

Preface

As the title suggests, the goal of this book is to give the reader a taste of the “unreasonable effectiveness” of Morse theory. The main idea behind this technique can be easily visualized.

Suppose M is a smooth, compact manifold, which for simplicity we assume is embedded in a Euclidean space E . We would like to understand basic topological invariants of M such as its homology, and we attempt a “slicing” technique.

We fix a unit vector \mathbf{u} in E and we start slicing M with the family of hyperplanes perpendicular to \mathbf{u} . Such a hyperplane will in general intersect M along a submanifold (slice). The manifold can be recovered by continuously stacking the slices on top of each other in the same order as they were cut out of M .

Think of the collection of slices as a deck of cards of various shapes. If we let these slices continuously pile up in the order they were produced, we notice an increasing stack of slices. As this stack grows, we observe that there are moments of time when its shape suffers a qualitative change. Morse theory is about extracting quantifiable information by studying the evolution of the shape of this growing stack of slices.

From a mathematical point of view, we have a smooth function

$$h : M \rightarrow \mathbb{R}, \quad h(x) = \langle \mathbf{u}, x \rangle.$$

The above slices are the level sets of h ,

$$\{x \in M; \quad h(x) = \text{const}\},$$

and the growing stack is the time-dependent sublevel set

$$\{x \in M; \quad h(x) \leq t\}, \quad t \in \mathbb{R}.$$

The moments of time when the pile changes its shape are called the *critical values* of h and correspond to the moments of time t when the corresponding hyperplane $\{\langle \mathbf{u}, x \rangle = t\}$ intersects M tangentially. Morse theory explains how to describe the shape change in terms of *local* invariants of h .

A related slicing technique was employed in the study of the topology of algebraic manifolds called the *Picard–Lefschetz theory*. This theory is back in fashion due mainly to Donaldson’s pioneering work on symplectic Lefschetz pencils.

The present book is divided into three conceptually distinct parts. In the first part, we lay the foundations of Morse theory (over the reals). The second part consists of applications of Morse theory over the reals, while the last part describes the basics and some applications of complex Morse theory, a.k.a. Picard–Lefschetz theory. Here is a more detailed presentation of the contents.

In Chap. 1, we introduce the basic notions of the theory and we describe the main properties of Morse functions: their rigid local structure (Morse lemma) and their abundance. (Morse functions are generic.) To aid the reader, we have sprinkled the presentation with many examples and figures. One recurring simple example that we use as a testing ground is that of a natural Morse function arising in the design of robot arms. We conclude this chapter with a simple but famous application of Morse theory. We show that the expected number of critical points of the restriction of a random linear map $\ell : \mathbb{R}^3 \rightarrow \mathbb{R}$ to a knot $K \hookrightarrow \mathbb{R}^3$ is described by the total curvature of the knot. As a consequence, we obtain Milnor’s celebrated result [M0] stating that if a closed curve in \mathbb{R}^3 is “not too curved,” then it is not knotted.

Chapter 2 is the technical core of the book. Here we prove the fundamental facts of Morse theory: crossing a critical level corresponds to attaching a handle and Morse inequalities. Inescapably, our approach was greatly influenced by the classical sources on this subject, more precisely Milnor’s beautiful books on Morse theory and h -cobordism [M3, M4].

The operation of handle addition is much more subtle than it first appears, and since it is *the* fundamental device for manifold (re)construction, we devoted an entire section to this operation and its relationship to cobordism and surgery. In particular, we discuss in some detail the topological effects of the operation of surgery on knots in S^3 and illustrate this in the case of the trefoil knot.

In Chap. 2, we also discuss in some detail dynamical aspects of Morse theory. More precisely, we present the techniques of Smale about modifying a Morse function so that it is self-indexing and its stable/unstable manifolds intersect transversally. This allows us to give a very simple description of an isomorphism between the singular homology of a compact smooth manifold and the (finite dimensional) Morse–Floer homology determined by a Morse function, that is, the homology of a complex whose chains are formal linear combinations of critical points and whose boundary is described by the connecting trajectories of the gradient flow. We have also included a brief section on Morse–Bott theory, since it comes in handy in many concrete situations.

We conclude this chapter with a section of a slightly different flavor. Whereas Morse theory tries to extract topological information from information about critical points of a function, min–max theory tries to achieve the opposite goal, namely, to transform topological knowledge into information about the critical points of a function. In particular, we discuss the Lusternik–Schnirelmann category of a space, which is a homotopy invariant particularly adept at detecting critical points.

Chapter 3 is devoted entirely to applications of Morse theory. We present relatively few examples, but we use them as pretexts for wandering in many parts of mathematics that are still active areas of research. We start by presenting a recent result of Farber and Schütz, [FaSch], on the Betti numbers of the space of planar polygons, or equivalently, the space of configurations of planar robot arms such that the end-point of the arm coincides with the initial joint. Besides its intrinsic interest, this application has an added academic bonus: it gives the reader the chance to witness Morse theory in action, in all its splendor. Additionally it exposes the reader to the concept of Bott–Samelson cycle which is useful in many other applications of Morse theory.

We next discuss two classical applications: the computation of the Poincaré polynomials of complex Grassmannians and an old result of Lefschetz concerning the topology of Stein manifolds.

The complex Grassmannians give us a pretext to discuss at length the Morse theory of moment maps of Hamiltonian torus actions. We prove that these moment maps are Morse–Bott functions. We then proceed to give a complete presentation of the equivariant localization theorem of Atiyah, Borel, and Bott (for S^1 -actions only), and we use this theorem to prove a result of Conner [Co]: the sum of the Betti numbers of a compact, oriented smooth manifold is greater than the sum of the Betti numbers of the fixed point set of any smooth S^1 -action. Conner’s theorem implies among other things that the moment maps of Hamiltonian torus actions are *perfect* Morse–Bott function. The (complex) Grassmannians are coadjoint orbits of unitary groups, and as such they are equipped with many Hamiltonian torus actions leading to many choices of perfect Morse functions on Grassmannians. We conclude with a section on the celebrated Duistermaat–Heckman formula.

Chapter 4 is more theoretical in nature but it opens the door to an active area of research, namely Floer homology. While still in the finite dimensional context, we take a closer look at the topological structure of a Morse–Smale flow. The main results are inspired by our recent investigations [Ni2] and, to the best of our knowledge, they seem to have never appeared in the Morse theoretic literature.

The key results of this chapter (Theorem 4.32 and 4.33) state that a Morse flow on a compact manifold satisfies the Smale transversality condition if and only if the stratification given by the unstable manifolds satisfies the Whitney regularity conditions. Because the theory of Whitney stratifications is not part a standard graduate curriculum, we devoted a large part of this chapter surveying this theory. Since the proofs of the main results in this area are notoriously complex, we decided to skip most of them opting instead for copious references and numerous illuminating examples.

These results provide a rigorous foundation to Thom’s original insight [Th]. One immediate consequence of Theorem 4.33 is a result of Laudénbach [Lau] on the nature of the singularities of the closure of an unstable manifold of a Morse–Smale flow.

In Sect. 4.4 we investigate the spaces of tunnelings between two critical points of a Morse–Smale flow. Using a recent idea of Kronheimer and Mrowka [KrMr], we show that these spaces admit natural compactifications as manifolds with corners.

We do not use this fact anywhere else in the book, but since it is part of the core of Morse theoretic facts available to the modern geometer, we thought we had to include a short proof.

In the last section of this chapter, we have a second look at the Morse–Floer complex, from a purely dynamic point of view. We define the boundary operator ∂ in terms of signed counts of tunnelings, and we give a purely dynamic proof of the equality $\partial^2 = 0$. Our proof is similar in spirit to the proof in [Lau], but we have deliberately avoided the usage of currents because the unstable manifolds may not have finite volume. Instead, we use the theory of Whitney stratifications to show that the equality $\partial^2 = 0$ is a consequence of the cobordism invariance of the degree of a map.

The application to the topology of Stein manifolds offered us a pretext for the last chapter of the book on the Picard–Lefschetz theory. Given a complex submanifold M of a complex projective space, we start slicing it using a (complex) one-dimensional family of projective hyperplanes. Most slices are smooth hypersurfaces of M , but a few of them are mild singularities (nodes). Such a slicing can be encoded by a holomorphic Morse map $M \rightarrow \mathbb{CP}^1$.

There is one significant difference between the real and the complex situations. In the real case, the set of regular values is *disconnected*, while in the complex case, this set is *connected* since it is a punctured sphere. In the complex case, we study not what happens as we cross a critical value, but what happens when we go once around it. This is the content of the Picard–Lefschetz theorem.

We give complete proofs of the local and global Picard–Lefschetz formulæ and we describe basic applications of these results to the topology of algebraic manifolds.

We conclude the book with a chapter containing a few exercises and solutions to (some of) them. Many of them are quite challenging and contain additional interesting information we did not include in the main body, since it may have been distracting. However, we strongly recommend to the reader to try solving as many of them as possible, since this is the most efficient way of grasping the subtleties of the concepts discussed in the book. The solutions of these more challenging problems are contained in the last section of the book.

Penetrating the inherently eclectic subject of Morse theory requires quite a varied background. The present book is addressed to a reader familiar with the basics of algebraic topology (fundamental group, singular (co)homology, Poincaré duality, e.g., Chaps. 0–3 of [Ha]) and the basics of differential geometry (vector fields and their flows, Lie and exterior derivative, integration on manifolds, basics of Lie groups and Riemannian geometry, e.g., Chaps. 1–4 in [Ni1]). In a very limited number of places we had to use less familiar technical facts, but we believe that the logic of the main arguments is not obscured by their presence.

Acknowledgements This book grew out of notes I wrote for a one-semester graduate course in topology at the University of Notre Dame in the fall of 2005. I want to thank the attending students, Eduard Balreira, Daniel Cibotaru, Stacy Hoehn, and Sasha Lyapina, for their comments, questions, and suggestions, which played an important role in smoothing out many rough patches

in presentation. While working on these notes I had many enlightening conversations on Morse theory with my colleague Richard Hind. I want to thank him for calmly tolerating my frequent incursions into his office, and especially for the several of his comments and examples I have incorporated in the book.

Last, but not the least, I want to thank my wife. Her support allowed me to ignore the “publish or perish” pressure of these fast times, and I could ruminate on the ideas in this book with joyous abandonment.

What’s new in the second edition

- I have included several immediate but useful consequences of the results proved in the first edition: Corollary 1.26 and Theorem 2.37.
- I have included several new sections of applications: Sects. 1.3, 3.1, 3.7.
- The whole of Chap. 4 is new.
- I have added several new exercises.
- I have fixed many typos and errors in the first edition. In this process I was aided by many readers. I would especially like to thank Professor *Steve Ferry* of Rutgers University for his numerous suggestions, corrections, and overall very useful critique. I would also like to thank Leonardo Biliotti and Alessandro Ghigi for drawing my attention to some problems in Theorems 3.45 and 3.48 of the first edition (Theorems 3.61 and 3.63 in the current edition.) I have addressed them in the current edition.

This work was partially supported by NSF grants DMS-0303601 and DMS-1005745.

An Invitation to Morse Theory

Nicolaescu, L.

2011, XVI, 353 p. 47 illus., Softcover

ISBN: 978-1-4614-1104-8