Chapter 7

Morse Theory

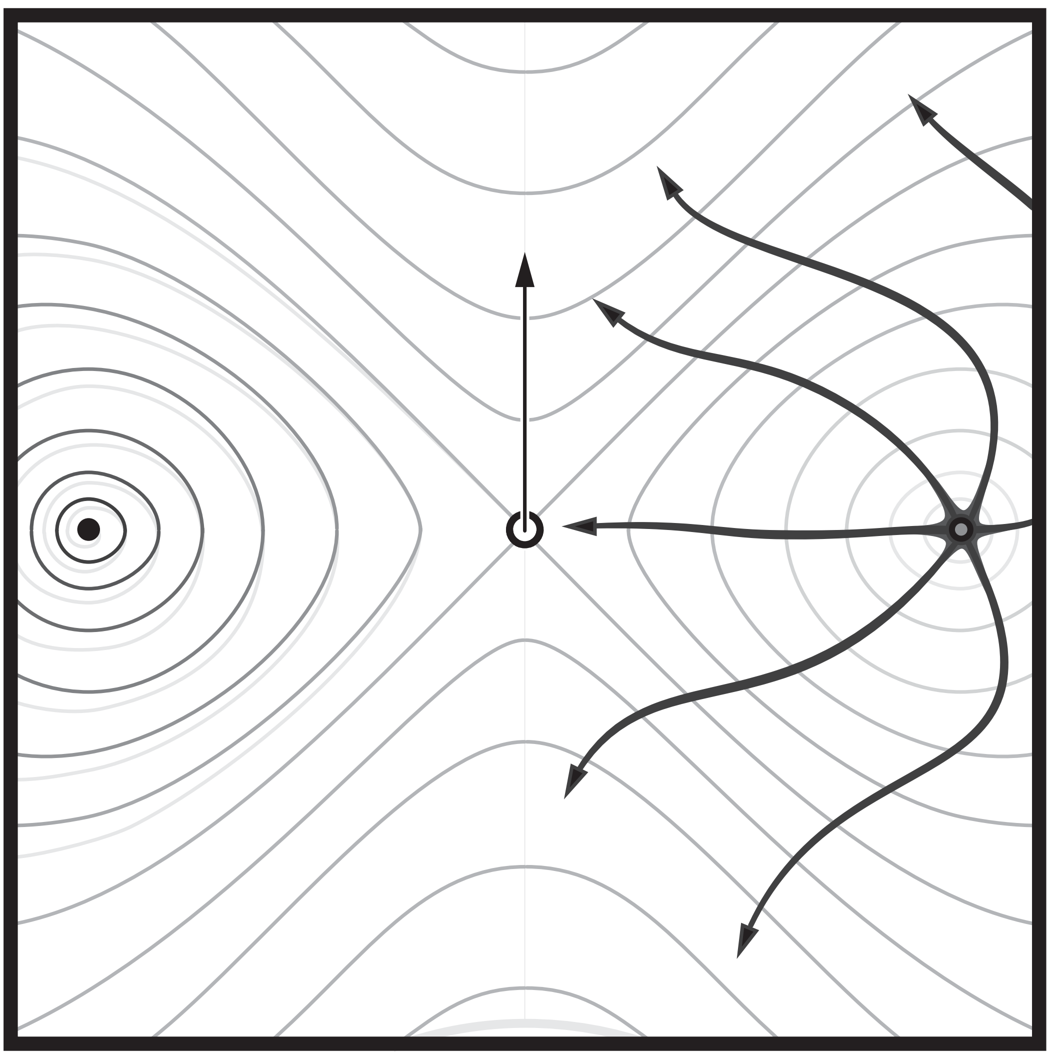

The two great themes of smooth manifolds and assembled complexes compete politely in topology. One locus of synthesis between the two lies in the eponymous theory of Morse. This chapter integrates all of the previous chapters into a suite of perspectives, tools, and applications connecting local behavior to global.

7.1 Critical points

In any homology theory, one counts certain objects with respect to an appropriate cancelation, usually with some ancillary structure imposed to keep counts finite. Morse theory uses a *height function* to facilitate homological counting.

Fix M a compact Riemannian manifold without boundary. Morse theory operates via a real-valued function and the dynamics of its gradient flow. Fix $h\colon M \to \mathbb{R}$ a smooth function and consider the gradient field $-\nabla h$ on M. The dynamics of this vector field are simple: solutions either are fixed points (critical points of h) or flow *downhill* from one fixed point to another. Let $\mathrm{Cr}(h)$ denote the set of critical points, and assume for the sake of simplicity that all such critical points are **nondegenerate** – the second derivative (or *Hessian*) is nondegenerate (has nonzero determinant) at these points. Equivalently, the gradient field $-\nabla h$, thought of as a section of the tangent bundle T_*M, is transverse to the zero section (*cf.* §1.6), whence it follows that nondegeneracy is generic. These nondegenerate critical points are the basis elements of Morse theory.

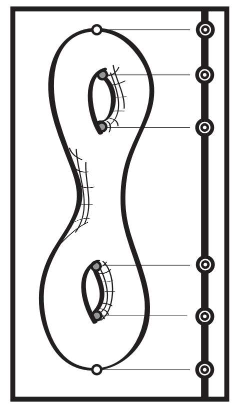

The critical points have a natural grading – the **Morse index**, $\mu(p)$, of $p \in \mathrm{Cr}(h)$ – the number of negative eigenvalues of the Hessian of second derivatives of h at p. This has the more topological interpretation as the dimension of the unstable manifold of the vector field $-\nabla h$ at p (recall Example 1.6): $\mu(p) = \dim W^u(p)$. The Morse index measures how unstable a critical point is: minima have the lowest Morse index; maxima the highest. Balancing a three-legged stool on k legs leads to an index $\mu = 3 - k$ equilibrium.

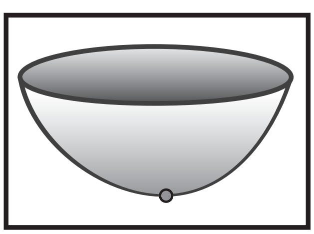

Classical Morse theory begins by observing the (lower) excursion sets $M_t := \{h \le t\}$ of a Morse function h on a compact manifold M. The story, in brief, is as follows. As $t \in \mathbb{R}$ increases, the family M_t gives a filtration of spaces that begins with the empty set and ends with M. In the beginning, a disc appears *ex nihilo* and evolves, pinching and branching as critical points are passed, ultimately capping at the maximum. The critical observation: M_t changes homeomorphism type *only* at critical values of h.

The local picture tells all. Morse theory asserts the following:

Lemma 7.1. *Consider a small compact ball B about a critical point $p \in \mathrm{Cr}(h)$ of Morse index μ. Denote by E the* lower *set $E = B \cap \{h \le h(p) - \epsilon\}$ for $\epsilon > 0$ small, and consider U, the closure of $B - E$, the complementary* upper *set.*

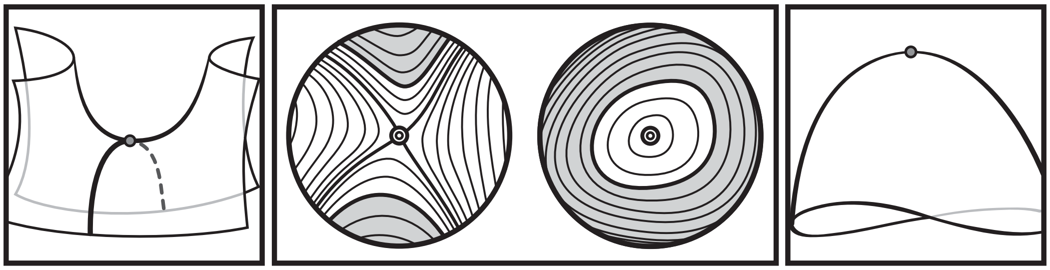

1. *U is homeomorphic to* $\mathbb{D}^n$ *for* $n = \dim M$;
2. *E is homeomorphic to* $\mathbb{D}^{n-\mu+1} \times \mathbb{S}^{\mu-1}$; *and*
3. $E \cap U$ *is homeomorphic to* $\mathbb{D}^{n-\mu} \times \mathbb{S}^{\mu-1}$.

One says that there is a **surgery** in the neighborhood of p that attaches a product disc $U \cong \mathbb{D}^{n-\mu} \times \mathbb{D}^{\mu}$ glued along $E \cap U \cong \mathbb{D}^{n-\mu} \times \partial\mathbb{D}^{\mu}$.

7.2 Excursion sets and persistence

The filtration of M by lower excursion sets M_t of a Morse function $h: M \to \mathbb{R}$ fits perfectly into the picture of persistence sketched in §5.13-5.15. In this setting (sometimes called **sublevel set persistence** [104]), one considers the persistence barcode of the filtration of M by subsets M_t, where the parameter t, is suitably discretized.

In keeping with the idea of Morse theory, the sublevel sets change their topological type only at critical points; hence barcodes for sublevel set persistence are tethered to the critical values. The meanings of the barcodes are as in §5.13-5.15, in that a long bar connotes significance; however, in sublevel set persistence, one is not trying to find an optimal cut-off t to approximate M; rather, one wants to know which topological features of a manifold are important when it is stretched out along a table ruled by h. Clearly a small bar in a barcode for excursion sets indicates something like a *wrinkle* in the Morse function: a transient hole. Long bars in the sublevel set barcode indicate a large-scale feature as seen by h. In the many applications of persistence, the lingering problem of noise is pertinent, since a wiggling of h leads to many small spurious bars. There is a large collection of stability-type results for persistence: see [70] for the first of these, which concerns sublevel set persistence and asserts an (*interleaving*) distance on barcodes with continuity guarantees under addition of noise to h, *cf.* §10.6.

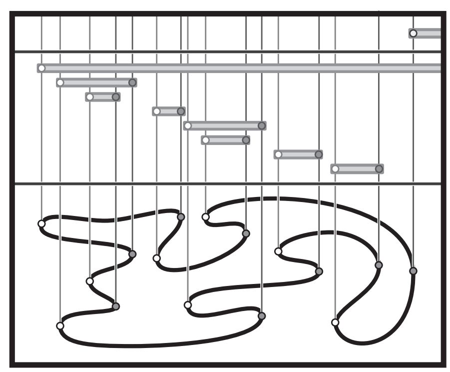

More salient to the themes of this chapter is the phenomenon of cancellation. The birth and death implicit in the barcodes for lower excursion sets reveal one of the great perspectives of Morse theory: critical points and the topological features they generate are naturally *paired* in a cancellative manner. Lemma 7.1 implies the following:

Corollary 7.2. *For a Morse function h on a compact manifold M, births and deaths in the sublevel set homology barcode of grading k implicate only critical points of Morse index k and k + 1 respectively.*

This cancellation of features foreshadows a self-contained homology theory for Morse functions.

7.3 Morse homology

Classical Morse theory concerns equivalence up to homeomorphism, based on the uniform behavior of nondegenerate critical points. By relaxing to a more homological view, the theory will connect better with the rest of the text and will naturally suggest extensions to non-Morse functions.

The constructs of the previous sections are perfect for homology. One has a natural set of objects ($\mathrm{Cr}(h)$) and a grading (μ). One obtains the **Morse complex**, $\mathcal{C}^h = (MC_\bullet, \partial)$, with MC_k the vector space with basis $\{p \in \mathrm{Cr}(h)\,;\, \mu(p) = k\}$. The boundary maps encode the global flow of the gradient field: ∂_k counts (modulo 2 in the case of $\mathbb{F}_2$ coefficients) the number of **connecting orbits** – flowlines from a critical point with $\mu = k$ to a critical point with $\mu = k - 1$. One hopes (or assumes) that this number is well-defined.

The difficult business is to demonstrate that $\partial^2 = 0$: this involves careful analysis of the connecting orbits, as in, e.g., [21, 274]. The use of $\mathbb{F}_2$ coefficients is highly recommended: dis-orientation is a plus. The ensuing **Morse homology**, $MH_\bullet(h)$, captures information about M.

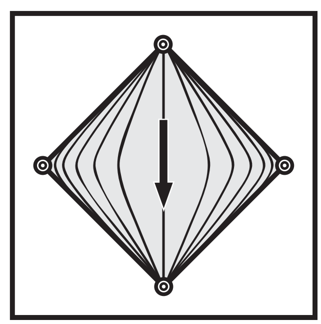

Theorem 7.3 (Morse Homology Theorem). *For M compact and $h\colon M \to \mathbb{R}$ Morse, $MH_\bullet(h;\mathbb{F}_2) \cong H_\bullet(M;\mathbb{F}_2)$, independent of h.*

The conceptually simplest proof involves an isomorphism to the cellular (CW) homology of M, where the cell structure is that given by the $W^u(p)$ for $p \in \mathrm{Cr}(h)$. The **Stable Manifold Theorem** from dynamical systems asserts that these unstable manifolds are all cells homeomorphic to $\mathbb{R}^{\mu(p)}$, and a transversality argument gives acceptable attaching maps. The proofs are clearest for **Morse-Smale functions**, for which all stable and unstable manifolds of critical points are transverse. Morse-Smale functions, like Morse functions, are generic, and for such, there is an isomorphism at the level of chain complexes: *cf.* Example 2.2.

Example 7.4 (Morse homology) ⊙

For the particularly simple height function h on $\mathbb{S}^2$ with two maxima, one minimum, and (of necessity) one saddle, the Morse complex in $\mathbb{F}_2$ coefficients is of the form,

$$0 \longrightarrow \mathbb{F}_2^2 \xrightarrow{[1\,|\,1]} \mathbb{F}_2 \xrightarrow{0} \mathbb{F}_2 \longrightarrow 0 \;;$$

thus, $MH_2(h) \cong \mathbb{F}_2$, $MH_1(h) \cong 0$, and $MH_0(h) \cong \mathbb{F}_2$; one notes that in this example $MH_\bullet(h) \cong H_\bullet(\mathbb{S}^2; \mathbb{F}_2)$. ⊚

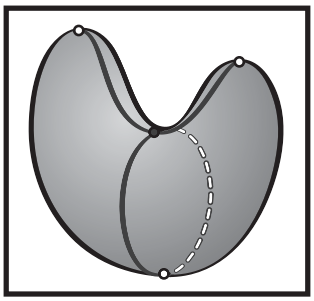

One efficient means of encoding all the critical point data of a Morse function is by means of the **Morse polynomial** of h, defined as $M_h(t) := \sum_{p\in \mathbf{Cr}(h)} t^{\mu(p)}$. This polynomial in the abstract variable t has as its coefficients $c_i \in \mathbb{N}$ the number of critical points of h with Morse index i. Theorem 7.3 implies a relationship between the Morse polynomial and the Poincaré polynomial, $P(t) = \sum_i \dim H_i(M)t^i$, which, recall, encodes the singular homology of the manifold M. At the very least, the i^{th} coefficient of $M_h(t)$ must be greater than or equal to that of $P(t)$. A stronger inequality uses the polynomial algebra explicitly.

Corollary 7.5 (Strong Morse Inequalities). *For $h\colon M \to \mathbb{R}$ Morse,*

$$M_h(t) = P(t) + (1+t)Q(t), \tag{7.1}$$

where $Q \in \mathbb{N}[t]$ is a polynomial with all coefficients in $\mathbb{N}$.

Corollary 7.6 (Euler characteristic). *For $h\colon M \to \mathbb{R}$ Morse on a compact manifold M,*

$$\chi(M) = \sum_{p\in \mathbf{Cr}(h)} (-1)^{\mu(p)}.$$

Proof. Use Corollary 7.5, Lemma 5.17, and $t = -1$. ⊙

Morse theory offers a painless demonstration of one manifestation of homological Poincaré duality:

Corollary 7.7 (Poincaré Duality). *The $\mathbb{F}_2$-homology of a compact n-manifold M is symmetric in its grading: $H_p(M;\mathbb{F}_2) \cong H_{n-p}(M;\mathbb{F}_2)$ for all $0 \le p \le n$.*

Proof. For h a Morse function, the function $-h$ is also Morse. Changing from h to $-h$ reverses the direction of the gradient flow, preserving the critical points and connecting orbits, but exchanging stable and unstable manifolds. At the chain complex level,

$$\begin{array}{ccccc} \xrightarrow{\partial} & MC_p(h) & \xrightarrow{\partial} & MC_{p-1}(h) & \xrightarrow{\partial} \\ & \cong\Big\downarrow & & \Big\downarrow\cong & \\ \xleftarrow[\partial]{} & MC_{n-p}(-h) & \xleftarrow[\partial]{} & MC_{n-p+1}(-h) & \xleftarrow[\partial]{} \end{array}$$

and thus $H_p(M) \cong MH_p(-h) \cong MH_{n-p}(h) \cong H_{n-p}(M)$. ⊙

7.4 Definable Euler integration

Corollary 7.6 hints at the role of Morse theory in integration with respect to Euler characteristic. The integral operator $\int d\chi \colon \mathrm{CF}(X) \to \mathbb{Z}$ of Chapter 3 does not readily extend to continuous real-valued integrands: there is an extension of the integral to real-valued integrands by Rota, then Chen, that vanishes on all continuous integrands [262, 65]. Recent work [25] has revealed a novel Euler calculus for $\mathbb{R}$-valued integrands, with interesting Morse-theoretic connections. Fix an o-minimal structure, as in §3.5, and denote by $\mathrm{Def}(X)$ the **definable functionals** $h \colon X \to \mathbb{R}$ – those (compactly supported) functions whose graphs in $X \times \mathbb{R}$ are definable sets. Recall, these are not necessarily continuous functions, as they include the constructible functions $\mathrm{CF}(X)$. Given $h \in \mathrm{Def}(X)$, define the integral of h as a limit of discretizations:

$$\int_X h\,\lfloor d\chi \rfloor := \lim_{n\to\infty} \frac{1}{n}\int_X \lfloor nh \rfloor\, d\chi \quad : \quad \int_X h\,\lceil d\chi \rceil := \lim_{n\to\infty} \frac{1}{n}\int_X \lceil nh \rceil\, d\chi\ . \tag{7.2}$$

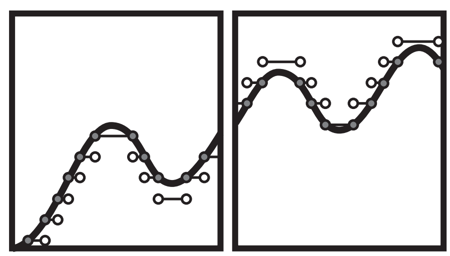

These limits exist and are well-defined, though *not equal* in general. The **triangulation theorem** for $\mathrm{Def}(X)$ [293] states that to any $h \in \mathrm{Def}(X)$, there is a definable triangulation (a definable bijection to a disjoint union of open affine simplices in some Euclidean space) on which h is affine on each open simplex. From this, one may reduce all questions about the integrals over $\mathrm{Def}(X)$ to questions of affine integrands over individual open simplices, using the additivity of the integral. Using this reduction technique, one proves the following analogue of Equation (3.10):

Proposition 7.8 ([25]). *For $h \in \mathrm{Def}(X)$,*

$$\int_X h\,\lfloor d\chi \rfloor = \int_{s=0}^{\infty} \chi\{h \geq s\} - \chi\{h < -s\}\, ds = -\int_X -h\,\lceil d\chi \rceil. \tag{7.3}$$

In its favor, the integral is coordinate-free, in the sense that $\int_X h \circ \phi \lfloor d\chi \rfloor = \int_X h \lfloor d\chi \rfloor$ for ϕ a homeomorphism of X. Less pleasant is that these integral operators are *not linear*; nor even homogeneous with respect to negative coefficients, by (7.3). The compelling feature of the valuations $\lfloor d\chi \rfloor$ and $\lceil d\chi \rceil$ is their relation to Morse theory. The following theorem has the effect of *concentrating* the "measure" $\lfloor d\chi \rfloor$ on the critical points of the integrand.

Proposition 7.9 ([25]). *If h is a Morse function on a compact n-manifold M, then:*

$$\int_M h\,\lfloor d\chi \rfloor = \sum_{p \in \mathrm{Cr}(h)} (-1)^{n-\mu(p)} h(p) \quad : \quad \int_M h\,\lceil d\chi \rceil = \sum_{p \in \mathrm{Cr}(h)} (-1)^{\mu(p)} h(p).$$

This result does not require having a Morse function: see Theorem 7.12. The theory of Euler integration on **CF** is the highly degenerate setting where the entire domain is critical.

Signal processing – with applications ranging from radar imagery to sensor networks – is fueled by integral transforms. The Euler integral of Chapter 3 yields some interesting novel examples of integral transforms with the limiting feature of being an integer-valued theory, and thus perhaps not directly applicable to, say, image processing. The extension of the Euler integral to $\mathbb{R}$-valued functions allows for a wider array of applications thanks to some novel integral transforms [79, 154].

Example 7.10 (Euler-Fourier transforms) ⦾

There is an integral transform $\mathcal{F}\colon \mathsf{CF}(\mathbb{R}^n) \times (\mathbb{R}^n)^\vee \to \mathsf{Def}(\mathbb{R}^n)$ that is best thought of as an Eulerian version of a Fourier transform. For $h \in \mathsf{CF}(\mathbb{R}^n)$ and $\xi \in (\mathbb{R}^n)^\vee$, define the **Euler-Fourier transform** of h in the direction ξ via

$$\mathcal{F}h(\xi) := \int_{-\infty}^{\infty} \int_{\xi^{-1}(s)} h \, d\chi \, ds.$$

It is clear that for A a compact convex subset of $\mathbb{R}^n$ and $\|\xi\| = 1$, $(\mathcal{F}\mathbb{1}_A)(\xi)$ equals the projected length of A along the ξ-axis. This points to the following theorem of [154]. Let A be a tame compact n-dimensional submanifold of $\mathbb{R}^n$ with boundary ∂A, decomposed into positive $\partial^+ A$ and negative $\partial^- A$ components (depending on whether the oriented normal to ∂A has positive or negative dot product with ξ). Then, the Euler-Fourier transform of A can be reduced to an integral over ∂A:

$$\mathcal{F}\mathbb{1}_A(\xi) = \int_{\partial^+ A} \xi \lfloor d\chi \rfloor - \int_{\partial^- A} \xi \lceil d\chi \rceil. \tag{7.4}$$

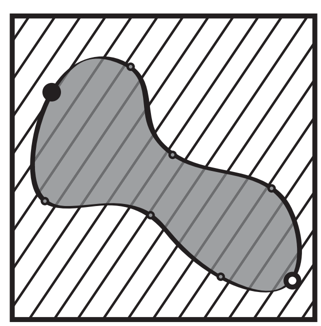

Note the resemblance to Stokes' theorem: an integral over the interior equals an integral of an "anti-derivative" over the boundary. One must be careful with orientations and signs, as with Stokes' theorem. Since ∂A is a manifold and $\xi\colon \partial A \to \mathbb{R}$ is generically a Morse function, Proposition 7.9 implies that $\mathcal{F}h(\xi)$ is an alternating sum of critical values, graded by Morse index, providing a vast generalization of the ξ-projected width observation for A convex. A polar version of this integral transform (called an **Euler-Bessel transform**) has a similar index formula and is useful in shape detection from enumerative sensor data [154]. ⦾

7.5 Stratified Morse theory

The initial emphasis of Morse theory on nondegenerate Morse functions and local coordinate representations leaves students with the impression that manifolds and nondegeneracy are a *sine qua non*. Though it is convenient to assume a Morse function, nature often interferes, necessitating a degenerate Morse theory. Morse Theory can be adapted to settings where the object of interest is not a manifold, but rather a stratified space, built from manifold pieces, assembled in a sufficiently tame manner. This leads to some complications, the consequence of increased generality.

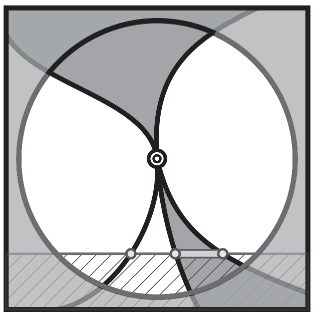

The theory of Goresky and MacPherson [163] recreates Morse theory for stratified spaces (see §1.8). This large and technical body of work cannot be summarized quickly and accurately: the following is elementary, at the expense of carefully-stated theorems. Instead of the usual (Whitney-type) stratified spaces, consider (*cf.* [272]) a fixed o-minimal structure. Let $Y \subset \mathbb{R}^n$ be tame and $h\colon \mathbb{R}^n \to \mathbb{R}$ a definable function – the tame analogue of a Morse function on Y. Stratified Morse theory delineates when the excursion sets $Y_t = \{h \leq t\}$ change their topological type. This occurs by defining a Morse index and classifying attachments in the manner of Lemma 7.1.

In classical Morse theory, the index of a critical point is a natural number, and attaching is by surgery of a disc along a sphere of dimension μ. In this stratified setting, one defines the **local Morse data** of Y at a point $p \in Y$ under h to be the *homeomorphism type* of the space

$$\mathrm{LMD}(p) = \mathrm{LMD}(Y, p, h) := \lim_{\epsilon' < \epsilon \to 0^+} \overline{B_\epsilon(p) \cap Y} \cap \{h > h(p) - \epsilon'\}\,.$$

The limit exists thanks to tameness [293] and is sometimes expressed in terms of a pair (B, E) of compact spaces: $B = \overline{B_\epsilon(p) \cap Y}$ and $E = B \cap \{h \leq h(p) - \epsilon'\}$ (sufficiently small), and $\mathrm{LMD} = B - E$. What is the appropriate Morse index in this setting? For a numerical value, it would be appropriate to call the **Euler-Morse index** of h the constructible function $\mathcal{I}_h \in \mathrm{CF}(Y)$ given by $\mathcal{I}_h(x) := \chi(\mathrm{LMD}(Y, x, h)) = \chi(B) - \chi(E)$. Note that one must be careful with open versus closed cells, as in Chapter 3. For a classical Morse function h, $\mathcal{I}_h$ is zero except at the critical points, at which the index takes on a value of $(-1)^\mu$, for μ the Morse index. A richer index would be the relative homology $H_\bullet(B, E)$, which, for a nondegenerate critical point of a Morse function on a manifold, would be concentrated in grading μ. For more degenerate critical points or critical sets, this index can capture some local topological behavior.

Stratified Morse theory was developed for applications in algebraic geometry that lie outside the bounds of this text, and only the very first step – the local Morse data – has been touched upon. Technical aspects of stratified spaces are numerous and crucial to the theory. The significant steps lie in a *tangential* versus *normal* decomposition of the local Morse data, with instructions as to how local changes in excursion sets are controlled by this data. All of this culminates in yet another homology theory – **intersection homology** – which requires the constructible sheaves of Chapter 9 to fully appreciate (but see [34]). Despite its rarefied origins, stratified Morse theory has found applications in several contexts, including problems in grasping and manipulation in robotics [249], in which potential functions on the relevant configuration spaces can lead to interesting critical sets, on which one cannot simply "compute derivatives" to determine stability.

Example 7.11 (Euler integration) ⊚

A blending of the Poincaré-Hopf Theorem 3.5 and Corollary 7.6 yields a link to Euler integration in that $\int_Y \mathcal{I}_h \, d\chi = \chi(Y)$. This, Proposition 7.9, and more follow from a connection between stratified Morse theory and Euler integration over $\mathsf{Def}(X)$.

Theorem 7.12 ([25]). *For $h \in \mathsf{Def}(X)$ continuous,*

$$\int_X h\lfloor d\chi \rfloor = \int_X h \, \mathcal{I}_{-h} \, d\chi \quad : \quad \int_X h\lceil d\chi \rceil = \int_X h \, \mathcal{I}_h \, d\chi$$

Thus, the appearance of not one but two valuations, $\lfloor d\chi \rfloor$ and $\lceil d\chi \rceil$, on $\mathsf{Def}(X)$, is not an anomaly, but rather another manifestation of Morse-theoretic Poincaré duality: $h \leftrightarrow -h$. This result, like so much of Morse theory done properly, does not require either a manifold or a nondegeneracy condition beyond tameness. ⊚

7.6 Conley index

Stratified Morse theory allows one to relax to non-manifolds and degenerate gradients: with the proper approach, it is also possible to ignore the Morse function altogether. One of the best approaches for doing generalized Morse theory is due to Conley [72], and has enjoyed great success in applications to mathematical biology [44, 181], rigorous numerics [292, 81, 228], experimental time-series inference [229], and more. What follows is a simple version of Conley's theory in the continuous-time setting.

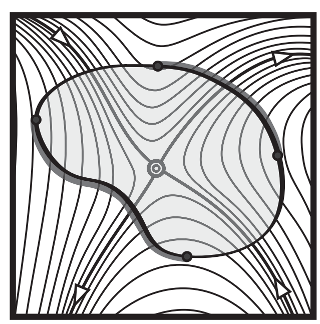

Consider the case of a gradient field $-\nabla h$ of a Morse function $h\colon M \to \mathbb{R}$ on a manifold M. Choose a fixed point $p \in \mathsf{Cr}(h)$ and consider a small ball B about p. The boundary ∂B is partitioned into an *exit set* on which the gradient field $-\nabla h$ points *out* of B; an *entrance set* on which the field points *in*; and the remaining points of tangency to ∂B. If p is a critical point of Morse index k, then the exit set is homotopic to a sphere of dimension $\mu(p) - 1$: Lemma 7.1. This example prompts a more general index. *The Conley index is not an integer, but a homotopy type of spaces.*

One simple approach is as follows. Let X be a complete locally compact metric space with a continuous flow $\varphi_t\colon X \to X$ (which may or may not come from a smooth vector field). The Conley index is associated to a suitable compact subset, $B \subset X$. The **invariant set** $S = \mathsf{Inv}(B; \varphi)$ of the flow on a set B is the set of all points $x \in B$ such that $\varphi_t(x) \in B$ for all t. One says a compact B is an **isolating block** if $\mathsf{Inv}(B; \varphi)$ lies strictly in the interior of B, and, for all $x \in \partial B$, the flowline through x exits B either in arbitrarily small forwards or backwards time or both. *No internal "tangencies" are permitted.* The isolating block condition is a loose type of transversality: small perturbations to B remain isolating blocks.

The Conley index of an invariant set S with isolating block B collates the topology of B relative to how the flow exits ∂B. The **exit set** of B is $E := \{x \in \partial B : \varphi_\epsilon(x) \notin B, \forall\, 0 < \epsilon \ll 1\}$. The **Conley index** of S is the pointed homotopy type $\mathrm{Con}(S) := \mathrm{h}[B/E, \{E\}]$. The index *is* the quotient space B/E (up to homotopy) with E remembered as an abstract basepoint. The **homological Conley index** of S is the relative singular homology $\mathrm{Con}_\bullet(B) := H_\bullet(B, E)$. The index is well-defined in that any two isolating blocks with the same invariant set have equivalent Conley indices.

Example 7.13 (Morse index) ⊙

The Conley index of a nondegenerate critical point of a Morse function with Morse index μ is the (homotopy type of the) sphere $\mathbb{S}^\mu$ (with basepoint). The basepoint, which initially seems extraneous, is vital when considering the case $\mu = 0$. There, the fixed point is a sink, and the exit set E is the empty set, which, when remembered as an abstract basepoint, gives a Conley index of $\mathbb{S}^0$, the disjoint union of two points. Notice that (1) the (local) topology of the unstable manifold of the invariant sets figures prominently in the Conley index; and (2) the Conley index depends only on the type of critical point B surrounds, not on B itself. ⊙

Isolating blocks are, unfortunately, not as abundant or flexible as one would like. The solution is to allow for a larger exit set than simply what lies on the boundary of B. Consider an invariant set S of the flow φ. One defines an **index pair** (B, E) of S to be compact subsets $E \subset B$ of X satisfying:

1. **Isolation:** $S = \mathrm{Inv}(\overline{B - E}; \varphi)$ lies in the interior of $B - E$;
2. **Invariance:** any flowline starting in E and staying in B lies within E; and
3. **Exit:** any flowline exiting B does so via E.

This B analogous to an isolating block, and E its exit set. The resulting Conley index is well-defined in its homotopical $\mathrm{Con}(S) := \mathrm{h}[B/E, \{E\}]$ and homological $\mathrm{Con}_\bullet(S) := H_\bullet(B, E)$ instantiations. Wonderful to tell, the Conley index generalizes the Morse index greatly:

1. Fixed points need not be nondegenerate, discrete, or stratified.
2. Vector fields need not be smooth or gradient.
3. Domains need not be manifolds or stratified/definable spaces.

Warning: some flows (*e.g.*, integrable Hamiltonian) admit few index pairs, and some minimal amount of local compactness is convenient. Computation of the Conley index is aided by the following:

Theorem 7.14 (Conley Index Theorem (see [266])). *The Conley index has the following properties:*

1. **[Invariance]** *The Conley index depends only on S, not on the choice of index pair (B, E) for S.*
2. **[Continuation]** *If (B_λ, E_λ) is a (Hausdorff-) continuous family of index pairs for a continuous family of flows, then $\mathrm{Con}_\lambda = \mathrm{Con}(B_\lambda, E_\lambda)$ is constant.*

3. **[Additivity]** *If* (B, E) *and* (B', E') *are disjoint index pairs for* S *and* S'*, then*

$$\mathrm{Con}(S \sqcup S') = \mathrm{Con}(S) \vee \mathrm{Con}(S').$$

Example 7.15 (Forcing theory) ⊚

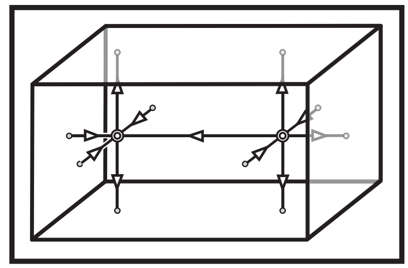

The signature application of index theory in dynamics is forcing the existence of invariant sets. The Poincaré-Hopf Theorem (Theorem 3.5), the Lefschetz Theorem (Theorem 5.19), and the Morse inequalities (Corollary 7.5) all can be used to force the existence of fixed points. The Conley index version is more general still: if $\mathsf{h}[B/E, \{E\}] \neq 0$ (*i.e.*, the pointed homotopy type is not that of a point), then $\mathsf{Inv}(B{-}E; \varphi) \neq \varnothing$. This invariant set may or may not be a fixed point – heteroclinic orbits, periodic orbits and even chaotic invariant sets are detectable [54, 141].

Consider the simple example of the 3-d system $\dot{x} = x^2 - x$, $\dot{y} = -y$, and $\dot{z} = z$. This gradient flow has two fixed points, which can be characterized by their Euler-Poincaré indices (-1 and $+1$), their Morse indices (1 and 2), or their Conley indices ($\mathbb{S}^1$ and $\mathbb{S}^2$), each computed using local information from small neighborhoods of the fixed points. If, instead of small neighborhoods of the fixed points, one chooses a rectangular prism B surrounding the pair of fixed points, then, clearly, the fixed-point index on B is $\mathcal{I}(B) = 1 - 1 = 0$; likewise, the Conley index of this isolating block is zero as well, since $E \subset \partial B$ is contractible. However, $\mathsf{Con}(B) = 0 \neq \mathbb{S}^1 \vee \mathbb{S}^2$ as would follow from Theorem 7.14 if there were no other invariant sets in B. One therefore concludes that there is another invariant set in B besides the fixed points: it is in fact the heteroclinic orbit connecting the two fixed points.

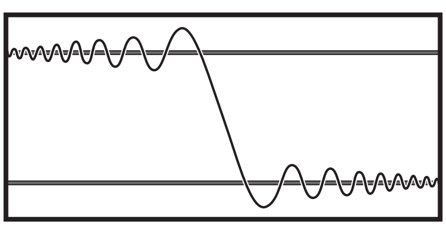

This simple example of an explicit gradient field is trivial, but it captures the spirit of deeper applications. One of the earliest uses of the Conley index was to show the existence of a **traveling wave** solution $u(x, t)$ to a class of reaction-diffusion partial differential equations of the form $u_t = u_{xx} + f(x, u, u_x)$ arising from mathematical biology (population dynamics [73], nerve impulses [279], and more). A pair of constant stationary solutions $u_0(x, t) = u_0$, $u_1(x, t) = u_1$ (*i.e.*, fixed points of the PDE flow) admits a traveling wave if there is a solution $u(x, t)$ that approaches u_1 as $t \to \infty$ and u_0 as $t \to -\infty$ (*i.e.*, a heteroclinic orbit of the PDE flow). Index pairs around each fixed point yield a nonzero Conley index, but a suitable index pair for the union of the two has trivial Conley index, as computed by 'deforming' the PDE and using continuation. An additivity argument *proves* that a travelling wave exists in cases where direct numerical simulation is inconclusive. Further extensions of these ideas have been very useful in finding stationary, time-periodic, and travelling-wave solutions to very general classes of PDEs [84, 156, 182, 292]. ⊚

Example 7.16 (Attractor-repeller pairs) ⊚

Rigorous arguments for the existence of a connecting orbit require a bit more machinery [129, 203, 222, 266]. The simplest example of such is as follows.

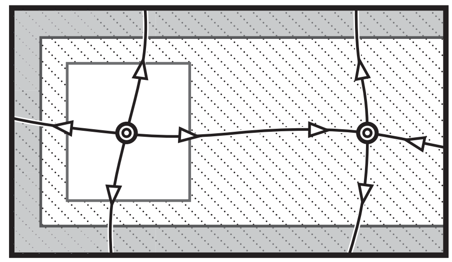

Consider an isolated invariant set $S \subset X$ for the flow φ. An **attractor-repeller pair** is a pair of disjoint compact sets (A, R) in S such that, for every $x \in S-(A \sqcup R)$, the flowline $\varphi_t(x) \to A$ as $t \to \infty$ and $\varphi_t(x) \to R$ as $t \to -\infty$. Thus, S decomposes as A, R, and $S-(A \cup R)$ the connecting orbits. Given any attractor-repeller pair, there exists an **index triple** $N_0 \supset N_1 \supset N_2$ of compact subsets of X, where (N_0, N_2) is an index pair for S, (N_1, N_2) is an index pair for A, and (N_0, N_1) is an index pair for R. The homological long exact sequence of the triple (N_0, N_1, N_2) (from the end of §5.3) becomes an exact sequence relating Conley indices:

$$\cdots \xrightarrow{\delta} \mathrm{Con}_k(A) \xrightarrow{H(i)} \mathrm{Con}_k(S) \xrightarrow{H(j)} \mathrm{Con}_k(R) \xrightarrow{\delta} \cdots,$$

This can be used for forcing connecting orbits from R to A.

Lemma 7.17. *If the connecting homomorphism δ on an index triple is nonvanishing, then there exists a connecting orbit from R to A in S.*

Proof. Assume that $S-(A \sqcup R) = \varnothing$. Then $S = A \sqcup R$ and, as per Theorem 7.14, the Conley index decomposes as a wedge sum. Passing to homology, one obtains $\mathrm{Con}_\bullet(S) = \mathrm{Con}_\bullet(A) \oplus \mathrm{Con}_\bullet(R)$. Exactness then implies that $\delta = 0$. ⊙⊚

Example 7.18 (Floer homology) ⊚

One of the great triumphs of the Conley's approach to Morse theory is in an infinite-dimensional version due to Floer [122] (see also [21, 224, 274]). It has been known since the original works in Morse theory that the methods are applicable to gradient flows on certain infinite-dimensional settings[1] when a Morse index can be defined. Unfortunately, those functionals for which Morse index is finite are rare. More common is the case of functionals whose critical points have Hessians with an infinite number of positive and negative eigenvalues. The breakthrough of Floer was to mimic the Conley approach in settings where the critical points whose linearizations yield Fredholm operators: the positive and negative eigenspaces are both infinite-dimensional, but the difference of their dimensions – the **Fredholm index** – is finite. By combining these insights with the pseudoholomorphic curve technology of Gromov [166] and an appropriate (nontrivial) array of analytic results, Floer and those who followed construct a chain complex over $\mathbb{F}_2$ freely generated by critical points of the functional and graded by the Fredholm index. The resulting **Floer homology** $\mathrm{Fl}_\bullet$ possesses similar properties as the Conley homology, including continuation. Most remarkable is the fact that,

[1] For example, Banach manifolds – spaces locally modeled on a fixed Banach space.

since the Fredholm index takes values in $\mathbb{Z}$ as opposed to $\mathbb{N}$, Floer homology is graded over the full integers: it is common to have $\mathsf{Fl}_k \neq 0$ for negative values of k. ⊙

Example 7.19 (Arnol'd Conjecture) ⊙

The first achievement of Floer's theory was a resolution of the **Arnol'd Conjecture**. Fix a compact $2n$-manifold M with a **symplectic form**, a closed nondegenerate[2] 2-form $\omega \in \Omega^2(M)$. A symplectic manifold allows for Hamiltonian dynamics as follows. Given a function $H\colon M \to \mathbb{R}$, let V_H be the (unique) vector field satisfying $\omega(V_H, \cdot) = -dH$. This **Hamiltonian field** V_H is a twisted analogue of the gradient $-dH$, and it follows from the Morse inequalities that V_H has at least $\sum_k \dim H_k(M; \mathbb{R})$ fixed points.

Choose a smooth $\mathbb{S}^1$-parameter family of functions H_t, $t \in \mathbb{S}^1$, and let $V_t = X_{H_t}$ be the corresponding t-dependent vector field. Arnold's conjecture – that the number of 1-periodic orbits of the family V_t is bounded below by $\sum_k \dim H_k(M)$ – is a deceptively simple-sounding analogue of the Strong Morse Inequalities. It was first proved using Floer-theoretic arguments as follows. Periodic orbits of V_t satisfy a variational principle with respect to *action*. This action yields a real-valued function on the space of loops in M, extrema of which correspond to 1-periodic orbits. The action functional is unbounded and the loop space of M is infinite dimensional, thus necessitating Floer-theoretic arguments to find critical points of action. An equivalence between the Floer homology of the action functional and the homology of M – a variant of Theorem 7.3 – provides the key to the Arnol'd Conjecture. Details are beyond the scope of this text (and are largely analytic), but the theme of counting certain invariant sets by means of a specialized homology theory and relating it to classical homology is fully in the spirit of Morse theory. ⊙

7.7 Lefschetz index, redux

As hinted at in §5.10, there is a way to compute the Lefschetz index of a self-map $f\colon X \to X$ that is localized at the fixed point set. This turns out to have a deep connection to both stratified Morse theory and the Conley index. The number and types of Lefschetz theorems are difficult to keep track of. Let the reader keep in mind that the utility of fixed-point theorems in economics, game theory, differential equations, and dynamical systems justifies the sometimes prickly technical machinery that arises.

Example 7.20 (Degree-theoretic Lefschetz) ⊙

The fixed point index $\mathcal{I}$ of a vector field (from §3.4) is intimately related to the Euler characteristic, thanks to Poincaré-Hopf (Theorem 3.5). It is also easily computed as a degree, as per Example 4.23. This perspective lifts to the Lefschetz theorem as well [98]. Consider first the case where $U \subset \mathbb{R}^n$ is open and $f\colon U \to \mathbb{R}^n$ has compact fixed

[2] Nondegenerate means that $\omega^n = \omega \wedge \cdots \wedge \omega \neq 0$ is a volume form.

point set $F = \mathrm{Fix}(f)$. Then the fixed point index, $\mathcal{I}_f(F)$, is defined to be the degree of $\mathrm{Id} - f$:

$$\mathcal{I}_f(F) := \deg\Big((U, U{-}F) \xrightarrow{\mathrm{Id}-f} (\mathbb{R}^n, \mathbb{R}^n{-}0)\Big). \tag{7.5}$$

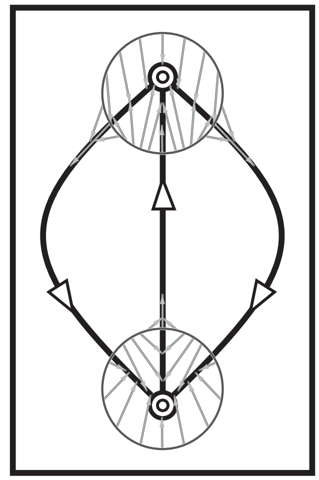

The domain and codomain local homologies (in $\mathbb{Z}$-coefficients) have rank one in dimension n; thus the degree is well-defined and has the usual properties, including homotopy invariance and additivity. Thus, the index can be computed as the sum $\mathcal{I}_f = \sum_F \mathcal{I}_f(F)$, where the sum is over all F disjoint connected components of $\mathrm{Fix}(f)$. When $f\colon X \to X$ is not defined on open neighborhoods in $\mathbb{R}^n$, but only on X a neighborhood retract (suitably small open neighborhoods of X in $\mathbb{R}^n$ retract to X), then, given a neighborhood U of a fixed point component F, $\mathcal{I}_f(F)$ can be computed as $\deg(\mathrm{Id} - f \circ r)$, where $r : U \to X$ is the retraction. In either setting, each term is determined by the local behavior near F. The deep result is that this sum of *local* fixed point indices is equal to the *global* Lefschetz index:

Theorem 7.21 (Lefschetz-Hopf Theorem). *For $X \subset \mathbb{R}^n$ a neighborhood retract and $f\colon X \to K \subset X$ a map to a compact subset K, the Lefschetz index (on $H_\bullet(X;\mathbb{R})$) equals the fixed point index:*

$$\mathcal{I}_f = \tau_f = \sum_k (-1)^k \mathrm{trace}\left(H(f)\colon H_kX \to H_kX\right).$$

The proof is beyond what can be reasonably done briefly: see, *e.g.*, [49, 98]. It is very instructive to draw some pictures of neighborhoods of fixed point sets and compute the degrees *by hand*. It is the best way to see the connection to Morse theory, since some fixed point components are *attracting*, some are *repelling*, and some have *mixed* behavior. ⊚

Example 7.22 (Stratified Morse-theoretic Lefschetz) ⊚

Equation (7.5) is at best of limited use given the requirement that X have an explicit embedding in $\mathbb{R}^n$: at the very least, it is poor form to work extrinsically. An intrinsic approach works with a Morse-type assumption on f, as pointed out in [164]. What follows is a slight reformulation. Assume that $f\colon X \to X$ is definable. Then $\mathrm{Fix}(f)$ is automatically compact and decomposed into connected components.

Following the index pair construction of §7.6, one could define for each connected component F of $\mathrm{Fix}(f)$ an **index pair** (B, E) of compact subsets $E \subset B$ of X satisfying:

1. **Isolation:** $F = \mathrm{Fix}(f|B - E)$ lies in the interior of $B - E$;
2. **Invariance:** $f(E) \cap (B - E) = \varnothing$; and

3. **Exit:** if $x \in B$ and $f(x) \notin B$, then $x \in E$.

For any such index pair, there is an induced map $f : (B/E, \{E\}) \to (B/E, \{E\})$ since points in $B - E$ either remain in $B - E$ or are sent to E. The choice of an index pair is by no means unique; nor is the pointed homotopy type of $(B/E, \{E\})$, in contradistinction to the Conley index. However, remarkably, the analogue of a localized relative Lefschetz index *is* well-defined:

$$\tau_f(F) = \sum_k (-1)^k \text{trace}\left(H(f)\colon H_k(B,E) \to H_k(B,E)\right). \tag{7.6}$$

In this setting, the Lefschetz-Hopf theorem becomes a theorem about the action of f on the local homologies of the fixed point components, relative to the appropriate exit sets:

$$\mathcal{I}_f(F) = \tau_f(F) \quad : \quad \mathcal{I}_f = \sum_F \tau_f(F)$$

There is yet a better version of the theorem that uses the tools of stratified Morse theory. For X compact and $f\colon X \to X$ definable, the fixed point set not only splits into a finite number of compact components; each of these is further stratified into disjoint open simplices on which the local behavior of f is well-defined and ‘constant’ in the sense that for each stratum F_α of $\text{Fix}(f)$, there is a well-defined local Lefschetz index $\tilde{\tau}_f(F_\alpha) \in \mathbb{Z}$. The definition of this index is similar in spirit to Equation (7.6), but for a localized index pair. By defining the local Lefschetz index to be zero off of $\text{Fix}(f)$, one can interpret $\tilde{\tau}_f$ as a constructible function on X. It was shown by Goresky and MacPherson [164] that

$$\tau_f = \int_X \tilde{\tau}_f \, d\chi.$$

This beautiful result is sadly under-appreciated, in part because the construction of $\tilde{\tau}$ (general enough to apply to multivalued mappings $F\colon X \rightrightarrows X$) requires techniques from the theory of sheaves that will only be hinted at in Chapters 9 and 10. ⦾

7.8 Discrete Morse theory

An idea as deep as Morse theory has emanations throughout all of Mathematics. This chapter has focused primarily on the smooth or continuous theory; however, there is a discrete version of Morse theory that has of late yielded powerful results in combinatorics [197, 198], braid groups [118], computational homology [230] and certain problems in computer science. The original papers of Forman [123, 126] are complemented by the recent book of Kozlov [198] and a growing literature.

Consider for concreteness a simplicial or cell complex X. The critical ingredient for Morse theory is *not* the Morse function but rather its gradient flow. A **discrete vector field** is a pairing V which partitions the cells of X (graded by dimension) into pairs $V_\alpha = (\sigma_\alpha \lhd \tau_\alpha)$ where σ_α is a codimension-1 face of τ_α. All leftover cells of X not paired by V are the **critical cells** of V, $\text{Cr}(V)$. A **discrete flowline** is a sequence

(V_i) of distinct paired cells with codimension-1 faces, arranged so that

$$\overbrace{\sigma_1 \lessdot \tau_1}^{V_1} \gtrdot \overbrace{\sigma_2 \lessdot \tau_2}^{V_2} \gtrdot \cdots \gtrdot \overbrace{\sigma_N \lessdot \tau_N}^{V_N} .$$

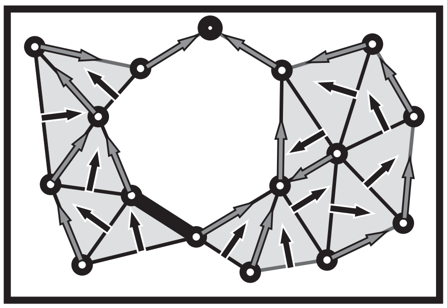

A flowline is **periodic** if $\tau_N \gtrdot \sigma_1$ for $N > 1$. A **discrete gradient field** is a discrete vector field devoid of periodic flowlines.

It is best to lift everything to algebraic actions on the chain complex $\mathcal{C} = (C^{\mathrm{cell}}_\bullet, \partial)$ associated to the cell complex X. By linearity, the vector field V induces a chain map $V : C_k \to C_{k+1}$ induced by the pairs $\sigma \lessdot \tau$ – one visualizes an arrow from the face σ to the cell τ. As with classical Morse homology, $\mathbb{F}_2$ coefficients is simplest; when oriented, one specifies $V : \sigma \mapsto [\sigma : \tau]\tau$ using incidence numbers. The **discrete flow** of V is generated by the degree-zero chain map Φ given by

$$\Phi := \mathrm{Id} + \partial V + V \partial,$$

with iterations of Φ describing how cells descend along the gradient field V. Unlike continuous time flows, the discrete flow has a limit: $\Phi^\infty = \lim_{n\to\infty} \Phi^n$ is constant for n sufficiently large. To every discrete gradient field is associated a discrete Morse complex, $\mathcal{C}^V = (MC_\bullet, \tilde{\partial})$ with MC_k the vector space (or module) with basis the critical cells $\{\sigma \in \mathrm{Cr}(V) \,;\, \dim(\sigma) = k\}$. Note that dimension plays the role of Morse index.

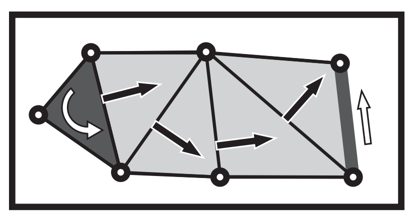

The boundary maps $\tilde{\partial}_k$ count (modulo 2 in the case of $\mathbb{F}_2$ coefficients; with a complicated induced orientation else) the number of discrete flowlines from a critical simplex of dimension k to a critical simplex of dimension $k-1$. Specifically, given τ a critical k-simplex and σ a critical $(k-1)$-simplex, the contribution of $\tilde{\partial}_k(\tau)$ to σ is the number of gradient paths from a face of τ to a coface of σ. In the case that $\sigma \lessdot \tau$, then this number is 1, ensuring that the trivial V for which all cells are critical yields $\mathcal{C}^V$ the usual cellular chain complex. It is not too hard to show that $\tilde{\partial}^2 = 0$ and that, therefore, the homology $MH_\bullet(V) = H_\bullet(\mathcal{C}^V)$ is well-defined. As usual, the difficulty lies in getting orientations right for $\mathbb{Z}$ coefficients.

Theorem 7.23 ([123]). *For any discrete gradient field V, $MH_\bullet(V) \cong H^{\mathrm{cell}}_\bullet(X)$.*

It follows from the proof that the strong Morse inequalities, Equation (7.1), hold with Morse polynomial $M_V(t) = \sum_{\sigma \in \mathrm{Cr}(V)} t^{\dim \sigma}$. Discrete Morse theory, like the Conley index theory, shows that the classical constraints – manifolds, smooth dynamics, nondegenerate critical points – are not necessary. Applications of discrete Morse theory are numerous and expansive, including to combinatorics [198], mesh simplification [208], image processing [252], configuration spaces of graphs [117, 118], and, most strikingly, efficient computation of homology of cell complexes [230].

Example 7.24 (20 questions) ⊚

One of the early applications by Forman was to a problem in decision theory related to evasiveness [124]. Let Δ be an abstract n-simplex on the vertex set $\{v_i\}_0^n$ and let $K \subset \Delta$ be a known subcomplex. There is a *hidden* (*i.e.*, unknown) simplex σ of Δ, and the goal of the *20 questions* game is to determine whether $\sigma \subset K$ by asking questions of the form *"Is v_i in σ?"* Clearly, one can win the game with $n+1$ questions by interrogating each vertex. One says that K is **nonevasive** if there is a strategy for determining whether $\sigma \subset K$ in strictly less than $n+1$ questions, independent of σ; else, K is **evasive**. For example, if $K = \partial\Delta$, then determining if $\sigma \subset K$ is clearly evasive, since one must check that all $v_i \in \sigma$. However, there is only one **evader** – only one simplex τ of Δ has the property that all $n+1$ questions must be asked in order to determine if $\tau \subset K$.

The insight of discrete Morse theory is that a guessing algorithm for determining if $\sigma \subset K$ induces a discrete gradient field on Δ, with the twist that one of the vertices is paired with the formal basepoint $\varnothing$, leading to reduced homology. This yields Morse-theoretic proofs of the following [124]:

1. If K is nonevasive, K collapses to a point.
2. If $\tilde{H}_\bullet(K) \neq 0$, then K is evasive.
3. The number of evaders is $\geq 2\sum_i \dim \tilde{H}_i(K)$.

One clever application of this result is to independence tests for random variables. Let $\mathcal{X} = \{X_i\}$ be a collection of random variables and recall from Example 2.1 the independence complex $\mathcal{I}_{\mathcal{X}} \subset \Delta^n$ of $\mathcal{X}$. Given an unknown subcollection $\sigma \subset \mathcal{X}$ of the random variables, how many trials of the form *"Is X_i a member of σ"* are required to determine if the collection is statistically independent? According to the results cited above, statistical independence is evasive if $\mathcal{I}_{\mathcal{X}}$ is not acyclic: any nontrivial homology class in $\tilde{H}_\bullet(\mathcal{I}_{\mathcal{X}})$ is an obstruction to evasiveness of statistical independence. How many such evasive collections of random variables are there? It is at least twice the total dimension of $\tilde{H}_\bullet(\mathcal{I}_{\mathcal{X}})$. ⊚

7.9 LS category

Given a space, how complicated is it? One means of characterizing topological complexity is co/homology. Another approach might involve critical points and other Morse-theoretic constructs. There is a more primal measure of topological complexity dating back to the work of Lusternik and Schnirelmann in the 1930s that goes under the (suboptimal) name of **category**. This classical measure of complexity for spaces, living in the shadow of Morse theory, informs various contemporary problems ranging from statistics to motion-planning in robotics.

Given X, the **LS category** of X, $\mathrm{LScat}(X)$, is the minimal number $\#\alpha$ of elements in an open cover $\{U_\alpha\}$ of X by sets which are nullhomotopic in X, meaning that

U_α is contractible in X, not that U_α is necessarily a contractible (or even connected) set. If a finite cover does not suffice, one sets $\mathsf{LScat} = \infty$. The **geometric category** of X, $\mathsf{gcat}(X)$, is the minimal such number $\#\alpha$ where each U_α is a contractible set. Invariance of LScat and gcat under, respectively, homotopy type and homeomorphism type, follows from standard results: see [74] for a comprehensive introduction.

Example 7.25 (LS category) ⊚

A sphere has $\mathsf{LScat}(\mathbb{S}^n) = \mathsf{gcat}(\mathbb{S}^n) = 2$ for all n. A compact surface S_g of genus $g > 0$ has $\mathsf{LScat}(S_g) = \mathsf{gcat}(S_g) = 3$. It is a nontrivial exercise to find a space on which LScat and gcat differ. One simple example is the space $X = \mathbb{S}^2 \vee \mathbb{S}^1$ obtained by gluing together $\mathbb{S}^2$ and $\mathbb{S}^1$ at a single point: $\mathsf{gcat}(X) = 3$, but $\mathsf{LScat}(X) = 2$. ⊚

The original motivation for investigating LS category was (in modern parlance) degenerate Morse theory. The Morse inequalities (Corollary 7.5) give a lower bound on the number of critical points of a smooth non-degenerate functional on a manifold M. In the case where the functional is not necessarily non-degenerate or M not necessarily a manifold, the LS category gives the correct lower bound.

Theorem 7.26 ([74]). *Any C^2 function $h\colon M \to \mathbb{R}$ on a compact manifold M must have at least* $\mathsf{LScat}(M)$ *critical points.*

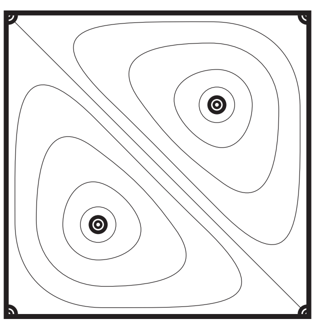

For example, any smooth functional on the 2-torus $\mathbb{T}^2$ must have at least three critical points. For a Morse function, the smallest number of critical points is

$$\sum_k \dim H_k(\mathbb{T}^2; \mathbb{R}) = 1 + 2 + 1 = 4.$$

Given a space X, it is typically difficult to compute the category of X, either geometric or LS. For example, the **Ganea conjecture**, open from 1971 until its disproval in 1998, was the deceptively simple statement that $\mathsf{LScat}(X \times \mathbb{S}^n) = \mathsf{LScat}(X) + 1$ for $n > 0$ and X a smooth closed manifold. Given such subtleties, one adopts a strategy of estimation, which, fortunately, has some reasonable steps.

Theorem 7.27 ([74]). *The LS category of a path-connected CW complex X is bounded by*

$$\mathsf{cup}(X) \le \mathsf{LScat}(X) - 1 \le \dim X.$$

The **cup length**, cup, is the smallest N such that there are N cohomology classes $\alpha_i \in H^\bullet(X)$ with nonzero grading and nonzero cup product $\alpha_1 \smile \cdots \smile \alpha_N \neq 0$. Cup length may depend on the coefficient ring used; the bound above does not. The theorem holds for more general (locally-contractible paracompact) spaces, at the cost of using *covering dimension* in the upper bound. These elementary bounds are the beginning of a rich theory of complexity for topological spaces. It complements (classical) Morse theory in its insensitivity to nondegeneracy.

7.10 Unimodal decomposition in statistics

LS category inspires definitions of topological complexity in several settings. Distributions on $\mathbb{R}^n$ form one excellent example relevant to statistics and mode-counting. Let $\mathfrak{D} = \mathfrak{D}(\mathbb{R}^n)$ denote the set of all compactly supported continuous functions $f\colon \mathbb{R}^n \to [0,\infty)$ and consider the statistical problem of **mode counting**. Given $f \in \mathfrak{D}$, assume it is the result of a sum of basis Gaussian distributions, or **modes**, of unknown mean, variance, and height. How many modes are there? This is an ill-defined question, but the minimal number of such modes is a reasonable measure of the distribution's complexity.

The problem becomes more topological in the coordinate-free setting where the distribution f is not known in terms of a fixed coordinate system, as might occur if the function values are sampled over a network of non-localized sensors. In this context, the following coordinate-free notion of a mode is relevant: $u \in \mathfrak{D}$ is **unimodal** if the non-empty upper excursion sets $u^c = u^{-1}([c,\infty))$ are contractible.

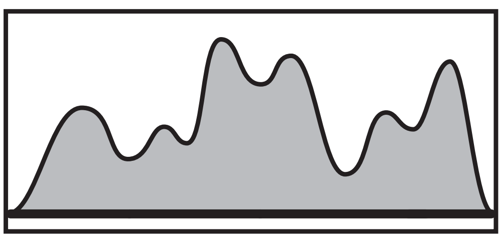

All Gaussians and other typical basis modes are, as the name connotes, unimodal. Following [26], define the **unimodal category** of a distribution $f \in \mathfrak{D}$ to be the minimal number ucat of unimodal distributions u_α, for $\alpha = 1, \ldots, \mathsf{ucat}$ such that f is a combination of unimodals:

$$f(x) = \sum_\alpha u_\alpha(x).$$

Unimodal category is invariant under changes of coordinates, as follows. For $u \in \mathfrak{D}$ unimodal and $\varphi\colon \mathbb{R}^n \to \mathbb{R}^n$ a homeomorphism, $(u \circ \varphi)^c = \varphi(u^c)$, which, being the homeomorphic image of a contractible set, is contractible. Thus, $\mathsf{ucat}(f)$ is a topological invariant of f. In the same way that one lifts the Euler characteristic from subsets of a space to distributions over subsets (integer or real valued) via the Euler integral, one lifts gcat from subsets of a space to distributions thereon. The unimodal category of the constant distribution $\mathbb{1}_U$ is, simply, $\mathsf{gcat}(U)$.

In general, the computation of unimodal category is, as with LScat or gcat, difficult. There is a simple algorithm [26] for the case of a univariate distribution: a greedy sweep of the distribution from left to right (or, via topological invariance, right to left) suffices. The correctness of this algorithm is based on the following result:

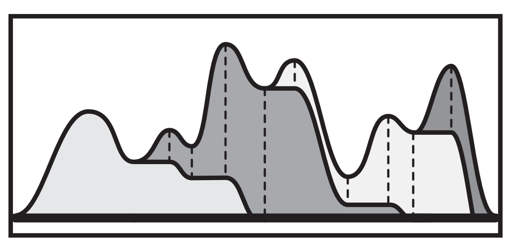

Proposition 7.28 ([26]). *For any $f \in \mathfrak{D}(\mathbb{R})$, $\mathsf{ucat}(f)$ is equal to the maximal number of closed intervals I_k covering the support of f such that*

$$\int_{I_k} f\lfloor d\chi \rfloor < 0$$

for all k.

From Theorem 7.12, this criterion can be translated to critical value data. For distributions over $\mathbb{R}^n$, an algorithmic solution is unknown and appears difficult.

Notes

1. One of the lessons of discrete Morse theory, Conley index theory, and nearly all modern variants of Morse theory is that it is not the *function* but the *dynamics* that matters. The initial emphasis on the Morse Lemma in classical texts was, in this author's opinion, an unfortunate distraction.
2. The subject of topological signal processing is embryonic. The Euler-Fourier transform of Example 7.4, the book of Robinson [256] and a few papers [27, 79, 154, 257, 253] are the starting points of deriving qualitative features of environments via low-fidelity signals. Already, Morse theory seems to play a prominent role (but this may be the author's bias).
3. The intrinsic volumes μ_k of §3.10 are likewise liftable to measures $d\mu_k$ on $\mathsf{CF}(\mathbb{E}^n)$ and then to $\lfloor d\mu_k \rfloor$ and $\lceil d\mu_k \rceil$ on $\mathsf{Def}(\mathbb{E}^n)$ via procedures analogous to those given here for $\lfloor d\chi \rfloor$ and $\lceil d\chi \rceil$. This requires extensive use of currents [28].
4. The definitions involved in the Conley index require more care than is given in the brief overview of §7.6, particularly in defining index pairs, as there are multiple formulations in the literature, with subtle differences in applicability. In the definition of an attractor-repeller pair, it is more proper to use the omega-limit set of the flow, see [167, 258]. The definition as given here is suitable for intuition only, and is but the beginning of a more refined Morse decomposition of the flow.
5. Conley index has been defined for maps (discrete-time dynamics) as well as flows. An **index pair** for f is a pair (B, E) of compact sets satisfying the three properties as listed in Example 7.22, with the exception that the isolation property requires $\mathsf{Inv}(\overline{B-E}; f)$ to be in the interior of $B-E$; invariance and exit properties are the same. The homotopy type of $(B/E, \{E\})$ is not unique. However, one can obtain a well-defined class by looking at the action on this homotopy type up to a certain equivalence [128, 286].
6. The work of Vandervorst *et al.* [155, 156] has adapted Conley and Floer indices to *braids* (see Example 1.8 and §8.3), using, in some cases, the flow of a parabolic PDE on $\mathbb{S}^1$ to set up a stratified Morse theory on the spaces of braids. This leads to some novel examples of forcing, where a single stationary solution to a PDE can force chaotic dynamics, complete with an infinite collection of forced stationary braided solutions.
7. Floer theory is at the moment multifarious, bubbling into many branches of topology, (symplectic, contact, and knot-theoretic). Though the perspective of this chapter is dynamical, much of the current work in Floer theory is symplectic in nature. Among topologists, Floer homology tends to be spoken of as a black box, unfortunately. It remains to incarnate Floer theory into a computational toolset for more directly applied problems. Several authors are progressing to this end [10, 90, 188, 261], but much work remains.
8. Stronger results on discrete Morse theory than presented here are proved in, *e.g.*, [197, 123]. In particular, homotopy-theoretic results about CW complexes are given. The analogue of Forman's discrete Morse theory for differential forms and cohomology is presented in [125]. Forman derived Morse inequalities for arbitrary (non-gradient) discrete fields by counting periodic flowlines properly. A discrete-Morse-theoretic analogue for the Conley index theory appears in the recent monograph of Nicaolescu [240]

(which is also an excellent source for Conley index theory).

9. Farley and Sabalka [117, 118] use discrete Morse theory to explicitly compute the cohomology ring of the (discretized) configuration space $\mathcal{UC}^n(T)$ of unlabeled points on a tree (a cycle-free graph). Their insight was to find a well-suited gradient flow which illuminates a small number of critical cells, the identification of which reveals not merely the homology, but the full cohomology product structure.
10. Mischaikow and Nanda have recently implemented a discrete Morse theory algorithm for computing homology and persistent homology quickly [230].
11. LS category is an abbreviation of Lusternik and Schnirelmann. The abbreviation is convenient vis-a-vis parsimony and frequent variations in spelling past the first letters. The tragic story of these two mathematicians present at the discovery of this invariant is told in [74]. The author apologizes for not using the normalization convention of [74]: the LS category there is one less than that here. The (excellent) motivations for normalizing in this way do not enter the picture in the elementary applications presented here.
12. One can generalize **ucat** to the **unimodal *p*-category** of f – the minimal number $\mathbf{ucat}^p$ of unimodal distributions $u_\alpha, \alpha = 1, \ldots, \mathbf{ucat}^p$ such that f is pointwise an ℓ^p combination of the unimodals:

$$f(x) = \left(\sum_\alpha (u_\alpha(x))^p\right)^{\frac{1}{p}} \quad \text{or} \quad f(x) = \max_\alpha \{u_\alpha(x)\} \text{ when } p = \infty.$$

The unimodal category $\mathbf{ucat} = \mathbf{ucat}^1$ adopts a simple additive model of interference between modes; $\mathbf{ucat}^2$ measures something akin to an *energy* of a distribution; and $\mathbf{ucat}^\infty$ is a natural 'tropical' measure for problems in which mode interference is negligible and the *strongest mode wins.*