}$$

and similarly $\lim_{t \to f(a)^-} \gamma(t) = a$.

Definition 2.1.1. Let φ be a (non-constant) gradient flow line of f, and define $h(t) = f(\varphi(t))$. The *height-reparameterized flow* is

$$\gamma = \varphi \circ h^{-1}\colon [f(b), f(a)] \to M.$$

If φ is a constant flow on a critical point a, we define its reparameterization $\gamma\colon \{f(a)\} \to M$ to be the constant flow $\gamma(f(a)) = a$. As before, we let γ_p denote the minimal unbroken flow-line through $p \in M$.

Remark 2.1.2. Note that this parametrization reverses the direction of the original flow line φ, since h is decreasing. So if φ is a flow from a to b, the height-parameterized γ is a flow from b to a. This is an aesthetic choice of the authors in [CJS95b] adopt as well. However, for the sake of conceptual consistency, we will still talk about some concepts in terms of the original φ parameterization. For instance, we will still say that γ "starts at" $a = s(\gamma)$ and "ends at" $b = e(\gamma)$, although strictly speaking the opposite is true; this well-intentioned untruth allows us to keep many of the same definitions from Chapter 1, and ultimately the moduli space $\mathcal{M}(a, b)$ does not depend on the parameterization of the flow.

We can find the new ordinary differential equation that such a reparameterized flow line will satisfy. Differentiating, we have

$$\begin{aligned}\frac{d\gamma}{dt} &= \frac{d}{dt}(\varphi \circ h^{-1}) \\ &= \frac{d\varphi}{dt}(h^{-1})\frac{dh^{-1}}{dt} \\ &= \nabla_{\varphi \circ h^{-1}} f \frac{1}{\frac{dh}{dt}(h^{-1})} \\ &= \nabla_\gamma f \frac{1}{|\nabla_\varphi f|^2 (h^{-1})} \\ &= \frac{\nabla_\gamma f}{|\nabla_\gamma f|^2},\end{aligned}$$

assuming $\nabla_\gamma f \neq 0$. Away from critical points, we may consider the vector field $X_p = \frac{\nabla_p f}{|\nabla_p f|^2}$. The integral curves of X, those that satisfy the differential equation

$$\frac{d\gamma}{dt} - \frac{\nabla_\gamma f}{|\nabla_\gamma f|^2} = 0, \tag{2.1.3}$$

are precisely the height-parameterized curves (see [CIN06, Lemma 4.8]), and so X and ∇f have the same integral curves under different parametrizations. The advantage of reparameterizating is that we can now "glue" flow lines together in a coherent way, as we will see in the following discussion.

2.1.2 The Space of Broken Flow Lines

There is a partial ordering on the critical points via flow lines. We say $a \succeq b$ if $W(a,b)$ is non-empty— or equivalently if there is a flow γ that starts at a and ends at b— and $a \succ b$ if $a \succeq b$ and $a \neq b$. We call a sequence of critical points $\mathbf{c} = \{c_1, \dots, c_k\}$ *ordered* if $c_i \succ c_{i+1}$ for all i. We will mostly be interested in ordered sequences of critical points from a to b, and so we let $\mathbf{c}(a,b)$ denote such a chain of critical points, that is, $\mathbf{c}(a,b) = \{a, c_1, \dots, c_k, b\}$ with $a \succ c_1 \succ \cdots \succ c_k \succ b$, and say that the *length* of this sequence is $l(\mathbf{c}(a,b)) = k$.[1] Define $l(a,b)$ to be the maximum length of such a sequence $\mathbf{c}(a,b)$. If $l(a,b) = 0$, meaning that there is *no* critical point c such that $a \succ c \succ b$, then we call b a *successor* of a. Finally, we define the moduli space of an ordered sequence $\mathbf{c} = \{c_1, \dots, c_k\}$ to be

$$\mathcal{M}(\mathbf{c}) = \mathcal{M}(c_1, c_2) \times \cdots \times \mathcal{M}(c_{k-1}, c_k).$$

Definition 2.1.4. Let $\mathcal{M}(a,b)$ be the moduli space of (height-parameterized) flows between critical points a, b, viewed as a subset of $\mathrm{Map}([f(b), f(a)], M)$. Define the

[1] Note that this definition of length results in the slightly awkward convention that an arbitrary chain $\mathbf{c} = \{c_1, \dots, c_k\}$ has length $k - 2$.

moduli space of broken flow lines from a to b by

$$\begin{aligned}\overline{\mathcal{M}}(a,b) &= \bigcup_{\mathbf{c}(a,b)} \mathcal{M}(\mathbf{c}(a,b)) \\ &= \bigcup_{\mathbf{c}(a,b)} \mathcal{M}(a,c_1) \times \cdots \times \mathcal{M}(c_k,b),\end{aligned}$$

where the union is over ordered sequences of critical points $\mathbf{c}(a,b) = \{a, c_1, \ldots, c_k, b\}$.

The curves in $\overline{\mathcal{M}}(a,b)$ are thus smooth on $M \setminus \mathrm{Crit}(f)$, hence referred to as *broken flow lines* (or sometimes *piecewise flow lines*). As is suggested by the notation, this space is meant to be the compactification of $\mathcal{M}(a,b)$ (asserted in [CJS95b] and proved in [Qin10]).

Theorem 2.1.5. *The space* $\overline{\mathcal{M}}(a,b)$ *is compact.*

Moreover, when (f,g) is Morse-Smale (Definition 1.2.9), the compactified moduli space carries the structure of a manifold with corners (cf. [Coh19, Proposition 2]).[2] This result, also in [Qin10], is crucial to the proof of the Morse-Smale case of the Cohen-Jones-Segal theorem which we will discuss in Chapter 4.

Remark 2.1.6. As discussed in Theorem 1.2.12, we could instead choose to view $\overline{\mathcal{M}}(a,b)$ with topology induced by the topology of M, and it turns out that compactness holds in this case as well (see [AD14, §3.2]).

There is a natural associative composition law

$$\overline{\mathcal{M}}(a,c) \times \overline{\mathcal{M}}(c,b) \to \overline{\mathcal{M}}(a,b)$$

given by the concatenation of curves. We denote the composition of two broken flows $\gamma_1 \in \overline{\mathcal{M}}(a,c)$ and $\gamma_2 \in \overline{\mathcal{M}}(c,b)$ by $\gamma_1 \circ \gamma_2 \in \overline{\mathcal{M}}(a,b)$.

Proposition 2.1.7. *The composition of broken flows* $(\gamma_1,\gamma_2) \to \gamma_1 \circ \gamma_2$ *is continuous.*

Proof. Consider two broken (height-parameterized) flow lines, $\gamma_1 \in \overline{\mathcal{M}}(a,c)$ and $\gamma_2 \in \overline{\mathcal{M}}(c,b)$. The composition is given by

$$(\gamma_1 \circ \gamma_2)(t) = \begin{cases} \gamma_2(t) & t \in [f(b), f(c)]; \\ \gamma_1(t) & t \in [f(c), f(a)]. \end{cases}$$

Note that $\gamma_1 \circ \gamma_2$ is indeed a well-defined broken flow, since $\gamma_2(f(c)) = c = \gamma_1(f(c))$ and the composition satisfies the defining differential equation (Equation (2.1.3)) away from critical points. Recall that the compact open topology on $\overline{\mathcal{M}}(a,b)$ is generated by the subbasis

$$D(K,U) = \{\gamma \in \overline{\mathcal{M}}(a,b) \mid \gamma(K) \subseteq U\}$$

for compact $K \subseteq [f(b), f(a)]$ and open $U \subseteq M$. We can decompose $K = K_1 \cup K_2$ with $K_2 \subseteq [f(b), f(c)]$ and $K_1 \subseteq [f(c), f(a)]$, so that the preimage is

$$\circ^{-1}(D(K,U)) = D(K_1,U) \times D(K_2,U),$$

that is, a product of subbasis elements. □

[2] A *manifold with corners* is a (second countable, Hausdorff) space such that each point has neighborhood U and homeomorphism $\phi\colon U \to \mathbb{R}^{n-k} \times [0,\infty)^k$ for some k, such that the transition maps are smooth.

2.2 Topology of the Flow Category

The following discussion introduces the flow category, as defined in [CJS95b]. We can view the flow category as a topological category (a category internal to **Top**) as in [ML71]. We give a general overview of the basic definitions for topological categories before turning to the flow category specifically.

2.2.1 Categories Internal to Top

When working with a small category, meaning that the collections of objects and morphisms are sets, we can equip the objects and morphisms with some additional structure. While the general theory merely requires that the fixed ambient category is finitely complete[3] (cf. [ML71, §XII.1]), we will restrict our attention to *categories internal to* **Top**, which we shall call *topological categories*.

Remark 2.2.1. The term 'topological category' is somewhat ambiguous. A category internal to **Top** should not be confused with the topologically enriched categories in [Rie14, §3] nor the topological concrete categories of [rAHS04, S VI.21]. In this thesis, topological categories shall be understood to be categories internal to **Top**.

Definition 2.2.2. A category $\mathscr{C}$ *internal to* **Top** consists of *space of objects* $\mathscr{C}_0$ and a *space of morphisms* $\mathscr{C}_1$ in **Top**, together with four continuous maps:

$$\mathrm{dom}, \mathrm{cod}\colon \mathscr{C}_1 \to \mathscr{C}_0, \quad i\colon \mathscr{C}_0 \to \mathscr{C}_1, \qquad \mathscr{C}_1 \times_{\mathscr{C}_0} \mathscr{C}_1 \xrightarrow{\ \circ\ } \mathscr{C}_1 .$$

- The *domain map* $\mathrm{dom}\colon (f\colon X \to Y) \mapsto X$,
- The *codomain map* $\mathrm{cod}\colon (f\colon X \to Y) \mapsto Y$,
- The *identity map* $i\colon X \mapsto \mathrm{id}_X$,
- The *composition map* $\circ$ sends a pair of morphisms (f, g) to their composite $g \circ f = gf$. Here $\circ$ is defined on the pullback of $\mathscr{C}_1 \times_{\mathscr{C}_0} \mathscr{C}_1$:

$$\begin{array}{ccc} \mathscr{C}_1 \times_{\mathscr{C}_0} \mathscr{C}_1 & \xrightarrow{\pi_2} & \mathscr{C}_1 \\ {\scriptstyle \pi_1}\downarrow & & \downarrow{\scriptstyle \mathrm{cod}} \\ \mathscr{C}_1 & \xrightarrow[\mathrm{dom}]{} & \mathscr{C}_0 \end{array} .$$

These maps must satisfy a variety of compatibility conditions, expressed as diagrams in **Top**:

[3] A *finitely complete* category has all finite products, pullbacks, and a terminal object.

$$\begin{array}{ccc} \mathscr{C}_0 & \xrightarrow{i} & \mathscr{C}_1 \\ & \searrow^{=} & \text{dom}\downarrow\downarrow\text{cod} \\ & & \mathscr{C}_0 \end{array}, \qquad \begin{array}{ccccc} \mathscr{C}_1 & \xleftarrow{\pi_1} & \mathscr{C}_1 \times_{\mathscr{C}_0} \mathscr{C}_1 & \xrightarrow{\pi_2} & \mathscr{C}_1 \\ {\scriptstyle\text{cod}}\downarrow & & \downarrow{\scriptstyle\circ} & & \downarrow{\scriptstyle\text{dom}} \\ \mathscr{C}_0 & \xleftarrow[\text{cod}]{} & \mathscr{C}_1 & \xrightarrow[\text{dom}]{} & \mathscr{C}_0 \end{array}.$$

These two diagrams specify the domain and codomain of the identity map i and the composition map $\circ$, respectively. The following two diagrams assert that composition is unital (with identity i) and associative:

$$\begin{array}{ccccc} \mathscr{C}_0 \times_{\mathscr{C}_0} \mathscr{C}_1 & \xrightarrow{i\times\text{id}} & \mathscr{C}_1 \times_{\mathscr{C}_0} \mathscr{C}_1 & \xleftarrow{\text{id}\times i} & \mathscr{C}_1 \times_{\mathscr{C}_0} \mathscr{C}_0 \\ {\scriptstyle\pi_2}\downarrow & & \downarrow{\scriptstyle\circ} & & \downarrow{\scriptstyle\pi_1} \\ \mathscr{C}_1 & = & \mathscr{C}_1 & = & \mathscr{C}_1 \end{array}$$

$$\begin{array}{ccc} \mathscr{C}_1 \times_{\mathscr{C}_0} \mathscr{C}_1 \times_{\mathscr{C}_0} \mathscr{C}_1 & \xrightarrow{\circ\times\text{id}} & \mathscr{C}_1 \times_{\mathscr{C}_0} \mathscr{C}_1 \\ {\scriptstyle\text{id}\times\circ}\downarrow & & \downarrow{\scriptstyle\circ} \\ \mathscr{C}_1 \times_{\mathscr{C}_0} \mathscr{C}_1 & \xrightarrow[\circ]{} & \mathscr{C}_1 \end{array}.$$

Of course, any small category is a topological category under the discrete topology, but there is often more than one way to do it. Many familiar notions from category theory have internal counterparts, such as functors and natural transformations.

Definition 2.2.3. A *continuous functor* is a map $F\colon \mathscr{C} \to \mathscr{D}$ between two topological categories that consists of two continuous maps,

$$F_0\colon \mathscr{C}_0 \to \mathscr{D}_0 \quad \text{and} \quad F_1\colon \mathscr{C}_1 \to \mathscr{D}_1,$$

which are compatible with the four structure maps. That is, such that the following diagrams commute:

$$\begin{array}{ccc} \mathscr{C}_1 \times_{\mathscr{C}_0} \mathscr{C}_1 & \xrightarrow{F_1\times F_1} & \mathscr{D}_1 \times_{\mathscr{D}_0} \mathscr{D}_1 \\ {\scriptstyle\circ_{\mathscr{C}}}\downarrow & & \downarrow{\scriptstyle\circ_{\mathscr{D}}} \\ \mathscr{C}_1 & \xrightarrow[F_1]{} & \mathscr{D}_1 \end{array} \qquad \begin{array}{ccccc} \mathscr{C}_1 & \overset{\text{dom}}{\underset{\text{cod}}{\rightrightarrows}} & \mathscr{C}_0 & \xrightarrow{i} & \mathscr{C}_1 \\ {\scriptstyle F_1}\downarrow & & \downarrow{\scriptstyle F_0} & & \downarrow{\scriptstyle F_1} \\ \mathscr{D}_1 & \overset{\text{dom}}{\underset{\text{cod}}{\rightrightarrows}} & \mathscr{D}_0 & \xrightarrow[i]{} & \mathscr{D}_1 \end{array}.$$

We then assemble the category **TopCat** whose objects are topological categories and whose morphisms are continuous functors.

Definition 2.2.4. A *continuous natural transformation* $\eta\colon F \to G$ between a pair of continuous functors $F, G\colon \mathscr{C} \rightrightarrows \mathscr{D}$ consists of a continuous map $\eta\colon \mathscr{C}_0 \to \mathscr{D}_1$ such that the following diagrams commute in **Top**:

$$\begin{array}{ccc} \mathscr{C}_0 & \xrightarrow{\eta} & \mathscr{D}_1 \\ & {\scriptstyle F}\searrow & \downarrow{\scriptstyle\text{dom}} \\ & & \mathscr{D}_0 \end{array} \qquad \begin{array}{ccc} \mathscr{C}_0 & \xrightarrow{\eta} & \mathscr{D}_1 \\ & {\scriptstyle G}\searrow & \downarrow{\scriptstyle\text{cod}} \\ & & \mathscr{D}_0 \end{array}.$$

That is, the map η assigns every $X \in \mathscr{C}_0$ a morphism $\eta_X \colon FX \to GX$. The naturality condition on η requires that the following diagram also commutes:

$$\begin{array}{ccc} \mathscr{C}_1 & \xrightarrow{(G,\ \eta\circ\mathrm{dom})} & \mathscr{D}_1 \times_{\mathscr{D}_0} \mathscr{D}_1 \\ {\scriptstyle(\eta\circ\mathrm{cod},F)}\downarrow & & \downarrow{\scriptstyle\circ} \\ \mathscr{D}_1 \times_{\mathscr{D}_0} \mathscr{D}_1 & \xrightarrow[\circ]{} & \mathscr{D}_1 \end{array} .$$

Note that the commutativity of this diagram implies the usual naturality condition. That is, if $f \colon X \to Y \in \mathscr{C}_1$, the diagram above maps

$$\begin{array}{ccc} f & \longmapsto & (Gf, \eta_X) \\ \downarrow & & \downarrow \\ (\eta_Y, Ff) & \longmapsto & \eta_Y \circ Ff = Gf \circ \eta_X \end{array}$$

which is precisely the desired condition.

Remark 2.2.5. It is equivalent to require that η is a continuous functor

$$\eta \colon \mathscr{C} \times [1] \to \mathscr{D}$$

such that $\eta(-,0) = F$ and $\eta(-,1) = G$. Here [1] is the poset category $0 < 1$.

Definition 2.2.6. A *continuous equivalence* of categories internal to **Top** consists of two continuous functors $F \colon \mathscr{C} \leftrightarrows \mathscr{D} : G$ together with two continuous natural isomorphisms $\eta \colon \mathrm{id}_{\mathscr{C}} \to GF$ and $\epsilon \colon FG \to \mathrm{id}_{\mathscr{D}}$.

2.2.2 The Flow Category as a Topological Category

We can now define the flow category of a Morse function $f \colon M \to \mathbb{R}$, which records information about critical points and the flows between them. Moreover, we can endow the spaces of objects and morphisms with topologies, giving the flow category the structure of a topological category. In Chapter 4, we will see that we can recover the topological structure of M from the topology of the flow category.

Definition 2.2.7. The *flow category* of f is the category $\mathscr{C}_f$ whose objects are the critical points of f and whose morphisms are broken flow lines between these critical points. That is, for each $a, b \in \mathrm{Ob}\,\mathscr{C}_f$,

$$\mathscr{C}_f(a,b) = \overline{\mathcal{M}}(a,b).$$

Composition in this category is given by composition of broken flows.

To endow $\mathscr{C}_f$ with topological structure, we consider the objects, $(\mathscr{C}_f)_0 = \mathrm{Crit}(f)$, under the subspace topology. However, since critical points of a Morse function are isolated (Corollary 1.1.8), this is equivalent to a discrete topology. Each homspace $\overline{\mathcal{M}}(a,b)$ is topologized as a subspace of $\mathrm{Map}([f(b), f(a)], M)$, continuous maps under the compact open topology. The space of morphisms $(\mathscr{C}_f)_1$ is the disjoint union of all the homspaces, over all pairs of critical points $a, b \in \mathrm{Crit}(f)$. Note that whenever $a = b$, the only possible flow is the steady solution, so $\mathscr{C}_f(a,a) = \{\mathrm{id}_a\}$. We claim that the flow category is indeed a topological category, and so must show the structure maps are continuous. Recall that the four maps are

- the identity map $i\colon a \mapsto \mathrm{id}_a \in \mathscr{C}_f(a,a)$ for $a \in (\mathscr{C}_f)_0$,
- the domain map $\mathrm{dom}\colon \gamma \mapsto a$ for $\gamma \in \mathscr{C}_f(a,b)$,
- the codomain map $\mathrm{cod}\colon \gamma \mapsto b$ for $\gamma \in \mathscr{C}_f(a,b)$,
- the composition map $\circ\colon (\gamma_1, \gamma_2) \to \gamma_1 \circ \gamma_2$ for $\gamma_1 \in \mathscr{C}_f(a,b)$ and $\gamma_2 \in \mathscr{C}_f(b,c)$.

Proposition 2.1.7 takes care of composition, and the identity is clearly continuous since its domain is a discrete space. To see that the domain and codomain maps are continuous, it suffices to consider the pre-image of a single critical point a, which will be disoint unions of homspaces and so open in $(\mathscr{C}_f)_1$.

Example 2.2.8. Let us determine the flow category for a simple but instructive example. Recall the alternative sphere from Example 1.3.10, whose height function has four critical points: two maxima p_1, p_2, a saddle point p_3, and a minimum p_4. Thus our flow category has four objects, $\mathrm{Ob}\,\mathscr{C}_f = \{p_1, p_2, p_3, p_4\}$, and the homspaces are as illustrated in Fig. 2.1.

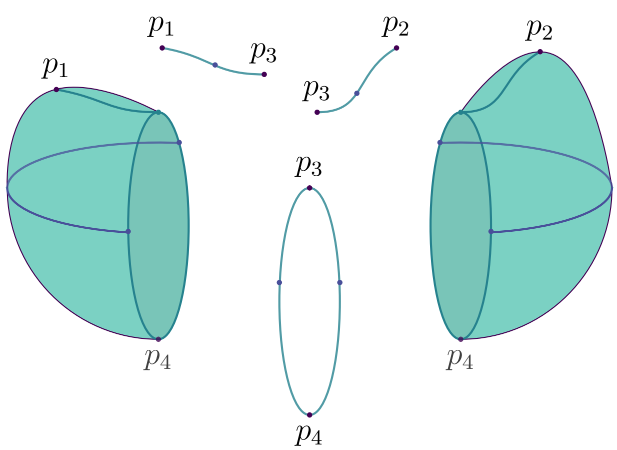

Figure 2.1: Homspaces of the flow category for the alternate sphere. The compactified moduli space $\overline{\mathcal{M}}(p_i, p_j)$ consists of all broken and unbroken flows connecting p_i and p_j. In the illustration above, we have shown the compactified moduli spaces (in a darker color) sitting inside the pieces of the alternate sphere that their flows cover. For instance, there are only two flow lines connecting p_3 and p_4 (up to composition by steady state flows), so $\overline{\mathcal{M}}(p_3, p_4)$ consists of two points.

We can see that both $\overline{\mathcal{M}}(p_1, p_4)$ and $\overline{\mathcal{M}}(p_2, p_4)$ consist of a one-parameter family of unbroken flowlines, in addition to the broken flows going through p_3. The other non-empty homspaces $\overline{\mathcal{M}}(p_1, p_3)$, $\overline{\mathcal{M}}(p_2, p_3)$, and $\overline{\mathcal{M}}(p_3, p_4)$ are just the usual moduli spaces of unbroken flows. Finally, we also have $\overline{\mathcal{M}}(p_i, p_i) = \{\mathrm{id}_{p_i}\}$ for each i. All other homspaces are empty.

Chapter 3

The Classifying Space of the Flow Category

Before examining the classifying space of the flow category $\mathscr{C}_f$, we first discuss the classifying space of an arbitrary small category $\mathscr{C}$. Despite its name, it is difficult to say precisely what a classifying space of a category "classifies" (but for one answer to this question, see [Wei05]). Intuitively, the classifying space records information about the way that we can "move" through the category via morphisms. This information about composable morphisms is recorded in a simplicial set known as the *nerve* of $\mathscr{C}$ (Definition 3.1.12). We can then turn the simplicial set into a topological space via the *geometric realization* (Definition 3.1.13); the resulting space is what is known as the *classifying space* of $\mathscr{C}$, denoted $B\mathscr{C}$. If we carefully adjust our definitions to respect the topological structure, we can form the classifying space for any topological category. Examples for particular flow categories are presented later in Section 4.1.

As is typical in algebraic topology, we are interested in this classifying space up to homotopy. It is perhaps unsurprising that an equivalence of categories induces a homotopy equivalence of their classifying spaces (Theorem 3.1.17). The second half of this chapter is devoted to developing and understanding different relationships between categories that yield an equivalence on their classifying spaces. One relatively well-known construction is the *twisted arrow category* of $\mathscr{C}$ (Definition 3.1.22), denoted $\mathrm{tw}(\mathscr{C})$. By relating the twisted arrow category with Segal's *edgewise subdivision* of a simplicial space (Definition 3.1.19), we get that $B\mathscr{C} \cong B\mathrm{tw}(\mathscr{C})$ (Corollary 3.1.24).

The final part of the chapter is dedicated to proving that a continuous functor between topological categories that is an levelwise homotopy equivalence induces a homotopy equivalence on the classifying spaces (Theorem 3.2.12) Fix. That is, we show that $\mathscr{C}_0 \simeq \mathscr{D}_0$ and $\mathscr{C}_1 \simeq \mathscr{D}_1$ implies $B\mathscr{C} \simeq B\mathscr{D}$. The homotopies on the objects and morphisms induce homotopy equivalences on every level of the nerve, and these levelwise homotopies assemble into a map of simplicial spaces (Theorem 3.2.7). The crucial piece of the argument is that the nerve is a *good* simplicial space,[1] which implies that the geometric realization preserves equivalence (Theorem 3.2.12). This work sets us up to prove the first part of the main theorem in Chapter 4, which

[1]Here, "good" is a technical term, see Definition 3.2.8.

establishes a homotopy equivalence $B\mathscr{C}_f \simeq M$ for an arbitrary Morse function f on M.

3.1 Some Simplicial Homotopy Theory

The following section covers some basic notions in simplicial homotopy theory, such as simplicial sets and their geometric realization, following classic references such as [May67] and [GJ99]. The discussion of simplicial sets is primarily based on [Rie08], and a more detailed exposition can be found in [Fri11]. The author also enjoyed reading [Bae18] for a less formal introduction to these concepts. Since we are particularly interested in categories internal to **Top**, we then generalize these definitions slightly to respect the topological structure, as in [ML71, §XII.1]. We can then define the classifying space for both categories and topological categories. Finally, we relate the edgewise subdivision of a simplicial set to the twisted arrow functor, as described in [Bar13, BR, Seg73].

3.1.1 Simplicial Sets

Simplicial sets generalize the geometric simplicial complexes found in algebraic topology. The following exposition introduces key notions and basic examples, the most important example being the nerve.

Let $\mathbf{\Delta}$ be the standard simplex category whose objects are finite, non-empty ordinals

$$[n] = \{0, 1, \dots, n\}$$

and whose morphisms are order-preserving maps. We can also regard the objects of $\mathbf{\Delta}$ as categories themselves, more specifically as posets, so $[n]$ is the category with $n + 1$ objects (the elements of the set above) and morphisms $i \to j$ whenever $i \leq j$.

Definition 3.1.1. A *simplicial set* is a (set-valued) presheaf on $\mathbf{\Delta}$, that is, a contravariant functor $\mathbf{\Delta} \to \mathbf{Set}$.

As is standard, we write X_n for the set $X[n]$, and call its elements n-simplicies. More generally, a *simplicial object* in a category $\mathscr{C}$ is a functor $X\colon \mathbf{\Delta}^{\mathrm{op}} \to \mathscr{C}$. We will be particularly concerned with simplicial objects in **Top**, which are called *simplicial spaces*.

The simplicial sets form a category, **sSet**, which is just the functor category $\mathbf{Set}^{\mathbf{\Delta}^{\mathrm{op}}}$.[2] More specifically, a map $X \to Y$ between simplicial sets is a natural transformation, and consists of maps $X_n \to Y_n$ that commute with the morphisms of $\mathbf{\Delta}$. In fact, as given by Proposition 3.1.5, it suffices to show that these maps between n-simplices commute with a smaller selection of maps, which we introduce presently.

For each $n \geq 0$, there are $n + 1$ injective *coface* maps $d^i\colon [n-1] \to [n]$, where the superscript indicates which object is not contained in the image. Similarly, there

[2]In general, the category of simplicial objects in a category $\mathscr{C}$ is denoted $\mathbf{s}\mathscr{C}$. For example, the category of simplicial spaces is **sTop**.

are $n+1$ surjective *codegeneracy* maps $s^j\colon [n+1] \to [n]$, where now the superscript indicates which object in the image is mapped onto twice. Explicitly,

$$d^i(k) = \begin{cases} k & k < i; \\ k+1 & k \geq i, \end{cases} \quad \text{and} \quad s^j(k) = \begin{cases} k & k \leq j; \\ k-1 & k > j, \end{cases}$$

for $0 \leq i, j \leq n$. It is straightforward to verify that these morphisms satisfy the following *cosimplicial relations*:

$$d^i d^j = d^j d^{i-1} \qquad i < j, \tag{3.1.2}$$

$$s^i s^j = s^j s^{i+1} \qquad i \leq j, \tag{3.1.3}$$

$$s^i d^j = \begin{cases} \mathrm{id} & i = j, j+1, \\ d^j s^{i-1} & i < j, \\ d^{j-1} s^i & i > j+1. \end{cases} \tag{3.1.4}$$

Proposition 3.1.5. *The morphisms of* $\mathbf{\Delta}$ *are generated by composing the coface and codegeneracy maps. In fact, any morphism* $f\colon [n] \to [m]$ *in* $\mathbf{\Delta}$ *can be expressed uniquely as a composite*

$$f = d^{i_k} \dots d^{i_1} s^{j_1} \dots s^{j_{k'}}$$

for $0 \leq i_1 < \cdots < i_k \leq m$ *and* $0 \leq j_1 < \cdots < j_{k'} \leq n$ *such that* $n + k - k' = m$.

Proof. An order-preserving function $f\colon [n] \to [m]$ is determined by its image in $[m]$ and those elements of $[n]$ on which f does not increase. Take $i_1, \dots, i_k \in [m]$ in unique increasing order to be those elements not in the image of f and $j_1, \dots, j_{k'} \in [n]$ (again in unique increasing order) to be the elements on which f does not increase. A quick verification proves the desired equality. □

The opposite category $\mathbf{\Delta}^{\mathrm{op}}$ has corresponding *face* maps d_i and *degeneracy* maps s_j. If X is a simplicial set, we have

$$d_i := Xd^i\colon X_n \to X_{n-1} \quad \text{and} \quad s_j := Xs^j\colon X_n \to X_{n+1},$$

for $0 \leq i, j \leq n$. Every morphism in $\mathbf{\Delta}^{\mathrm{op}}$ can similarly be expressed as a composition of face and degeneracy maps. These maps satisfy the dual relations to those given above, namely

$$d_j d_i = d_{i-1} d_j \qquad j < i, \tag{3.1.6}$$

$$s_j s_i = s_{i+1} s_j \qquad j \leq i, \tag{3.1.7}$$

$$d_j s_i = \begin{cases} \mathrm{id} & i = j, j+1, \\ s_{i-1} d_j & i < j, \\ s_i d_{j-1} & i > j+1. \end{cases} \tag{3.1.8}$$

These relations are called the *simplicial relations*.

The standard way to write down the data of a simplicial set is to provide the sets of n-simplicies X_n and the face and degeneracy maps that satisfy the necessary relations. In this way, the data of a simplicial set is entirely specified by the sets X_n and the d_i, s_j maps, prompting the following, alternative definition.

Definition 3.1.9. A *simplicial set* X is a collection of sets X_n for each integer $n \geq 0$ together with maps $d_i\colon X_n \to X_{n-1}$ and $s_j\colon X_{n-1} \to X_n$ for $1 \leq i, j \leq n$ that satisfy the simplicial relations (3.1.6), (3.1.7), and (3.1.8).

Given an n-simplex $x \in X_n$, we can visualize it as an n-dimensional tetrahedron whose $n+1$ vertices are ordered by $0, 1, \ldots, n$ and whose faces are labeled by simplices of the appropriate dimension. The image $d_i(x)$ of x under the i^{th} face map is the $(n-1)$-simplex that does not include the i^{th} vertex of x. Each of the $(n+1)$-simplices $s_0(x), s_1(x), \ldots, s_n(x)$ represent the same simplex geometrically, each with a different degeneracy; the image $s_j(x)$ is the simplex such that collapsing the edge between the j^{th} and $(j-1)^{th}$ vertices to a single point gives the n-simplex x. Accordingly, a simplex is called *degenerate* if it is the image of some s_j, and is *non-degenerate* otherwise. Unlike in a simplicial complex, we allow simplices to be degenerate.

Example 3.1.10 (The standard n-simplex)**.** The simplicial set called the *standard n-simplex* is the functor represented by $[n] \in \mathbf{\Delta}$. Letting $y\colon \mathbf{\Delta} \hookrightarrow \mathbf{sSet}$ denote the Yoneda embedding (see Appendix A.3.2), the standard n-simplex is just the image of $[n]$. That is,

$$\Delta^n := y[n] = \mathbf{\Delta}(-, [n]),$$

so $\Delta^n_k = \mathbf{\Delta}([k], [n])$ by definition. The face and degeneracy maps are given by pre-composition in $\mathbf{\Delta}$ by d^i and s^j, so

$$\begin{array}{cc} d_i\colon \Delta^n_k \to \Delta^n_{k-1} & s_j\colon \Delta^n_k \to \Delta^n_{k+1} \\ ([k] \xrightarrow{f} [n]) \mapsto ([k-1] \xrightarrow{d^i} [k] \xrightarrow{f} [n]) & ([k] \xrightarrow{f} [n]) \mapsto ([k+1] \xrightarrow{s^j} [k] \xrightarrow{f} [n]). \end{array}$$

Non-degenerate k-simplices correspond to the injective maps $[k] \to [n]$ in $\mathbf{\Delta}$; there is a unique non-degenerate n-simplex in Δ^n corresponding to the identity on $[n]$. There are many degenerate simplicies in this data as well: for instance, Δ^0 contains one element in each Δ^0_k, the zero function $[k] \to [0]$, which is degenerate for $k > 0$.

This perspective allows us to better understand the key role that Δ^n plays in $\mathbf{sSet}$. Since the Yoneda embedding is full and faithful, the maps $f\colon \Delta^n \to \Delta^m$ of simplicial sets are in bijection with the maps $f\colon [n] \to [m]$ in $\mathbf{\Delta}$. The maps $f_k\colon \Delta^n_k \to \Delta^m_k$ are given by post-composition by f. The Yoneda Lemma implies that simplicial maps $\Delta^n \to X$ correspond bijectively to the n-simplices in X, which is to say

$$\mathbf{sSet}(\Delta^n, X) \cong X_n.$$

An n-simplex $x \in X$ can thus be regarded as a map $x\colon \Delta^n \to X$ that sends the unique non-degenerate n-simplex in Δ^n to x. Lower-dimensional simplicies in X can be seen as a composition of maps in Δ^n, post-composed by x.

Example 3.1.11 (Total singular complex)**.** Perhaps unsurpisingly, one standard example of a simplicial set is related to the topological notion of a simplex. Letting $|\Delta^n|$ denote the standard n-simplex in $\mathbf{Top}$,

$$|\Delta^n| = \left\{ (x_0, \ldots, x_n) \in \mathbb{R}^{n+1} \,\middle|\, \sum_{i=0}^{n} x_i = 1,\ x_i \geq 0 \right\} \subseteq \mathbb{R}^{n+1},$$

there is a natural covariant functor $\mathbf{\Delta} \to \mathbf{Top}$ given by $[n] \mapsto |\Delta^n|$. A map $f\colon [n] \to [m]$ induces a map $f_*\colon\ |\Delta^n| \to |\Delta^m|$ given by $(x_0, \dots, x_n) \mapsto (y_0, \dots, y_n)$ where

$$y_i = \begin{cases} 0 & f^{-1}(i) = \emptyset; \\ \sum_{j \in f^{-1}(i)} x_i & \text{otherwise.} \end{cases}$$

Thus the i^{th} coface map inserts a 0 in the i^{th} coordinate and the j^{th} codegeneracy map adds the x_j and x_{j+1} coordinates. Geometrically, the former inserts $|\Delta^{n-1}|$ as the i^{th} face of $|\Delta^n|$ and the latter projects $|\Delta^{n+1}|$ onto the topological n-simplex orthogonal to its j^{th} face.

Given a topological space Y, the *total singular complex* (or *singular set*) is the simplicial set SY given by $[n] \mapsto \mathbf{Top}(|\Delta^n|, Y)$. Elements of SY_n are the singular n-simplices of Y familiar to algebraic topologists. The face and degeneracy maps are given by pre-composition by d^i and s^j. This functor S is essential to the definition of the singular homology of the space Y.

One might be confused by the notation $|\Delta^n|$ for the topological simplex, since this object is more commonly denoted by Δ^n or Δ_n. We have chosen this notation in order to distinguish the standard n-simplex *as a simplicial set* from the standard n-simplex *as a topological space*. We implore the reader who remains unconvinced to reassess after we define the geometric realization functor (Definition 3.1.13); the observation

$$|\Delta^n| = |\Delta^n|,$$

while both amusing and infuriating, is perhaps also reassuring that our notation is well-chosen.[3]

3.1.2 Classifying Spaces

The most crucial example of a simplicial set, for our purposes, is that of the nerve of a category (which is sometimes itself called the classifying space, as in [GJ99]). Ultimately, we wish to understand the nerve of a category internal to **Top** as a simplicial space, rather than a simplicial set, but first we lay out the more traditional definitions.

Definition 3.1.12. The *nerve* of $\mathscr{C}$ is the simplicial set $N\mathscr{C}$ given by

$$N\mathscr{C}_n = \mathbf{Cat}([n], \mathscr{C}),$$

where **Cat** denotes the category of small categories and functors between them.

In other words, an n-simplex is a string of n composable arrows in $\mathscr{C}$. Thus $N\mathscr{C}_0 = \mathrm{Ob}\,\mathscr{C}$, $N\mathscr{C}_1 = \mathrm{Mor}\,\mathscr{C}$, and $N\mathscr{C}_n = \{\text{strings of } n \text{ composable arrows in } \mathscr{C}\}$ more generally. In this sense, we can think of the n-simplices of the nerve as the commutative diagrams in the category that "look like" a standard n-simplex.

[3] The expression on the left side of the equality is the geometric realization of the standard n-simplex as a simplicial set, and the expression on the right is the topological n-simplex.

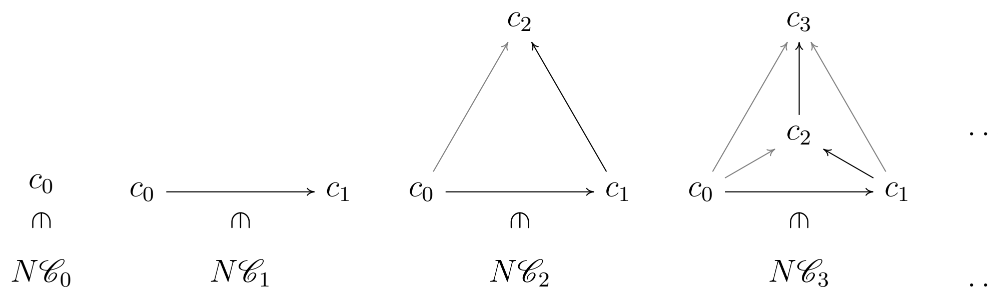

Figure 3.1: Some n-simplices of the nerve. The black arrows indicate the string of n morphisms, while the gray arrows are the morphisms induced by composition.

Given a string of n composable arrows,

$$c_0 \to c_1 \to \cdots \to c_{i-1} \to c_i \to c_{i+1} \to \cdots \to c_n,$$

the face map $d_i \colon N\mathscr{C}_n \to N\mathscr{C}_{n-1}$ returns the string of $n-1$ composable arrows

$$c_0 \to c_1 \to \cdots \to c_{i-1} \longrightarrow c_{i+1} \to \cdots \to c_n,$$

where the arrow $c_{i-1} \to c_{i+1}$ is the composition of the i^{th} and $(i+1)^{th}$ arrows. In the cases that $i = 0, n$, we instead omit that i^{th} arrow. Similarly, the degeneracy map $s_j \colon N\mathscr{C}_n \to N\mathscr{C}_{n+1}$ returns the string of $n+1$ composable arrows

$$c_0 \to c_1 \to \cdots \to c_{j-1} \to c_j \xrightarrow{\text{id}} c_j \to c_{j+1} \to \cdots \to c_n.$$

From the definition above, we can see that the nerve is a functor $N \colon \mathbf{Cat} \to \mathbf{sSet}$. Given a map $\mathscr{C} \to \mathscr{D}$, the induced map $N\mathscr{C} \to N\mathscr{D}$ is given by post-composition. This functor can also be seen as the right adjoint of an embedding $\mathbf{\Delta} \hookrightarrow \mathbf{Cat}$, and for more details on this perspective, we refer the reader to [Rie08, §4].

To build a topological space out of a simplicial set, we use the geometric realization functor $|-| : \mathbf{sSet} \to \mathbf{Top}$.[4] This space is constructed by interpreting each n-simplex as a copy of $|\Delta_n|$ and using information from the face and degeneracy maps to obtain gluing and collapsing instructions. The resulting space is fairly well-behaved, and the geometric realization of any simplicial set is a CW complex (cf. [GJ99, Proposition I.2.3]).

Definition 3.1.13. The *geometric realization* of a simplicial set X is

$$|X| = \operatorname{colim} \left(\coprod_{f \colon [n] \to [m]} X_m \times |\Delta^n| \underset{f^*}{\overset{f_*}{\rightrightarrows}} \coprod_{[n]} X_n \times |\Delta^n| \right)$$

[4]Technically, the geometric realization lands in **CGHaus**, the category of compactly generated Hausdorff spaces, as discussed in [GJ99, §I.2].

The map $f_*\colon X_m \times |\Delta^n| \to X_m \times |\Delta^m|$ includes the faces as described in Example 3.1.11, and the map $f^*\colon X_m \times |\Delta^n| \to X_n \times |\Delta^n|$ collapses degeneracies via the simplicial structure $Xf\colon X_m \to X_n$. In practice, it is often useful to use a more concrete description of the geometric realization, which is given by

$$\left(\coprod_{n\geq 0} X_n \times |\Delta^n|\right) / \sim,$$

where $(x, s^j y) \sim (s_j x, y)$ and $(x, d^i y) \sim (d_i x, y)$. Here, the equivalence relation does the work of the commutativity of the cone diagram. The first relation ensures that the degeneracies are "glued in" in a compatible way, and the second relation does the same for the faces.

Finally, composing the nerve and geometric realization functors, we get the classifying space of a category. This trick, first due to Segal [Seg68], generalizes the idea of classifying spaces for topological groups (see [May99, §16.5]).

Definition 3.1.14. The *classifying space* of a small category $\mathscr{C}$ is $B\mathscr{C} = |N\mathscr{C}|$.

In the context of a topological category, the nerve is a simplicial space rather than a simplicial set. Recall that a category $\mathscr{C}$ internal to **Top** consists of a space of objects $\mathscr{C}_0$ and a space of morphisms $\mathscr{C}_1$. The n^{th} level of the nerve is given inductively by the iterated pullback with n factors,

$$N\mathscr{C}_n = \mathscr{C}_1 \times_{\mathscr{C}_0} \mathscr{C}_1 \times_{\mathscr{C}_0} \cdots \times_{\mathscr{C}_0} \mathscr{C}_1,$$

where each $\mathscr{C}_1 \times_{\mathscr{C}_0} \mathscr{C}_1$ is the limit of $\mathscr{C}_1 \xrightarrow{\text{cod}} \mathscr{C}_0 \xleftarrow{\text{dom}} \mathscr{C}_1$. As before, the 0^{th} level of the nerve is the objects, $N\mathscr{C}_0 = \mathscr{C}_0$. The degeneracy map s_j for $N\mathscr{C}_n$ inserts the appropriate identity map in the j^{th} position, and d_i omits the outermost morphism when $i = 0, n$ and otherwise composes the i^{th} and $(i+1)^{th}$ morphisms. The definition of geometric realization does not change for a topological category, except now we must also take the topology on the X_n into consideration. We would hope that the resulting classifying space map would retain its functorial properties, and indeed we may rest easy thanks to the following result from [Seg68].

Theorem 3.1.15. *A continuous functor $F\colon \mathscr{C} \to \mathscr{D}$ induces a continuous map $BF\colon B\mathscr{C} \to B\mathscr{D}$, hence $B\colon$ **TopCat** $\to$ **Top** is a functor.*

To conclude this section, we present a few more results about classifying spaces for topological categories. We point the interested reader to Segal's original paper [Seg68] for proofs, or [Lai13] for more in-depth treatments.

Proposition 3.1.16. *Given topological categories $\mathscr{C}$, $\mathscr{D}$, at least one of which has finitely many objects, there is a homeomorphism $B(\mathscr{C} \times \mathscr{D}) \cong B\mathscr{C} \times B\mathscr{D}$.*

Applying this proposition to $\mathscr{D} = [1]$ (the poset category) and pondering Remark 2.2.5, we get the following theorem.

Theorem 3.1.17. *Let $F, G\colon \mathscr{C} \to \mathscr{D}$ be continuous functors with a continuous natural transformation $\eta\colon F \to G$. Then the induced maps $BF, BG\colon B\mathscr{C} \to B\mathscr{D}$ are homotopic.*

Remark 3.1.18. Note that this theorem implies that if $\mathscr{C}$ has an initial object, then the classifying space $B\mathscr{C}$ is contractible, induced by the continuous natural transformation between $\mathrm{id}_{\mathscr{C}}$ and the constant functor on the initial object (cf. Example A.3.7).

Another corollary of the theorem is that equivalent topological categories will have homotopy equivalent classifying spaces. A primary concern of this thesis is understanding when maps between topological categories induce homotopy equivalences on the classifying spaces of those categories, but the requirements of Theorem 3.1.17 are rather stringent. Later work, in particular Section 3.2.3, will establish slightly weaker conditions where we can reach the same conclusion.

3.1.3 Edgewise Subdivision

The category $\mathbf{\Delta}$ has a (non-symmetric) monoidal structure via the join $\star$. Given linearly ordered sets I, J, their *join* $I \star J$ is the set $I \amalg J$ with the original orderings on I and J, along with the additional condition that $i < j$ for all $i \in I$, $j \in J$. For example, we can think of $[n] \star [m]$ as

$$\overline{0} < \overline{1} < \cdots < \overline{n} < 0 < 1 < \cdots < m,$$

where the overline is merely meant to distinguish between the elements of $[n]$ and those of $[m]$. Now, let $\varepsilon\colon \mathbf{\Delta} \to \mathbf{\Delta}$ be given by $\mathrm{op} \star \mathrm{id}$, so $\varepsilon([n]) = [n]^{\mathrm{op}} \star [n] \cong [2n+1]$, where $[n]^{\mathrm{op}}$ is meant to indicate $[n]$ with reversed ordering. Hence we can think of $\varepsilon([n])$ as

$$\overline{n} < \overline{n-1} < \cdots < \overline{1} < \overline{0} < 0 < 1 < \cdots < n-1 < n.$$

Definition 3.1.19. Given a simplicial set X, the *edgewise subdivision* of X is the simplicial set $\mathrm{sd}(X) = X \circ \varepsilon$, with each component $\mathrm{sd}(X)_n \cong X_{2n+1}$. The vertices of $\mathrm{sd}(X)$ are the edges of X, and an edge of $\mathrm{sd}(X)$ from $a \to b$ to $c \to d$ can be viewed as a commutative diagram,

$$\begin{array}{ccc} a & \longleftarrow & c \\ \downarrow & & \downarrow \\ b & \longrightarrow & d \end{array}.$$

The edgewise subdivision is thus a functor $\mathrm{sd}\colon \mathbf{sSet} \to \mathbf{sSet}$ specified by $\mathrm{sd}(X)_n = X_{2n+1}$ and the structure maps $\mathrm{sd}(d_i) = d_{n-i} \circ d_{n+1+i}\colon \mathrm{sd}(X)_n \to \mathrm{sd}(X)_{n-1}$ and $\mathrm{sd}(s_j) = s_{n-j} \circ s_{n+1+j}\colon \mathrm{sd}(X)_n \to \mathrm{sd}(X)_{n+1}$.

Example 3.1.20. The edgewise subdivision of a standard k-simplex $\mathrm{sd}(\Delta^k)$ will divide Δ^k into 2^k non-degenerate k-simplicies. For $k = 2, 3$, we get the pictures in Fig. 3.2. A quick check verifies that we do indeed have the correct number of simplicies.

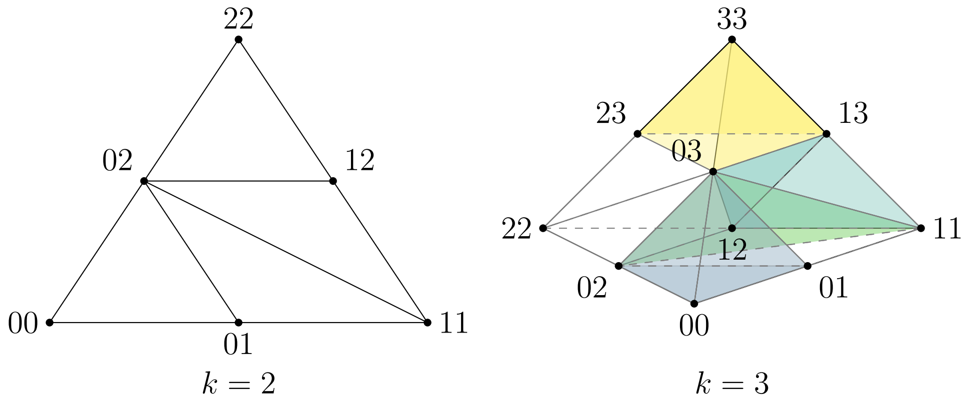

Figure 3.2: The edgewise subdivision of the standard k-simplex, for $k = 2, 3$. A collection of vertices $\{v_{i_0,j_0}, \ldots, v_{i_{k+1},j_{k+1}}\}$ determines a simplex when $i_0 \geq \cdots \geq i_{k+1}$ and $j_0 \leq \cdots \leq j_{k+1}$ (and each of the numbers $0, \ldots, k$ appears at least once).

As the illustration above indicates, there is a homeomorphism between the geometric realization of a simplicial set and its edgewise subdivision. This result is due to Segal in [Seg73, Appendix 1], and a more in-depth proof is given in [Ber09, §5.3].

Theorem 3.1.21 (Segal)**.** *For any simplicial space X, we have $|X| \cong |\mathrm{sd}(X)|$.*

In the case where X is the nerve of some category, the edgewise subdivision is related to a construction known as the twisted arrow category.

Definition 3.1.22. Given a small category $\mathscr{C}$, the *twisted arrow category* of $\mathscr{C}$, denoted $\mathrm{tw}(\mathscr{C})$, the category whose objects are morphisms of $\mathscr{C}$, written vertically, and whose morphisms are commutative squares. That is, a morphism from $a \to b$ to $c \to d$ is given by

$$\begin{array}{ccc} a & \longleftarrow & c \\ \downarrow & & \downarrow \\ b & \longrightarrow & d \end{array}$$

For example, the twisted arrow category of the poset category $[n]$ has objects (i, j) for every $0 \leq i \leq j \leq n$, and morphisms $(i \to j) \leftrightarrows (i' \to j')$ whenever $i \geq i'$ and $j \leq j'$.

This definition strongly resembles the edgewise subdivision of a simplicial set, except in this case we consider a category. To link the two concepts together, we turn a category into a simplicial set via the nerve functor.

Proposition 3.1.23. *If $\mathscr{C}$ is a small category, then $\mathrm{sd}(N\mathscr{C}) \cong N\mathrm{tw}(\mathscr{C})$.*

Proof. Starting small, we can see that $\mathrm{sd}(N\mathscr{C})_0 \cong N\mathscr{C}_1 = \mathrm{Mor}\,\mathscr{C} = \mathrm{Ob}\,\mathrm{tw}(\mathscr{C}) = N\mathrm{tw}(\mathscr{C})_0$. Similarly,

$$\mathrm{sd}(N\mathscr{C})_1 \cong N\mathscr{C}_3 = \{\cdot \to \cdot \to \cdot \to \cdot\} = \left\{ \begin{array}{ccc} \cdot & \longleftarrow & \cdot \\ \downarrow & & \downarrow \\ \cdot & \longrightarrow & \cdot \end{array} \right\} = N\mathrm{tw}(\mathscr{C})_1$$

and so on. It suffices to show that $N\mathrm{tw}(\mathscr{C})_n \cong N\mathscr{C}_{2n+1}$. An element of $N\mathscr{C}_{2n+1}$ looks like a diagram of $2n+1$ composable morphisms between $2n$ objects of $\mathscr{C}$

$$\begin{array}{ccccccccc} \overline{0} & \longleftarrow & \overline{1} & \longleftarrow & \dots & \longleftarrow & \overline{n-1} & \longleftarrow & \overline{n} \\ \big\downarrow & & & & & & & & \\ 0 & \longrightarrow & 1 & \longrightarrow & \dots & \longrightarrow & n-1 & \longrightarrow & n \end{array}$$

but, composing arrows, this is the same as the diagram

$$\begin{array}{ccccccccc} \overline{0} & \longleftarrow & \overline{1} & \longleftarrow & \dots & \longleftarrow & \overline{n-1} & \longleftarrow & \overline{n} \\ \big\downarrow & & \big\downarrow & & & & \big\downarrow & & \big\downarrow \\ 0 & \longrightarrow & 1 & \longrightarrow & \dots & \longrightarrow & n-1 & \longrightarrow & n \end{array}$$

which is just an element of $N\mathrm{tw}(\mathscr{C})_n$. A quick inspection shows that this correspondence is compatible with the face and degeneracy maps as well. □

Thus we get a commutative diagram of functors,

$$\begin{array}{ccc} \mathrm{Cat} & \xrightarrow{\mathrm{tw}} & \mathrm{Cat} \\ {\scriptstyle N}\big\downarrow & & \big\downarrow{\scriptstyle N} \\ \mathrm{sSet} & \xrightarrow[\mathrm{sd}]{} & \mathrm{sSet} \end{array}$$

Stringing together Theorem 3.1.21 and Proposition 3.1.23 yields the following corollary.

Corollary 3.1.24. *$B\mathscr{C} \cong B\mathrm{tw}(\mathscr{C})$.*

Note that when $\mathscr{C}$ is a topological category, $\mathrm{tw}(\mathscr{C})$ inherits topological structure as well. The objects of $\mathrm{tw}(\mathscr{C})$ are the space $\mathscr{C}_1$ and the morphisms are homeomorphic to the iterated pullback $\mathscr{C}_1 \times_{\mathscr{C}_0} \mathscr{C}_1 \times_{\mathscr{C}_0} \mathscr{C}_1$, as explained above in the proof of Proposition 3.1.23.

In particular, when considering the flow category $\mathscr{C}_f$, the twisted arrow category records the ways we can "break up" a broken flow line. The objects are all broken flow lines, and there is a morphism $\gamma \to \gamma'$ precisely when the image of γ is contained in the image of γ'; that is, there are other broken flow lines α, β such that $\gamma' = \beta \circ \gamma \circ \alpha$. It is straightforward to show that such a decomposition, if it exists, is unique (up to composing with the identities on critical points). Consequently, if there is a morphism $\gamma \to \gamma'$, then that morphism is unique.

3.2 Homotopy Invariance

This section is concerned with establishing conditions between topological categories that induce a homotopy equivalence on their classifying spaces. After recalling the basics of fibrations and cofibrations (available in classics such as [Hat02, May99]), we establish some crucial properties of the nerve using homotopy pullbacks, as described in [BK72, May72, Seg74]. The author also enjoyed reading selections from [Dug08] for an additional, more modern perspective.

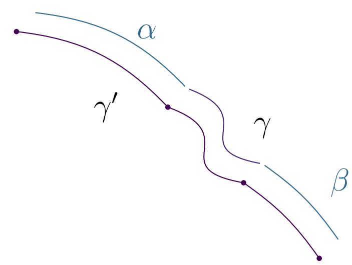

Figure 3.3: Decomposing flow lines. There is a morphism $\gamma \to \gamma'$ in $\mathrm{tw}(\mathscr{C}_f)$ if and only if $\gamma' = \alpha \circ \gamma \circ \beta$ for some $\alpha, \beta \in \mathscr{C}_f$.

3.2.1 Fibrations and Cofibrations

Recall (or see Appendix A.1) that a fiber bundle is a structure $F \to E \xrightarrow{\pi} B$ where π locally behaves like the projection map from $F \times B$. Fibrations essentially behave like fiber bundles from the point of view of homotopy theory, where now the fibers may merely be homotopy equivalent rather than homeomorphic (when B is connected).

Definition 3.2.1. A map $p\colon E \to B$ is said to have the *homotopy lifting property* with respect to a space X if for every homotopy $h_t\colon X \to B$ and map $\tilde{h}_0\colon X \to E$ such that $p\tilde{h}_0 = h_0$, there exists a homotopy $\tilde{h}_t\colon X \to E$ such that $p\tilde{h}_t = h_t$. The homotopy $\tilde{h}_t$ is said to *lift* h_t. A map is a *(Hurewicz) fibration* if it has the homotopy lifting property with respect to all spaces.

Fibrations are particularly well-known for their relation to the long exact sequence of homotopy groups (see, for instance, [May99, §9.3]), which in turn yields a nice test for homotopy equivalence between well-behaved spaces.

Theorem 3.2.2. *A fibration between CW complexes with contractible fibers is a homotopy equivalence.*

Idea of proof. Let $p\colon E \to B$ be a fibration such that each fibre F_b is contractible. By the long exact sequence of homotopy groups, we have

$$\cdots \to \pi_n(F) \to \pi_n(E) \to \pi_n(B) \to \pi_{n-1}(F) \to \cdots \to \pi_0(F) \to \pi_0(E).$$

But each fiber is contractible, so $\pi_k(F) = 0$ for all $k \geq 0$. Thus we get a series of very short exact sequences

$$0 \to \pi_k(E) \to \pi_k(B) \to 0$$

which tells us that the homotopy groups are isomorphic, hence we have a weak homotopy equivalence. If E and B are CW complexes, then Whitehead's theorem (cf. [Hat02, Theorem 4.5]) gives the desired result. □

Dualizing fibrations, we arrive at the notion of cofibrations. Cofibrations can be thought of as a "nice" sort of inclusion, where the subspace has some room to wiggle.

Definition 3.2.3. The dual notion of the homotopy lifting property is called the *homotopy extension property*, where now a map $i\colon B \to E$ has this property if for every homotopy $h_t\colon B \to X$ and map $\tilde{h}_0\colon E \to X$ with $\tilde{h}_0 i = h_0$, there is a homotopy $\tilde{h}_t\colon E \to X$ such that $\tilde{h}_t i = h_t$. A map is a *(Hurewicz) cofibration* if it has the homotopy extension property with respect to all spaces.

In practice, it may be easier to recognize cofibrations using different criteria. While there are multiple equivalent conditions, given in [May99, §6.4], we will primarily use the following one.

Proposition 3.2.4. *An inclusion $B \hookrightarrow E$ is a cofibration precisely when $B \times I \cup E \times \{0\}$ is a retract of $E \times I$.*

The retract requirement means there is a map $r\colon E \times I \to B \times I \cup E \times \{0\}$ with a section $i\colon B \times I \cup E \times \{0\} \to E \times I$. The map r is called a *retraction*.

Fibrations and cofibrations feature heavily in the study of homotopy, (co)homology, and model categories. We have presented only the most basic definitions and properties that will be crucial in the following subsections, but we encourage the interested reader to pursue these topics in [Hat02], [May99], and others.

3.2.2 Homotopy Pullbacks

This subsection is concerned with developing sufficient conditions on topological categories that will yield a homotopy equivalence on the nerves. Since the nerve $N\mathscr{C}$ is the simplicial space with $N\mathscr{C}_n = \mathscr{C}_1 \times_{\mathscr{C}_0} \cdots \times_{\mathscr{C}_0} \mathscr{C}_1$, an n-fold iterated pullback, it makes sense that we should first seek to understand the homotopy theory surrounding pullbacks.

We are familiar with the pullback $X \times_Z Y$ as the limit of $X \xrightarrow{f} Z \xleftarrow{g} Y$, however, these strict pullbacks do not necessarily preserve homotopy equivalences. To address this issue, we look to the homotopy pullback. Homotopy pullbacks are an example of homotopy limits as described in [BK72, §XI.3], and are dual to the homotopy pushout squares discussed in [Lur09, §A.2.4].

Definition 3.2.5. A *homotopy pullback* consists of a diagram

$$\begin{array}{ccc} W & \longrightarrow & Y \\ \downarrow & & \downarrow{\scriptstyle g} \\ X & \xrightarrow[f]{} & Z \end{array}$$

which commutes up to homotopy, and such that for any other diagram

$$\begin{array}{ccc} \tilde{W} & \longrightarrow & Y \\ \downarrow & & \downarrow{\scriptstyle g} \\ X & \xrightarrow[f]{} & Z \end{array}$$

which commutes up to homotopy, there is a morphism $\tilde{W} \to W$ (unique up to homotopy) such that the triangles in the following diagram commute up to homotopy:

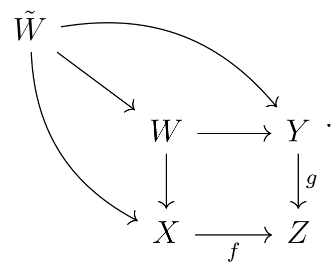

A concrete example of a homotopy pullback is given by

$$X \times^h_Z Y := X \times_Z Z^I \times_Z Y = \{(x, \gamma, y) \mid f(x) = \gamma(0), \gamma(1) = g(y)\}$$

(see [Dug08, Example 5.1]). We can think of an element of the homotopy pullback as a point $x \in X$, a point $y \in Y$, and a path γ connecting $f(x)$ and $g(y)$ in Z.

This construction preserves homotopy equivalence in the sense that if we have a map of diagrams that is an objectwise-homotopy equivalence

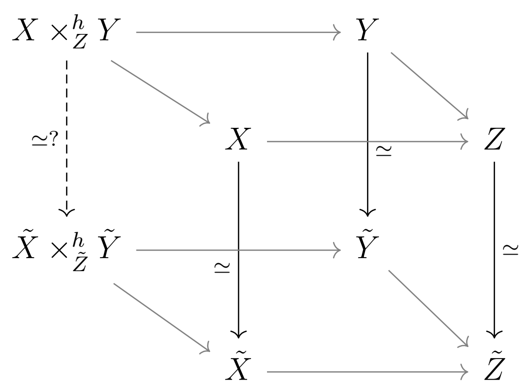

then the induced map in question is also a homotopy equivalence, meaning $X \times^h_Z Y \simeq \tilde{X} \times^h_{\tilde{Z}} \tilde{Y}$ completes the cube above.

Of course, we are interested in filling out the cube for $N\mathscr{C}_n \xrightarrow{\text{cod}} \mathscr{C}_0 \xleftarrow{\text{dom}} \mathscr{C}_1$ and $N\mathscr{D}_n \xrightarrow{\text{cod}} \mathscr{C}_0 \xleftarrow{\text{dom}} \mathscr{D}_1$, which we thus far understand as strict pullbacks.[5] The following proposition establishes key conditions under which this strict pullback is in fact a homotopy pullback.

Proposition 3.2.6. *If f or g is a fibration, then the strict pullback diagram is a homotopy pullback.*

Thus, for our purposes, it suffices to show that either the dom (source) or cod (target) map is a fibration. The argument we give below only considers dom, but the argument for cod is almost identical.

[5]Technically, we mean $N\mathscr{C}_n \xrightarrow{\text{cod}} \mathscr{C}_0$ to be $N\mathscr{C}_n \xrightarrow{\circ} \mathscr{C}_1 \xrightarrow{\text{cod}} \mathscr{C}_0$ where $\circ$ denotes the composition of the n morphisms from $N\mathscr{C}_n$. We will omit this finer detail in the following discussion for the sake of simplicity.

Suppose we have a homotopy $h_t\colon X \to \mathscr{C}_0$ and a map $\tilde{h}_0\colon X \to \mathscr{C}_1$ such that $\mathrm{dom}\tilde{h}_0 = h_0$. Then, since the identity map $i\colon \mathscr{C}_1 \to \mathscr{C}_0$ is a section of the source map, we get a homotopy $\tilde{h}_t\colon X \to \mathscr{C}_1$ given by $\tilde{h}_t = ih_t$. Diagrammatically,

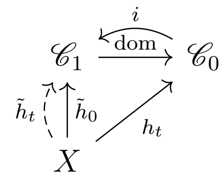

Clearly $\tilde{h}_t$ is continuous as a composition of two continuous maps, and moreover $\mathrm{dom}\tilde{h}_t = h_t$ by construction. Hence the source map is a fibration, and so the following proposition is immediate.

Theorem 3.2.7. *Suppose $F\colon \mathscr{C} \to \mathscr{D}$ is a continuous functor. If $F_0\colon \mathscr{C}_0 \simeq \mathscr{D}_0$ and $F_1\colon \mathscr{C}_1 \simeq \mathscr{D}_1$, then $NF\colon N\mathscr{C} \to N\mathscr{D}$ is a map of simplicial spaces that is a levelwise homotopy equivalence, meaning that $NF_n\colon N\mathscr{C}_n \simeq N\mathscr{D}_n$ for all $n \geq 0$.*

Proof. Since $N\mathscr{C}_n$ is the pullback over $N\mathscr{C}_{n-1} \xrightarrow{\mathrm{cod}} \mathscr{C}_0 \xleftarrow{\mathrm{dom}} \mathscr{C}_1$ for $n \geq 2$, the inductive claim follows by the argument above. We know NF is a map of simplicial spaces since $N\colon \mathbf{TopCat} \to \mathbf{sTop}$ is a functor. □

3.2.3 Inducing Equivalence on Classifying Spaces

Recall that a map of simplicial spaces $X \to Y$ is a continuous natural transformation of functors. We call such a map an *levelwise homotopy equivalence* if each $X_n \to Y_n$ is a homotopy equivalence. Unfortunately, it is not always the case that an objectwise homotopy equivalence $X \to Y$ yields a homotopy equivalence $|X| \to |Y|$. We discuss certain conditions where this *is* the case, as developed by [May72, Chapter 11] and [Seg74, Appendix A].

Definition 3.2.8. A simplicial space X is *good* if every degeneracy map $s_j\colon X_n \to X_{n+1}$ is a closed cofibration, and is *Reedy cofibrant* (or *proper*) if every latching map $L_nX \hookrightarrow X$ is a cofibration, where

$$L_nX = \bigcup_{i=0}^{n-1} s_i(X_{n-1})$$

is the n^{th} *latching object.*

We can think of $L_nX \subseteq X_n$ as the set of degenerate n-simplices, which gives a natural inclusion $L_nX \hookrightarrow X_n$ (this is the n^{th} latching map in the definition above). It is well known that every good simplicial space is Reedy cofibrant (this follows from the fact that cofibrations are preserved under pushouts— namely, unions). In either case, we ensure that geometric realization preserves homotopy equivalence.

Theorem 3.2.9. *Let $f\colon X \to Y$ be a map of Reedy cofibrant simplicial spaces. If $f_n\colon X_n \to Y_n$ is a homotopy equivalence for all n, then $|f|\colon |X| \to |Y|$ is a homotopy equivalence.*

Our statement of this result follows [May72, Theorem 11.13], although the reader familiar with model categories may be interested in the approach of [RV14] (see their Corollary 10.6 for a statement of the theorem above).

Having established necessary conditions on topological categories to induce a homotopy equivalence on the nerves in Theorem 3.2.7, we hope that geometric realization preserves such a homotopy. The following lemma describes a sufficient condition.

Lemma 3.2.10. *If the image of the identity map $i\colon \mathscr{C}_0 \to \mathscr{C}_1$ is closed, then the nerve is a good simplicial space.*

Proof. Consider $s_j\colon N\mathscr{C}_n \to N\mathscr{C}_{n+1}$ for some n and some $0 \le j \le n$. Recall that this map sends a string of n composable morphisms

$$c_0 \to c_1 \to \cdots \to c_{j-1} \to c_j \to c_{j+1} \to \cdots \to c_n$$

to the string of $n+1$ composable morphisms

$$c_0 \to c_1 \to \cdots \to c_{j-1} \to c_j \xrightarrow{\mathrm{id}_{c_j}} c_j \to c_{j+1} \to \cdots \to c_n.$$

Thus the image of s_j is

$$\mathscr{C}_1 \times_{\mathscr{C}_0} \cdots \times_{\mathscr{C}_0} \{\mathrm{id}_c \mid c \in \mathscr{C}_0\} \times_{\mathscr{C}_0} \cdots \times_{\mathscr{C}_0} \mathscr{C}_1 \subseteq N\mathscr{C}_{n+1}$$

where the funny business happens in the $(j+1)^{th}$ factor. Thus $s_j(N\mathscr{C}_n)$ and $N\mathscr{C}_{n+1}$ agree as products on all but one of the factors, namely the $(j+1)^{th}$. So to conclude that $s_j(N\mathscr{C}_n)$ is closed, it suffices to check that

$$\{\mathrm{id}_c \mid c \in \mathscr{C}_0\} \subseteq \mathscr{C}_1$$

is a closed subspace, which follows by our assumption.

To see that s_j is a cofibration, we observe that $\mathrm{im}(s_j) \times I \cup N\mathscr{C}_{n+1} \times \{0\}$ is a retract of $N\mathscr{C}_{n+1} \times I$, since the j^{th} face map d_j is a section of s_j. □

Remark 3.2.11. Suppose we have $\mathscr{C}_1 = \coprod_{c,d\in\mathscr{C}_0} \mathscr{C}(c,d)$ (as holds, for example, for the flow category). Then

$$\{\mathrm{id}_c \mid c \in \mathscr{C}_0\} = \coprod_{c\in\mathscr{C}_0} \mathscr{C}(c,c)$$

is a disjoint union of homspaces, and each homspace is closed in $\mathscr{C}_1$, implying that the nerve is good.

It is well-known that an equivalence of categories induces a homotopy equivalence on their classifying spaces. By combining Lemma 3.2.10 and Theorem 3.2.12, we see that it is sufficient in some cases to merely require that the spaces of objects and morphisms are homotopy equivalent.

Theorem 3.2.12. *Suppose that $F\colon \mathscr{C} \to \mathscr{D}$ is a continuous functor between topological categories such that $F_0\colon \mathscr{C}_0 \simeq \mathscr{D}_0$ and $F_1\colon \mathscr{C}_1 \simeq \mathscr{D}_1$. Then, if the images of the identity maps i of $\mathscr{C}$ and $\mathscr{D}$ are closed, the classifying spaces are homotopy equivalent. That is, the map $|NF| : B\mathscr{C} \to B\mathscr{D}$ is a homotopy equivalence.*

Proof. By Lemma 3.2.10, the nerve is a good simplicial space under the given conditions, and hence Reedy cofibrant. Thus Theorem 3.2.9 applies, yielding a homotopy equivalence on the geometric realization of the nerves. □

Chapter 4

The Cohen-Jones-Segal Theorem

The goal of this thesis is to understand and prove the main results appearing in [CJS95b], and the work of the previous chapters gives us the language to do so. These results relate the flow category of a Morse function to the underlying manifold via its classifying space.

Theorem 4.0.1. *Let f be a Morse function on a closed Riemannian manifold (M, g). Then*

(1) there is a homotopy equivalence $M \simeq B\mathscr{C}_f$,

(2) if (f, g) is Morse-Smale, then there is a homeomorphism $M \cong B\mathscr{C}_f$.

Although widely cited, the original preprint by Cohen, Jones, and Segal was never published. As explained by Cohen, quoted below, the key piece of missing information was the "folk theorem" that (in the Morse-Smale case) the compactified moduli spaces are manifolds with corners and that the gluing maps between these spaces are associative.

> The fact that [CJS95b] was never submitted for publication was due to the fact that the "folk theorem" mentioned above, as well as the associativity of gluing, both of which the authors of [CJS95b] assumed were "well known to the experts," were indeed not in the literature, and their proofs which were eventually provided in [Qin11], were analytically more complicated than the authors imagined. [Coh19, p.16]

In any case, the efforts of [Qin11] and later [Weh12] remedy the issue, and the proof of part (2) of the theorem can now be completed using these results. Indeed, such a proof is supplied in [CIN06, §12.2], and will be discussed in Section 4.2.2.

The work in Section 4.2.1 provides a fix to another, more minor error in the original proof of part (1). In [CJS95b, §6], the authors produce two functors, $\Theta\colon \tilde{\mathscr{C}}_f \leftrightarrows \mathscr{M} : \Gamma$ (where $\tilde{\mathscr{C}}_f$ and $\mathscr{M}$ are variants on the flow category and the manifold M, respectively) whose induced maps $B\Theta$, $B\Gamma$ are inverse homotopy equivalences. Unfortunately, Γ is not a *continuous* functor, and this failure of continuity essentially boils down to the fact that the "assignment" map $p \to \gamma_p$ is not continuous.[1] However, if we expand the flow category to a slightly larger category (Definition 4.2.1), then we can extend Γ to be a continuous local section of Θ, and the desired result follows. This adjustment was suggested in an email correspondence between Cohen and Segal, which was generously shared with the author by Cohen.

Before getting into the nitty-gritty aspects of the proofs, we provide some simple examples to see the theorem in action. While these examples have appeared in other places (see [CIN06, HD10, Rot10] and even the original [CJS95b]), we hope that these examples will help the reader gain some intuition for how everything pieces together.

4.1 Examples

We present a series of examples that illustrate Theorem 4.0.1, many of which are continuations of examples from previous chapters. In particular, we will examine the (Morse-Smale) height function on the sphere as well as few different functions on the torus. The reader is also invited to revisit the alternate sphere, whose classifying space we illustrated in the Introduction.

4.1.1 The Sphere

Continuing Example 1.1.5, Example 1.2.2, and Example 1.3.10, we consider the n-sphere S^n embedded in $\mathbb{R}^{n+1}$ under the height function $f(x_0, \ldots, x_n) = x_n$. This function has two critical points— namely the north and south poles, N and S— and under the standard flat metric in $\mathbb{R}^{n+1}$ restricted to S^n, the pair $(f, -\nabla f)$ is Morse-Smale. Note that $\mathcal{M}(N,S) = W(N,S)/\mathbb{R} \cong S^{n-1}$, and moreover $\overline{\mathcal{M}}(N,S) = \mathcal{M}(N,S)$ since there are only the two critical points.

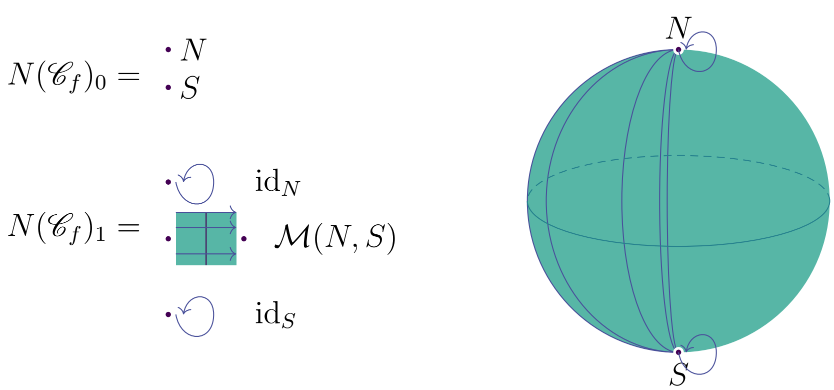

Figure 4.1: The classifying space of S^n (when $n = 2$). The non-degenerate simplices of the nerve are displayed on the left and are "glued in" on the right. The identity maps are sent to their respective critical points and the non-trivial morphisms $\mathcal{M}(N,S)$ are glued in around the equator.

[1] For instance, if some $\mathscr{C}_f(a,b)$ (which is open in $\operatorname{Mor}\mathscr{C}_f$) contains unbroken flow lines, then its preimage under this assignment map is $W(a,b)$, which is not necessarily open in M.

The flow category has the two objects N, S and morphisms $\mathcal{M}(N,S)$ as well as the identities on N and S. Hence the only non-degenerate simplices in the nerve are the 0-simplices N, S and the 1-simplices that are the elements of $\mathcal{M}(N,S)$. Thus the classifying space looks like

$$B\mathscr{C}_f \cong (\{N,S\} \times \{*\} \amalg \mathcal{M}(N,S) \times I) \,/ \sim$$

with the gluing instructions $(N,*) \sim (\varphi, 0)$ and $(S,*) \sim (\varphi, 1)$. (We also glue the identities $(N,*) \sim (t, \mathrm{id}_N)$ and $(S,*) \sim (t, \mathrm{id}_S)$ for all $t \in I$.) This describes $B\mathscr{C}_f$ as a suspension[2] of $\mathcal{M}(N,S) \cong S^{n-1}$, hence $B\mathscr{C}_f \cong S^n$.

4.1.2 The Vertical Torus

The torus is often the prototypical example used to illustrate the ideas of Morse theory, as it is a topologically interesting space that is simple enough to visualize. Recall from Example 1.2.13 that we view the vertical torus as embedded in ordinary three-space, standing on one end. The function f is the ordinary height function.

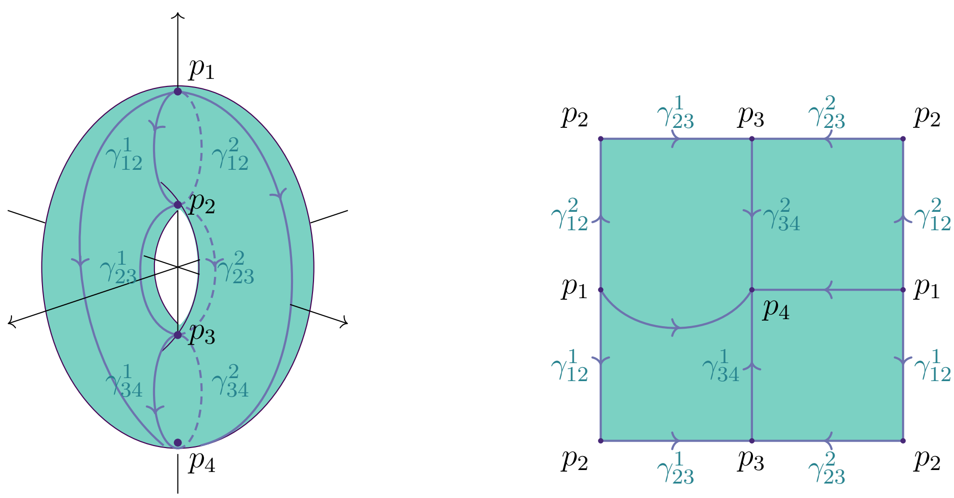

Figure 4.2: Flow lines on the vertical torus, seen as the quotient $\mathbb{R}^2/\mathbb{Z}^2$. On the left we show the vertical torus embedded in $\mathbb{R}^3$. There are four critical points of the height function (p_1, p_2, p_3, p_4). Any point on the torus lies on a flow line between critical points: either a constant flow at a critical point, a distinguished flow γ_{ij}^k, or a flow from p_1 to p_4. There are two open intervals worth of flows from p_1 to p_4. On the right, we have illustrated these flows on the torus, seen as the standard quotient of $\mathbb{R}^2$ by the lattice $\mathbb{Z}^2$.

There are four critical points of f: p_1 of index 2, p_2 and p_3 of index 1, and p_4 of index 0. As illustrated in Fig. 4.2, the moduli spaces $\mathcal{M}(p_1,p_2)$, $\mathcal{M}(p_2,p_3)$, and $\mathcal{M}(p_3,p_4)$ each consist of two distinct points which we denote by $\gamma_{ij}^* \in \mathcal{M}(p_i,p_j)$ (for

[2]For a space X, the *suspension* SX of X is the quotient of $X \times I$ obtained by collapsing $X \times \{0\}$ to one point and $X \times \{1\}$ to another. It is well known (cf. [Hat02, Chapter 0]) that $SS^{n-1} = S^n$, with the two suspension points being precisely the north and south poles.

$* = 1, 2$). Any point on the vertical torus not in the image of one of these γ_{ij}^* lies on a flow in $\mathcal{M}(p_1, p_4)$. The moduli space $\mathcal{M}(p_1, p_4)$ is one dimensional, and is the disjoint union of two open intervals; its compactification $\overline{\mathcal{M}}(p_1, p_4)$ is the disjoint union of two closed intervals.

We saw in Example 1.2.13 that the the height function on the vertical torus is *not* Morse-Smale, namely because of the γ_{23}^k (recalling that the Morse-Smale condition prohibits flows between critical points of the same index), and Fig. 1.5 shows how the presence of these flows prevented the decomposition into unstable manifolds from being a CW complex. Similarly, we expect that the classifying space to reflect the absence of the Morse-Smale condition. Namely, we should expect the classifying space to be homotopic to the torus, but not homeomorphic. In fact, if we consider the simplicial description of the classifying space, we see that $B\mathscr{C}_f$ will be three-dimensional (given the existence of the triples of composable flows $(\gamma_{12}^{*_1}, \gamma_{23}^{*_2}, \gamma_{34}^{*_3})$). Therefore we cannot hope for a homeomorphism.

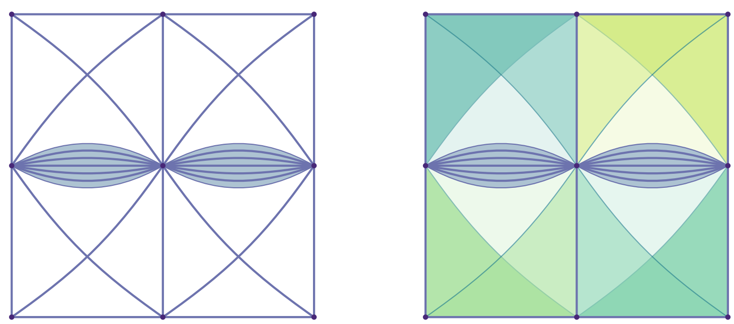

Figure 4.3: Simplicial description of the classifying space for the vertical torus, as the quotient $\mathbb{R}^2/\mathbb{Z}^2$. This figure illustrates three stages of the simplicial description of $B\mathscr{C}_f$. In the upper left square (labeled like the square in Fig. 4.2), we have added one-simplices for every morphism of $\mathscr{C}_f$, attached to the appropriate vertices. When we glue in the two-simplices, illustrated in the upper right square, each quadrant contains four overlapping triangles.

The vertices of the classifying space correspond to the objects of $\mathscr{C}_f$, the critical points of f, so there are four vertices. There is an edge (a one-simplex) for every morphism (flow) of $\mathscr{C}_f$, where the endpoints of the edge are glued to the start and end of the flow. The morphisms of $\overline{\mathcal{M}}(p_1, p_4)$ will index two one-parameter families of edges attached to the vertices for p_1 and p_4. Even at this first stage, the classifying space can no longer be embedded in two-dimensions— the one-simplices in the left illustration of Fig. 4.3 have four intersection points which are not vertices. Now, we add in the two-simplices for every pair of twice-composable flows. There are sixteen of these pairs, eight of which come from the choices of pairs $(\gamma_{ij}^{*_1}, \gamma_{jk}^{*_2})$ for $1 \leq i < j < k \leq 4$, and the other eight of which come from the composition of some broken flow and some γ_{ij}^*. A two-simplex associated with a pair of flows, say (γ, γ'), will have its three faces (edges) identified with the one-simplices labeled by γ, γ', and $\gamma \circ \gamma'$.

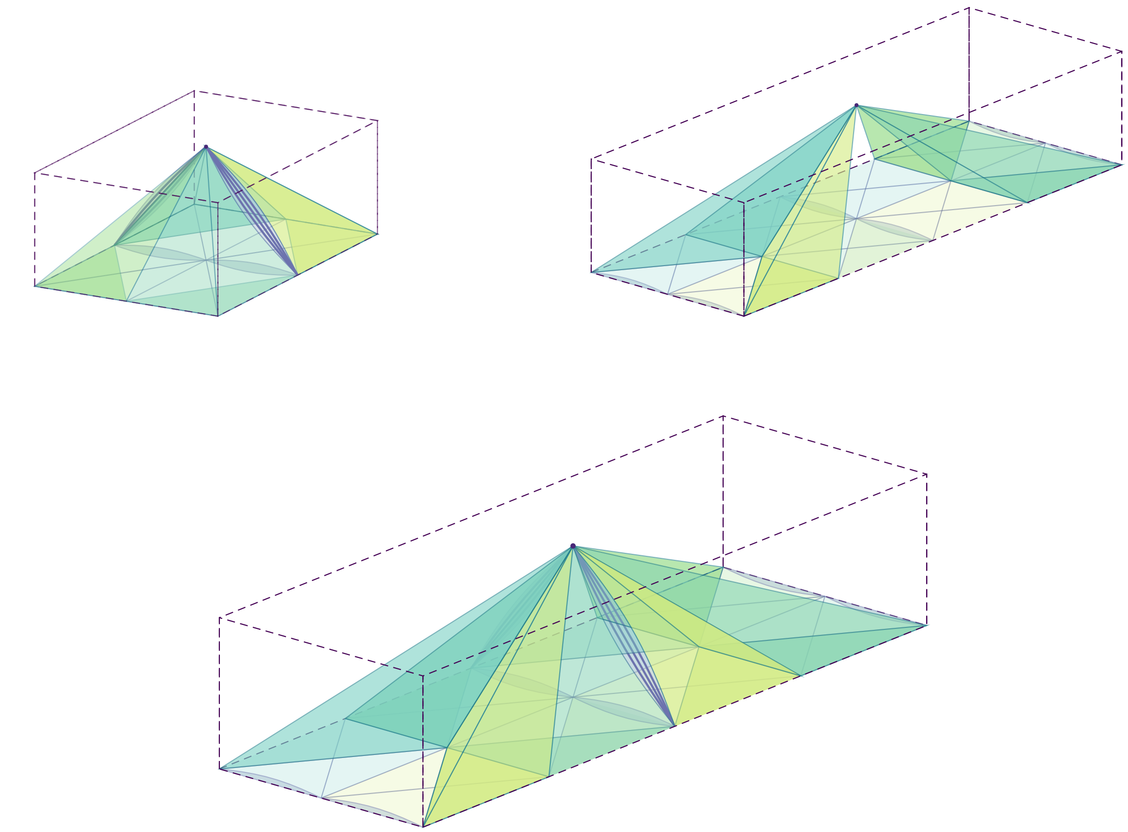

Figure 4.4: Classifying space of the vertical torus, embedded in $\mathbb{R}^2/\mathbb{Z}^2 \times \mathbb{R}/\mathbb{Z}$. We glue in eight tetrahedra for the eight thrice-composable flows. To be able to draw this space, we have raised p_4 to the top of the box. All the simplices connect to p_4 because all the thrice-composable flows end at p_4. The floor of the box, tiled by the square from Fig. 4.3, is the two-dimensional shadow of this space. In an attempt to make the illustration clearer, we have separated the eight simplices to groups of four: the "inner" simplices are displayed on the lower left (along with the one-parameter families of edges from $\mathcal{M}(p_1, p_4)$) and the four "outer" simplices on the right. The bottom illustration puts the two pieces together. The color of the triangle corresponds to the quadrant that the base of the simplex is on; this color is determined by the choice of two flows connecting p_1 to p_3. For example, the inner bright yellow-green tetrahedron corresponds to the triple $(\gamma_{12}^2, \gamma_{23}^2, \gamma_{34}^2)$, while the outer bright yellow-green triangle corresponds to $(\gamma_{12}^2, \gamma_{23}^2, \gamma_{34}^2)$. Note that there is no way for us to draw all eight simplices over one square from Fig. 4.3 without self-intersections. The additional presence of the gluing instructions (inherited from the original torus) means that we cannot embed the classifying space in $\mathbb{R}^3$.

The final step is to add in the three-simplices associated with triples of composable morphisms. Any higher-dimensional simplex in the nerve $N\mathscr{C}_f$ is degenerate, and so does not contribute to the classifying space. The triples must come from the product of moduli spaces $\mathcal{M}(p_1, p_2) \times \mathcal{M}(p_2, p_3) \times \mathcal{M}(p_3, p_4)$, and there are eight possible choices of $(\gamma_{12}^{*_1}, \gamma_{23}^{*_2}, \gamma_{34}^{*_3})$ given by the 2^3 possible choices for the $*$'s. Fig. 4.4

illustrates how we might glue these eight tetrahedra together along the appropriate two-dimensional faces. We leave it up to the reader to imagine the effect of the identifications and gluings that would turn our illustration into the "true" classifying space.

4.1.3 The Tilted Torus

By tilting the torus slightly, we change the gradient flow so that p_2 and p_3 are no longer connected by any flow lines, as shown in Fig. 4.5. We saw in Example 1.2.13 that this adjustment puts us in the Morse-Smale case. In fact, the tilted torus is the example in [CJS95b] used to illustrate the Morse-Smale case of the main theorem.

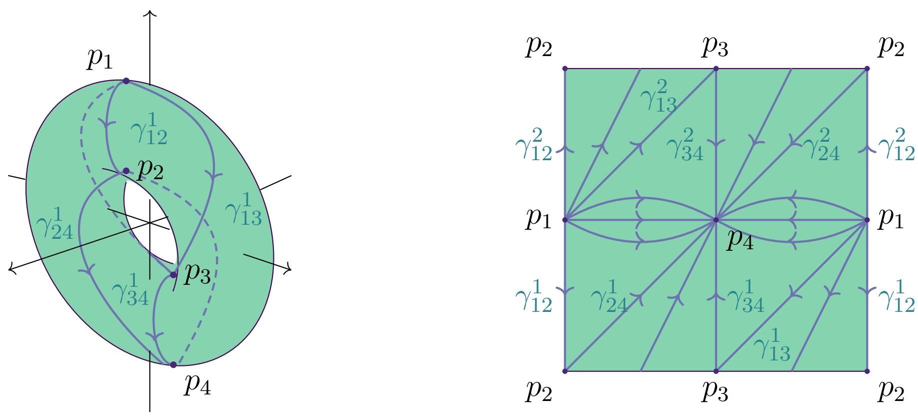

Figure 4.5: Flow lines on the tilted torus, seen as the quotient $\mathbb{R}^2/\mathbb{Z}^2$. The height function on the tilted torus still has four critical points (p_1, p_2, p_3, p_4), but now there are no flows connecting p_2 and p_3. In addition to the γ_{12}^* and γ_{34}^* from the vertical torus example, there are now flows γ_{13}^* and γ_{24}^* as well as four open intervals worth of flows from p_1 to p_4. On the right, we have illustrated these flows on the torus, seen as the standard quotient of $\mathbb{R}^2$ by the lattice $\mathbb{Z}^2$.

As for vertical torus, the height function f on the tilted torus has four critical points: p_1 of index 2, p_2 and p_3 of index 1, and p_4 of index 0. However, $\mathcal{M}(p_2, p_3)$ is now empty. Instead we have two distinct points in each of $\mathcal{M}(p_1, p_2)$, $\mathcal{M}(p_1, p_3)$, $\mathcal{M}(p_2, p_4)$, and $\mathcal{M}(p_3, p_4)$, which we will denote by γ_{ij}^* (for $* = 1, 2$) as before. Any point of the tilted torus not on one of the γ_{ij}^* must be on a flow in $\mathcal{M}(p_1, p_4)$. Now the compactification $\overline{\mathcal{M}}(p_1, p_4)$ consists of four disjoint closed intervals.

We go through the same process as before to build the classifying space for the tilted torus. Since any composition in $\mathscr{C}_f$ involving three or more morphisms will be degenerate, we only need to worry about the zero-, one-, and two-dimensional simplices. Once again, there are four vertices, corresponding to the four objects of $\mathscr{C}_f$, and one edge for every morphism in $\mathscr{C}_f$, including four one-parameter families connecting p_1 and p_4. The last step is gluing in triangles for every pair of composable morphisms. There are eight pairs, coming from the four points in the products $\mathcal{M}(p_1, p_i) \times \mathcal{M}(p_i, p_4)$ for $i = 2, 3$. The left illustration in Fig. 4.6 shows the resulting simplicial structure, which fits nicely into the quotient $\mathbb{R}^2/\mathbb{Z}^2$. In this case, we can

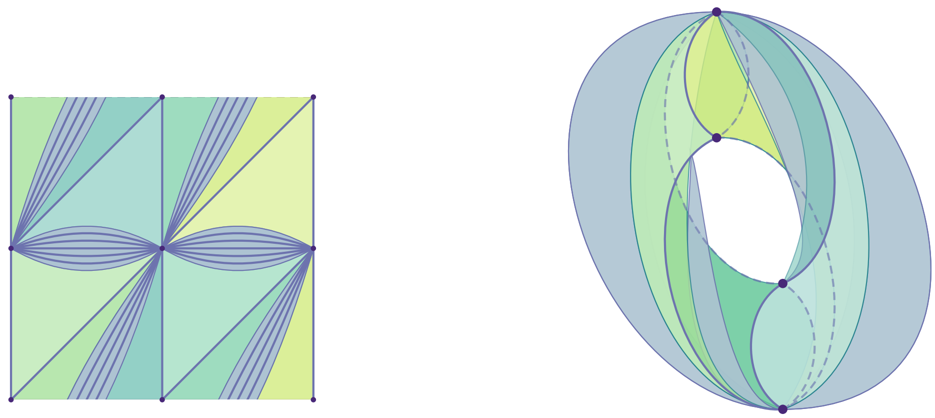

Figure 4.6: Classifying space of the tilted torus. The left portion of the figure illustrates the classifying space in the quotient $\mathbb{R}^2/\mathbb{Z}^2$, and the right portion follows the gluing instructions to reconstruct the torus. There are four one-parameter families of edges between p_1 and p_4; two of which form the "outer sides" of the torus, and the other two of which wrap around the torus by passing through the center hole. There are eight triangles corresponding to the eight (non-degenerate) pairs of composable morphisms in $\mathscr{C}_f$.

follow the gluing instructions to reconstruct the torus; the resulting figure illustrates the Morse-Smale case of Theorem 4.0.1 that the classifying space is homeomorphic to the underlying manifold.

4.2 Proof of the Theorem

The remainder of this chapter is dedicated to proving Theorem 4.0.1. We divide the work into two subsections, dealing first with the general case and second with the Morse-Smale case.

4.2.1 Part 1: The General Case

Given any Morse function f on a compact Riemannian manifold M, we can show that there is a homotopy equivalence $M \simeq B\mathscr{C}_f$. We record information about M in a constant topological category $\mathscr{M}$, so that $B\mathscr{M} \cong M$. Instead of comparing $\mathscr{M}$ to $\mathscr{C}_f$ directly, we introduce an enlarged version of the flow category C_f that allows gradient flows to have endpoints nearby critical points. This increased flexibility allows us to define a continuous functor Θ between a variation on the twisted arrow category of C_f and $\mathscr{M}$, and this functor admits a continuous local section Γ_p. We can then show that the induced map $B\Theta$ is a fibration with contractible fibers, which yields a homotopy equivalence on the classifying spaces. Roughly, the work of this section is to establish the following chain:

$$\mathscr{C}_f \hookrightarrow C_f \xrightarrow{\text{tw}} \text{tw}C_f \xrightarrow{\simeq} \tilde{C}_f \underset{\Theta}{\overset{\Gamma_p}{\leftarrow\!\!\!\to}} \mathscr{M}$$

Here, the $\simeq$ and $\cong$ relations are meant to be taken on the spaces of objects and morphisms separately. The work of the previous chapters then tells us that we have

$$M \cong B\mathscr{M} \simeq B\tilde{C}_f \simeq B\text{tw}C_f \cong BC_f \simeq B\mathscr{C}_f$$

whenever $f\colon M \to \mathbb{R}$ is a Morse function on M.

In more detail now, define $\mathscr{M}$ to be the topological category whose objects are the points of M and whose morphisms are only the identity functions on the points. That is,

$$\mathscr{M}_0 = M \text{ and } \mathscr{M}_1 = \{\text{id}_p \mid p \in M\} \cong M,$$

where the topology is inherited from M. It is straightforward to verify that $B\mathscr{M} \cong M$.[3]

To reach the desired result, we will show that there is a homotopy equivalence between $B\mathscr{M}$ and the classifying space of another category that is homotopy equivalent to the flow category.

Definition 4.2.1. The *almost-flow category* C_f is an enlarged version of the flow category. Suppose $\text{Crit}(f) = \{a_1, \dots, a_m\}$ and choose disjoint neighborhoods $U_1, \dots, U_m$ of these critical points. Since (M, g) is a Riemannian manifold, we may take U_i to be an open ball $B_{R_i}(a_i)$ around the critical point a_i of radius R_i (with respect to the metric induced by g, see Appendix A.1.5). Moreover, by the compactness of M, we may choose one radius R for all the neighborhoods. Our space of objects for C_f is the union $U = \bigcup_{i=1}^m U_i$, under the subspace topology.

The morphisms are broken gradient flow lines with endpoints in U which are height-parameterized flows (satisfying Equation (2.1.3)) away from critical points. That is, a morphism in C_f looks like a broken flow line whose endpoints may not be critical points but are within R of some elements of $\text{Crit}(f)$. In order to give this morphism space a nice topology, we reparameterize these almost-flow lines to share a common domain. Since M is compact, the image of f is compact and moreover is of the form $[f(a_i), f(a_j)]$ for some $a_i, a_j \in \text{Crit}(f)$. We will denote this interval by I_f. If $\gamma\colon [f(q), f(p)] \to M$ is an almost-flow from p to q, we reparameterize

$$\tilde{\gamma}(t) = \begin{cases} p & t \in [f(p), f(a_j)]; \\ \gamma(t) & t \in [f(q), f(p)]; \\ q & t \in [f(a_i), f(q)]. \end{cases}$$

[3]In general, given a topological space S and defining $\mathscr{S}$ to be the topological category whose objects $\mathscr{S}_0$ are the points of S and whose morphisms $\mathscr{S}_1$ are the identity maps on those points, we get a homeomorphism $B\mathscr{S} \cong S$. This result follows because all non-degenerate sequences have length 0 and so

$$B\mathscr{S} \cong \Delta^0 \times N_0\mathscr{S} \cong * \times \mathscr{S}_0 \cong S.$$

To extend the domain of γ to I_f, the newly parameterized $\tilde{\gamma}$ only "moves" in the interval $[f(q), f(p)] \subseteq I_f$, staying constant at the endpoints for the appropriate amount of time. If $\tilde{\gamma}_1 \in C_f(p,q)$ and $\tilde{\gamma}_2 \in C_f(q,r)$, their composition is given by

$$\tilde{\gamma}_1 \circ \tilde{\gamma}_2(t) = \begin{cases} \tilde{\gamma}_1(t) & t \in [f(q), f(a_j)]; \\ \tilde{\gamma}_2(t) & t \in [f(a_i), f(q)]. \end{cases}$$

We can then topologize the morphisms of C_f as a subspace of $\mathrm{Map}(I_f, M)$ under the compact open topology.

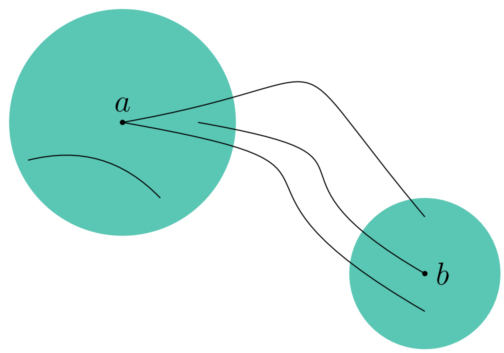

Figure 4.7: The almost-flow category C_f. The morphisms in C_f are height-parameterized gradient (broken) flows that share a common domain I_f. These "almost" flow lines may have their endpoints anywhere in the neighborhood, not necessarily at the critical points a or b.

Note that if $\tilde{\gamma}_1$ and $\tilde{\gamma}_2$ are composable, then $\tilde{\gamma}_1 \circ \tilde{\gamma}_2$ is another genuine almost-flow line. That is, we will not have any flow lines that are "broken" at any points other than critical points. As in the flow category, an almost-flow $\tilde{\gamma}$ from p to q is said to start at $s(\tilde{\gamma}) = p$ and end at $e(\tilde{\gamma}) = q$, although in fact the parameterization has $\tilde{\gamma}$ flowing in the opposite direction.

Proposition 4.2.2. *The almost-flow category C_f is a topological category.*

Proof. We need to check that the four structure maps are continuous (as discussed in Definition 2.2.2). The argument for composition is almost identical to the one given in Proposition 2.1.7. To see that the identity map $i\colon (C_f)_0 \to (C_f)_1$ is continuous, observe that the pre-image of any subbasis element

$$D(K, V) = \{\tilde{\gamma} \in (C_f)_1 \mid \gamma(K) \subseteq V\}$$

for compact $K \subseteq I_f$ and open $V \subseteq M$ will be precisely the intersection $V \cap U$, which is open in M. As for the domain map $\mathrm{dom}\colon (C_f)_1 \to (C_f)_0$ which sends a flow $\tilde{\gamma} \in C_f(p,q)$ to p, given any open set $V \subseteq M$, we can write $\mathrm{dom}^{-1}(V) = D(\{f(a_i)\}, V)$. Since every almost-flow "starts" at $f(a_i)$ (the right endpoint of I_f), this subbasis element is precisely the collection of flows that start at some point in V. Continuity of the codomain map follows similarly. $\square$

We work in a larger category for technical reasons, as suggested in the correspondence between Cohen and Segal that was generously shared with the author by Cohen. The method of proof is similar to the one in the original preprint [CJS95b], but working in this larger category fixes the issue of the continuity of the functor Γ discussed at the beginning of this chapter.

Lemma 4.2.3. *$BC_f \simeq B\mathscr{C}_f$.*

Proof. The inclusion of the flow category as a subcategory of C_f induces homotopy equivalences on the objects and morphisms. On the object level, we can contract each U_i to the critical point a_i. On the morphism level, the idea is that we will send a flow in C_f that passes through the neighborhoods of some sequence of critical points $a_1 \succ \cdots \succ a_k$ to a broken flow that connects $a_1, \ldots, a_k$.

Since f is strictly decreasing along the non-constant gradient flow lines (Proposition 1.1.10), once the flow leaves some U_i it cannot return. Note that if the almost-flow stays entirely within some U_i, then the contraction of U_i to a_i will take this flow to the steady state solution. It suffices to consider an almost-flow $\tilde{\gamma}$ that has its end points in some U_i and U_j (for $i \neq j$) and passes through no other U_i neighborhoods, since we can repeat the procedure for each piece of the longer flow. By the construction of C_f, there is some genuine gradient flow that $\tilde{\gamma}$ follows. The continuity of the flow guarantees that there is some nearby γ which starts at a_i and ends at a_j, and we can continuously deform $\tilde{\gamma}$ to this flow, yielding a levelwise homotopy equivalence of topological categories.

In order to conclude the classifying spaces are equivalent using Theorem 3.2.12, we must verify that the image of $i\colon (C_f)_0 \to (C_f)_1$ is closed. Equivalently, we will show that its complement, the collection of non-constant flows, is open in $(C_f)_1$. It suffices to produce a neighborhood for an arbitrary non-constant flow $\tilde{\gamma} \in C_f(p,q)$. We can exclude constant flows by considering an intersection of subbasis elements containing $\tilde{\gamma}$ whose open sets are disjoint. For instance, if $d_g(p,q) = \varepsilon$, we could take the neighborhood $D(\{f(p)\}, B_{\varepsilon/2}(p)) \cap D(\{f(q)\}, B_{\varepsilon/2}(q))$. This shows the image of i is closed, and so combining Lemma 3.2.10 and Theorem 3.2.12, we get the desired result. □

Remark 4.2.4. Note that the argument we have used above (and will use again in Lemma 4.2.6) requires that the spaces of objects and morphisms of both $\mathscr{C}_f$ and C_f are sufficiently well-behaved, that is, homotopy equivalent to CW complexes. While the spaces of objects clearly satisfy this condition, we could not find anything about the spaces of morphisms. We suspect that this fact is one of those notorious things that does not appear in the literature but is "known to the experts."

To establish a connection between C_f and M, we turn to the twisted arrow category of C_f (Definition 3.1.22). However, we extend to a slightly larger category to keep track of a bit more information.

Definition 4.2.5. The *pointed twisted almost-flow category* $\tilde{C}_f$ is the twisted arrow category $\mathrm{tw}(C_f)$ with some additional structure. Namely, the objects of $\tilde{C}_f$ are pairs $(\tilde{\gamma}, p)$ where $\tilde{\gamma} \in \mathrm{Ob}\,\mathrm{tw}(C_f)$ and $\tilde{\gamma}$ flows through $p \in M$. The morphisms of $\tilde{C}_f$ are commutative squares

$$\begin{array}{ccc} \bullet & \xleftarrow{\alpha} & \bullet \\ {\scriptstyle (\tilde\gamma, p)}\big\downarrow & & \big\downarrow{\scriptstyle (\tilde\gamma', q)} \\ \bullet & \xrightarrow[\beta]{} & \bullet \end{array}$$

where there is a morphism between $(\tilde{\gamma}, p)$ and $(\tilde{\gamma}', q)$ if and only if $p = q \in M$. We denote such a morphism by $(\alpha, \beta)_p$. This topological category inherits its topology as a subspace of $\mathrm{tw}(C_f) \times \mathscr{M}$.

The original [CJS95b] considers the similarly "pointed" version of $\mathrm{tw}(\mathscr{C}_f)$. The authors then define Γ to be the functor from $\mathscr{M}$ into this pointed flow category that maps $p \mapsto (\gamma_p, p)$. As explained in the beginning of this chapter, it turns out that this functor is not continuous, requiring us to expand to $\tilde{C}_f$.

Lemma 4.2.6. $B\tilde{C}_f \simeq B\mathrm{tw}(C_f)$.

Proof. The forgetful functor induces a fibration with contractible fibers between the objects and morphisms of the two topological categories. The proof of this fact follows from the observation that $\tilde{C}_f$ can be seen as a subcategory of $\mathrm{tw}(C_f) \times \mathscr{M}$ and that the image of every morphism $\tilde{\gamma} \in C_f$ is contractible (as a subset of M). □

Define $\Theta \colon \tilde{C}_f \to \mathscr{M}$ to be the functor that maps

$$(\tilde{\gamma}, p) \mapsto p \text{ and } (\alpha, \beta)_p \mapsto \mathrm{id}_p.$$

Note that Θ is just the projection map on $\mathrm{Ob}\,\tilde{C}_f$ and the composition of a projection with the domain (or codomain) map on morphisms, and so is continuous on both levels. A quick check against the functorial diagrams in Definition 2.2.3 verifies that Θ is indeed a continuous functor.

The benefit of working in the somewhat *ad hoc* category $\tilde{C}_f$ is that we can now find a local section of Θ. For $p \in M$, we denote this section by Γ_p, in the spirit of the original preprint. The idea is that Γ_p will take a point q nearby p to an almost-flow line that stays "close" to the minimal flow γ_p for all times $t \in I_f$. We claim that, for q sufficiently close, we can use the flow line γ_q which is allowed to be non-constant only on $[f(b), f(a)]$. We denote this restricted version of γ_q by γ_q^p.

Lemma 4.2.7. *Let $p \in M$, and suppose the minimal flow γ_p goes between critical points $a \in U_a$ and $b \in U_b$. Then there is a neighborhood V_p of p such that for all $q \in V_p$, the truncated flow line $\gamma_q^p \colon I_f \to M$ given by*

$$\gamma_q^p(t) = \begin{cases} \gamma_q(f(a)) & t \in [f(a), f(a_j)] \\ \gamma_q(t) & t \in [f(b), f(a)] \\ \gamma_q(f(b)) & t \in [f(a_i), f(b)], \end{cases}$$

also flows between U_a and U_b, meaning that $\gamma_q^p(f(b)) \in U_b$ and $\gamma_q^p(f(a)) \in U_a$. The local assignment $\tau_p \colon q \mapsto \gamma_q^p$ is continuous on V_p.

Proof. Recall that we have $U_a = B_R(a)$ and $U_b = B_R(b)$ by the construction of C_f. Since the flow assignment $\Phi(p,t) = \gamma_p(t)$ is smooth, there is a $\delta > 0$ such that

$$d_g(\gamma_p(t), \gamma_q(t')) < R \text{ whenever } d_g(p,q) + |t - t'| < \delta.$$

In particular, if $t = t'$, it is sufficient to require $d_g(p,q) < \delta$. Set $V_p = \{q \in M \mid d_g(p,q) < \delta\} \subseteq M$. Then at time $f(a)$, we have $d_g(\gamma_p(f(a)), \gamma_q(f(a))) = d_g(a, \gamma_q^p(f(a))) < R$. In other words, $\gamma_q^p(f(a)) \in U_a$, and similarly $\gamma_q^p(f(b)) \in U_b$.

Now we will show that the assignment $q \mapsto \gamma_q^p$ is continuous. Since I_f is compact and M is equipped with a metric d_g, the compact-open topology is metrizable, with the metric

$$\tilde{d}(\gamma, \gamma') = \sup_{t \in I_f}\{d_g(\gamma(t), \gamma'(t))\}.$$

Now consider some neighborhood of $\tau_p(q) = \gamma_q^p$, without loss of generality of the form $B_\varepsilon(\gamma_q^p)$ for some $\varepsilon > 0$ (where $B_\varepsilon(\gamma_q^p)$ is given with respect to $\tilde{d}$). Once again, the continuity of the flow gives us a $\delta > 0$ such that $d_g(\gamma_p(t), \gamma_q(t)) < \varepsilon$ whenever $d_g(p,q) < \delta$, for all times t that make sense. Take the neighborhood of q given by $B_\delta(q)$ and consider its image under the assignment map. We can see that $\tau_p(B_\delta(q)) \subseteq B_\varepsilon(\gamma_q^p)$ since

$$\begin{aligned}
\tau_p(B_\delta(q)) &= \left\{\gamma_q^p \mid d_g(p,q) < \delta\right\} \\
&\subseteq \left\{\gamma_q^p \mid d(\gamma_p(t), \gamma_q(t)) < \varepsilon \text{ for all } t \in I_f\right\} \\
&\subseteq \left\{\gamma_q^p \mid \tilde{d}(\gamma_p, \gamma_q) < \varepsilon\right\},
\end{aligned}$$

which shows that τ_p is continuous on V_p. □

We let $\mathscr{V}_p$ denote the topological subcategory of $\mathscr{M}$ which is given by the constant category on V_p. Now define $\Gamma_p \colon \mathscr{V}_p \to \tilde{C}_f$ to be the functor that takes

$$q \mapsto (\gamma_q^p, q) \quad \text{and} \quad \mathrm{id}_q \mapsto \mathrm{id}_{(\gamma_q^p, p)}.$$

It is straightforward to verify that Γ_p is both a functor and a local section of Θ.

Proposition 4.2.8. *The section Γ_p is continuous.*

Proof. Note that Γ_p on morphisms is the composition of $i \circ (\Gamma_p)_0 \circ \mathrm{dom}$ (where $(\Gamma_p)_0$ is Γ_p on objects), so it suffices to show that Γ_p is continuous on objects. The fact that Γ_p is continuous on objects follows from the observation that $\Gamma_p = (\tau_p, \mathrm{id}) \circ \Delta$, where Δ is the continuous diagonal map $q \mapsto (q,q)$. □

Since this section is only local, we cannot produce a homotopy inverse to Θ as the original authors claimed in [CJS95b]. However, we can use Γ_p to show that the induced map $B\Theta$ is a homotopy equivalence.

Proposition 4.2.9. *The induced map $B\Theta$ is a fibration with contractible fibers.*

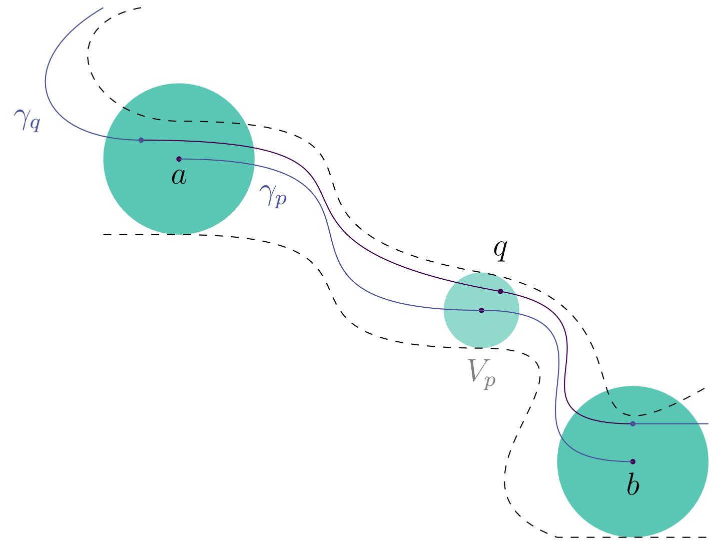

Figure 4.8: The truncated flow γ_q^p. The minimal flow γ_q of a point $q \in V_p$ passes through the neighborhoods of the critical points associated with γ_p. Restricting the flow γ_q to be non-constant only on the domain $[f(b), f(a)]$ gives the truncated flow γ_q^p.

Proof. That fact that $B\Theta$ is a fibration follows from the observation that we can lift any homotopy into $B\mathscr{M}$ to a homotopy into $B\Theta$ using the map induced by the continuous section Γ_p. By the functorial properties of B, the induced map $B\Gamma_p$ will be a local section of $B\Theta$.

It remains to show that the fibers $(B\Theta)^{-1}(p)$ are contractible. Consider $\Theta^{-1}(p)$, the subcategory of $\tilde{C}_f$ which consists of flows through p, pointed by p. By the argument in Lemma 4.2.3, this subcategory is homotopy equivalent to its restriction to $\mathrm{tw}(\mathscr{C}_f)$. Note that this identification entails that whenever p is in one of the distinguished neighborhoods U_i, we have $\Theta^{-1}(p) \simeq \Theta^{-1}_{\mathrm{tw}\mathscr{C}_f}(a_i)$, where the fiber on the left lives in $\tilde{C}_f$ and the fiber on the right lives in $\mathrm{tw}\mathscr{C}_f$. Otherwise, $\Theta^{-1}(p) \simeq \Theta^{-1}_{\mathrm{tw}\mathscr{C}_f}(p)$, where both subcategories are pointed by the same point p. In any case, the subcategory $\Theta^{-1}_{\mathrm{tw}\mathscr{C}_f}(p)$ has an initial object (the constant flow id_{a_i} if $p \in U_i$, and the minimal flow γ_p otherwise), which implies that $B(\Theta^{-1}_{\mathrm{tw}\mathscr{C}_f}(p)) \simeq B(\Theta^{-1}(p)) = (B\Theta)^{-1}(p)$ is contractible (see Remark 3.1.18). □

Assuming that the homotopy type of $B\mathscr{C}_f$ is sufficiently nice (see Remark 4.2.4), this proposition implies that $B\Theta\colon B\mathscr{C}_f \xrightarrow{\simeq} M$ is a homotopy equivalence, which finishes our proof of the general case.

4.2.2 Part 2: The Morse-Smale Case

In the special case that f is Morse-Smale, there is a homeomorphism $M \cong B\mathscr{C}_f$. The technical heart of the proof is the existence of an associative gluing map

$$\mu\colon (0, \varepsilon] \times \mathcal{M}(a, c) \times \mathcal{M}(c, b) \to \mathcal{M}(a, b).$$

This gluing map can be thought of as deforming broken flow lines $\gamma_1 \circ \gamma_2$ into smooth ones $\gamma_1 \circ_s \gamma_2$ that "stay s away from" the critical point c, for any $s \in (0, \varepsilon]$. We then consider the subspace $\mathcal{K}(a, b) \subseteq \mathcal{M}(a, b)$ that consists of all smooth flow lines from a to b that stay at least ε far away from all other critical points. It turns out that we can redefine the homspaces of the flow category in terms of these $\mathcal{K}(a, b)$ spaces. We can build an intermediary space between M and the flow category, denoted by $\mathcal{R}_f$, that is built out of products of $\mathcal{K}(a, b)$ with cubes that describe how to glue the flow lines together. The link between $\mathcal{R}_f$ and M is given by the evaluation map

$$\mathcal{R}_f \ni (t, (s_1, \ldots, s_m), (\gamma_0, \ldots, \gamma_m)) \mapsto (\gamma_0 \circ_{s_1} \cdots \circ_{s_m} \gamma_m)(t) \in M,$$

along with certain equivalence relations. Turning the cubes in $\mathcal{R}_f$ into simplices, we get another space $\tilde{\mathcal{R}}_f$ which is homeomorphic to the classifying space $B\mathscr{C}_f$. We can then establish a chain of homeomorphisms

$$B\mathscr{C}_f \cong \tilde{\mathcal{R}}_f \cong \mathcal{R}_f \cong M.$$

We follow the original proof from [CJS95b] (replicated in [CIN06, §12.2]), at times skimming over some of the more technical details. Our exposition should serve to familiarize the reader with the general method of proof, and convince the reader that this approach works.

Recall from Section 2.1.2 that we can define a partial ordering on critical points, where $a \succ b$ when $\mathcal{M}(a, b)$ is non-empty for $a \neq b$. In general, the moduli spaces $\mathcal{M}(a, b)$ are not compact, but have a canonical compactification given by

$$\overline{\mathcal{M}}(a, b) = \bigcup_{a \succ c_1 \succ \cdots \succ c_k \succ b} \mathcal{M}(a, c_1) \times \cdots \times \mathcal{M}(c_k, b),$$

where the union is taken over all ordered sequences of critical points from a to b. Essentially, we form $\overline{\mathcal{M}}(a, b)$ from $\mathcal{M}(a, b)$ by formally adjoining broken flow lines. The composition map

$$\overline{\mathcal{M}}(a, c) \times \overline{\mathcal{M}}(c, b) \to \overline{\mathcal{M}}(a, b)$$

is just the concatenation of curves. The issue is that these broken flows are not themselves flow lines, and so we must rethink how we can associate a pair of composable flows (γ_1, γ_2) with an actual flow line. In the Morse-Smale case, this can be done by deforming "patched curves" $\gamma_1 \#_\varepsilon \gamma_2$ to genuine flows $\gamma_1 \circ_\varepsilon \gamma_2$.

First, we choose $\varepsilon > 0$ so that for any two points $p, q \in M$ with geodesic distance $\delta < \varepsilon$, there is a unique geodesic $g_{p,q} \colon [-\frac{\delta}{2}, \frac{\delta}{2}] \to M$ joining p and q, parameterized by arclength. Now, suppose we have a sequence of critical points $a \succ c \succ b$ connected by (unbroken) flow lines γ_1, γ_2. For $0 < s < \frac{\varepsilon}{3}$, let p be the last point on the curve γ_1 whose distance from c is s and q be the first point on γ_2 whose distance from c is s. We define the *patched curve* by

$$(\gamma_1 \#_s \gamma_2)(t) = \begin{cases} \gamma_1(t) & t \leq -\frac{\delta}{2}; \\ g_{p,q}(t) & t \in [-\frac{\delta}{2}, \frac{\delta}{2}]; \\ \gamma_2(t) & t \geq \frac{\delta}{2}. \end{cases}$$

While these patched curves are not exactly flow lines, they are "close" to flow lines in such a way that every patched flow $\gamma_1 \#_s \gamma_2$ can be suitably associated with a bona fide flow $\gamma_1 \circ_s \gamma_2$. The specifics of this "closeness" is discussed in more detail in [CIN06, §9.2]. The parameter s can be thought of as a measure of how close the flow $\gamma_1 \circ_s \gamma_2$ comes to the critical point c (where γ_1 ends and γ_2 begins). Note that we do not yet allow $s = 0$, although we can think of $\gamma_1 \circ_0 \gamma_2$ as the familiar broken flow from Chapter 2; these flows will be formally adjoined later in the proof.

The crucial piece of the proof is the existence of certain associative gluing maps, and this is the moment when the Morse-Smale condition becomes necessary. To get these maps, the authors of [CJS95b] relied on a certain "folk theorem" regarding the structure of the compactified moduli spaces of a Morse-Smale flow, as discussed at the beginning of this chapter. The exact nature of the interplay between the Morse-Smale condition and the gluing maps is beyond our scope, but we point the interested reader to [Qin10, Qin11, Weh12].

Theorem 4.2.10. *There is some $\varepsilon > 0$ with a gluing map*

$$\mu\colon (0, \varepsilon] \times \mathcal{M}(a, c) \times \mathcal{M}(c, b) \to \mathcal{M}(a, b)$$

given by $(s, \gamma_1, \gamma_2) \mapsto \gamma_1 \circ_s \gamma_2$ that is a diffeomorphism onto its image. Moreover, μ is associative, in the sense that

$$(\gamma_1 \circ_s \gamma_2) \circ_{s'} \gamma_3 = \gamma_1 \circ_s (\gamma_2 \circ_{s'} \gamma_3)$$

for all $s, s' \leq \varepsilon$.

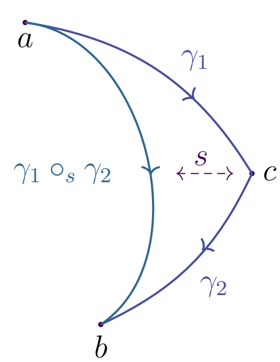

Figure 4.9: The glued flow $\gamma_1 \circ_s \gamma_2$. The gluing map μ takes (s, γ_1, γ_2) to the flow line that stays at least s away from c.

These gluing maps give us a way to associate a pair of flows (whose direct concatenation is a broken flow) with a genuine flow line. Recall that we use $\mathbf{c}(a, b)$ to denote an ordered sequence of critical points $a \succ c_1 \succ \cdots \succ c_k \succ b$, and say the length of $\mathbf{c}(a, b)$ is k. To save ourselves from notational headaches, we shall now denote such a sequence by merely $\mathbf{c}$. The associativity condition given above allows us to extend the gluing maps to products of moduli spaces

$$\mathcal{M}(\mathbf{c}) = \mathcal{M}(a, c_1) \times \cdots \times \mathcal{M}(c_k, b).$$

Corollary 4.2.11. *For any ordered sequence* $\mathbf{c}$ *of length* $k \geq 1$*, the gluing map*

$$\mu_{\mathbf{c}} \colon (0, \varepsilon]^k \times \mathcal{M}(\mathbf{c}) \to \mathcal{M}(a, b)$$

given by $(s_1, \ldots, s_k; \gamma_0, \ldots, \gamma_k) \mapsto \gamma_0 \circ_{s_1} \cdots \circ_{s_k} \gamma_k$ *is a diffeomorphism onto its image.*

By rescaling the metric, we can take $\varepsilon = 1$ (cf. [CIN06, §12.2]). The image of $\mu_{\mathbf{c}}$ consists of flow lines from a to b that come within 1 of each of the critical points $c_1, \ldots, c_k$. The space we are interested in is $\mathcal{M}(a, b)$ with these flows that "come too close" removed.

Definition 4.2.12. Define $\mathcal{K}(a, b) \subseteq \mathcal{M}(a, b)$ by

$$\mathcal{K}(a, b) := \mathcal{M}(a, b) - \bigcup_{\mathbf{c}} \mu_{\mathbf{c}} \left((0, 1)^k \times \mathcal{M}(\mathbf{c}) \right).$$

Since the gluing maps are diffeomorphisms onto their images, we can think of $\mathcal{K}(a, b)$ as the space of flow lines connecting a and b that stay at least 1 away from all other critical points. As it turns out, there is a homeomorphism between $\mathcal{K}(a, b)$ and $\overline{\mathcal{M}}(a, b)$. We can see that the ends of the moduli space $\mathcal{M}(a, b)$ consist of spaces of half-open cubes $(0, 1]^k$ parameterized by the composable sequences of flow lines in $\mathcal{M}(\mathbf{c})$; if we remove the *open* cubes $(0, 1)^k$, we arrive at $\mathcal{K}(a, b)$, and if we formally close the cubes, we arrive at $\overline{\mathcal{M}}(a, b)$.

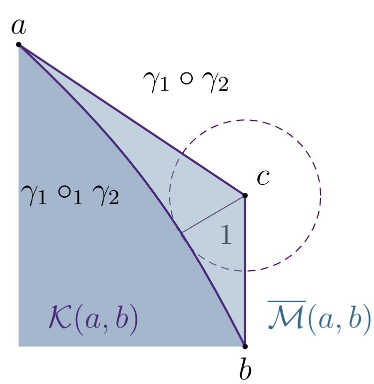

Figure 4.10: Comparing the spaces $\mathcal{K}(a, b)$ **and** $\overline{\mathcal{M}}(a, b)$**.** The space $\overline{\mathcal{M}}(a, b)$ is formed by adjoining the broken flow $\gamma_1 \circ \gamma_2$ that passes through the critical point c centered in the ball of radius 1. The space $\mathcal{K}(a, b)$ is formed from $\mathcal{M}(a, b)$ by removing all the flows that pass through the interior of that ball. The flow $\gamma_1 \circ_1 \gamma_2$ lies on the boundary of $\mathcal{K}(a, b)$. Both $\overline{\mathcal{M}}(a, b)$ and $\mathcal{K}(a, b)$ are compact, and we can imagine continuously deforming one to the other. The triangular region between $\gamma_1 \circ_1 \gamma_2$ and $\gamma_1 \circ \gamma_2$ (including the former but not the latter) is the image of the gluing map μ.

Theorem 4.2.13. *The space* $\mathcal{K}(a, b)$ *is compact and homeomorphic to*

$$\overline{\mathcal{M}}(a, b) := \operatorname{colim}_{\mathbf{c}} \left(\mathcal{M}(a, b) \cup_{\mu_{\mathbf{c}}} [0, 1]^k \times \mathcal{M}(\mathbf{c}) \right).$$

The expression in the parentheses is the pushout of $[0,1]^k \times \mathcal{M}(\mathbf{c}) \hookleftarrow (0,1]^k \times \mathcal{M}(\mathbf{c}) \xrightarrow{\mu_{\mathbf{c}}} \mathcal{M}(a,b)$, and the colimit is taken over all ordered sequences $\mathbf{c}$ from a to b. This theorem allows us to redefine the homspaces of the flow category using $\mathcal{K}(a,b)$. The benefit of this perspective is that $\mathcal{K}(a,b) \subseteq \mathcal{M}(a,b)$ is a collection of bona fide flow lines, whereas $\overline{\mathcal{M}}(a,b)$ includes the piecewise smooth broken flows. Composition in the flow category (previously given by $(\gamma_1,\gamma_2) \mapsto \gamma_1 \circ \gamma_2$) can now be given by $(\gamma_1,\gamma_2) \mapsto \gamma_1 \circ_1 \gamma_2$.

By defining $\gamma_1 \circ_0 \gamma_2$ to be the broken flow line $\gamma_1 \circ \gamma_2$, we extend $\mu_{\mathbf{c}}$ to a map

$$\mu_{\mathbf{c}} \colon [0,1]^k \times \mathcal{M}(\mathbf{c}) \to \overline{\mathcal{M}}(a,b).$$

The space $\overline{\mathcal{M}}(a,b)$ has a filtration given by

$$\mathcal{K}(a,b) = \mathcal{K}^{(0)}(a,b) \subseteq \cdots \subseteq \mathcal{K}^{(i-1)}(a,b) \subseteq \mathcal{K}^{(i)}(a,b) \subseteq \cdots \subseteq \overline{\mathcal{M}}(a,b) = \bigcup_i \mathcal{K}^{(i)}(a,b)$$

where $\mathcal{K}^{(i)}(a,b)$ consists of those flows γ that come within distance 1 of at most i intermediary critical points. In other words, γ can be decomposed as

$$\gamma = \gamma_0 \circ_{s_1} \cdots \circ_{s_k} \gamma_k$$

for some $k \leq i$, with the $s_j \in [0,1]$ and the $\gamma_j \in \mathcal{K}(c_j, c_{j+1})$ (with the understanding that $c_0 = a$ and $c_{k+1} = b$). More precisely,

$$\mathcal{K}^{(i)}(a,b) = \bigcup_{k \leq i} \bigcup_{\mathbf{c}} \mu\left([0,1]^k \times \mathcal{K}(\mathbf{c})\right)$$

where $\mathcal{K}(\mathbf{c})$ is the product

$$\mathcal{K}(\mathbf{c}) = \mathcal{K}(a, c_1) \times \cdots \times \mathcal{K}(c_k, b).$$

At each step, the piece we "add on" to $\mathcal{K}^{(i)}(a,b)$ are those flows that come within 1 of i intermediary critical points. That is,

$$\mathcal{K}^{(i)}(a,b) - \mathcal{K}^{(i-1)}(a,b) = \coprod_{\mathbf{c}} \mu_{\mathbf{c}}\left([0,1)^i \times \mathcal{K}(\mathbf{c})\right),$$

where the disjoint union is taken across all ordered sequences $\mathbf{c}$ of length i connecting a and b. Since $\mu_{\mathbf{c}}$ is a diffeomorphism onto its image and $\overline{\mathcal{M}}(a,b) = \bigcup_i \mathcal{K}^{(i)}(a,b)$, the map

$$\coprod_k \coprod_{\substack{\mathbf{c} \\ l(\mathbf{c})=k}} [0,1]^k \times \mathcal{K}(\mathbf{c}) \to \overline{\mathcal{M}}(a,b)$$
$$(s_1, \ldots, s_k; \gamma_0, \ldots, \gamma_k) \mapsto \gamma_0 \circ_{s_1} \cdots \circ_{s_k} \gamma_k$$

is surjective. In order to turn this map into a homeomorphism, we just need to define the appropriate equivalence relations. Now, the only places where the map may not be injective are those places where some $\gamma_{i-1} \circ_{s_i} \circ \gamma_i$ is already in some $\mathcal{K}(c_j, c_{j+1})$, which can happen if and only if $s_i = 1$. Thus if we declare

$$(s_1, \ldots, s_{i-1}, 1, s_{i+1}, \ldots, s_k; \gamma_0, \ldots, \gamma_k) \\ \sim (s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_k; \gamma_0, \ldots, \gamma_{i-1} \circ_1 \gamma_i, \ldots, \gamma_k),$$

we will get the desired homeomorphism.

Theorem 4.2.14. *The compactified moduli space of broken flows $\overline{\mathcal{M}}(a, b)$ is homeomorphic to*

$$\coprod_k \coprod_{\substack{\mathbf{c} \\ l(\mathbf{c})=k}} [0,1]^k \times \mathcal{K}(\mathbf{c}) \Big/ \sim,$$

where $\sim$ is the relation defined above.

We will make use of this homeomorphism soon to build the intermediary space $\mathcal{R}_f$, but first we establish a connection between $\mathcal{M}(a, b)$ and M using evaluation maps of flow lines, namely, the maps

$$[f(b), f(a)] \times \overline{\mathcal{M}}(a, b) \to M$$

which send $(t, \gamma) \mapsto \gamma(t)$. Note that the image of this map is the closure of $W(a, b)$, since the images of the broken height-parameterized flows γ include all the critical points that γ passes through, as discussed in Section 2.1.1. By considering all ordered sequences of critical points $\mathbf{c} = \{c_0, c_1, \ldots, c_k, c_{k+1}\}$, we get a map

$$\coprod_{\mathbf{c}} [f(c_{k+1}), f(c_0)] \times [0,1]^k \times \mathcal{K}(\mathbf{c}) \to M$$

which evaluates $(t; s_1, \ldots, s_k; \gamma_0, \ldots, \gamma_k) \mapsto (\gamma_0 \circ_{s_1} \cdots \circ_{s_k} \gamma_k)(t)$. It is clear that this map is onto, since every point $p \in M$ is in the image of its minimal flow γ_p. We will use the same trick as before to turn this map into a homeomorphism, namely, by defining the appropriate equivalence relation.

Definition 4.2.15. Let $\mathbf{c} = \{c_0, \ldots, c_{k+1}\}$ denote an arbitrary ordered sequence of critical points, with length $l(\mathbf{c}) = k$. Let $J_{\mathbf{c}} = [f(c_{k+1}), f(c_0)]$ and $I^{\mathbf{c}} = [0,1]^k$. Then define

$$\mathcal{R}_f = \coprod_{\mathbf{c}} J_{\mathbf{c}} \times I^{\mathbf{c}} \times \mathcal{K}(\mathbf{c}) \Big/ \sim,$$

where $\sim$ is the equivalence relation given by an adjustment of the previous relation

$$(t; s_1, \ldots, s_{i-1}, 1, s_i, \ldots, s_k; \gamma_0, \ldots, \gamma_k) \\ \sim (t; s_1, \ldots, s_{i-1}, s_i, \ldots, s_k; \gamma_0, \ldots, \gamma_{i-1} \circ_1 \gamma_i, \ldots, \gamma_k)$$

in addition to

$$(t; s_1, \ldots, s_{i-1}, 0, s_i, \ldots, s_k; \gamma_0, \ldots, \gamma_k) \\ \sim \begin{cases} (t; s_1, \ldots, s_{i-1}; \gamma_0, \ldots, \gamma_{i-1}) & t \in [f(c_i), f(c_0)]; \\ (t; s_i, \ldots, s_k; \gamma_i, \ldots, \gamma_k) & t \in [f(c_{k+1}), f(c_i)]. \end{cases}$$

The first relation takes care of the overlap between $I^{\mathbf{c}}$ and $\mathcal{K}(\mathbf{c})$, as before. The second relation takes care of the broken flow lines, allowing us to identify a piecewise flow $\gamma_1 \circ_0 \gamma_2$ with one of its pieces, γ_1 or γ_2, depending on where $(\gamma_1 \circ_0 \gamma_2)(t)$ lands. The point $(\gamma_1 \circ_0 \gamma_2)(t)$ is either a joining point (that is, when $t = f(c_i)$ in Definition 4.2.15) and so we may associate with either segment, or else $(\gamma_1 \circ_0 \gamma_2)(t)$ lies in the image of either γ_1 or γ_2. In any case, the only identifications that may take place when some s_i is 0 can be reached by that set of relations. Imposing these relations turns the evaluation map $(t; s_1, \ldots, s_k; \gamma_0, \ldots, \gamma_k) \mapsto (\gamma_0 \circ_{s_1} \cdots \circ_{s_k} \gamma_k)(t)$ into a homeomorphism.

Theorem 4.2.16. *The evaluation map $\mathcal{R}_f \to M$ is a homeomorphism.*

Finally, we need to show $\mathcal{R}_f \cong B\mathscr{C}_f$. Comparing the definitions of $\mathcal{R}_f$ and $B\mathscr{C}_f$ (particularly considering the concrete description of the geometric realization after Definition 3.1.13), we see that there is quite a bit of similarity, the essential difference being that $\mathcal{R}_f$ is built out of cubes while $B\mathscr{C}_f$ is built out of simplices. We can establish a homeomorphism between the two spaces via an intermediary space $\tilde{\mathcal{R}}_f$ which is built out of $\mathcal{R}_f$ in two steps: we first collapse the cubes into simplices, and then we impose relations that glue the simplices together to make $B\mathscr{C}_f$.

We want to turn the cubes $J_{\mathbf{c}} \times I^{\mathbf{c}}$ into a simplex without changing their image in the quotient space $\mathcal{R}_f$. Looking at the identifications made in $\mathcal{R}_f$ under $\sim$, we are motivated to give a name to the ordered sequence $\mathbf{c}$ with a critical point c_i removed; define $\bar{\mathbf{c}}_i = \{c_0, \ldots, c_{i-1}, c_{i+1}, \ldots, c_{k+1}\}$ to be this sequence. Now we will define a relation $\sim$ so that the space

$$J_{\mathbf{c}} \times I^{\mathbf{c}} / \sim$$

is naturally homeomorphic to a simplex. If $s_i = 0$ for some $1 \le i \le k$ (implying that $J_{\mathbf{c}} = J_{\bar{\mathbf{c}}_i}$), then we say

$$(t; s_1, \ldots, s_{i-1}, 0, s_{i+1}, \ldots, s_k) \sim \begin{cases} (t; s_1, \ldots, s_{i-1}, 0, s'_{i+1}, \ldots, s'_k) & t \in [f(c_i), f(c_0)]; \\ (t; s'_1, \ldots, s'_{i-1}, 0, s_{i+1}, \ldots, s_k) & t \in [f(c_{k+1}), f(c_i)], \end{cases}$$

for any $s'_j \in [0, 1]$. If $i = 0, k+1$, then we have

$$\begin{aligned} (t; 0, s_2, \ldots, s_k) &\sim (t; 0, s'_2, \ldots, s'_k) && t \in [f(c_1), f(c_0)], \\ (t; s_1, \ldots, s_{k-1}, 0) &\sim (t; s'_1, \ldots, s'_{k-1}, 0) && t \in [f(c_{k+1}), f(c_k)]; \end{aligned}$$

and finally we collapse $\{f(c_i)\} \times I^{\mathbf{c}}$ (for $i = 0, k+1$) to a point, imposing

$$\begin{aligned} (f(c_{k+1}); s_1, \ldots, s_k) &\sim (f(c_{k+1}); s'_1, \ldots, s'_k), \\ (f(c_0); s_1, \ldots, s_k) &\sim (f(c_0); s'_1, \ldots, s'_k). \end{aligned}$$

Note that if two points are identified by this relation, then they have the same image in $\mathcal{R}_f$. The point is that these quotiented cubes are homeomorphic to simplices, and so we can build $\mathcal{R}_f$ (or at least something homeomorphic to it) out of simplices rather than cubes. To define the appropriate gluing together of these simplices, we need the "face inclusion maps"

$$\delta_i \colon J_{\overline{\mathbf{c}}_i} \times I^{\overline{\mathbf{c}}_i} \to J_{\mathbf{c}} \times I^{\mathbf{c}},$$

for $0 \leq i \leq k+1$, defined by

$$\delta_i(t; s_1, \ldots, s_{k-1}) = \begin{cases} (t; 0, s_1, \ldots, s_{k-1}) & i = 0; \\ (t; s_1, \ldots, s_{i-1}, 1, s_i, \ldots, s_{k-1}) & 1 \leq i \leq k; \\ (t; s_1, \ldots, s_{k-1}, 0) & i = k+1. \end{cases}$$

Lemma 4.2.17. *There are homeomophisms* $\phi_{\mathbf{c}} \colon J_{\mathbf{c}} \times I^{\mathbf{c}} / \sim \to \left|\Delta^{k+1}\right|$ *such that the following diagrams commute:*

$$\begin{array}{ccc} J_{\mathbf{c}} \times I^k / \sim & \xrightarrow{\phi_{\mathbf{c}}} & \left|\Delta^{k+1}\right| \\ {\scriptstyle \delta_i}\uparrow & & \uparrow{\scriptstyle d_i} \\ J_{\overline{\mathbf{c}}_i} \times I^{k-1} / \sim & \xrightarrow[\phi_{\overline{\mathbf{c}}_i}]{} & \left|\Delta^{k}\right| \end{array} .$$

Up to this point, we have used the second set of relations in Definition 4.2.15. Using this lemma and imposing the first set of relations on $\mathcal{R}_f$, we get a homeomorphism between $\mathcal{R}_f$ and the following space.

Definition 4.2.18. Define

$$\tilde{\mathcal{R}}_f = \coprod_{\mathbf{c}} \left|\Delta^{k+1}\right| \times \mathcal{K}(\mathbf{c}) \Big/ \sim,$$

where $\sim$ is the relation given by

$$\begin{aligned} &(s_0, \ldots, s_{i-1}, 0, s_{i+1}, \ldots, s_k; \gamma_0, \ldots, \gamma_k) \\ &\qquad \sim \begin{cases} (s_1, \ldots, s_k; \gamma_0, \ldots, \gamma_k) & i = 0; \\ (s_0, \ldots, s_{i-1}, s_{i+1}, \ldots, s_k; \gamma_0, \ldots, \gamma_{i-1} \circ_1 \gamma_i, \ldots, \gamma_k) & 1 \leq i \leq k-1; \\ (s_0, \ldots, s_{k-1}; \gamma_0, \ldots, \gamma_{k-1}) & i = k. \end{cases} \end{aligned}$$

The definition of $\tilde{\mathcal{R}}_f$ is almost identical to that of the classifying space $B\mathscr{C}_f$, except $\tilde{\mathcal{R}}_f$ only uses the space of *non-degenerate* simplices, rather than the space of *all* simplices. Thus we have

$$\mathcal{R}_f \cong \tilde{\mathcal{R}}_f \cong B\mathscr{C}_f,$$

which implies that $B\mathscr{C}_f \cong M$.

Suggested Further Reading

There is only so much one can say in a thesis without being in danger of writing a book. While we have attempted to give a substantive overview of classical Morse theory, there are undoubtedly areas of the field that we have unfairly skimmed over or skipped entirely. The aim of this final portion of the thesis is to briefly outline a few areas of mathematics related to Morse theory that the interested reader may pursue.

More Morse Theory

One interesting way to extend Morse theory is to loosen our requirements on the functions we consider. For instance, *Morse-Bott theory* considers functions whose critical set may be a closed submanifold of M, such as the height function on the horizontal torus we saw in Example 1.2.13.[4] Morse functions become a special case of Morse-Bott functions, where the critical sets are zero dimensional. Many of the ideas of Morse theory can be extended to Morse-Bott theory, including Morse homology (see [AB95, §3] or [Hut02, §6]).

Rather than changes the types of functions under consideration, we might change the types of spaces. This is precisely the approach of *stratified Morse theory*, which generalizes Morse theory to certain spaces with singularities; the classical reference for this theory is the work of Mark Goresky and Robert MacPherson, particularly [GM88]. In particular, results about the topology of sublevel sets (Section 1.3.1) can be extended to (Whitney) stratified spaces, that is, spaces X that admit a finite filtration

$$\emptyset \subseteq X_0 \subseteq X_1 \subseteq \cdots \subseteq X_n = X$$

which satisfies certain nice properties.

A more combinatorial generalization of Morse theory is *discrete Morse theory*, which considers simplicial complexes under so-called discrete Morse functions that (roughly) assign higher numbers to higher-dimensional simplicies. This field, developed by Robin Foreman [For02], has found many practical uses in computational homology, topological data analysis, and other areas of applied math and computer science. Many of the basic results of standard Morse theory generalize to discrete Morse theory, such as the Morse inequalities, the Morse complex, and Morse homology. In fact, there is a result for the flow category for a discrete Morse function

[4] In fact, this example belongs to a special class of Morse-Bott functions, known as *round functions*.

analogous to the main theorem of this thesis, formulated and proved in [NTT16].

In addition to these topics we have mentioned, there is plenty more to say about Morse theory, its spin-offs, and its relations to other areas of mathematics. The expositions [Bot80, Bot88] and [Gue01] give good overviews of some aspects of Morse theory that escape the scope of this thesis.

Floer Homology and Homotopy Theory

Floer homology is a tool used in modern symplectic geometry and low-dimensional topology, and can be seen as an infinite-dimensional analogue of Morse homology. Andreas Floer introduced the first version of Floer homology in his proof of the Arnold conjecture (see [AD14, Part II] for a thorough treatment). Roughly, the Arnold conjecture proposes that the number of periodic trajectories of a Hamiltonian vector field on a symplectic manifold W is bounded below by

$$\sum_k \dim HM_k(W;\mathbb{Z}/2).$$

The connection between the Arnold conjecture and the Morse inequalities (Theorem 1.3.15) was key to Floer's proof strategy. Variations of Floer homology have played important roles in many other areas of mathematics and physics.

In finite-dimensional Morse theory, it is geometrically clear why the homotopy type of a manifold M is reflected in the structure of critical points and gradient flows between them, as captured by Morse homology. A program of Cohen, Jones, and Segal [CJS95a] seeks to develop a comparable Floer homotopy theory. The recent paper of Cohen [Coh19] summarizes some recent applications of ideas from Floer homotopy theory, including the work of Lipshitz and Sarkar related to Khovanov homotopy theory [LS11]. With many unanswered questions, this area of mathematics provides a fertile ground for further research in the intersection of homotopy theory, symplectic geometry, and low-dimensional topology.

Appendix A

Some Things to Know

When the author began working on this thesis, she quickly realized how much she did not know. The hope is that this appendix will provide you, the reader, with a relatively succinct account of Some Things to Know in order to understand this thesis. The three main areas discussed are differential topology (Appendix A.1), algebraic topology (Appendix A.2), and category theory (Appendix A.3), appearing roughly in the order that the material is used in the main body of the thesis.

A.1 A Bit of Differential Topology

We first cover some background material from differential geometry and differential topology. This section is essentially a retelling of [AD14, Appendix A], with some bits skipped and some bits expanded upon. While the reader could likely get away with imagining everything taking place in familiar Euclidean space, replacing "manifold" with "$\mathbb{R}^n$" and invoking the notions of tangent vector and derivative from multivariable calculus, we encourage examining the definition of tangent vectors and the tangent space (Definition A.1.8), the definition of the tangent map (Definition A.1.9), and the discussion of vector bundles (Appendix A.1.3). Both [Mor01] and [Küh15] provide a more detailed exposition of these basic concepts, the classic [Lee71] is an excellent reference for all things related to topological manifolds, and the author found [Hir76, War71] helpful for the more technical aspects of calculus on manifolds.

A.1.1 Manifolds

Intuitively, manifolds are abstract topological spaces that locally resemble $\mathbb{R}^n$. The concept of manifold allows us to describe complicated spaces in terms of the simpler and well-understood structure of Euclidean space.

Definition A.1.1. A topological *manifold* M of dimension n is a Hausdorff space such that every point has a neighborhood U that is homeomorphic to an open subset of $\mathbb{R}^n$ via a map ϕ. The pair (U, ϕ) is called a *(coordinate) chart.* An *atlas* is a family $\{U_i, \phi_i\}_{i \in I}$ for which the U_i constitute a covering of M.

Remark A.1.2. In some cases, a second countability condition[1] is also included in the definition of a topological manifold. However, we will primarily consider compact manifolds, for which this property automatically holds.

Given a chart $\{U, \phi\}$, a point $p \in U$ is determined by $\phi(p)$, and we often identify the two. The components of $\phi(p) \in \mathbb{R}^n$ are called the *local coordinates* of p.

Definition A.1.3. Two charts (U_i, ϕ_i), (U_j, ϕ_j) are *compatible* if the chart transition

$$\phi_j \circ \phi_i^{-1} \colon \phi_i(U_i \cap U_j) \to \phi_j(U_i \cap U_j)$$

is differentiable of class C^∞ (just in case $U_i \cap U_j$ is non-empty). An atlas is *differentiable* if all its chart transitions are compatible, and we say that such an atlas determines a smooth (or differentiable) structure on M. A manifold with a smooth structure is aptly called a *smooth* (or *differentiable*) manifold.

Two differentiable atlases are said to be *equivalent* if their union is also differentiable. The *maximal atlas* determined by a differentiable atlas $\{U_i, \phi_i\}_{i\in I}$ is the union of all differentiable atlases equivalent to $\{U_i, \phi_i\}_{i\in I}$. Hence two atlases are equivalent if and only if they determine the same maximal atlas.

Example A.1.4 (Examples of differentiable manifolds)**.** We present some examples of manifolds that may be familiar from other areas of mathematics.

- Euclidean n-space is of course an example of a smooth manifold.
- The n-sphere S^n admits a smooth structure via stereographic projection.
- The n-dimensional torus $T^n = S^1 \times \cdots \times S^1$ can also be seen as the quotient of $\mathbb{R}^n$ by the lattice $\mathbb{Z}^n$ acting by vector addition.
- Projective real space,
$$\mathbb{RP}^n = S^n/(x \sim -x)$$
has a smooth structure given by the collection of open sets $U_i \subseteq S^n$ whose points have a non-zero i^{th} coordinate. Each U_i is homeomorphic to an open ball in $\mathbb{R}^n$. We can also think of projective space as the collection of lines passing through the origin in $\mathbb{R}^{n+1}$.
- Similarly, we have complex projective space
$$\mathbb{CP}^n = (\mathbb{C}^{n+1} \setminus \{0\})/\mathbb{C}^*,$$
where $\mathbb{C}^* = \mathbb{C} \setminus \{0\}$ acts by multiplication on each coordinate of $\mathbb{C}^{n+1}$. This space is covered by $U_i = \{z \mid z_i \neq 0\}$, where now each U_i is homeomorphic to a complex open ball of dimension n, and so has real dimension $2n$. The chart transitions are given by linear fractional transformations.

[1] This condition asserts that there is a base of countably many elements for the system of open sets.

Manifolds can have many different qualities and characteristics. A manifold is *orientable* if it admits an orientation, meaning that (in the case of a differentiable manifold) there is an atlas on M whose transition functions have positive Jacobians. For instance, the circle S^1 is orientable but the Möbius strip is not. This thesis primarily considers *closed* manifolds, that is, manifolds which are both compact (as topological spaces) and without boundary.[2] All of the manifolds in the example above are closed and orientable.

Definition A.1.5. A map $f\colon M \to N$ between differentiable manifolds with charts $\{U_i, \phi_i\}$ and $\{V_j, \psi_j\}$ is *smooth* (or *differentiable*) if all the maps $\psi_j \circ f \circ \phi_i^{-1}$ are differentiable (where defined). A smooth bijective map with a smooth inverse is said to be a *diffeomorphism*, and the manifolds are then *diffeomorphic*.

In the case where $N = \mathbb{R}$, the map f is called a *function* on M. We let $C^\infty(M)$ denote the set of all smooth functions on M, and we note that $C^\infty(M)$ has the structure of an algebra[3] over $\mathbb{R}$.

Certain special types of maps get their own names, such as immersion, embedding, and submersion (we will define these qualities more precisely in Definition A.1.10 after introducing tangent maps). Roughly, immersions and embeddings can be thought of as inclusions of spaces, and the prototypical example of a submersion is a projection map.

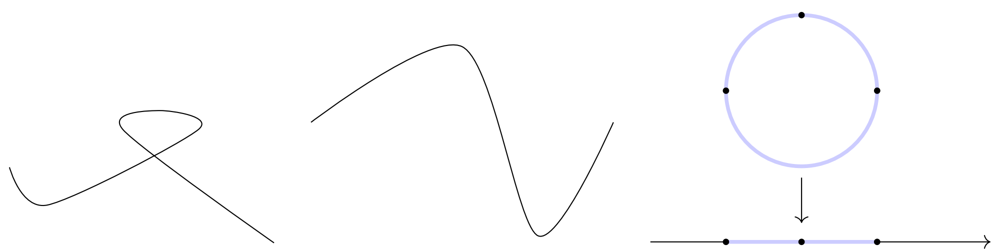

Definition A.1.6. Let M be an n-dimensional smooth manifold. A subset $N \subseteq M$ is a *submanifold* of M if the inclusion map $i\colon N \hookrightarrow M$ is an immersion (see Definition A.1.10).

Intuitively, this means we may allow N to self-intersect within M, but N cannot be "squished" in such a way that we lose information about the tangent space of a

[2]Our definition of manifold excludes many objects that we might naturally wish to study, particularly spaces with "edges" such as the closed ball $D^n \subseteq \mathbb{R}^n$. To extend our studies to such objects, we allow charts to map U to an open subset of the half-space

$$\mathbb{H}^n = \{x \in \mathbb{R}^n \mid x_n \geq 0\}, \text{ with } \partial\mathbb{H} = \{x \in \mathbb{H}^n \mid x_n = 0\}.$$

The *boundary* of M, denored ∂M, is the collection of all points in $\phi^{-1}(\partial\mathbb{H})$ for some ϕ. If $\partial M \neq \emptyset$, then we call M a *manifold with boundary*, otherwise M is called a *manifold without boundary.*

[3]An *algebra* over a field is a vector space over that field along with a product operation that plays nice with addition and scaling (see, for example, [Mor01, §1.3]).

point. Aside from the definition above, there are other equivalent characterizations of a submanifold (see, for example, [AD14, Theorem A.1.1]). Somewhat surprisingly, every compact manifold can be embedded in Euclidean space of a sufficiently large dimension (see [Hir76, Theorem 3.5]).

Theorem A.1.7 (Whitney Embedding Theorem)**.** *Every n-dimensional compact manifold M can be embedded as a submanifold in $\mathbb{R}^{2n+1}$.*

A.1.2 Tangent Vectors

In Euclidean space, a vector $v = (v_1, \dots v_n)$ at $p \in \mathbb{R}^n$ can be thought of as an operator on differential functions. Given a function $f : \mathbb{R}^n \to \mathbb{R}$ that is differentiable on a neighborhood of p, $v\colon f \mapsto v(f)$ "returns" the directional derivative of f along v at p. That is,

$$v(f) = \sum_{i=1}^{n} v_i \left.\frac{\partial f}{\partial x_i}\right|_p \in \mathbb{R},$$

where x_i is understood to be the i^{th} standard coordinate vector. This operator satisfies two properties:

(i) *v is linear: $v(f + \lambda g) = v(f) + \lambda v(g)$,*

(ii) *v is a derivation: $v(fg) = f(p)v(g) + g(p)v(f)$,*

for all $f, g \in C^\infty(\mathbb{R}^n)$ and $\lambda \in \mathbb{R}$. It is with these properties in mind that we define tangent vectors on a manifold, where a tangent vector should be independent of a choice of particular chart.

Definition A.1.8. Given a smooth manifold M and a point $p \in M$, a *tangent vector* to M at p is any map $v_p : C^\infty(M) \to \mathbb{R}$ that satisfies conditions (i) and (ii) above. The *tangent space* at p is just the collection of all such v_p, and is denoted T_pM.

By defining $(v_p+w_p)(f) = v_p(f)+w_p(f)$ and $(\lambda v_p)(f) = \lambda(v_p(f))$ for $v_p, w_p \in T_pM$ and $\lambda \in \mathbb{R}$, we give T_pM the structure of a real vector space.

There is another perspective from which we can view tangent vectors, namely as an equivalence class of curves, which is related to the more intuitive notion of a tangent vector as a velocity vector of a curve. From this point of view, a tangent vector at p is the collection of curves $c\colon (-\varepsilon, \varepsilon) \to M$ with $c(0) = p$ under the relation $c_1 \sim c_2$ if for a chart ϕ centered at p, we have

$$(\phi \circ c_1)'(0) = (\phi \circ c_2)'(0).$$

Given such a class of curves, we can associate it to the derivation $v \in T_pM$ with

$$v(f) = \frac{d}{dt} f(c(t))|_{t=0},$$

where c is a representative of the equivalence class.

Definition A.1.9. The *tangent map at* p is the map $T_p f\colon T_pM \to T_{f(p)}N$ associated with $f\colon M \to N$, and sends the equivalence class of a curve γ to that of the curve $f \circ \gamma$. The *differential at* p is the special case when $N = \mathbb{R}$, and is denoted by $(df)_p$.

From the perspective of derivations, the tangent map is the linear map $f_*\colon v \mapsto f_*v$ where f_*v is the derivation that maps $h \mapsto v(h \circ f)$ for $h \in C^\infty(M)$. The tangent map can also reveal information about the behavior of f.

Definition A.1.10. Consider a smooth map $f\colon M \to N$.

(i) We call f an *immersion* if $T_pf\colon T_pM \to T_pN$ is injective for all $p \in M$,

(ii) We call f an *embedding* if f is an immersion and also a homeomorphism onto its image,

(iii) We call f a *submersion* if f is surjective and T_pf is surjective for all $p \in M$.

A point $q \in N$ is called a *regular value* of f if f is submersive— meaning the tangent map T_pf is surjective— for all $p \in f^{-1}(q)$. The following theorem, which will be generalized in Proposition 1.2.8, says that the preimage of such points is in fact a submanifold of M.

Theorem A.1.11 (Regular Value Theorem)**.** *If* $q \in N$ *is a regular value of* $f\colon M \to N$*, then* $f^{-1}(q)$ *is a submanifold of* M *of dimension* $\dim(M) - \dim(N)$*.*

A.1.3 Vector Bundles

A vector bundle is a formal way of thinking of a family of vector spaces as being parameterized by a base space B, and is an example of a more general structure known as a fiber bundle.

Definition A.1.12. A *fiber bundle* $F \to E \xrightarrow{\pi} B$ consists of a *total space* E, a *base space* B, a *fiber* F, and a projection map $\pi\colon E \to B$ such that for all $p \in B$ there is a neighborhood U and a homeomorphism $\phi\colon \pi^{-1}(U) \to U \times F$ that make the following diagram commute:

$$\begin{array}{ccc} \pi^{-1}(U) & \xrightarrow{\ \phi\ } & U \times F \\ & \searrow^{\pi} & \downarrow \\ & & U \end{array}$$

The unlabeled arrow indicates projection onto the first coordinate. Note that commutivity of the diagram implies that each fiber $F_p := \pi^{-1}(p)$ is homeomorphic to $\{p\} \times F$ via ϕ, and so the map ϕ is called a *local trivialization.*

A fiber bundle is determined by the map $\pi\colon E \to B$, but may be denoted by (E, B, π), the short exact sequence notation used above, or even just the total space E. When E is isomorphic to $B \times F$, the fiber bundle is *trivial.*

Definition A.1.13. A *(differentiable) vector bundle* of rank n is a fiber bundle such that $\pi\colon E \to B$ is a differentiable map between differentiable manifolds, each fiber F_p is an n-dimensional vector space over $\mathbb{R}$, and each local trivialization ϕ is a diffeomorphism.

The word 'differentiable' is often omitted when clear from context. More intuitively, we can think of a vector bundle as a family of vector spaces paramatrized by the manifold B. The following vector bundles are two prime motivational examples from differential topology.

Definition A.1.14. The *tangent bundle* of M is the vector bundle TM whose fibres are the tangent spaces T_pM. The natural projection map $\pi\colon TM \to M$ sends a vector $v_p \in T_pM$ to the point p.

The tangent bundle is equipped with a natural topology and smooth structure. As a set,

$$TM = \bigcup_{p\in M} T_pM = \{(p,v) \mid p \in M, v \in T_pM\}.$$

A local coordinate system (U,ϕ) on M induces a chart $(TU, T\phi)$ where

$$T\phi\colon TU \to \phi(U) \times \mathbb{R}^n$$

maps $(p,u) \in TU$ to $(\phi(p), d\phi_p(v)) \in \mathbb{R}^{2n}$. This notation can be a bit confusing at first, so we will take a second to unpack what is going on. If $\phi = (x_1, \dots x_n)$ is a local coordinate map, we can write a tangent vector $v \in T_pM$ as $\sum_{i=1}^n (dx_i)_p(v) \left.\frac{\partial}{\partial x_i}\right|_p$, where each $(dx_i)_p(v) =: v_i \in \mathbb{R}$ is like "the component of v pointing in the x_i direction." The $d\phi_p(v)$ is meant to indicate the n-tuple $(v_1, \dots, v_n)$. It is straightforward to show that the induced $T\phi$ is bijective. The topology on TM is given by the collection $\{(\pi^{-1}(U), T\phi) \mid (U,\phi)\}$, after declaring each $\pi^{-1}(U)$ to be open and each $T\phi$ to be a homeomorphism. All in all, if M is an n-dimensional manifold, then TM is a $2n$-dimensional manifold.

Definition A.1.15. The vector space dual to T_pM is called the *contangent space* of M at p and is denoted by T_p^*M; its elements are called *cotangent vectors*. The vector bundle over M whose fibers are the contangent spaces is called the *cotangent bundle* of M, denoted by T^*M. The projection map $\pi\colon T^*M \to M$ sends T_p^*M to p.

Like the tangent bundle, the cotangent bundle admits the structure of a $2n$-dimensional manifold. As a set,

$$T^*M = \bigcup_{p\in M} T_p^*M = \{(p,\xi) \mid p \in M,\ \xi \in T_p^*M\},$$

and a chart (U,ϕ) on M induces the natural chart on T^*M via the map

$$T^*\phi\colon T^*U \to \phi(U) \times (\mathbb{R}^n)^*$$

that sends $\xi \in T_p^*M$ to $(\phi(p), \xi d\phi_p^{-1})$. As with the tangent bundle, the topology is given by these induced charts, giving the cotangent bundle the structure of a $2n$-dimensional manifold.

A.1.4 Vector Fields

A vector field can be thought of as an assignment of a tangent vector at each point of M. A more formal definition approaches vector fields from the perspective of vector bundles.

Definition A.1.16. A *vector field* is a section of the tangent bundle TM.

Recall that a *section* of $\pi\colon TM \to M$ is a (differentiable) map $s\colon M \to TM$ such that $\pi \circ s = \mathrm{id}_M$. The collection of all vector fields on M is a module both over $\mathbb{R}$ and over $C^\infty(M)$.

Remark A.1.17. Dually, a 1-*form* is a section of the cotangent bundle T^*M, that is, a map $s\colon M \to T^*M$ such that $\pi \circ s = \mathrm{id}_M$ (where now π is the natural projection $T^*M \to M$). The *differential*

$$df\colon M \to T^*M$$

which sends $p \mapsto (df)_p$ is a 1-form, since the natural projection sends $(df)_p \mapsto p$. The differential also gives an embedding of M into T^*M (cf. [AD14, Exercise 2]).

Definition A.1.18. The *directional derivative* of f along a vector field V is $Vf = V(f)\colon M \to \mathbb{R}$, which is the function $(Vf)(p) = V_p(f)$.

The directional derivative at p can be thought of as representing the instantaneous rate of change of f, moving through p in the direction of V. The function Vf is not to be confused with fV, which denotes the vector field with $(fV)_p = f(p)V_p$.

Definition A.1.19. The *Lie bracket* for vector fields V, W is the vector field correponding ot the derivation $f \mapsto V(Wf) - W(Vf)$, denoted $[V, W]$. That is,

$$[V, W]_p(f) = V_p(Wf) - W_p(Vf).$$

The bracket satisfies several key properties:

(i) $\mathbb{R}$-*bilinearity*: $[aV + b\tilde{V}, W] = a[V, W] + b[\tilde{V}, W]$ for all $a, b \in \mathbb{R}$; similarly for the second slot;

(ii) *anti-symmetry*: $[V, W] = -[W, V]$;

(iii) *Jacobi identity*: $[[V, \tilde{V}], W] + [[\tilde{V}, W], V] + [[W, V], \tilde{V}] = 0$.

Definition A.1.20. A *Lie algebra* is a vector space with an operation $[\ ,\ ]$ satisfying the three characterizing properties above.

The collection of vector fields on M under the Lie bracket forms a Lie algebra, and is naturally isomorphic to the set of all derivations of $C^\infty(M)$ as Lie algebras.

A.1.5 Riemannian Metrics

One of the most crucial vector fields that we will discuss is the gradient of a function $f\colon M \to \mathbb{R}$. In Euclidean coordinates, the gradient at a point p is dual to the Jacobian at p,

$$\nabla_p f = \left[\frac{\partial f}{\partial x_1}(p), \ldots, \frac{\partial f}{\partial x_n}(p)\right].$$

The vector field ∇f can be thought of as pointing in the direction of "steepest increase," perpendicular to the level sets of f. In the more general setting of manifolds, we need to do a little bit more work to develop a similar type of vector field. In particular, the gradient will depend on a choice of Riemannian metric, which will allow us to make sense of things like angle and distance on a manifold. The idea is that we equip each tangent space T_pM with an inner product g_p, so that the inner products vary smoothly as p varies across M.

Definition A.1.21. A *Riemannian structure* is a "smoothly varying" association $g\colon p \mapsto g_p$ of a point p to a symmetric, positive-definite bilinear form

$$g_p\colon T_pM \times T_pM \to \mathbb{R}.$$

Here, "smoothly varying" means that for any pair of smooth vector fields V, W, the function $p \mapsto g_p(V_p, W_p)$ is smooth.[4] Such a g is called *Riemannian metric* (or *metric tensor*), and a manifold M with such a metric is called a *Riemannian manifold.*

It is well-known that every smooth manifold admits a Riemannian structure (cf. [Mor01, §4.1]). Since g_p is an inner product on T_pM, we often use the notation $\langle V, W\rangle_p$ instead of $g_p(V, W)$. Strictly speaking, a Riemannian metric g is not a metric on M in the usual sense, however g does induce a genuine metric on M.

Definition A.1.22. Let $\gamma\colon [a, b] \to M$ be a piece-wise smooth curve. The *length* of γ is

$$L(\gamma) = \int_a^b \sqrt{\langle \gamma'(t), \gamma'(t)\rangle_{\gamma(t)}}\,dt.$$

We therefore have a well-defined notion of distance, via the *metric induced by g*,

$$d_g(p, q) = \inf\{L(\gamma)\},$$

where the infimum is taken over piece-wise smooth curves $\gamma\colon [a, b] \to M$ from p to q, that is γ such that $\gamma(a) = p, \gamma(b) = q$.

We will often drop the specific reference to the metric from our notation. Equipped with this distance function, our manifold carries the structure of a metric space. Furthermore, on a connected manifold M, the metric topology on M coincides with the original manifold topology (cf. [Lee71, Proposition 8.19]).

[4]More formally, we require g to be a smooth section of the tensor product $T^*M \otimes T^*M$.

A.1.6 Flow

Let V be a vector field on M. A curve $c\colon (a,b) \to M$ is an *integral curve* of V if $V_{c(t)} = \frac{dc}{ds}(t)$ for all $t \in (a,b)$. The vector $\frac{dc}{ds}(t) \in T_{c(t)}M$ is called the *velocity vector* and is sometimes denoted $\dot{c}(t)$. Without loss of generality, we may assume $0 \in (a,b)$. For every point $p \in M$, there is a uniquely determined smooth curve

$$c_p\colon (a,b) \to M, \quad \text{with} \quad c_p(0) = p \quad \text{and} \quad \frac{dc_p}{ds}(t) = V_{c_p(t)}.$$

The existence and uniqueness of this curve is a consequence of the existence and uniqueness theorem for ordinary differential equations. We can use these curves to describe the local flow of the vector field V, which describes how our buoyant grain will travel after time t after being placed at any point p on the manifold. However, this local flow will only "make sense" for certain t, depending on the domain of the curve c_p.

Definition A.1.23. Let $M_t = \{p \in M \mid t \text{ is in the domain of } c_p\}$. The *local flow* of V at p is $\varphi_V^t : M_t \to M$ with $\varphi_V^t(p) = c_p(t)$.

If we start at the point p, and flow along the integral curve c_p, then $\varphi_V^t(p)$ describes the point at which we arrive after time t. Note that $\varphi_V^0 = \mathrm{id}_{M_t}$.

Remark A.1.24. When ϕ_V^t is defined for all $t \in \mathbb{R}$, we say that the vector field (or flow) is *complete*. In particular, any vector field on a compact manifold is complete (see [Mil63, p.10]).

When the vector field is complete, we have a *one-parameter group of diffeomorphisms* generated by V, which is a smooth map $\varphi_V : \mathbb{R} \times M \to M$ such that $\varphi_V(t,p) = \varphi_V^t(p)$, with $\varphi_V^{t+s} = \varphi_V^t \circ \varphi_V^s$. To see that φ_V^t is a diffeomorphism, observe

$$\varphi_V^t \circ \varphi_V^{-t} = \varphi_V^{-t} \circ \varphi_V^t = \varphi_V^0 = \mathrm{id}_M.$$

Alternatively, given a one-parameter group of diffeomorphisms $\varphi\colon \mathbb{R} \times M \to M$, we can define a vector field V on M. For every smooth $f\colon M \to \mathbb{R}$, let

$$V_p(f) = \lim_{h \to 0} \frac{f(\varphi(h,p)) - f(p)}{h}.$$

The assignment $V : p \mapsto V_p$ is a vector field.

A.2 A Bit of Algebraic Topology

This section provides only the most essential definitions in homotopy theory and algebraic homology. We recommend that the reader unfamiliar with these concepts to at least become comfortable with homotopy (Definition A.2.1), homotopy equivalence (Definition A.2.2), chain complexes (Definition A.2.6), and homology groups (Definition A.2.7), although they may skip the examples of singular homology and cellular homology which are presented as a supplement to Section 1.3.2. The reader who is interested in more than our pithy refresher of this material will be pleased to know that there are many excellent books to turn to, such as [Hat02, Mun61, May99].

A.2.1 Homotopy

Homotopy provides a way to think of spaces as being "similar" without requiring something as strict as a homeomorphism.

Definition A.2.1. A *homotopy* between continuous maps $f, g\colon X \to Y$ is a continuous map $H\colon X \times I \to Y$ such that $H(x,0) = f(x)$ and $H(x,1) = g(x)$, where $I = [0,1]$. We say that f is *homotopic* to g and write $f \simeq g$. For a fixed time $t \in I$, we write $H_t = H(-,t)$.

We can think of a homotopy as a one-parameter family of continuous maps $X \to Y$. Imagining the parameter as representing time, then the homotopy "deforms" one map into the other, continuously, as time goes from 0 to 1.

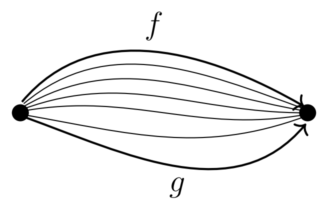

Definition A.2.2. A map $f\colon X \to Y$ is a *homotopy equivalence* if there is a map $g\colon Y \to X$ such that $gf \simeq \mathrm{id}_X$ and $fg \simeq \mathrm{id}_Y$. The spaces X, Y are said to have the same *homotopy type*, denoted $X \simeq Y$.

If $f \simeq c$ for some constant map c, we say that f is *nullhomotopic*. When the identity map on the space is nullhomotopic, the space is said to be *contractible*, meaning it has the same homotopy type as a point.

Definition A.2.3. Let $A \subseteq X$. A *retraction* of X onto A is a map $r\colon X \to X$ such that $r(X) = A$ and $r|_A = \mathrm{id}$. A *deformation retraction* is a homotopy from the identity on X to a retraction of X onto A.

In general, two spaces are homotopy equivalent if and only if there is a third space containing them both as deformation retracts. Homotopy equivalence provides a good criterion for the algebraic topologist to talk about "sameness" of topological spaces, without restricting to homeomorphic spaces. One such way to classify topological spaces is via their homotopy groups, which record information about the "holes" in the space.

Definition A.2.4. For a space X and basepoint $x_0 \in X$, define the n^{th} *homotopy group* $\pi_n(X, x_0)$ to be the set of homotopy classes of maps $f\colon (I^n, \partial I^n) \to (X, x_0)$ where homotopies f_t are required to satisfy $f_t(\partial I^n) = x_0$ for all t.

Here I^n denotes the n-product of I, with boundary ∂I^n consisting of points with at least one coordinate equal to 0 or 1. When $n = 0$, we define I^0 to be just a point and ∂I^0 to be empty. Equivalently, we could consider homotopy classes of maps from $S^n \to X$ that preserve a chosen base point. In this way, the homotopy groups describe the n-dimensional holes (the places we could fit an n-dimensional sphere) in X. For example, $\pi_1(X, x_0)$ records information about homotopies of loops in X (based at x_0); this group is called the *fundamental group* of X. These homotopy groups give us a slightly weaker notion of "sameness" between spaces.

Definition A.2.5. A map $f: X \to Y$ is a *weak homotopy equivalence* if it induces isomorphisms $\pi_n(X, x_0) \to \pi_n(Y, f(x_0))$ for $n \geq 0$ and all choice of basepoints x_0.

When the spaces X and Y are nice enough to have the homotopy type of a CW complex (see Definition A.2.11), the classic theorem of Whitehead tells us that the weak homotopy equivalence f is in fact a homotopy equivalence [Hat02, Theorem 4.5].

A.2.2 Homology

Homology, like homotopy, provides a way to think of spaces as "the same" by associating a sequence of algebraic objects to a topological space, the idea being that similar spaces will be associated with similar objects.

Definition A.2.6. A *chain complex* is a sequence C_* of homomorphisms between Abelian groups

$$\cdots \to C_{k+1} \xrightarrow{\partial_{k+1}} C_k \xrightarrow{\partial_k} C_{k-1} \to \cdots \to C_1 \xrightarrow{\partial_1} C_0 \xrightarrow{\partial_0} 0$$

such that $\partial_k \circ \partial_{k+1} = 0$ for each k. The ∂ maps are typically referred to as *boundary operators*, and in many cases the subscript is omitted.

The nilpotency of ∂ is equivalent to the requirement that $\operatorname{im} \partial_{k+1} \subseteq \ker \partial_k$. An element of the kernel of ∂ is called a *cycle* and an element of the image is called a *boundary*.

Definition A.2.7. The k^{th} *homology group* of a chain complex C_* is the quotient group

$$H_k(C_*) = \ker \partial_k / \operatorname{im} \partial_{k+1}$$

An element of $H_k(X)$ is a coset of $\operatorname{im} \partial_{k+1}$ called a *homology class*. Two cycles are *homologous* if their difference is a boundary and so belong to the same homology class. The chain complex is exact (meaning $\operatorname{im} \partial_{k+1} = \ker \partial_k$) if and only if $H_k(C_*) = 0$ for each k. Thus the homology groups measure the failure of exactness.

Definition A.2.8. A map $f: C_* \to D_*$ of chain complexes is a sequence of homomorphisms $f_k: C_k \to D_k$ such that the following diagram commutes for all k:

$$\begin{array}{ccc} C_k & \xrightarrow{f_k} & D_k \\ \partial_k \downarrow & & \downarrow \partial_k \\ C_{k-1} & \xrightarrow[f_{k-1}]{} & D_{k-1} \end{array}$$

A map of chain complexes induces a map between the homology groups, $H_k(f) := f_*: H_k(C_*) \to H_k(D_*)$.

There are many different types of homologies for a given space, depending on the assigned chain complex, and it is often enlightening to compare different homology theories. The first example we will look at, singular homology, is one of the simpler homology theories, being built on relatively concrete constructions.

Example A.2.9 (Singular homology)**.** Given a topological space X, a *singular n-simplex in X* is a continuous map $\sigma\colon |\Delta^n| \to X$ where $|\Delta^n| \subseteq \mathbb{R}^{n+1}$ is the standard topological n-simplex. The word 'singular' expresses the idea that $\mathrm{im}(\sigma)$ might have 'singularities' where the image does not look like a simplex.[5] So a singular 0-simplex is a map from one-point space into X, which we may identify with just a point of X, and a singular 1-simplex gives a path in X.

To define singular homology, we take $C_n(X)$ to be the free Abelian group on the set of all singular n-simplicies in X. This group is called the *singular chain group in dimension n*, and an element of $C_n(X)$ is called a *singular n-chain* in X. We define the boundary operator $\partial_n\colon C_n(X) \to C_{n-1}(X)$ by specifying its values on the basis elements: For any singular n-simplex $\sigma\colon |\Delta^n| \to X$, define the $(n-1)$-chain $\partial_n\sigma$ by

$$\partial_n\sigma = \sum_{i=0}^{n} (-1)^i \sigma \circ d_i,$$

where $d_i\colon |\Delta^{n-1}| \to |\Delta^n|$ is the *ith face map* which sends $|\Delta^{n-1}|$ to the face of $|\Delta^n|$ opposite the i^{th} vertex.[6] The boundary of a 0-chain is defined to be zero. After checking that $\partial^2 = 0$ (as is done in [Hat02, Chapter 2]), we obtain well-defined singular homology groups.

Singular homology is particularly nice because $f \simeq g$ implies that the induced homology maps will be equal, so any two homotopy equivalent spaces will yield the same singular homology groups. In the language of category theory, singular homology group defines a homotopy invariant covariant functor from the category of topological spaces to the category of (graded) Abelian groups. However, calculations using singular homology can be quite complicated, which motivates the development of other homology theories, such as cellular homology.

Example A.2.10 (Cellular homology)**.** This homology is defined for a certain kind of topological space called a CW complex. In order to define cellular homology, we must first understand CW complexes. Essentially, a CW complex is a type of space constructed from building blocks called *cells*. An *open (closed) n-cell* is a topological space that is homeomorphic to the n-dimensional open (closed) unit ball.

Definition A.2.11. A *CW complex* is a cell complex X built out of cells in the following manner: starting with a discrete set X^0 (whose points are 0-cells), inductively construct the *n-skeleton* X^n from X^{n-1} by attaching n-cells via attaching maps $\varphi_\alpha\colon S^{n-1} \to X^{n-1}$. We can write X_n as the pushout

[5] Placing appropriate restrictions on σ to get a nice embedding yields *simplicial homology*.

[6] The reader familiar with simplicial sets can note the connection here with Example 3.1.11.

$$\begin{array}{ccc} \coprod_\alpha S^{n-1}_\alpha & \xrightarrow{\coprod_\alpha \varphi_\alpha} & X^{n-1} \\ \big\downarrow & & \big\downarrow \\ \coprod_\alpha D^n_\alpha & \longrightarrow & X^n \end{array}$$

Here, S^n denotes the topological n-sphere, which can be seen as the boundary of the n-disk D^n. The diagram above just means that X^n is the quotient space $X^{n-1} \amalg \coprod_\alpha D^n_\alpha / \sim$ where $x \sim \varphi_\alpha(x)$ for $x \in S^{n-1}_\alpha = \partial D^n_\alpha$. If this inductive process stops at a finite stage, setting $X = X^n$ for some $n < \infty$, then the complex is said to have dimension n. Otherwise, setting $X = \bigcup_n X^n$, the dimension is said to be infinite.

Given a CW complex X, the *cellular complex* $K_*(X)$ is defined by taking $K_n(X)$ to be the free module over $\mathbb{Z}$ (or some other group, like $\mathbb{Z}/2\mathbb{Z}$,) generated by the n-cells of X. The differential is given by

$$\partial_n c = \sum_{c' \in K_{n-1}(X)} N(c, c')c'$$

where c and c' are cells of dimension n and $n-1$, respectively. The coefficient $N(c, c')$ is determined as follows: Let φ_c be the attachment map of the cell c, and consider the composition

$$S^{n-1} \xrightarrow{\varphi_c|_{S^{n-1}}} X^{n-1} \xrightarrow{\psi_{c'}} S^{n-1}$$

where $\psi_{c'}$ is the map that sends $X^{n-1} \setminus c'$ to a single point. Then $N(c, c')$ is the degree of this composition, meaning the integer d so that the induced map $(\psi_{c'} \circ \varphi_c|_{S^{n-1}})(\alpha) \simeq d\alpha$ (see the beginning of [Hat02, §2.2]). It can be shown that $\partial^2 = 0$ and that the homology of this complex does not depend on the cellular decomposition.

Given two homology theories, it can be enlightening to compare them to each other. It is well known that the two examples we have presented are isomorphic (a proof is given in [Hat02, Theorem 2.35]). Consequently, we are able to work with whichever theory is better suited to the task at hand, without worry of losing valuable topological information.

A.3 A Bit of Category Theory

The following section covers some basic ideas in category theory that will make this thesis easier to understand. The first subsection covers the most essential definitions, such as categories, functors, and natural transformations. The second subsection discusses the Yoneda lemma and embedding that is referenced in Section 3.1.1. We then move to limits and colimits, which are used in various places in Chapter 3. Finally, we define adjunctions, which are not dicussed very much in this thesis, but are still good to know about. The reader who is in a hurry could likely get away with only reading Appendix A.3.1 and skimming Appendix A.3.3. Since this exposition is meant to be supplementary, we do not offer much detail beyond the most basic definitions; the interested reader should look to [Rie16] or the more classic references [ML71, May99].

A.3.1 Basic Notions

The most basic notion is a category, which is specified by a collection of objects and morphisms between them.

Definition A.3.1. A *category* $\mathscr{C}$ is a collection of objects $X, Y, Z, \ldots$, denoted $\operatorname{Ob}\mathscr{C}$, and a collection of morphisms $f, g, h, \ldots$, denoted $\operatorname{Mor}\mathscr{C}$, such that

- Each morphism has a *domain* (or *source*) and a *codomain* (or *target*), which are objects of $\mathscr{C}$. If a morphism f has domain X and codomain Y, we use the notation $f\colon X \to Y$ or $X \xrightarrow{f} Y$. We can assemble all morphisms with domain X and codomain Y into the collection $\mathscr{C}(X,Y)$, sometimes written $\operatorname{Mor}(X,Y)$.
- Each object X in $\mathscr{C}$ has an *identity morphism*, $\mathrm{id}_X\colon X \to X$.
- Any pair of morphisms f, g with $f\colon X \to Y$ and $g\colon Y \to Z$ has a *composite morphism* $gf\colon X \to Z$. We require that $\mathrm{id}_Y f = f = f\mathrm{id}_X$ and that composition is associative, so for any morphism $h\colon Z \to W$, we have $h(gf) = (hg)f$, which we denote by hgf.

We will oftentimes write '$X \in \mathscr{C}$' or '$f \in \mathscr{C}$' when it is clear from context whether we are referring to objects or morphisms in $\mathscr{C}$. By an abuse of notation, we use the symbol '$\in$' (typically reserved for set-membership) without requiring that $\mathscr{C}$, $\operatorname{Ob}\mathscr{C}$, or $\operatorname{Mor}\mathscr{C}$ are sets.

If $\operatorname{Ob}\mathscr{C}$ and $\operatorname{Mor}\mathscr{C}$ are in fact sets, we say that the category is *small*. Alternatively, when $\mathscr{C}(X,Y)$ is a set for each $X, Y \in \mathscr{C}$, we say $\mathscr{C}$ is *locally small*. These $\mathscr{C}(X,Y)$ are called *homsets*. A *concrete* category is one whose objects have underlying sets and whose morphisms are so-called 'structure-preserving' function between these sets.

Example A.3.2 (Examples of categories)**.** Many familiar classes of mathematical objects can be assembled into a category. We list a few key examples here that are topical to this thesis.

Set	sets and functions,
Top	topological spaces and continuous maps,
Diff	smooth manifolds and smooth maps,
Htpy	topological spaces and homotopy classes of continuous maps,

In any category, a morphism $f\colon X \to Y$ is an *isomorphism* $\cong$ if there is another morphism $g\colon Y \to X$ that is a two-sided inverse for f, meaning that $fg = \mathrm{id}_Y$ and $gf = \mathrm{id}_X$. Isomorphisms in a particular category might be given a different name, for instance 'bijection' in **Set** or 'homeomorphism' in **Top**. In general, isomorphisms are the strongest type of equivalence (other than identity) in the category. A category is a *groupoid* if every morphism is an isomorphism and *discrete* if every morphism is the identity.

Definition A.3.3. Given a category $\mathscr{C}$, we define its *opposite category* $\mathscr{C}^{\mathrm{op}}$ to be the category with the same objects as $\mathscr{C}$ but with reversed morphisms. That is, we have $X \xrightarrow{f} Y \in \mathscr{C}$ if and only if $Y \xrightarrow{f^{\text{op}}} X \in \mathscr{C}^{\text{op}}$. The identity of X in $\mathscr{C}^{\text{op}}$ is the same identity map id_X (with the arrow reversed), and composition is given by $(gf)^{\text{op}} = f^{\text{op}} g^{\text{op}}$.

This process of 'flipping the arrows' creates a duality between a category and its opposite— we can learn about one by examining the other. To see an example of this phenomenon, we can turn to certain 'special' objects in a category, known as the initial and terminal objects. An object $X \in \mathscr{C}$ is *initial* if for every $Y \in \mathscr{C}$ there is a unique morphism $X \to Y$, and *terminal* if there is a unique morphism $Y \to X$.[7] Thus an object is initial in $\mathscr{C}$ if and only if it is terminal in $\mathscr{C}^{\text{op}}$. In this spirit, any theorem proved about an arbitrary category $\mathscr{C}$ can also be interpreted in the dual context, where the arrows in the argument are reversed, thus yielding a *dual theorem*. In the words of Riehl, "The result is a two-for-one deal: any proof in category theory simultaneously proves two theorems, the original statement and its dual" [Rie16, p.10]. Mathematicians are fond of calling the dual notion of some concept its *co*concept (see, for instance, limits and colimits in Appendix A.3.3 or fibrations and cofibrations in Section 3.2.1). This affinity can lend some light-heartedness to abstract mathematics, such as the infamous pants and copants of cobordism theory, or the joke that every nut is a coconut.

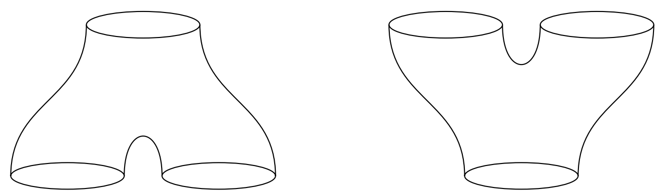

pants copants

Having defined and explored various categories, it is only fitting that we next investigate the morphisms between them.

Definition A.3.4. A morphism between categories is called a *functor*. More specifically, a (covariant) functor $F\colon \mathscr{C} \to \mathscr{D}$ is specified by

- an object $FX \in \text{Ob}\,\mathscr{D}$ for every $X \in \text{Ob}\,\mathscr{C}$,
- a morphism $Ff \in \mathscr{D}(FX, FY)$ for every morphism $f \in \mathscr{C}(X, Y)$,

such that the *functoriality axioms* are satisfied: (i) $F(gf) = Fg \circ Ff$ for every composable pair f, g in $\mathscr{C}$, and (ii) $F\text{id}_X = \text{id}_{FX}$ for every $X \in \text{Ob}\,\mathscr{C}$. Similarly, a *contravariant functor* from $\mathscr{C}$ to $\mathscr{D}$ is a morphism $F\colon \mathscr{C}^{\text{op}} \to \mathscr{D}$ specified by

- an object $FX \in \text{Ob}\,\mathscr{D}$ for every object $X \in \mathscr{C}$,
- a morphism $Ff \in \mathscr{D}(FY, FX)$ for every morphism $f \in \mathscr{C}(X, Y)$,

that satisfy (i) $F(gf) = Ff \circ Fg$ for every composable pair f, g in $\mathscr{C}$, and (ii) $F\text{id}_X = \text{id}_{FX}$ for every $X \in \text{Ob}\,\mathscr{C}$.

[7] An object that is both initial and terminal is called a *zero object*.

Given a functor F between (locally) small categories, we say F is *full* if (for every X, Y in the domain) the map $\mathscr{C}(X,Y) \to \mathscr{D}(FX, FY)$ is surjective and *faithful* if the map is injective. Since functors preserve the structure of categories (such as isomorphisms and commutative diagrams), we get the category of categories **Cat** whose objects are small categories and whose morphisms are functors.

Example A.3.5 (Examples of functors)**.** There are many mathematical notions that naturally exhibit the structure of a functor, and we present a few examples.

- The *constant functor* on an object $X \in \mathscr{C}$ is the functor $c_X\colon \mathscr{C} \to \mathscr{C}$ that takes every object to X and every morphism to id_X.
- There is a type of functor known as a "forgetful functor" that lands in **Set** and 'forgets' structure. For example, the forgetful functor **Top** → **Set** sends a space to its set of points and a continuous map to its underlying function.
- There is a functor **Top** → **Htpy** that acts as the identity on objects and sends a continuous morphism to its homotopy class,
- A simplicial set is a contravariant functor $\mathbf{\Delta} \to \mathbf{Set}$, as treated in Section 3.1.1.

Definition A.3.6. A *natural transformation* is a functor between functors. That is, if $F\colon \mathscr{C} \to \mathscr{D}$ and $G\colon \mathscr{C} \to \mathscr{D}$ are functors, $N\colon F \to G$ assigns to each object $X \in \mathrm{Ob}\,\mathscr{C}$ a morphism $N_X \in \mathrm{Mor}\,\mathscr{D}$ such that the following diagram commutes:

$$\begin{array}{ccc} F(X) & \xrightarrow{N_X} & G(X) \\ {\scriptstyle Ff}\downarrow & & \downarrow{\scriptstyle Gf} \\ F(Y) & \xrightarrow[N_Y]{} & G(Y) \end{array}$$

The map N_X is called the *component* of N at X. A *natural isomorphism* is a natural transformation whose components are all isomorphisms.

Example A.3.7. If $\mathscr{C}$ has an initial object $\emptyset$ (meaning there is a unique map $\emptyset \to X$ for every $X \in \mathscr{C}$), then there is a natural transformation between the constant functor $c_\emptyset$ and $\mathrm{id}_\mathscr{C}$. Each component N_X of this natural transformation is given by the unique map $\emptyset \to X$.

We write $\mathscr{D}^\mathscr{C}$ for the category of functors $\mathscr{C} \to \mathscr{D}$ and natural transformations between them, called a *functor category.* We can also use natural transformations to develop a notion of when categories are "essentially the same," or *equivalent.*

Definition A.3.8. An *equivalence of categories* consists of functors $F\colon \mathscr{C} \rightleftarrows \mathscr{D} : G$ together with natural isomorphisms $\eta\colon \mathrm{id}_\mathscr{C} \to GF$ and $\epsilon\colon FG \to \mathrm{id}_\mathscr{D}$.

The reader may note that this definition of categorical equivalence is strikingly similar to that of homotopy equivalence (Definition A.2.2), where homotopies are now replaced with natural isomorphisms.

A.3.2 The Yoneda Lemma

The Yoneda Lemma (and embedding) provides a way to understand set-valued functors in terms of a certain class of functors, known as representable functors. For the sake of brevity, we merely state the lemma and subsequent embedding, without providing proofs. The interested reader may look to [ML71, §III.2] or [Rie16, §2.2] for further explanation. We assume that $\mathscr{C}$ is a locally small category throughout.

Definition A.3.9. For any object $c \in \mathscr{C}$, we define two functors $\mathscr{C} \to \mathbf{Set}$, called the *functors represented by c*. A functor $\mathscr{C} \to \mathbf{Set}$ is called *representable* if it can be represented by some object $c \in \mathscr{C}$. The *covariant functor represented by c*, $\mathscr{C}(c, -)$, maps an object X to the homset $\mathscr{C}(c, X)$. A morphism $f\colon X \to Y$ is sent to the post-composition $f_*\colon \mathscr{C}(c, X) \to \mathscr{C}(c, Y)$. The *contravariant functor represented by c*, $\mathscr{C}(-, c)$, takes X to $\mathscr{C}(X, c)$ and f to the pre-composition $f^*\colon \mathscr{C}(Y, c) \to \mathscr{C}(X, c)$. Diagramically, we have

$$\begin{array}{ccc} X & \xmapsto{\mathscr{C}(c,-)} & \mathscr{C}(c,X) \\ {\scriptstyle f}\big\downarrow & \longmapsto & \big\downarrow{\scriptstyle f_*} \\ Y & \longmapsto & \mathscr{C}(c,Y) \end{array} \qquad \begin{array}{ccc} X & \xmapsto{\mathscr{C}(-,c)} & \mathscr{C}(X,c) \\ {\scriptstyle f}\big\downarrow & \longmapsto & \big\uparrow{\scriptstyle f^*} \\ Y & \longmapsto & \mathscr{C}(Y,c) \end{array}$$

In general, we use f_* and f^* to denote post- and pre-composition by f, respectively. Note that post-composition defines a covariant action on homsets, while pre-composition defines a contravariant one.

Theorem A.3.10 (The Yoneda Lemma)**.** *For any functor $F\colon \mathscr{C} \to \mathbf{Set}$ and any object $c \in \mathscr{C}$, there is a bijection*

$$\mathbf{Cat}(\mathscr{C}(c, -), F) \cong Fc$$

which maps a natural transformation $N\colon \mathscr{C}(c, -) \to F$ to $N_c \mathrm{id}_c$, the image of the identity under the component N_c.

Applying the Yoneda Lemma allows us to characterize natural transformations between representable functors, via the Yoneda embeddings.

Corollary A.3.11 (Yoneda embeddings)**.** *The covariant and contravariant Yoneda functors given by*

$$\mathscr{C} \xhookrightarrow{\;y\;} \mathbf{Set}^{\mathscr{C}^{\mathrm{op}}} \qquad\qquad \mathscr{C}^{\mathrm{op}} \xhookrightarrow{\;y\;} \mathbf{Set}^{\mathscr{C}}$$

$$\begin{array}{ccc} c & \longmapsto & \mathscr{C}(-,c) \\ {\scriptstyle f}\big\downarrow & \longmapsto & \big\downarrow{\scriptstyle f_*} \\ d & \longmapsto & \mathscr{C}(-,d) \end{array} \qquad \begin{array}{ccc} c & \longmapsto & \mathscr{C}(c,-) \\ {\scriptstyle f}\big\downarrow & \longmapsto & \big\uparrow{\scriptstyle f^*} \\ d & \longmapsto & \mathscr{C}(d,-) \end{array}$$

define full and faithful embeddings.

This says that any locally small category is isomorphic to the subcategory of $\mathbf{Set}^{\mathscr{C}^{\mathrm{op}}}$ spanned by the contravariant represented functors, and the dual statement holds of $\mathscr{C}^{\mathrm{op}}$.

A.3.3 Limits and Colimits

Limits and colimits can be found in any category, and include under their umbrella familiar constructions like the infimum and supremum, cartesian products, direct sums, kernels, cokernels, and unions, among others. We will only brush the surface of this rich area of category theory, providing a basic definition of a categorical (co)limit, as is used to define the realization functor in Section 3.1.2 and the homotopy pullback in Section 3.2.2. We first introduce some basic terminology with which to build these definitions.

A diagram in a category $\mathscr{C}$ is typically thought of as a directed graph of morphisms in $\mathscr{C}$, but we can make this notion more precise. Specifically, a *diagram* of shape (or type) J in $\mathscr{C}$ is a functor $D\colon J \to \mathscr{C}$, where the domain is a small category. The domain J is called the *indexing category* (or *scheme*) of the diagram D. For example, for any category J and object $c \in \mathscr{C}$, the *constant diagram* $\underline{c}$ is the functor $\underline{c}\colon J \to \mathscr{C}$ that sends every object of J to c and all morphisms in J to id_c. A *morphism of diagrams* is just a natural transformation in the functor category $\mathscr{C}^J$.

Definition A.3.12. A *cone over* a diagram $D\colon J \to \mathscr{C}$ with *apex* (or *summit* or *vertex*) $c \in \mathscr{C}$ is a natural transformation $\underline{c} \to D$. The components of this natural transformation are called the *legs* of the cone.

A cone over D can be specified by a collection of morphisms $\lambda_j\colon \underline{c} \to F_j$, indexed by the objects $j \in J$ such that the following triangle commutes in $\mathscr{C}$

$$\begin{array}{ccc} & c & \\ {\scriptstyle\lambda_i}\swarrow & & \searrow{\scriptstyle\lambda_j} \\ Di & \xrightarrow[Df]{} & Dj \end{array},$$

whenever $f\colon i \to j$ in J. For example, if D is a diagram indexed by $J = [n]$, the poset category with $n+1$ objects $0, 1, \ldots, n$ and n non-identity generating morphisms ($i \to j$ just in case $i \leq j$), then the cone over D with apex c looks like

$$\begin{array}{ccccccccc} & & & & c & & & & \\ & {\scriptstyle\lambda_0}\swarrow & & \swarrow{\scriptstyle\lambda_1} & \ldots & \searrow{\scriptstyle\lambda_{n-1}} & & \searrow{\scriptstyle\lambda_n} & \\ D0 & \longrightarrow & D1 & \longrightarrow & \ldots & \longrightarrow & D(n-1) & \longrightarrow & Dn \end{array}.$$

Definition A.3.13. The dual notion to the cone over D is called the *cone under* (or *cocone*), which is given by flipping the arrows in the definition of the cone over D. That is, a *cone under* D with *nadir* c is a natural transformation $D \to \underline{c}$.

In this case, the cone is specified by morphisms λ_j with domain F_j and codomain the constant diagram of c, again indexed by $j \in J$, such that the following triangle commutes:

$$\begin{tikzcd} Di \arrow[rr, "Df"] \arrow[dr, "\lambda_i"'] & & Dj \arrow[dl, "\lambda_j"] \\ & c & \end{tikzcd}.$$

Returning to the example above, when $J = [n]$, the cone under D looks as we would expect:

$$\begin{tikzcd} D0 \arrow[r] \arrow[drr, "\lambda_0"'] & D1 \arrow[r] \arrow[dr, "\lambda_1"] & \dots \arrow[r] & D(n-1) \arrow[r] \arrow[dl, "\lambda_{n-1}"'] & Dn \arrow[dll, "\lambda_n"] \\ & & c & & \end{tikzcd}.$$

Put simply, the colimit of D is the universal cone under the diagram D, whereas the limit is the universal cone over D. Like many categorical notions, there are multiple perspectives from which to view limits and colimits. The definition we present here is taken from [May99, §6.2], but the interested reader can look to [Rie16, §3.1] for more detailed exposition.

Definition A.3.14. The *colimit* of a J-shaped diagram D, $\operatorname{colim} D$, is an object of $\mathscr{C}$ along with a morphism of diagrams $\iota\colon D \to \underline{\operatorname{colim} D}$ that is initial among all such morphisms. (Recall that here $\underline{\operatorname{colim} D}$ denotes the constant diagram on $\operatorname{colim} D$.) That is, any morphism of diagrams $\eta\colon D \to \underline{X}$ factors uniquely through ι. In terms of diagrams, for any map $f\colon i \to j$ in J, we have a commutative diagram

$$\begin{tikzcd} Di \arrow[rr, "Df"] \arrow[dr, "\iota"] \arrow[ddr, "\eta"'] & & Dj \arrow[dl, "\iota"'] \arrow[ddl, "\eta"] \\ & \operatorname{colim} D \arrow[d, "\exists!"'] & \\ & X & \end{tikzcd}.$$

Dually, the *limit* of D, $\lim D$, is an object of $\mathscr{C}$ with morphisms $\pi\colon \underline{\lim D} \to D$ that is terminal among all such morphisms. That is, any morphism $\varepsilon\colon D \to \underline{X}$ factors uniquely through π and a map $X \to \lim D$ in $\mathscr{C}$, giving the commutative diagram

$$\begin{tikzcd} Di \arrow[rr, "Df"] & & Dj \\ & \lim D \arrow[ul, "\pi"'] \arrow[ur, "\pi"] & \\ & X \arrow[u, "\exists!"] \arrow[uul, "\varepsilon"] \arrow[uur, "\varepsilon"'] & \end{tikzcd}.$$

Many limits and colimits are given special names; for example, a *(co)equalizer* is the (co)limit of a diagram indexed by $\bullet \rightrightarrows \bullet$, a *pullback* is the limit of a diagram indexed by $\bullet \to \bullet \leftarrow \bullet$, and a *pushout* is the colimit of a diagram indexed by the dual $\bullet \leftarrow \bullet \to \bullet$. A category might have some colimits and limits, but not others. A category is called *(co)complete* if it has all (co)limits.

A.3.4 Adjunctions

An adjunction consists of a pair of functors standing in a particular relation to one another.

Definition A.3.15. An *adjunction* from $\mathscr{C}$ to $\mathscr{D}$ consists of two functors, $F\colon \mathscr{C} \to \mathscr{D}$ and $G\colon \mathscr{D} \to \mathscr{C}$, along with isomorphisms

$$\mathscr{D}(Fc, d) \cong \mathscr{C}(c, Gd)$$

for each $c \in \mathscr{C}$ and $d \in \mathscr{D}$ that are natural in both c and d. We say that F is *left adjoint* to G and G is *right adjoint* to F. The morphisms corresponding under the isomorphism are said to be *adjunct*.

Following convention, we use '$\vdash$' to indicate when a pair of functors are adjoint. The expressions $G \vdash F$ and $F \dashv G$ assert that $F\colon \mathscr{C} \to \mathscr{D}$ is left-adjoint to $G\colon \mathscr{D} \to \mathscr{C}$.

When $\mathscr{C}$ and $\mathscr{D}$ are locally small, the naturality condition amounts to the assertion that the isomorphisms assemble into a natural isomorphism

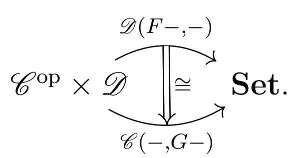

Here $\mathscr{C}^{\mathrm{op}} \times \mathscr{D}$ denotes the *product* of the two categories, whose objects are ordered pairs (c, d) for $c \in \mathscr{C}$ and $d \in \mathscr{D}$ and whose morphisms are ordered pairs $(f, g)\colon (c, d) \to (c', d')$ for $f\colon c \to c' \in \mathscr{C}$ and $g\colon d \to d' \in \mathscr{D}$. Composition and identities in the product are defined componentwise.

Intuitively, we can think of an adjunction as giving us a way to move between the two categories. An equivalent definition (cf. [Rie16, Proposition 4.2.6]) of adjoint functors says that two functors are adjoint if and only if there exist a pair of natural transformations $\mathrm{id}_{\mathscr{C}} \Rightarrow GF$ and $FG \Rightarrow \mathrm{id}_{\mathscr{D}}$ that satisfy certain identities given in [Rie16, Definition 4.2.5]. In this sense, we can think of adjoint functors as giving a weak form of an equivalence of categories.

References

[AB95] D.M. Austin and P.J. Braam. Morse-Bott theory and equivariant cohomology. In *The Floer Memorial Volume*, volume 133 of *Progress in Mathematics*, pages 123–183. Birkhauser, 1995.

[AD14] Michèle Audin and Mihai Damian. *Morse Theory and Floer Homology*. Universitext. Springer, London, 2014.

[Bae18] John Baez. A Quick Tour of Basic Concepts in Simplicial Homotopy Theory, 2018. `http://math.ucr.edu/home/baez/calgary/homotopy.html` (access date: Jan. 2020).

[Bar13] Clark Barwick. On the Q-construction for exact ∞-categories, 2013. arXiv:1301.4725 [math.KT].

[Ber09] Knut Berg. Edgewise subdivision and simple maps. Master's thesis, University of Oslo, Norway, 2009.

[BK72] A.K. Bousfield and Daniel M. Kan. *Homotopy Limits, Completions and Localizations*. Number 304 in Lecture Notes in Mathematics. Springer, 1972.

[BOO+20] Julia Bergner, Angélica Osorno, Viktoriya Ozornova, Martina Rovelli, and Claudia Scheimbauer. The edgewise subdivision criterion for 2-Segal objects. *Proceedings of the American Mathematical Society*, 148(1):71–82, 2020.

[Bot80] Raoul Bott. Marston Morse and his mathematical works. *Bulletin of the American Mathematical Society*, 3:907–950, 1980.

[Bot88] Raoul Bott. Morse theory indomitable. *Publications mathématiques de l'I.H.É.S*, 68:99–114, 1988.

[BR] Clark Barwick and John Rognes. On the Q-construction for exact ∞-categories. `https://www.maths.ed.ac.uk/~cbarwick/papers/qconstr.pdf` (access date: Jan. 2020).

[CIN06] Ralph Cohen, Kevin Iga, and Paul Norbury. Topics in Morse theory: Lecture notes, 2006.

[CJS95a] Ralph Cohen, John Jones, and Graeme Segal. Floer's infinite dimensional Morse theory and homotopy theory. In *The Floer Memorial Volume*, volume 133 of *Progress in Mathematics*, pages 297–325. Birkhauser, 1995.

[CJS95b] Ralph Cohen, John Jones, and Graeme Segal. Morse theory and classifying spaces, 1995. preprint, `http://math.stanford.edu/ralph/morse.ps`.

[Coh19] Ralph Cohen. Floer homotopy theory, revisited, 2019. arXiv.1901.08694v1 [math.AT].

[Coh20] Ralph Cohen. *Bundles, Homotopy, and Manifolds*. Unpublished draft, Jan. 2020.

[Dug08] Daniel Dugger. A primer on homotopy colimits, 2008. `https://pages.uoregon.edu/ddugger/hocolim.pdf` (access date: Feb. 2020).

[Flo89] Andreas Floer. Witten's complex and infinite dimensional Morse theory. *Journal of Differential Geometry*, 30:207–221, 1989.

[FOOO09] Kenji Fukaya, Yong-Geun Oh, Hiroshi Ohta, and Kaoru Ono. *Lagrangian Intersection Floer Theory: Anomaly and Obstruction*, volume 46 of *AMS/IP Studies in Advanced Mathematics*. American Mathematical Society and International Press of Boston, Inc., 2009.

[For02] Robin Foreman. A user's guide to discrete Morse theory. *Séminaire Lotharingien de Combinatoire*, 2002.

[Fre11] Daniel S. Freed. Commentary on "Lectures on Morse theory, old and new". *Bulletin of the American Mathematical Society*, 48:517–523, 2011.

[Fri11] Greg Friedman. An elementary illustrated introduction to simplicial sets, 2011. arXiv.0809.4221 [math.AT].

[GJ99] Paul G. Goerss and John F. Jardine. *Simplicial Homotopy Theory*, volume 174 of *Progress in Mathematics*. Birkhauser Verlag, Basel, 1999.

[GM88] Mark Goresky and Robert MacPherson. *Stratified Morse Theory*, volume 14 of *Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)]*. Springer Verlag, Berlin, 1988.

[Gue01] Martin Guest. Morse theory in the 1990's, 2001. arXiv:math/0104155 [math.DG].

[Hat02] Allen Hatcher. *Algebraic Topology*. Cambridge University Press, Cambridge, 2002.

[HD10] Brad Hannigan-Daley. The flow category of a Morse function, 2010. `http://individual.utoronto.ca/hannigandaley/flowcategory.pdf` (access date: Mar. 2020).

[Hir76] Morris W. Hirsch. *Differential Topology*. Number 33 in Graduate Texts in Mathematics. Springer, 1976.

[Hut02] Michael Hutchings. Lecture notes on Morse homology (with an eye towards Floer theory and pseudoholomorphic curves), 2002.

[Küh15] Wolfgang Kühnel. *Differential Geometry: Curves–Surfaces–Manifolds*, volume 77 of *Student Mathematical Library*. American Mathematical Society, Providence, RI, 3rd edition, 2015.

[Lai13] Johan Laine. Simplicial spaces and homotopy colimits, 2013. `http://web.math.ku.dk/~jg/students/laine.ms-project.2013.pdf` (access date: Mar. 2020).

[Lee71] John M. Lee. *Introduction to Topological Manifolds*. Number 202 in Graduate Texts in Mathematics. Springer, 2nd edition, 1971.

[LS11] Robert Lipshitz and Sucharit Sarkar. A Khovanov homotopy type, 2011. arXiv.1112.3932 [math.GT].

[Lur09] Jacob Lurie. *Higher Topos Theory*, volume 170 of *Annals of Mathematics Studies*. Princeton University Press, 2009. Accessed updated version at `https://www.math.ias.edu/~lurie/papers/HTT.pdf`, Mar. 2020.

[Mat02] Yukio Matsumoto. *An Introduction to Morse Theory*, volume 208 of *Translations of Mathematical Monographs*. American Mathematical Scoeity, Providence, RI, 2002.

[May67] J. P. May. *Simplicial Objects in Algebraic Topology*. Chicago lectures in mathematics. The University of Chicago Press, Chicago, IL, 1967.

[May72] J.P. May. *The geometry of iterated loop spaces*, volume 271 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin-New York, 1972.

[May74] J. P. May. E_∞ spaces, group completions, and permutive categories. *New Advancements in Topology*, (11):61–93, 1974. London Mathematical Society Lecture Notes.

[May99] J. P. May. *A Concise Course in Algebraic Topology*. Chicago lectures in mathematics. The University of Chicago Press, Chicago, IL, 1999.

[Mil56] John Milnor. On manifolds homeomorphic to the 7-sphere. *Annals of Mathematics*, 64:399–405, 1956.

[Mil63] John Milnor. *Morse Theory*. Number 51 in Annals of Mathematics Studies. Princeton University Press, Princeton, NJ, 1963.

[Mil65] John Milnor. *Topology from the Differentiable Viewpoint*. The University Press of Virginia, Charlottesville, VA, 1965.

[ML71] Saunders Mac Lane. *Categories for the Working Mathematician*. Springer, 1971.

[Mor01] Shigeyuki Morita. *Geometry of Differential Forms*, volume 201 of *Translations of Mathematical Monographs*. American Mathematical Society, Providence, RI, 2001.

[Mun61] James Munkres. *Elementary Differential Topology*. Number 54 in Annals of Mathematics Studies. Princeton University Press, Princeton, NJ, 1961.

[Mun18] James Munkres. *Topology*. Pearson, New York, NY, (2018 reissue) 2nd edition, 2018.

[NTT16] Vidit Nanda, Dai Tamaki, and Kohei Tanaka. Discrete Morse theory and classifying spaces, 2016. arXiv.1612.08429 [math.AT].

[Qin10] Lizhen Qin. On moduli spaces and cw structures arising from morse theory on hilbert manifolds. *Journal of Topology and Analysis*, 2:469–526, 2010.

[Qin11] Lizhen Qin. On the associativity of gluing, 2011. arXiv:1107.5527 [math.GT].

[rAHS04] Jiří Adámek, Horst Herrlich, and George E. Strecker. *Abstract and Concrete Categories: The Joy of Cats*. Online edition, 2004. `http://katmat.math.uni-bremen.de/acc/acc.pdf` (access date: Spring 2020).

[Rie08] Emily Riehl. A leisurely introduction to simplicial sets. 2008. `http://www.math.jhu.edu/~eriehl/ssets.pdf` (access date: Nov. 2019).

[Rie14] Emily Riehl. *Categorical Homotopy Theory*. New Mathematical Monographs. Cambridge University Press, 2014.

[Rie16] Emily Riehl. *Category Theory in Context*. Aurora: Modern Math Originals. Dover Publications, Chicago, IL, 2016.

[Rot10] T.O. Rot. Morse theory via the flow category. Master's thesis, Utrecht University, Netherlands, 2010.

[RV14] Emily Riehl and Dominic Verity. The theory and practice of Reedy categories. *Theory & Applications of Categories*, 9(29), 2014.

[Seg68] Graeme Segal. Classifying spaces and spectral sequences. *Publications Mathématiques de l'IHÉS*, 34:105–112, 1968.

[Seg73] Graeme Segal. Configuration-spaces and iterated loop spaces. *Inventiones math.*, 21:231–221, 1973.

[Seg74] Graeme Segal. Categories and cohomology theories. *Topology*, 13:293–312, 1974.

[Sma61] Stephan Smale. On gradient dynamical systems. *Annals of Mathematics*, 74:199–206, 1961.

[War71] Frank Warner. *Foundations of Differentiable Manifolds and Lie Groups*. Scott, Foresman and Company, Glenview, IL, 1971.

[Weh12] Katrin Wehrheim. Smooth structures on Morse trajectory spaces, featuring finite ends and associative gluing, 2012. arXiv:1205.0713 [math.DG].

[Wei05] Michael Weiss. What does the classifying space of a category classify? *Homology, Homotopy and Applications*, 7(1):185–195, 2005.