Differential Topology in Higher Categories

From regular values and cobordism
to derived smooth families

Anonymous

November 2026

Contents

I Smooth manifolds, maps, and jets

1 Smooth manifolds and their probes **1**
- 1.1 Coordinates and gluing . . . 1
- 1.2 Smooth families and full faithfulness . . . 2
- 1.3 Atlases and interval homotopies . . . 3

2 Tangent functors and smooth maps **5**
- 2.1 First-order probes . . . 5
- 2.2 Rank and normal forms . . . 6
- 2.3 Transverse pullbacks . . . 7

3 Function spaces and smooth families **9**
- 3.1 Compact control and global control . . . 9
- 3.2 Approximation with fixed data . . . 10
- 3.3 The object representing jointly smooth families . . . 12

4 Jets and differential constraints **13**
- 4.1 Finite Taylor data . . . 13
- 4.2 Composition, formal data, and rank . . . 14
- 4.3 Jets at distinct source points . . . 15

II Transversality and normal geometry

5 Critical values and smooth pullbacks **17**
- 5.1 The fiber of a regular value . . . 17
- 5.2 Sard's theorem in the smooth case . . . 18
- 5.3 Boundary strata and openness . . . 19

6 Transversality in parameter families **21**
- 6.1 The projection of an incidence manifold . . . 21
- 6.2 Genericity in a space of maps . . . 22
- 6.3 Boundary and relative constraints . . . 23

7 Bundles, frames, and classifying objects **25**
7.1 Linear gluing and pullback . . . 25
7.2 Finite Grassmannians and stabilization . . . 26
7.3 Classifying the whole groupoid . . . 27
7.4 Orientation as reduction of structure group . . . 28

8 Tubular neighborhoods and collars **29**
8.1 Normal coordinates near an embedding . . . 29
8.2 Isotopy of normalized germs . . . 30
8.3 Collars and neat intersections . . . 31

III Degree, intersections, and Morse theory

9 Degree and oriented fibers **33**
9.1 Local orientations and the signed count . . . 33
9.2 Homotopy and the boundary of a fiber . . . 34
9.3 Families of regular values . . . 35
9.4 Parity and examples . . . 36

10 Intersections and the Euler class **37**
10.1 Oriented transverse pullbacks . . . 37
10.2 Thom classes and zero sets . . . 38
10.3 The excess bundle in a clean intersection . . . 39
10.4 Self-intersection and tangent zeros . . . 40

11 Morse functions and handle filtrations **41**
11.1 The Hessian and the local form . . . 41
11.2 Regular bands . . . 42
11.3 Crossing one critical level . . . 42

12 Trajectory spaces and Morse homology **45**
12.1 Flow lines and the boundary operator . . . 45
12.2 Continuation and homotopies of continuations . . . 46
12.3 The enriched flow category . . . 47
12.4 Two explicit flow categories . . . 47

IV Cobordism, isotopy, and surfaces

13 Geometric bordism and the Thom construction **49**
13.1 Bordisms and their elementary groups 49
13.2 Collapsing a normal neighborhood 50
13.3 Transversality gives the inverse 51
13.4 What stabilization remembers 52

14 Bordism as a higher category **53**
14.1 Gluing and its space of choices 53
14.2 The spectrum of an embedded cobordism category 54
14.3 Corners, duals, and field theories 55

15 Isotopies and ambient transformations **57**
15.1 Extending velocities 57
15.2 Restriction and the stabilizer 58
15.3 Changing a seam by isotopy 60

16 Surfaces and families of models **61**
16.1 Disk and band presentations 61
16.2 Cutting along simple curves 61
16.3 Models in a family 63

V Smooth stacks and parametrized families

17 Smooth spaces and quotient stacks **65**
17.1 Smooth test manifolds 65
17.2 Locally represented geometry 66
17.3 An action groupoid as a quotient 66
17.4 The reflection quotient 67

18 Mapping objects and families over a base **68**
18.1 Smooth mapping objects 68
18.2 Evaluation over a base 69
18.3 Transport over a homotopy type 70

19 Orientation and characteristic classes in families **71**
19.1 Determinants and orientation covers 71
19.2 An exact sequence of relative tangent spaces 72

19.3 An oriented count over a parameter space 72
19.4 Euler classes and pullback 73

20 Discriminants and the topology of families **75**
20.1 A regular total space with singular fibers 75
20.2 Fold and cusp calculations 76
20.3 Transverse jet conditions in a family 77

VI Derived geometry and stratified topology

21 Derived smooth spaces **79**
21.1 The local model of a zero set 79
21.2 The tangent complex 80
21.3 Compatibility as a mapping property 81

22 Derived intersections and excess **83**
22.1 Clean intersections 83
22.2 Local functions in the clean case 84
22.3 The Euler factor 84
22.4 Tangencies and families 86

23 Stratifications and exit paths **87**
23.1 Conical neighborhoods and stratified simplices 87
23.2 Constructible sheaves and their transport 88
23.3 A real fold and a complex branch 88
23.4 Ordinary transport on a regular region 89
23.5 Specialization between local systems 90

24 Spaces of immersions and embeddings **91**
24.1 Formal immersions and regular homotopy 91
24.2 Tests on pairs of points 92
24.3 The configuration stages 94

VII Equivariance, factorization, and persistence

25 Compact group actions and orbit types **95**
25.1 Invariant metrics and fixed points 95
25.2 A slice from the normal bundle 96
25.3 The quotient stack in slice coordinates 97

26 Equivariant transversality and fixed-locus constraints **99**
26.1 The normal map at a stabilizer 99
26.2 Recovery for free source actions 100
26.3 Pullbacks of quotient stacks 102

27 Factorization through an intermediate manifold **103**
27.1 Differential constraints 103
27.2 Obstructions preserved by homotopy 104
27.3 A sharp uniform obstruction on the sphere 105
27.4 Spaces of factorizations 105

28 Persistence of sublevel families **107**
28.1 The filtered diagram 107
28.2 When persistence diagrams are available 108
28.3 Transport at finitely many regular thresholds 109

A Notation **111**

B Subject index **112**

Bibliography **113**

Part I Smooth manifolds, maps, and jets

CHAPTER 1

Smooth manifolds and their probes

A smooth manifold may be specified by an atlas, reconstructed by gluing its charts, or characterized by the smooth families of points that it admits. The three descriptions support different operations. Charts give access to calculus. Gluing exposes compatibility conditions. Smooth families provide a test for morphisms that continues to make sense even when a construction has no manifold representative. We begin with the ordinary smooth objects of Hirsch's first chapter [15, Ch. 1], and pass from these to their functors of points.

1.1 Coordinates and gluing

A smooth n-manifold is a Hausdorff, second-countable topological space with coordinate neighborhoods in $\mathbb{R}^n$ and smooth transition diffeomorphisms. The local dimension is constant within a connected component. For a manifold with boundary, charts instead take values in relatively open subsets of the half-space $[0,\infty)\times\mathbb{R}^{n-1}$, with the usual zero-dimensional convention. Smoothness in these coordinates means local extension to smooth maps on open subsets of the ambient Euclidean space. The boundary is the locus whose charts have first coordinate zero. It is intrinsic, as each transition and its inverse preserve the cone of inward velocities. At a boundary point this cone is a half-space, whereas at an interior point every direction occurs. The derivatives of smooth extensions agree on overlaps by continuity from the interior. Coordinate transitions therefore make the same smoothness test in every chart [18].

A map $f : M \to N$ is smooth if its coordinate expressions are smooth on the open sets where they are defined. Composition is smooth by the coordinate chain rule. The resulting category has a terminal point, and maps out of the point recover the underlying sets. A map of manifolds with boundary need not preserve the boundary. For example, a map from a closed interval that is constant at an interior point is smooth. When an argument requires a map of pairs $(M,\partial M)\to(N,\partial N)$, the boundary condition is part of the structure.

One can construct manifolds from coordinates given in advance. Let U_i be chart domains, let $U_{ij}\subset U_i$ be open subsets, and let $\phi_{ji} : U_{ij}\to U_{ji}$ be smooth diffeomorphisms. Assume the overlap domains and these maps obey the identity,

inverse, and triple-overlap conditions. The quotient of the disjoint union $\coprod_i U_i$ by the resulting identifications has each U_i as an open chart. To see this, the saturate of an open subset of U_i is a union of open sets obtained by the transition diffeomorphisms. The overlap identities ensure that the quotient map is injective on each U_i. When the quotient is Hausdorff and second countable, these charts give a smooth manifold, with the manifold dimension and boundary behavior specified on each chart.

Hausdorffness is a condition on this gluing. Take two copies of the real line and identify them on $\mathbb{R} \setminus \{0\}$. The resulting space has two origins. A neighborhood of either contains the image of some punctured interval, so two neighborhoods of the different origins cannot be disjoint. The coordinate changes are smooth, but the result is outside the category of Hausdorff manifolds. A setting for smooth families can still retain the gluing diagram and its compatibility data.

1.2 Smooth families and full faithfulness

Fix a small site U of coordinate domains, with open covers as its coverings. For boundary arguments we allow relatively open subsets of quadrants $[0,\infty)^k \times \mathbb{R}^{n-k}$ with the same convention of local smooth extension. Morphisms are all smooth maps between the domains. Including quadrants allows finite products of tests. The manifolds in a particular argument need only use the boundary or boundaryless charts required there. Let $\mathcal{E}$ denote the ∞-category of sheaves of spaces on this site.

Associate to a manifold M the discrete-valued functor

$$\underline{M}(U) = C^\infty(U, M), \qquad U \in \mathsf{U}.$$

A map $U \to M$ can be recovered from compatible restrictions to an open cover of U. Its assembled function is continuous because continuity is local, and it is smooth because every point has a smooth local expression. This proves the sheaf condition on $\underline{M}$. If a test U has coordinates $(u_1, \dots, u_d)$, a section of $\underline{M}(U)$ is a smooth d-parameter family, possibly with parameters on the faces.

Proposition 1.1 (Reconstruction from test families)**.** *The association $M \mapsto \underline{M}$ is fully faithful. A natural transformation $\underline{M} \to \underline{N}$ is composition with a unique smooth map $M \to N$. As maps into $\underline{N}$ in $\mathcal{E}$, these form a discrete mapping space.*

Proof. Let $a_i : U_i \to M$ be open chart maps from coordinate domains whose images cover M. A natural transformation τ assigns the smooth map $\tau_{U_i}(a_i) : U_i \to N$ to each of these families. On an overlap, use an open subdomain of U_i and its coordinate change to U_j. Naturality identifies the two restrictions. The maps thus glue to a smooth $f : M \to N$. Given any $u : U \to M$, restrict its

domain to open pieces on which it factors through one of the a_i. Naturality on each piece shows that $\tau_U(u)$ restricts to $f \circ u$. The sheaf condition then gives $\tau_U(u) = f \circ u$. This also proves uniqueness. The representable smooth objects under consideration are 0-truncated sheaves. By definition all mapping spaces into a 0-truncated object are 0-truncated, giving the last assertion. □

The value of $\underline{M}$ at the point remembers only its underlying set. The extra naturality requirement is substantial. For instance, $x \mapsto x^3$ is bijective on points, but the identity plot $u \mapsto u$ of its target has no smooth lift on a neighborhood of 0. The inverse is therefore absent even as a local map of smooth sheaves. Test families detect a singularity of a set-theoretic bijection.

1.3 Atlases and interval homotopies

An atlas takes a universal form in the category of sheaves. Set $A = \coprod_i \underline{U_i}$ in $\mathcal{E}$ and let $A \to \underline{M}$ be induced by the chart maps. Every test-section of $\underline{M}$ locally lifts to A, by pulling back the coordinate cover. In this sense the atlas is an effective epimorphism, meaning that $\underline{M}$ is reconstructed as the colimit of its Čech nerve. The first stages are

$$A \times_{\underline{M}} A \rightrightarrows A \longrightarrow \underline{M}.$$

The iterated fiber products of A in this nerve record all multiple overlaps. The reconstruction can be checked by testing against a sheaf F of spaces. A map $\underline{M} \to F$ is given on plots into M by sections of F. A compatible system on the nerve restricts, for each plot, to compatible sections on its pulled-back open cover, with the higher homotopies supplied on repeated intersections. Sheaf descent glues these into a section over the plot. This gluing is compatible with changes of plot and has a contractible space of choices once the descent data are fixed. Equivalently, mapping out of $\underline{M}$ gives the limit of the mapping spaces out of the nerve.

Two different levels of smooth comparison are relevant. Proposition 1.1 compares maps in the unlocalized sheaf category. A smooth homotopy, on the other hand, is a map $H : M \times [0,1] \to N$ restricting to specified maps at the two endpoints. It is a family parametrized by the smooth interval. For example, the maps $x \mapsto 0$ and $x \mapsto x$ from $\mathbb{R}$ to $\mathbb{R}$ are joined by $(x,t) \mapsto tx$. They are still distinct points of the discrete mapping space of the corresponding unlocalized sheaves. If one subsequently localizes at the projections $X \times [0,1] \to X$, the two endpoint inclusions become equivalent, since both are sections of an inverted arrow. Smooth homotopies then induce equal maps in the homotopy category of the localization. This explains which extra operation makes an interval family

into a homotopical comparison. The shape of a smooth object and its sheaf of plots serve different purposes.

Finally, objects representing plot families can exist without being manifolds of the chosen kind. For instance, the product $\underline{[0,\infty)} \times \underline{[0,\infty)}$ is represented on the chosen site by the quadrant. Its corner has an inward cone with zero-dimensional lineality space, whereas an inward half-plane has one-dimensional lineality and $\mathbb{R}^2$ has two-dimensional lineality. Here lineality means the largest vector subspace contained in the cone. An isomorphism of smooth coordinate models would induce an invertible derivative carrying one cone onto the other. Thus at the corner this product has no local model by an ordinary two-dimensional manifold with boundary. The sheaf product, with maps given by pairs of smooth maps, supplies the expected universal property. Function objects and singular pullbacks will use this enlargement in a similar way [2].

CHAPTER 2

Tangent functors and smooth maps

Differentiation associates a linear map to a smooth map at each point. Its functoriality comes from the fact that each linear map describes the effect on first-order probes. Pointwise linear algebra detects local immersion, submersion, and transversality conditions. Global properties such as embedding also require control of the topology of the image.

2.1 First-order probes

At a point x of a smooth manifold M, let $C^\infty_{M,x}$ be the real algebra of germs of smooth functions at x. A derivation at x is an $\mathbb{R}$-linear map $v : C^\infty_{M,x} \to \mathbb{R}$ satisfying

$$v(ab) = a(x)v(b) + v(a)b(x).$$

Equivalently, it is a unital $\mathbb{R}$-algebra map from $C^\infty_{M,x}$ to $\mathbb{R}[\epsilon]/(\epsilon^2)$ reducing to evaluation at x. The map is $a \mapsto a(x) + \epsilon v(a)$. Multiplication of such expressions gives the derivation rule exactly because $\epsilon^2 = 0$. This is a description by infinitesimal algebra. The ordinary site of Chapter 1 uses smooth coordinate domains as probes. To use the dual numbers as an actual probe object one enlarges the site or otherwise specifies a theory of infinitesimals.

In a chart centered at x, write the coordinate germs as $t_1, \dots, t_n$. The identity

$$a(t) - a(0) = \sum_i t_i \int_0^1 \partial_i a(st)\, ds$$

holds in a sufficiently small coordinate neighborhood, using a smooth extension if x is on the boundary. A derivation thus has $v(a) = \sum_i v(t_i)\partial_i a(0)$. Conversely every choice of the numbers $v(t_i)$ gives a derivation by this formula. This identifies the tangent space T_xM with $\mathbb{R}^n$ in coordinates. At an interior point the same coordinates identify tangent vectors with first derivatives of smooth curves. On the boundary the full tangent space has dimension n. Velocities of curves defined only for nonnegative times form its inward half-space, and two-sided curves through the point have velocities tangent to the boundary if their images lie in the manifold.

For smooth $f : M \to N$, define $D_xf(v)$ by applying v after pulling back the germ from N to M. Composition of pullbacks gives $D_x(gf) = D_{f(x)}g \circ D_xf$.

For the identity the derivative is the identity. Under coordinate changes the derivation coordinates transform by the Jacobian matrices, so gluing $U \times \mathbb{R}^n$ gives the smooth tangent bundle TM and a smooth map $Tf : TM \to TN$. In bundle notation its defining square is

$$\begin{array}{ccc} TM & \xrightarrow{Tf} & TN \\ \downarrow & & \downarrow \\ M & \xrightarrow{f} & N. \end{array}$$

This construction preserves products, allowing corner charts if necessary. A product chart separates the two lists of variables, giving the canonical identification $T(M \times N) \cong TM \times TN$. For example, the tangent map of multiplication $m : \mathbb{R}^2 \to \mathbb{R}$ satisfies $D_{(x,y)}m(v, w) = yv + xw$. This formula follows as well by multiplying $x + \epsilon v$ and $y + \epsilon w$ in the dual numbers.

2.2 Rank and normal forms

An immersion is a smooth map with injective derivative at every point. A submersion has surjective derivative. We use the constant-rank theorem for boundaryless coordinate neighborhoods in the following form [15, Ch. 1].

Proposition 2.1 (Local normal form)**.** *Suppose f is a smooth map between boundaryless manifolds, of constant rank r near x. There are source and target coordinates in which $f(u, v) = (u, 0)$, with $u \in \mathbb{R}^r$.*

Proof. Work near the origin and choose coordinates so the first r components have an invertible minor in the first r variables. Use $(f_1, \dots, f_r, t_{r+1}, \dots, t_m)$ as a change of source coordinates, by the inverse function theorem. The map is now $(u, v) \mapsto (u, h(u, v))$. Since the derivative has rank exactly r, its columns in the v variables, whose first r entries vanish, must themselves vanish. Restricting to a coordinate box shows that $h(u, v) = h(u, 0)$ by integration along the v-segments. Subtracting $h(u, 0)$ from the last target coordinates gives the stated form. Rank zero is the same argument with an empty first minor. □

Full rank is open, so the hypothesis applies near every point of an immersion or a submersion in the interior. For an immersion the form is a coordinate inclusion. For a submersion it is a coordinate projection. Boundary versions require hypotheses on the boundary behavior. For instance $t \mapsto t^2$ on the half-line has a critical boundary point even though it is a submersion on the interior. Requiring a boundary-preserving map between half-spaces does not by itself ensure a smooth inverse up to the boundary.

An embedding is an immersion that is a homeomorphism onto its image with the subspace topology. The irrational-slope map $\mathbb{R} \to \mathbb{R}^2/\mathbb{Z}^2$ given by

$t \mapsto (t, \alpha t) \bmod \mathbb{Z}^2$ illustrates the second condition. For irrational α it is an injective immersion. There are unbounded integers n_j with αn_j tending to zero modulo $\mathbb{Z}$. Indeed the pigeonhole principle applied to the fractional parts of $0, \alpha, \ldots, N\alpha$ gives a nonzero integer of size at most N whose phase is within $1/N$ of zero. A fixed finite set of nonzero integers cannot do this as N grows. The images of n_j tend to the image of 0, so the inverse on the image is discontinuous.

A proper injective immersion between the manifolds here is an embedding. To check that it is closed, take y in the closure of the image of a closed set. In the first-countable target there is a sequence $y_j \to y$ in that image. The sequence together with its limit is compact, so the preimages of its terms lie in a compact set by properness. A convergent subsequence of the preimages has limit in the closed set and maps to y. Thus the immersion is closed, and injectivity makes the inverse on its image continuous. For an embedded $S \subset M$, the quotient bundle $TM|_S/TS$ is its normal bundle. A metric identifies it with the orthogonal complement, whereas the quotient comes directly from the embedding.

At the boundary a useful class is given by *neat embeddings*. In this chapter this means an embedding $f : M \to N$ with $f^{-1}(\partial N) = \partial M$ and with derivative transverse to ∂N at each point of ∂M. The boundary normal form is a coordinate inclusion preserving the first coordinate. In fact, let t and s denote source and target boundary coordinates, with $t, s \geq 0$. In a neighborhood of a boundary point the first component of f is $tb(t, u)$ for a smooth b by integrating its derivative in t, since the component vanishes at $t = 0$. Transversality and nonnegativity give $b > 0$ after shrinking. The other target coordinates can be linearly recombined along ∂N so that the first $\dim M$ components, including s, have an invertible derivative. They become source coordinates by the inverse function theorem on smooth extensions. The sign of $tb(t, u)$ ensures this coordinate change respects the half-spaces. Subtracting the graph of the remaining components gives the coordinate inclusion by a change that fixes s. This proves the claimed local form and exhibits the boundary hypothesis needed for it.

2.3 Transverse pullbacks

Let $f : M \to Z$ and $g : N \to Z$ be smooth maps of boundaryless manifolds. At a coincidence (m, n) with $f(m) = g(n) = z$, their derivatives define the linear map

$$(u, v) \longmapsto D_m f(u) - D_n g(v) \quad \text{from } T_m M \oplus T_n N \text{ to } T_z Z.$$

When this is surjective at every coincidence, call the pair transverse. In a target chart the coincidence condition is the zero set of the difference of the coordinate expressions. Its derivative on the zero set is precisely the displayed map, in coordinates. The local submersion normal form gives the solution set a smooth structure and its dimension $\dim M + \dim N - \dim Z$ when nonempty. The change

of target chart does not change this embedded submanifold structure, since it is the subspace preimage of the diagonal with the induced charts. A smooth map into $M \times N$ meeting this set factors through it smoothly by its coordinate inclusion. Hence this is the smooth pullback $M \times_Z N$. Its tangent exact sequence at (m, n) is

$$0 \longrightarrow T_{(m,n)}(M \times_Z N) \longrightarrow T_m M \oplus T_n N \xrightarrow{Df - Dg} T_z Z \longrightarrow 0.$$

The initial tangent term is the fiber product of the tangent spaces over $T_z Z$. This explains the precise sense in which differentiation preserves this pullback.

For $f(s, t) = s^2 + t^2$ and g the inclusion of 1 in $\mathbb{R}$, the pullback is the unit circle, with tangent line $su + tv = 0$ at (s, t). Replacing 1 by 0 leaves the singleton $\{(0, 0)\}$, while the linear equation given by the derivative has the whole plane as its kernel. Both pullbacks exist in manifolds in this particular example. At the singular value the ordinary tangent to the solution set loses the first-order equation. Testing with an infinitesimal point sends (s, t) to $(a\epsilon, b\epsilon)$ and makes this difference explicit, since $s^2 + t^2 = 0$ in the dual numbers for all a, b. Later intersection constructions will distinguish the agreement of smooth families from the infinitesimal information of a singular equation.

CHAPTER 3

Function spaces and smooth families

Convergence of smooth maps can be imposed on one compact set at a time or simultaneously over the entire source. A smooth family of maps has another kind of structure: it is evaluated jointly in its parameter and source variables. Distinguishing these structures is especially important for noncompact sources. We follow Hirsch's distinction between the weak and strong topologies [15, Ch. 2], then describe the object that represents jointly smooth families.

3.1 Compact control and global control

For a fixed nonnegative integer r, write $J^r(M,N)$ for the bundle of order-r jet data. A jet at a point consists of the value and the coordinate derivatives through order r, interpreted in all charts by the chain rule. Chapter 4 proves the basic properties of this bundle. The weak C^r topology is the compact-open topology on jet prolongations. In particular

$$\{\, g : j^r g(K) \subset O \,\}, \qquad K \subset M \text{ compact}, \quad O \subset J^r(M,N) \text{ open},$$

is a subbasic set. A neighborhood of f in this topology requires such a condition on finitely many compact sets. After finitely many charts are chosen around the compact jet data, these neighborhoods amount to uniform coordinate control through order r on compact subsets.

For the strong C^r topology we allow a locally finite family (K_i) of compact subsets of M, with an open jet neighborhood O_i for each. The basic condition is $j^r g(K_i) \subset O_i$ for every i. This is the same topology as that given by tubes about jet graphs, namely sets of maps whose entire jet graph lies in one open set of $J^r(M,N)$. To see one direction, let π be the source map of the jet bundle. The complement of $\bigcup_i (\pi^{-1}(K_i) \setminus O_i)$ is open, since the union is locally finite and each term is closed. Requiring a graph to lie in that complement imposes exactly the compact-set conditions. For the other direction, use a locally finite compact cover of M and impose the tube condition on each compact member. Such a cover is available from a compact exhaustion or from a locally finite refinement by precompact coordinate domains. Throughout this discussion the order r is fixed. A smooth map can be tested at each finite order without committing to a single topology that controls indefinitely many orders at infinity.

Consider a smooth bump function $b : \mathbb{R} \to [0,1]$ supported in $(-1,1)$ with $b(0) = 1$. The functions $g_j(x) = b(x - j)$ converge weakly to zero through every finite order, since their supports eventually miss any compact set. They miss the strong neighborhood of zero given by the condition $|g(x)| < 1/2$ at every x. On a compact source every locally finite family of nonempty compact sets is finite, so the two topologies agree. The bump sequence shows why the compactness hypothesis cannot simply be dropped.

Baire category arguments commonly use spaces in which an exhaustion allows a completeness proof. For instance, on $C^r(\mathbb{R}^m, \mathbb{R}^n)$ take the seminorm p_j of uniform control of all partial derivatives through order r on the closed ball of radius j. A compatible complete metric for the weak topology is

$$d(f,g) = \sum_{j \geq 1} 2^{-j} \min\{1, p_j(f - g)\}.$$

For a Cauchy sequence the functions and each derivative converge uniformly on each ball. The fundamental theorem of calculus along coordinate segments shows that the limits of the derivatives are the corresponding derivatives of the function limit. Induction on the order gives a global C^r limit. This proves completeness and the Baire property in this setting. On smooth Euclidean-valued maps one can use the same construction with p_j controlling orders up to j, proving completeness in the weak smooth topology by the same argument. Hirsch proves more generally that the weak topology on $C^r(M,N)$ admits a complete metric, and every weakly closed subset is a Baire space for the induced strong topology, using closure of holonomic jets and completeness on the target of jet maps [15, Ch. 2, §4]. In later genericity arguments one has to justify the topology and the Baire property of the actual class of maps being perturbed.

3.2 Approximation with fixed data

Here is a useful approximation statement in a form that separates the local analytic argument from the geometric target.

Proposition 3.1 (Relative smoothing)**.** *Let M, N be ordinary smooth manifolds without boundary and let $f : M \to N$ be C^r for finite r. If f is smooth on a neighborhood of a closed set $A \subset M$, it can be approximated in any strong C^r neighborhood by a smooth g agreeing with f on a possibly smaller neighborhood of A. The approximation can be joined to f by a homotopy fixed near A.*

Proof. For an $\mathbb{R}^L$-valued map F take a locally finite smooth partition of unity by functions ψ_i with compact support in coordinate domains. In each coordinate chart mollify $\psi_i F$ with sufficiently small support radius. The mollified functions $\widetilde{F}_i$ are smooth and supported in slightly larger compact subsets of the same

coordinate domains, chosen locally finite from the start. Since convolution commutes with derivatives through order r, their differences from $\psi_i F$ approach zero in C^r in the coordinate chart. Extend by zero and sum.

To meet a given strong neighborhood, use locally finite compact controls on the sum. Each enlarged support meets only finitely many of these control sets. On each meeting, derivatives in the control charts are estimated by the chart derivatives using the chain rule on compact subsets. Choose the smoothing radius for the ith summand so its contribution to each relevant constraint is at most 2^{-i-1} times the allowed error there. The sum then lies in the prescribed neighborhood. If F is smooth near A, first choose a smooth cutoff χ with $\chi = 1$ near A and with support inside that smooth locus. Replace the sum $\widetilde{F}$ by $\chi F + (1-\chi)\widetilde{F}$. Multiplication by $1-\chi$ has a bounded derivative estimate on each compact constraint. Choosing smaller local errors gives the same approximation with the exact agreement required.

For a manifold target use a proper smooth embedding $i : N \hookrightarrow \mathbb{R}^L$ and a tubular retraction $\rho : V \to i(N)$ from an open neighborhood [15, Chs. 2, 4]. Apply the Euclidean argument to $F = i \circ f$. Before smoothing, tighten the compact controls so that the straight segments between F and its approximation remain in V. This is possible on each compact set since F maps it into a compact subset of the open set V. The chain rule and compact control also allow the approximation to be chosen small enough so that $\rho \circ \widetilde{F}$ meets the original jet conditions. Composing the straight-line homotopy with $i^{-1}\rho$ gives the assertion, including the fixed neighborhood. □

Real-analytic approximation has another source. Local partitions of unity cannot serve as analytic cutoffs, since an analytic function on a connected domain that vanishes on a nonempty open set vanishes identically. A compact interval illustrates a different mechanism. If f is C^r on $[a, b]$, approximate $f^{(r)}$ uniformly by a polynomial q. For $r \geq 1$ integrate q successively r times, choosing the constants to match f and its lower derivatives at a. The resulting polynomial approximates the jth derivative for $0 \leq j < r$ with error bounded by

$$\|q - f^{(r)}\|_\infty \frac{(b-a)^{r-j}}{(r-j)!}.$$

This follows by integrating the derivative error at each step. For $r = 0$ it is the original polynomial approximation. On general real-analytic manifolds the global approximation theorem uses complex analytic methods. Hirsch states the density of analytic maps in the strong topology for targets and sources equipped with analytic structures [15, Ch. 2, §5]. Density does not impose exact agreement with arbitrary prescribed smooth data.

3.3 The object representing jointly smooth families

Let $\mathcal{E}$ be the category of sheaves of spaces from Chapter 1. It has an exponential $[\underline{M}, \underline{N}]$, characterized by a natural equivalence

$$\mathrm{Map}(S, [\underline{M}, \underline{N}]) \simeq \mathrm{Map}(S \times \underline{M}, \underline{N}).$$

If $S = \underline{U}$ is a coordinate probe, the right-hand side is the discrete space of jointly smooth maps $U \times M \to N$. To see this, the product is represented by the usual product manifold, allowing corners for boundary probes, and test-family reconstruction applies to it. The exponential therefore treats a family and its evaluation under a single universal property. Unlike a topology alone, it specifies which parameter tests have smooth evaluation [2].

A jointly smooth family $F : U \times M \to N$ is continuous into the weak C^r topology. For a point u_0 and compact set K, joint continuity of the partial derivatives in the M variables and compactness of K yield one parameter neighborhood for each required compact control. This proof does not supply infinitely many controls at once. For $F(s, x) = sx$ on $\mathbb{R} \times \mathbb{R}$, the parameter family is not continuous at $s = 0$ into the strong topology: the strong neighborhood $|g(x)| < 1$ on the whole source contains $F(0, -)$ and no other member. Thus the smooth-family structure does not implicitly choose the strong topology.

When M is compact, familiar infinite-dimensional charts explain the first-order geometry of this function object. Choose a Riemannian metric on N. For g sufficiently close to f in C^0, the inverse to the exponential map near the zero section is defined for all pairs $(f(x), g(x))$, by compactness of $f(M)$. The resulting section of f^*TN is a local coordinate for g, and its inverse sends a small section v to $x \mapsto \exp_{f(x)} v(x)$. For a smoothly evaluated curve through f, differentiating in its curve parameter gives a smooth section of f^*TN. Every such section occurs as a velocity, since compactness permits one common small time interval for this exponential construction. This gives the first-order tangent model $\Gamma(M, f^*TN)$. A full manifold structure on the space of maps requires an infinite-dimensional calculus. For the compact-source convenient model and its exponential law see [17].

CHAPTER 4

Jets and differential constraints

A map and its first derivative determine whether it is an immersion or a submersion at a point. A second derivative can distinguish a nondegenerate critical point from a more singular one. A jet is the portion of a map needed for such finite-order questions. Its construction turns differential conditions into geometric conditions on a bundle, which can then be tested on local families.

4.1 Finite Taylor data

Use boundaryless manifolds for the constructions here. Two germs of smooth maps $M \to N$ at x have the same r-jet if their values are the same and their coordinate expressions have equal partial derivatives through order r at x. In charts centered at x and at that common value y, this is equality of Taylor polynomials through order r. The chain rule for compositions of Taylor polynomials, discarding terms of degree $> r$, shows that the equivalence relation is independent of the charts. The same calculation shows that changing charts gives a smooth change of jet coordinates. Such a change uses derivatives of the coordinate transitions evaluated at x and y, and polynomial expressions in the given coefficients. Thus the classes form a smooth bundle $J^r(M, N)$ over $M \times N$. At order zero the bundle is $M \times N$. The first-jet bundle has fiber $\mathrm{Hom}(T_xM, T_yN)$ at (x, y).

For $r \geq 1$ the truncation $J^r(M, N) \to J^{r-1}(M, N)$ is an affine bundle. If two jets lie over the same lower jet at (x, y), the difference of the rth-order coefficients defines an element of

$$\mathrm{Sym}^r(T_x^*M) \otimes T_yN.$$

Indeed, in a target coordinate change the linear term of the change acts on the difference by its derivative at y. Any nonlinear term that would contain a difference of order r has another factor of positive order and contributes only in order $> r$. Source coordinate change acts on the difference by the rth symmetric power of the derivative of the source change. Terms depending only on lower derivatives cancel. The tensor difference is therefore intrinsic. In coordinates one can add any homogeneous degree-r term to a fixed Taylor polynomial, and that polynomial is the germ of a smooth map to the target chart after restricting the source. This gives the free and transitive action on each fiber that defines the affine structure. The resulting tower and its applications to singularities are developed in [11].

A map $f : M \to N$ has a prolongation $j^r f : M \to J^r(M, N)$, given by its jet at each point. Over $M \times N$ it lies over the graph of f:

$$\begin{array}{ccc} M & \xrightarrow{j^r f} & J^r(M,N) \\ & \searrow^{(\mathrm{id}_M, f)} & \downarrow \\ & & M \times N. \end{array}$$

If $F : P \times M \to N$ is smooth, differentiating only in the M variables gives a smooth family $(p, x) \mapsto j^r_x F_p$. It commutes with every smooth parameter change $P' \to P$. This is the relative prolongation used when the parameters sample different maps from the same source.

4.2 Composition, formal data, and rank

Write $j^r_x g$ for the jet of a germ with specified source point x. For a composable pair of germs, truncated composition gives

$$j^r_{f(x)} g \circ j^r_x f = j^r_x (g \circ f).$$

This is well-defined for arbitrary representatives of the two classes, by the Taylor argument above. Associativity follows from composition of germs. Consequently postcomposition by a fixed smooth map gives a map of jet bundles, also changing the target value. Precomposition by a fixed $h : M' \to M$ first pulls back jet data by its source map h and then uses the jet of h at each point. The two operations commute. Between points of a fixed manifold, the composable jets of germs of local diffeomorphisms, for $r \geq 1$, form a groupoid. Inverse jets come from smooth local inverses. In coordinates, taking the inverse of an invertible linear term and solving successively in each higher degree gives smooth inversion on this groupoid.

A smooth section of $J^r(M, N) \to M$ is called formal differential data. It is holonomic when it equals $j^r f$. For $M = N = \mathbb{R}$, coordinates (x, y, a) on J^1 designate source, value, and derivative. A section $(x, y(x), a(x))$ is holonomic exactly when $dy - a\,dx$ pulls back to zero. The choice $y(x) = 0$ and $a(x) = 1$ fails this condition. Thus a bundle of possible derivatives does not by itself supply a space of solutions by ordinary maps. A global example is the first-jet section over S^1 with value identically zero and derivative the angular 1-form $d\theta$. It specifies an injective linear map on each tangent line. There can be no smooth immersion $S^1 \to \mathbb{R}$, since a global smooth real function on the circle has a critical point. In particular this everywhere regular formal data cannot be deformed through formal immersions to a holonomic one between these fixed manifolds.

The rank strata of first jets provide useful smooth submanifolds. For a linear map between spaces of dimensions m, n, the locus of rank ℓ has codimension

$(m-\ell)(n-\ell)$ in the space of matrices. To prove this, choose a nonvanishing $\ell \times \ell$ minor A and decompose a nearby matrix into blocks

$$\begin{pmatrix} A & B \\ C & D \end{pmatrix}.$$

The rank is ℓ exactly when $D - CA^{-1}B = 0$. In these coordinates $(A, B, C, D - CA^{-1}B)$ is a smooth coordinate change, so the displayed condition imposes $(m-\ell)(n-\ell)$ independent equations. Every rank-ℓ matrix belongs to one of these charts. At rank zero the condition is that all entries vanish. This proof applies in local bundle frames, giving the strata in $J^1(M, N)$.

For example, the family $f_a(x) = x^3 - ax$ has derivative $3x^2 - a$. The total parameter-source family meets the zero-derivative stratum transversely, since differentiating that equation in a gives -1. Its critical-point locus is therefore the smooth parabola $a = 3x^2$. The projection of this locus to parameters has a critical point at the origin. For $a > 0$ the fixed map has two critical points, for $a < 0$ it has none, and for $a = 0$ the two branches meet at a degenerate critical point. Regularity of a total family and regularity of each parameter fiber are separate questions even in this elementary example.

The construction admits a description by smooth families. Let $\mathcal{F}_r(U)$ consist of the smooth maps $U \times M \to J^r(M, N)$ over the projection to M. These sets are a sheaf in U, and relative prolongation is a natural map $[\underline{M}, \underline{N}] \to \mathcal{F}_r$. Taking the value coordinate gives a retraction, since every relative jet has the original map as its zeroth component. The subfunctor of relative holonomic data consists precisely of the image. Indeed, if formal data are holonomic over an open cover of a parameter-source product, their local realizing maps must be the value coordinate there. These maps consequently glue, and their prolongations agree on overlaps. Holonomicity is a local compatibility of derivatives with values, rather than a restriction on the topology of the parameter alone.

4.3 Jets at distinct source points

For a positive integer q let $M^{(q)}$ be the ordered configuration space of q distinct points of M. Over it the order-r multijet bundle has fiber at $(x_1, \dots, x_q)$ equal to the product of the jet fibers at the specified source points, with values free in N. The multijet $j^r_q f$ is evaluation of $j^r f$ in all the slots. Its domain excludes repeated source points. The purpose of this exclusion is that local perturbations can be varied independently in the slots.

Here is the local independence argument. Fix a map f and distinct points x_i. Take disjoint coordinate neighborhoods of the x_i with f mapping each into a target chart. In each source coordinate centered at x_i, choose a bump function equal to one near the center and supported compactly within the neighborhood.

Add a polynomial of degree at most r times this bump to the target coordinates of f. For small polynomial coefficients all perturbed points still lie in the target chart, since the change is made on a compact support. Continue by f outside the supports. At x_i the derivatives of the change are exactly the derivatives of the chosen polynomial, and at all other chosen points the jet is unchanged. Therefore the derivative of multijet evaluation in the polynomial coefficients surjects onto the product of the vertical tangent spaces of the jet bundle at the configuration. This finite-dimensional family is the input to multijet transversality arguments [11, Ch. II].

At order zero and $q = 2$, the pair of values $(f(x_1), f(x_2))$ can be tested against the diagonal of $N \times N$. If this test is transverse, the double-point locus away from the source diagonal is a submanifold of dimension $2 \dim M - \dim N$ when nonempty, by the transverse pullback calculation of Chapter 2. The rank condition is a one-point test. Accidental equality of values is a two-point test. The distinction explains why injectivity questions require multijets even when the immersion condition is already settled.

CHAPTER 5

Critical values and smooth pullbacks

A level equation $f(x) = y$ asks for a pullback by the point $y : * \to N$. To identify this pullback with a smooth manifold of the expected dimension, the crucial hypothesis is a rank condition on the derivative along the fiber. Sard's theorem makes that condition available for almost every choice of the point. The point of view in this chapter is to keep the rank test and the genericity argument distinct, so that each can subsequently be applied to families.

5.1 The fiber of a regular value

Let M, N be finite-dimensional second-countable smooth manifolds without boundary, with $\dim M = m$ and $\dim N = n$. A point $x \in M$ is *critical* for a smooth map f if $df_x : T_xM \to T_{f(x)}N$ fails to be surjective. The *critical values* form the image of the critical points. A regular value is a point outside this image, so a point outside $f(M)$ is a regular value with empty fiber.

Theorem 5.1 (Regular preimages)**.** *If y is a regular value, then $f^{-1}(y)$ is an embedded submanifold, of dimension $m-n$ when nonempty. At a point $x \in f^{-1}(y)$ its tangent space is $\ker df_x$. More generally, for an embedded submanifold $W \subseteq N$ of codimension c, suppose $df_x(T_xM) + T_{f(x)}W = T_{f(x)}N$ whenever $f(x) \in W$. Then $f^{-1}(W)$ is an embedded submanifold of codimension c when nonempty, with tangent space $df_x^{-1}(T_{f(x)}W)$.*

Proof. For the general statement choose a local submersion $q : U \to \mathbb{R}^c$ defining W near $f(x)$, so that $W \cap U = q^{-1}(0)$. The displayed hypothesis says exactly that $d(qf)_x$ is surjective. In source coordinates whose first c columns give an invertible minor, augment qf by the remaining source coordinates. The derivative of the augmented map is invertible, so by the inverse function theorem its components are local coordinates. In these coordinates the preimage is the zero set of the first c coordinates. Differentiating the equations identifies its tangent space with $\ker d(qf)_x = df_x^{-1}(T_{f(x)}W)$. The case $W = \{y\}$ has $c = n$. The local descriptions use the subspace topology and thus assemble to the asserted embedded structures. □

The square in the theorem is a smooth pullback.

$$\begin{array}{ccc} f^{-1}(W) & \longrightarrow & W \\ \downarrow & & \downarrow \\ M & \xrightarrow{f} & N. \end{array}$$

Indeed, a smooth test map into M whose image lands in $f^{-1}(W)$ factors smoothly through it. This follows from the adapted coordinates in the proof. Testing against all smooth manifolds gives the same pullback square of representable sheaves. Such sheaves are 0-truncated, so their pullback in sheaves of spaces is also computed by the ordinary fiber-product condition. Representability of that sheaf pullback requires a separate argument. For $f(x) = x^2$ and $y = 0$, the fiber sheaf is represented by a point, but $df_0 = 0$ and its kernel is one-dimensional. A derived enhancement of a smooth equation can record this failure of the differential rank test.

5.2 Sard's theorem in the smooth case

A subset of a manifold is *null* if its intersection with each coordinate chart has Lebesgue outer measure zero in that chart. This definition is invariant under smooth change of coordinates because the change is locally Lipschitz on a countable cover by compact coordinate boxes. A smooth positive density induces the same local class of null sets.

Theorem 5.2 (Sard)**.** *For a smooth map between finite-dimensional second-countable manifolds without boundary, the critical values are null. In particular the regular values are dense.*

Proof. The statement is local in source and target. Use countable charts with the image of each source chart in a target chart, and work with $f : U \to \mathbb{R}^n$ for $U \subseteq \mathbb{R}^m$ open. The case $n = 0$ has no critical points. For $n \geq 1$, induct on m, with the case $m = 0$ giving at most a single image point on each chart. Write C for the critical set and C_j for the subset on which every partial derivative of every component of f, of orders 1 through j, vanishes. Thus $C_1 \subseteq C$.

Positive derivative rank. Near a point of $C \setminus C_1$ there is a nonzero first partial of some component. After permuting coordinates, use that component in place of the first source coordinate. The inverse function theorem makes this a local coordinate change and the new expression for f has the form

$$f(t, u) = (t, F(t, u)).$$

If $n = 1$, there are no critical points in this chart. Otherwise the derivative has rank $1 + \operatorname{rank} d_u F(t, u)$. At each t, the corresponding section of the image of the

critical set lies among the critical values of $u \mapsto F(t, u)$. It is null by induction. To apply Fubini, exhaust the chart by countably many compact subsets and intersect the critical set with each one. Since criticality is a closed condition, the images of these intersections are compact, and their slices are measurable. Fubini gives measure zero on each compact image. Countably many such charts cover $C \setminus C_1$.

A first nonvanishing higher derivative. Near any $x \in C_j \setminus C_{j+1}$, for $j \geq 1$, some scalar partial derivative g of order j has nonzero differential. It vanishes on C_j, so the portion of C_j in this neighborhood lies in the smooth hypersurface $g^{-1}(0)$ through x. At every point of C_j there, $df = 0$, and therefore the restriction of f to the hypersurface is critical at each of these points. Its source dimension is $m - 1$, so the induction hypothesis gives a null image of this part of C_j. A countable collection of hypersurface charts covers $C_j \setminus C_{j+1}$.

High-order vanishing. Choose j with $(j + 1)n > m$. On a closed coordinate box $K \subset U$, fix a bound for all derivatives of order $j + 1$. Subdivide the box into $O(\epsilon^{-m})$ pieces of diameter at most a constant times ϵ. For each piece meeting C_j, expand at one such point. Taylor's formula on the segment inside K bounds the oscillation from that point by $O(\epsilon^{j+1})$. Thus $f(C_j \cap K)$ is covered by $O(\epsilon^{-m})$ balls of radius $O(\epsilon^{j+1})$. Their total n-volume is $O(\epsilon^{(j+1)n-m})$, which tends to zero. A countable box cover handles all of C_j. The three cases exhaust C. □

This smooth proof follows the induction underlying [24, Section 3]. Hirsch also treats finite differentiability. For C^r maps, the sufficient threshold is $r > \max\{m - n, 0\}$ [15, Chapter 3]. In the smooth argument above, choosing a high order of vanishing helps obtain the last covering estimate. The sharper finite-differentiability threshold uses a different argument. Sard's conclusion concerns the image of the critical set. The set itself can contain open regions, as for any map constant on an open set.

5.3 Boundary strata and openness

Suppose M has boundary and N has none. If a value is regular for both $f|_{\operatorname{int} M}$ and $f|_{\partial M}$, the fiber is a manifold with boundary, with boundary equal to the fiber of the second restriction. At a boundary point write a source chart as (u, t) with $t \geq 0$. The surjectivity on the boundary stratum lets us choose the invertible minor in the u-directions. The coordinate change in the regular preimage proof then leaves t unchanged. Apply Sard's theorem to the two strata and take a union of null sets to see that almost every value has this property. Testing only the ambient tangent space at a boundary point is insufficient. For the inclusion of $[0, \infty)$ in $\mathbb{R}$, the ambient derivative at 0 is surjective, but the boundary restriction is critical there.

For a proper f without boundary, the critical set is closed and its image is closed, so the regular values are open as well as dense. Properness cannot be discarded in this statement. For a connected example, let $f : \mathbb{R} \to \mathbb{R}$ be $f(x) = e^{-x}(2 + \sin(2x))$. It is strictly positive, so 0 is a regular value with empty fiber. At $x = j\pi$ the derivative

$$f'(x) = e^{-x}\big(2\cos(2x) - 2 - \sin(2x)\big)$$

vanishes, and the resulting critical values $2e^{-j\pi}$ approach zero as $j \to \infty$. This illustrates why openness statements for families typically include compactness or control of the behavior at infinity.

Example 5.3 (Completing the tangent exact sequence)**.** For the function $f(x_1, \ldots, x_m) = \sum_i x_i^2$, every positive real number a is regular and its fiber is the sphere of radius $\sqrt{a}$. At a point x of the sphere,

$$0 \longrightarrow T_x f^{-1}(a) \longrightarrow \mathbb{R}^m \xrightarrow{v \mapsto 2\langle x, v\rangle} \mathbb{R} \longrightarrow 0$$

is exact. Dividing by $\|x\|^2$ constructs a splitting in the radial direction, which varies smoothly for $a > 0$. At $a = 0$ the ordinary fiber contracts to the origin and the derivative has nontrivial cokernel. The comparison exhibits the information in the linearized equation that its set of solutions cannot retain.

CHAPTER 6

Transversality in parameter families

The preceding chapter used regular values to vary a point in the target. A family of maps allows another kind of variation. The parameter can repair a normal derivative that fails to be onto in the input variables. Parametric transversality makes precise when this collective freedom gives individual members with well-behaved pullbacks.

6.1 The projection of an incidence manifold

Let P, M, N be finite-dimensional second-countable smooth manifolds without boundary, and let $W \subset N$ be an embedded submanifold. Say $f \pitchfork W$ when f satisfies the normal surjectivity condition in Theorem 5.1. Write $F_p(x) = F(p, x)$ for the members of a smooth family $F : P \times M \to N$.

Theorem 6.1 (Parametric transversality)**.** *If $F \pitchfork W$, the incidence space $Z = F^{-1}(W)$ is an embedded submanifold. For $\pi : Z \to P$ the restricted projection, p is a regular value of π exactly when $F_p \pitchfork W$. The set of such parameters is conull, dense, and residual. If π is proper, it is also open.*

Proof. Theorem 5.1 constructs Z. At $(p, x) \in Z$, write $Q = T_{F(p,x)}N / T_{F(p,x)}W$. The derivative followed by projection to Q gives a pair of linear maps $A : T_pP \to Q$, $B : T_xM \to Q$, with $\operatorname{im} A + \operatorname{im} B = Q$. The tangent space of Z consists of (u, v) such that $Au + Bv = 0$. The map $d\pi$ sends this pair to u, so it is onto if and only if $\operatorname{im} A \subseteq \operatorname{im} B$. Joint surjectivity makes this equivalent to surjectivity of B. That is the transverse condition for F_p at x. The criterion is vacuous when $\pi^{-1}(p)$ is empty.

By Sard's theorem, the bad parameters form a null set. The critical set of π is closed in a second-countable, locally compact manifold Z, hence is a countable union of compact subsets. Each compact subset has compact, null image in P, which has empty interior. The bad set is therefore a countable union of closed nowhere dense sets. If π is proper its whole critical image is closed. □

The diagram compares the fibers with the total incidence space.

$$\begin{array}{ccccc} F_p^{-1}(W) & \hookrightarrow & Z & \hookrightarrow & P \times M \\ \downarrow & & \downarrow{\scriptstyle \pi} & & \downarrow \\ * & \xrightarrow{\ p\ } & P & = & P. \end{array}$$

The left square gives a smooth manifold pullback at every regular parameter. If M is compact and W is closed in N, the incidence space is a closed subset of $P \times M$, which makes its projection proper. Its restriction to the open regular locus is a proper submersion and is smoothly locally trivial. Over a coordinate ball choose a metric and lift radial paths from the center horizontally. Properness keeps their lifts over a closed smaller ball in a compact inverse image. The lifts can therefore be integrated all the way along each path, and can be reversed. Smooth dependence on the endpoints and initial values gives the local trivialization. In particular, after any smooth change of parameter into this regular locus, the incidence manifold pulls back as an ordinary smooth manifold. This is a representability result for that family pullback, proved by the normal derivative criterion together with local transport.

Example 6.2 (A constrained family on the circle)**.** Consider $F((a, b), t) = \sin t + a \cos t + b$ on $\mathbb{R}^2 \times S^1$, with $W = \{0\}$. The derivative with respect to b is 1, so the total incidence manifold is smooth. The trigonometric summand $\sin t + a \cos t$ has amplitude $R = \sqrt{1 + a^2}$. Its derivative vanishes exactly at its two extrema. The parameters with a nonregular zero are therefore the two curves $b = \pm R$. In the region $|b| < R$ each fiber has two points. In the region $|b| > R$ it is empty. On a discriminant curve the vanished input derivative has a nonzero second derivative. Although the full family is transverse, freezing b and a at such a point gives a nontransverse member. A different family $F(a, x) = ax^2$ fails even joint transversality along $x = 0$ because both partial derivatives vanish there.

6.2 Genericity in a space of maps

Parametric transversality can also be used to construct dense sets of transverse maps. Here the family is invented locally for a given map, so the argument needs no global finite dimensional chart on the space of maps.

Proposition 6.3 (Transversality on a compact set)**.** *Let $W \subset N$ be closed and embedded, and let $K \subset M$ be compact. Smooth maps that are transverse to W at all points of K form an open dense subset of $C^\infty(M, N)$ in the weak Whitney topology. With a compact exhaustion of M, global transversality is residual in this topology, and dense by its Baire property.*

Proof. For openness, cover K by finitely many relatively compact source charts on which the original map either avoids W or is tested against a target chart adapted to W. Avoidance on a compact subset persists because W is closed. Where the map meets W, the relevant minor of the normal derivative is nonzero and persists after shrinking the source chart and restricting to small C^1 changes. Together finitely many such conditions work on K.

For density, start with a smooth $f : M \to N$ and a compact neighborhood K' of K. Pick finitely many smooth sections v_i of f^*TN, supported in a compact neighborhood of K', whose values span at every point of K'. To construct them, work in local bundle frames, multiply frame vectors by cutoff functions that are nonzero on smaller neighborhoods, and use a finite cover. For small $a = (a_1, \dots, a_s)$ define

$$F(a, x) = \exp_{f(x)} \left(\sum_{i=1}^{s} a_i v_i(x) \right).$$

Use a Riemannian metric on N. Compact support lets all these exponentials be defined uniformly for small a. At $a = 0$ the parameter derivative is surjective at every point of K', and openness and compactness make it surjective on K' for all small parameters. Consequently the family is transverse on the open manifold $\operatorname{int} K'$, and Theorem 6.1 gives arbitrarily small parameters whose maps are transverse on K. The maps converge smoothly on compacta to f as $a \to 0$, so every weak neighborhood of f contains one of them. Intersect the open dense tests for the members of a countable compact exhaustion. The weak smooth mapping space is Baire [15, Chapter 2], proving the last assertion. □

The same local mechanism applies to jet conditions from Chapter 4. In coordinate neighborhoods of a point x, perturb a map by vector-valued polynomials of degree at most r multiplied by a cutoff equal to one near x. Varying their coefficients spans all derivatives through order r at any point of this smaller neighborhood: translation of a polynomial is an invertible change of basis of the finite-dimensional polynomial space. The source coordinate of the jet is supplied by the input variable. The evaluation of prolonged jets is thus a submersion there. Applying parametric transversality to a smooth submanifold in the jet bundle gives the local density step of the jet transversality theorem [15, Chapter 3]. In particular, maps into a jet bundle arising from actual smooth functions can be varied through such maps. Arbitrary perturbations of a section need not remain a prolongation.

6.3 Boundary and relative constraints

When M has boundary, test transversality of the full family on both $P \times \operatorname{int} M$ and $P \times \partial M$, with P, N, W as above. Their incidence spaces are the interior and

boundary strata of Z. At a boundary point one can choose a nonsingular normal minor in the boundary variables, leaving the collar parameter unchanged. Apply Sard to the two projections into P. Outside their union of critical images, both stratumwise restrictions of F_p are transverse, and the fiber is a manifold with boundary as in Chapter 5. When the parameter manifold itself has boundary, the same argument can be applied separately to the parameter strata, using the appropriate total family on each product stratum.

A relative problem can require the map to stay fixed on a closed subset $A \subset M$. If it is already transverse on a neighborhood of A, compact-set perturbations can be supported away from A. For the proof of Proposition 6.3, choose an open neighborhood V of A with closure inside the region where f is transverse. Then $K \setminus V$ is compact and disjoint from A. Build the spanning sections using supports in $M \setminus A$ and apply the parametric argument to an open neighborhood of $K \setminus V$. The existing transversality on $K \cap \overline{V}$ persists for sufficiently small perturbations, since this set is compact. At a point constrained to remain fixed on an open region, any preexisting failure of transversality in that region persists as well. The support of the allowed perturbation is therefore part of the geometric hypothesis.

CHAPTER 7

Bundles, frames, and classifying objects

The tangent exact sequence of a transverse preimage is a sequence of vector bundles rather than a list of separate linear maps. Bundles make linear operations vary with the base, and frame changes record how a local description varies. They provide a useful first example of the distinction between classifying objects up to existence of an isomorphism and retaining the groupoid of their presentations and symmetries.

7.1 Linear gluing and pullback

Let X be a smooth manifold. A smooth rank-r real vector bundle $E \to X$ has smooth local trivializations that are linear on fibers. Choose smooth frames s_i over open sets $U_i \subset X$, viewing $s_i(x)$ as a linear isomorphism $\mathbb{R}^r \to E_x$. On an overlap write $s_j = s_i g_{ij}$, where $g_{ij} : U_i \cap U_j \to \mathrm{GL}_r(\mathbb{R})$ is smooth. On triple overlaps these matrices satisfy $g_{ij}g_{jk} = g_{ik}$. Conversely, such a cocycle, with $g_{ii} = 1$, glues the products $U_i \times \mathbb{R}^r$ by $(x, v)_j \sim (x, g_{ij}(x)v)_i$. The cocycle makes the relations consistent on all overlaps, and the inclusion of each product gives a chart of the glued total space. Changing frames by $s'_i = s_i h_i$ replaces g_{ij} by $h_i^{-1} g_{ij} h_j$. A section is a family of smooth functions v_i satisfying $v_i = g_{ij} v_j$.

A map $a : Y \to X$ pulls the bundle back by pulling back its frames and transitions. The diagram describes the pullback on total spaces.

$$\begin{array}{ccc} a^*E & \longrightarrow & E \\ \downarrow & & \downarrow \\ Y & \xrightarrow{\ a\ } & X. \end{array}$$

If bundles E, F have transitions g_{ij}, h_{ij} on a common cover, the direct sum and tensor product have transitions $g_{ij} \oplus h_{ij}$ and $g_{ij} \otimes h_{ij}$. The bundle of linear maps $E \to F$ has transition $A \mapsto h_{ij} A g_{ij}^{-1}$. Thus a global linear map of bundles is a section of a bundle defined by this action on matrices. This local transformation law matters whenever a model takes a feature in one bundle to a feature in another, because all coordinate matrices must satisfy the overlap conditions.

A smooth positive definite fiber metric exists on any bundle over a paracompact smooth manifold. Choose local metrics and average them with a smooth partition

of unity. Each nonzero vector has positive length in every metric with nonzero coefficient, so the average is again positive. Orthonormalizing local frames for this metric gives $\mathrm{O}(r)$-valued transitions. Orthogonal complements then split smooth subbundles. For example, a short exact sequence of smooth bundles splits by choosing a metric in the middle term, but the splitting depends on that choice.

7.2 Finite Grassmannians and stabilization

Let $\mathrm{Gr}_r(\mathbb{R}^N)$ be the Grassmannian of r-planes with tautological bundle $\gamma_{r,N}$. Its fiber over a plane is the plane itself.

Theorem 7.1 (A classifying map for a compact smooth base)**.** *Let X be a compact smooth manifold. A smooth rank-r bundle over X is isomorphic to $c^*\gamma_{r,N}$ for some smooth map $c : X \to \mathrm{Gr}_r(\mathbb{R}^N)$ and some finite N. Two such maps give isomorphic bundles if and only if they are smoothly homotopic after inclusion in a common larger Grassmannian.*

Proof. Take a finite cover by bundle frames and smooth cutoff functions ϕ_i supported in the corresponding frame domains, with at least one $\phi_i(x) > 0$ at each x. For $v \in E_x$ define a linear embedding J_x by concatenating the vectors $\phi_i(x)s_i(x)^{-1}v$, extending by zero outside their domains. One nonzero coefficient suffices for injectivity. These embeddings are smooth over X, and their image planes form a smooth map c to the Grassmannian. Locally the image is spanned by a smooth full-rank matrix $C(x)$. Its orthogonal projection is $C(x)(C(x)^{\mathsf{T}}C(x))^{-1}C(x)^{\mathsf{T}}$, which also shows that the planes vary smoothly. The maps J_x give a bundle isomorphism onto the tautological pullback.

If two maps classify isomorphic bundles, identify their pullbacks with the same E and regard the resulting fiberwise injections J_0, J_1 as injections into the same $\mathbb{R}^N$ by padding with zeros. In the direct sum of two copies of $\mathbb{R}^N$, the family

$$v \longmapsto (\cos t\, J_0 v,\ \sin t\, J_1 v), \qquad 0 \le t \le \pi/2,$$

stays injective, since at least one coefficient is nonzero. Its endpoints are J_0 in the first block and J_1 in the second. A rotation of the two blocks moves the latter to J_1 in the first. Reparametrizing the two paths to be stationary near their concatenation gives a smooth homotopy of the stabilized classifying maps.

Conversely, a smooth homotopy pulls the tautological bundle back to a bundle on $X \times [0,1]$. A connection gives parallel transport along each interval $\{x\} \times [0,1]$. One may take the connection induced by orthogonal projection of the trivial connection in the ambient trivial bundle. The resulting transport is linear, invertible by reverse transport, and smooth in x by smooth dependence of linear

ordinary differential equations. It identifies the endpoint bundles. Stabilizing a plane by inclusion of its ambient vector space pulls the tautological bundle back to the original one, so this is also an isomorphism before stabilization. □

Hirsch gives finite-dimensional bounds from the dimension of the base [15, Chapter 4]. The construction here uses a finite cover and lets the ambient dimension grow. The topological theorem for paracompact bases classifies numerable bundles by homotopy classes of continuous maps to an infinite Grassmannian, with the relevant topology on that Grassmannian [14, Chapter 1]. For a compact smooth base, a continuous map into a finite Grassmannian can be smoothed up to homotopy, and continuous homotopies between smooth maps into one can be smoothed relative to their ends after making them stationary near the ends. The smooth formulation retains a smooth bundle and smooth changes of frames throughout its proof.

7.3 Classifying the whole groupoid

Consider the site of smooth test manifolds with covers by open sets, using sheaves of spaces for the ambient category. Let $\mathcal{G}_r$ be the group sheaf $\mathcal{G}_r(U) = C^\infty(U, \mathrm{GL}_r(\mathbb{R}))$. The geometric realization of the simplicial bar construction of $\mathcal{G}_r$, taken in sheaves, is its classifying stack $B\mathcal{G}_r$. Its groupoid of points on a test manifold X is the groupoid of smooth rank-r bundles on X and smooth bundle isomorphisms.

To see the classification, the frames of E form a right $\mathcal{G}_r$-torsor. They are locally nonempty, and any two frames differ by a unique smooth pointwise matrix. Conversely a torsor P gives the associated bundle with fibers $P_x \times_{\mathrm{GL}_r(\mathbb{R})} \mathbb{R}^r$. Local sections of P reduce this construction to the gluing equations above. These operations are inverse on objects and on isomorphisms. The bar construction presents the groupoid with one object and automorphism sheaf $\mathcal{G}_r$. Descent adds the torsors obtained by gluing this object, with the changes of local frames giving their isomorphisms. This computes the stack interpretation.

For example, the trivial line over a point has automorphism group $\mathbb{R}^\times$ as an ordinary group in this groupoid. Its single isomorphism class says nothing about these automorphisms. The topological classifying space uses $\mathrm{GL}_r(\mathbb{R})$ with its Lie group topology and classifies bundles through homotopy classes of maps. It answers a corresponding isomorphism-class question in a category with continuous homotopies. A map to the smooth stack, in contrast, records actual smooth changes of frame over each test manifold. The classifying map of Theorem 7.1 concerns only the isomorphism class, and its equivalence relation is stabilized homotopy.

7.4 Orientation as reduction of structure group

For $r \geq 1$, an orientation on a real rank-r bundle is a coherent choice of orientation on its fibers. The functions $\operatorname{sgn}\det g_{ij}$ define a double cover with fiber the two possible orientations. A section amounts to choosing local signs that agree under these transitions, equivalently to reducing the frame group to the positive determinant subgroup. With a metric this is a reduction from $\mathrm{O}(r)$ to $\mathrm{SO}(r)$. On stacks the sign of determinant acts as a map to the classifying stack of a two-element group, and an orientation is a trivialization of the resulting torsor. This interpretation also shows why the orientation pulls back with the bundle.

Example 7.2 (Twisted coefficients on a circle)**.** Over S^1 identify $(2\pi, v) \sim (0, -v)$ in $[0, 2\pi] \times \mathbb{R}$ to obtain the Möbius line L. A section is represented by an antiperiodic real function on the real line. A nowhere-zero section would have constant sign by continuity, contradicting antiperiodicity. Thus L has no orientation or global frame. The tensor-square line has transition $(-1)^2 = 1$ and is trivial. A local feature that is linear in a vector of L inherits the sign change. A quadratic scalar in the same vector can be invariant and global even when the frame for the original feature does not exist.

CHAPTER 8

Tubular neighborhoods and collars

An embedded submanifold determines normal directions by a quotient of tangent bundles. Turning the quotient into coordinates in the ambient manifold requires a choice. Tubular neighborhoods provide this realization, and their isotopy properties explain how subsequent constructions can use these coordinates and still be independent up to homotopy. Boundary collars play the same role for an inward half-line.

8.1 Normal coordinates near an embedding

Let $i : A \hookrightarrow M$ be an embedded submanifold, with A, M both without boundary, and set $\nu_i = i^*TM/TA$. A *normalized tubular embedding* is an embedding $e : U \to M$ from an open neighborhood U of the zero section in ν_i, with open image, such that $e(0_x) = i(x)$ and the induced derivative in the normal quotient at 0_x is the identity of $(\nu_i)_x$. The quotient property of ν_i makes the normalization meaningful without a metric. The source projection determines a retraction on the image.

$$\begin{array}{ccc} U & \xrightarrow{\ e\ } & e(U) \\ \downarrow{\scriptstyle p} & & \downarrow{\scriptstyle pe^{-1}} \\ A & = & A. \end{array}$$

Theorem 8.1 (Existence near a compact submanifold)**.** *If A is compact, a normalized tubular embedding exists on an open disk bundle of some positive radius.*

Proof. Choose a Riemannian metric on M, identifying ν_i with the orthogonal complement of TA in i^*TM. Restrict the Riemannian exponential to this complement. At 0_x, the derivative on the tangent of the zero section is di_x, and the derivative along a fiber is the inclusion of the normal complement into T_xM. Their sum is an isomorphism. It follows that the restriction is a local diffeomorphism on some neighborhood of the zero section. Because A is compact, a single small positive radius works for definition of the exponential and invertibility of its derivative on the resulting disk bundle.

If injectivity failed on every smaller disk bundle, there would be pairs u_j, v_j of distinct vectors of norms tending to zero with equal images. By compactness

their footpoints have convergent subsequences, with limits x and y. Continuity and injectivity of i force $x = y$. Both sequences then approach 0_x, contradicting injectivity in a coordinate neighborhood where the map is a diffeomorphism. An injective local diffeomorphism has open image and smooth inverse on it. The derivative calculation provides the normalization. □

For a closed noncompact embedded submanifold one can use a positive radius depending on the footpoint [15, Chapter 4]. Uniform radius can fail at infinity. For example, embed a countable discrete set in $\mathbb{R}^2$ as the points $(j, 0)$ and $(j, 1/j)$ for $j \geq 1$. This is a closed zero-dimensional submanifold. Under the Euclidean exponential its normal disks are ambient disks, and no uniform positive radius makes all of them disjoint. Choosing radii less than one third of the distance to the nearest other point remedies the problem.

8.2 Isotopy of normalized germs

Proposition 8.2 (Independence after shrinking)**.** *Let A be compact, with ambient manifold without boundary. Any two normalized tubular embeddings e_0, e_1 of the same embedding i, defined on possibly different neighborhoods, are joined by a smooth family of normalized tubular embeddings on a common smaller neighborhood. The family is constant as a germ wherever the two original germs agree.*

Proof. After shrinking, form the transition $g = e_0^{-1} e_1$ near the zero section of ν_i. It fixes the zero section, and its normal derivative is the identity. For $s > 0$ let $g_s = \delta_{1/s} \circ g \circ \delta_s$, where δ_s scales the fibers by s. In a local trivialization g has components $(b(x, v), w(x, v))$ and $w(x, 0) = 0$. The base component of g_s tends to x, and its fiber component has the formula

$$\frac{w(x, sv)}{s} = \int_0^1 D_v w(x, \tau s v)\, v \, d\tau.$$

The equality supplies a smooth extension at $s = 0$, with fiber component v by normalization. It also shows that sufficiently small vectors, uniformly over the compact zero section and $s \in [0, 1]$, have images in the domain of e_0. For positive s the transition is an embedding on this smaller domain, by its definition as a composition of embeddings. At zero it is the identity. Its derivative on the quotient remains the identity, and $e_0 \circ g_s$ gives the required family. On a region where the transition is the identity as a germ, every rescaled transition is the identity as a germ there. □

The formula uses compatible transition maps of the normal bundle, so the extension shown in each trivialization glues globally. This conclusion concerns

tubular germs. An isotopy on a prescribed radius without shrinkage requires control away from the zero section. There is also a precise base change statement for the quotient normal bundle. If $f : X \to M$ is transverse to A and $B = f^{-1}(A)$, then df followed by projection to ν_i is onto along B with kernel TB. Consequently

$$TX|_B/TB \;\cong\; (f|_B)^*\nu_i.$$

This identification is natural for repeated transverse base change. An exponential map uses additional geometric data. Comparing two chosen exponentials requires a corresponding comparison of these choices.

8.3 Collars and neat intersections

A collar of a manifold with boundary M is an embedding $c : \partial M \times [0, \epsilon) \to M$ onto a neighborhood of the boundary with $c(x, 0) = x$. For compact boundary such a collar is obtained by choosing a smooth vector field near it pointing strictly inward there. Use inward coordinate vectors in local boundary charts, and glue with a partition of unity. Positivity of the derivative of an inward boundary coordinate is invariant under changing the defining coordinate, and is preserved by positive sums of vector fields. Thus the glued field remains inward. Its positive-time flow on the boundary is defined for a uniform short interval by compactness. At time zero its derivative has image the sum of the boundary tangent space and the transverse inward direction. The local inverse function theorem in half-space charts gives a collar near each point. The same compactness argument as for tubular neighborhoods gives a common interval on which this flow is injective and has smooth inverse onto its image.

For a compact manifold M with boundary, suppose the compact submanifold A is *neat*. This means $\partial A = A \cap \partial M$ and near each of those points there are half-space charts with coordinates (u, v, t), $t \geq 0$, in which A is defined by $v = 0$. One can construct a collar of M whose restriction to ∂A is a collar of A. In the neat charts the vector ∂_t is tangent to A along A. Choose all local inward fields meeting ∂A in this form, and choose fields on the remaining patches disjoint from A near those patches of the boundary. A partition of unity preserves tangency to A and inwardness. Uniqueness of solutions of the flow equation then gives the compatible collar near ∂A.

Use product metrics on a smaller common collar and double M and A across their boundaries. The double is smooth in the signed collar coordinate, and the doubles form an embedded pair without boundary. Apply the normal exponential construction to this pair. On the product collar, the normal vectors to A have zero component in the collar direction. Their short normal geodesics have constant collar coordinate. Away from a smaller collar the feet of the vectors are a compact set separated from the boundary. After shrinking further, their

geodesics remain in the original half of the double. The resulting restriction is a tubular neighborhood for the original neat pair, compatible with its boundary product structure [15, Chapter 4]. The product condition avoids treating the inward half-line as an unrestricted normal direction.

Example 8.3 (Collapse to a Thom space)**.** Let A be compact inside a compact manifold M without boundary, and choose a closed radius-r disk bundle $D_r(\nu_i)$ inside a tubular domain. The Thom space here is $D_r(\nu_i)/S_r(\nu_i)$, with the sphere bundle collapsed to a base point. For positive normal rank, the tubular embedding defines a map $M \to D_r(\nu_i)/S_r(\nu_i)$: on the embedded disk bundle use the inverse embedding followed by the quotient, and on the complement of its interior use the base point. The definitions agree on the embedded sphere bundle, giving continuity by gluing over two closed subsets. Shrinking the radius identifies Thom spaces by fiberwise rescaling and gives homotopic collapse maps by varying the radius continuously. Likewise, a family of tubular embeddings on a common disk bundle gives a homotopy of collapse maps. The inverse embedding on that family is continuous jointly with the parameter by compactness of the total closed disk bundle over the compact parameter interval. Thus the collapse construction can use a tubular choice and still define a homotopy class independent of it.

CHAPTER 9

Degree and oriented fibers

A regular fiber of a map between manifolds of the same dimension is a finite set when the source is compact. Orientation gives one sign to each point of this set. The degree is the sum of these signs, and its invariance comes from two compatible constructions. The first uses the map on fundamental classes. The second compares regular fibers by an oriented one-dimensional preimage. Both constructions are useful in families. Here a closed manifold is a compact manifold without boundary.

9.1 Local orientations and the signed count

Let M and N be closed oriented n-manifolds with N connected. For a smooth map $f : M \to N$, a value y is regular if $D_x f$ is surjective at every point of $f^{-1}(y)$. At such a point the tangent map is an isomorphism. Define $\epsilon_x(f)$ to be 1 or -1 according as it preserves or reverses orientation. The inverse function theorem gives a discrete fiber. Its closedness in a compact space makes it finite, including the possibility of an empty fiber.

The orientation at x is equivalently a choice of generator of $H_n(M, M \setminus \{x\}; \mathbb{Z})$. This follows by excision to a coordinate ball and the relative homology of that ball punctured at its center. A global fundamental class $[M]$ restricts to this generator at each point. If M is disconnected, $[M]$ sums the fundamental classes of the oriented components. In particular, if f is surjective on some components and misses a given value on others, every component can still be included in the same formula.

Theorem 9.1 (Regular-fiber degree)**.** *The homomorphism on top homology satisfies*

$$f_*[M] = d[N], \qquad d = \sum_{x \in f^{-1}(y)} \epsilon_x(f)$$

at every regular value y. Thus the signed count is independent of the regular value and depends only on the homotopy class of f.

Proof. Since $H_n(N; \mathbb{Z})$ is generated freely by $[N]$, the first equation uniquely defines d. Put $Q = f^{-1}(y)$. Apply homology to the maps from absolute to

relative pairs.

$$\begin{array}{ccc} H_n(M) & \xrightarrow{\quad f_* \quad} & H_n(N) \\ \downarrow & & \downarrow \\ H_n(M, M \setminus Q) & \xrightarrow{f_*} & H_n(N, N \setminus \{y\}) \end{array}$$

All groups in the square use integer coefficients. Disjoint coordinate balls about the finitely many points of Q and excision identify the lower left group with the direct sum of the local orientation groups. The image of $[M]$ is the tuple of their chosen generators. If Q is empty, the relative group is zero. On each summand the bottom arrow sends the local generator to $\epsilon_x(f)$ times the local orientation at y. Indeed, in small charts the map is a diffeomorphism, and its action on local homology is its orientation sign. The right vertical arrow sends $[N]$ to the local orientation at y. The square proves the signed count. Homotopy invariance follows from that of the induced homology map. □

This argument also explains why the signs must be computed relative to both orientations. Reversing the orientation of M changes $[M]$ by a minus sign, and the differential signs all change accordingly. Reversing both orientations leaves the degree unchanged. If $f : M \to N$ and $g : N \to L$ have connected closed oriented targets of the same dimension, functoriality gives $\deg(gf) = \deg(g)\deg(f)$. Disjoint union adds degrees for maps to a common target. A map of nonzero degree is surjective, since an omitted value would be regular with empty fiber. The formula can apply even when the critical points occupy part of the source. Only the chosen fiber needs an isomorphism of tangent spaces at each of its points.

9.2 Homotopy and the boundary of a fiber

Suppose $f_0, f_1 : M \to N$ have a common regular value y and are joined by a smooth homotopy F. Reparametrize near the ends so that F agrees with the end maps on their collars. Relative transversality permits replacing the interior of F by a nearby homotopy transverse to y without altering these collars [15, Chs. 3 and 5]. Its fiber over y is a compact one-manifold W with boundary. Orient $[0,1] \times M$ using the time coordinate followed by the orientation on M. The preimage orientation on W is given by

$$0 \longrightarrow TW \longrightarrow T([0,1] \times M)|_W \xrightarrow{DF} T_yN \longrightarrow 0$$

using the kernel first. On each end collar a positive tangent vector of W points in the increasing-time direction exactly when the corresponding local sign of f_i is positive. The boundary convention therefore gives

$$\partial W = f_1^{-1}(y) \amalg -f_0^{-1}(y)$$

as oriented zero-manifolds. Each compact oriented interval has one positive and one negative endpoint, and circle components contribute none. This proves equality of the signed cardinalities with a geometric comparison. Points of opposite signs can therefore be the two boundary points of one oriented component of a transverse interpolating fiber.

The same reasoning for a homotopy of targets can be expressed as transversality to a graph. For a smooth path $\gamma : [0,1] \to N$, one forms the preimage of its graph in $[0,1] \times N$ under $(t,x) \mapsto (t, f(x))$. If this is transverse, oriented by the normal orientation pulled back from TN, it has the corresponding endpoint fibers. The cylinder argument and the graph argument are instances of the same normal exact sequence. The top-homology proof covers paths through arbitrary singular values without choosing such a transverse comparison.

9.3 Families of regular values

Let P be a smooth parameter manifold and M compact. Suppose $F : P \times M \to N$ is smooth and a smooth choice $y : P \to N$ is regular for each $F_p = F(p,-)$, with $\dim M = \dim N$. Consider the diagram

$$\begin{array}{ccc} Z & \longrightarrow & P \times M \\ \big\downarrow{\scriptstyle \pi} & & \big\downarrow{\scriptstyle (p,x)\mapsto(p,F(p,x))} \\ P & \xrightarrow{(1,y)} & P \times N. \end{array}$$

The diagram is a transverse pullback. In target coordinates the equation is $F(p,x) = y(p)$, and its derivative in x is invertible at every solution. The implicit function theorem makes π a local diffeomorphism. It is proper because Z is closed in $P \times M$ and M is compact. On each component with nonempty fibers it is a finite covering. To see the local triviality, choose disjoint inverse branches about the points of a fiber. Properness allows shrinking the base neighborhood until no points remain outside these branches. Otherwise points outside the branches over a convergent sequence would have a limit in the original fiber.

For oriented M, N, the signs are locally constant on Z. Transport around a loop can permute points, while preserving their signs. Thus the signed fiber is a locally constant family of finite sets with a map to $\{-1,1\}$. Forgetting transport and taking signed cardinality yields a locally constant integer. For $P = M = N = S^1$, let $F(p,z) = z^k$ and $y(p) = p$ with $k > 0$. The regular fiber has k positive points. A single turn in P cyclically permutes them. Its degree is k, and the permutation is additional data of the family.

9.4 Parity and examples

Over $\mathbb{F}_2$, each closed n-manifold has a canonical fundamental class even if it is nonorientable. The same excision argument gives the parity formula

$$f_*[M]_2 = (\#f^{-1}(y) \bmod 2)[N]_2$$

for closed manifolds of dimension n and connected N. The mod-two degree is homotopy invariant. An unoriented compact one-manifold has an even number of boundary points, so the geometric comparison works with parity as well. If integral orientations exist, coefficient reduction sends the integer degree to the mod-two degree. An unbranched two-sheeted covering, for instance the antipodal covering $S^2 \to \mathbb{RP}^2$, has mod-two degree zero, although the pullback of a point contains two geometrically distinct points.

For a different example let $0 < a < \pi$ and $f(e^{it}) = \exp(ia \sin t)$. The value 1 has two preimages, at $t = 0$ and $t = \pi$, since $a \sin t$ lies in $(-\pi, \pi)$. Their local angular derivatives are a and $-a$. The count is zero, in agreement with the homotopy $\exp(isa \sin t)$ from the constant map to f. By contrast the angle map $z \mapsto z^k$ has degree k for every integer k. For $k < 0$ all $|k|$ regular preimages have negative local derivative. For $k = 0$ the map misses a value.

There is also a relative version for compact oriented manifolds with boundary. If $f : (M, \partial M) \to (N, \partial N)$ is smooth and N is connected, use the map on $H_n(-, \partial -)$ and its relative fundamental classes. A regular value in the interior of N has no preimage in the boundary of M. Excision gives the same local formula by mapping the relative fundamental classes into the local pairs. This version makes the behavior on the boundary part of the hypotheses. Without it the absolute top homology used for closed manifolds does not supply the desired comparison.

CHAPTER 10

Intersections and the Euler class

The sign in degree theory compares two vector spaces of equal dimension. For intersections the relevant target is a normal space. A transverse preimage is oriented using its tangent and normal exact sequence. The Thom class expresses the same comparison in cohomology, and it gives a formula when the intersection has larger dimension than transversality predicts.

10.1 Oriented transverse pullbacks

Work with smooth manifolds without boundary in this chapter. Let $A \hookrightarrow W$ be a closed embedded submanifold of codimension d with oriented normal bundle $\nu_{A/W}$. Let $f : B \to W$ be transverse to A and put $Z = f^{-1}(A)$. The normal quotient of Df is surjective on Z and the transverse preimage theorem yields the diagram

$$\begin{array}{ccc} Z & \longrightarrow & A \\ \downarrow{\scriptstyle j} & & \downarrow \\ B & \xrightarrow{\ f\ } & W. \end{array}$$

It is a pullback of smooth manifolds. Smoothness of a factorization through Z follows from the defining local submersion coordinates, so the square has the same universal property on all smooth test manifolds. The differential identifies the normal bundle of Z in B with the pullback of $\nu_{A/W}$, by the exact sequence

$$0 \longrightarrow TZ \longrightarrow TB|_Z \longrightarrow (f|_Z)^*\nu_{A/W} \longrightarrow 0.$$

If B is oriented, we orient TZ so that an oriented basis of TZ, followed by lifts of a positive normal basis, gives the orientation on B. When $\dim B = d$, each point of Z is oriented by the determinant sign of the normal differential. For compact B the intersection number is the sum of these signs. Taking A to be a point in an oriented target recovers the regular-value formula of Chapter 9.

If f_t is a homotopy between transverse maps with compact source without boundary, a transverse replacement relative to end collars has oriented preimage with boundary the two end intersections, taken with opposite signs. The preimage is compact by closedness of A. This proves invariance of the intersection number and gives an oriented bordism in higher dimension. The tangent sequence provides

the sign convention throughout. Specifying normal orientations also shows why the order of intersection can matter. Interchanging two independent oriented normal blocks of ranks a and b changes their combined orientation by $(-1)^{ab}$.

10.2 Thom classes and zero sets

For an oriented rank-d real vector bundle $E \to B$ with zero section z, the Thom class is the class

$$U_E \in H^d(E, E \setminus z(B); \mathbb{Z})$$

which restricts on each fiber to the chosen orientation generator. Fiber restriction and the Thom isomorphism give its uniqueness. One construction glues the local generator on each oriented product chart, using the trivializations and the orientation compatibility in the Thom isomorphism [26, §9]. Its pullback along z, after forgetting the relative condition, is the Euler class $e(E) \in H^d(B; \mathbb{Z})$.

Let s be a section transverse to z and let $Z = s^{-1}(0)$. The class s^*U_E belongs to $H^d(B, B \setminus Z; \mathbb{Z})$. Its absolute image is $e(E)$, by the homotopy $t \mapsto ts$ of sections after forgetting the complements. In a local bundle chart the derivative of the coordinate vector of s identifies $\nu_{Z/B}$ with $E|_Z$. This identification preserves the induced normal orientation. The Thom class s^*U_E is consequently the normal Thom class of Z. For compact oriented B with $\dim B = d$, evaluation on its fundamental class gives

$$\langle e(E), [B] \rangle = \sum_{x \in Z} \operatorname{sgn} \det(D_x^{\mathrm{vert}} s).$$

Indeed, excise to disjoint oriented coordinate balls and use the local fiber generators as in the degree proof. This also proves independence of the signed zero count from the transverse section. If s is nowhere zero, the pullback of the relative class is zero, so $e(E)$ vanishes. The Euler class detects a necessary condition for a nonvanishing section.

More generally, an oriented closed embedding $j : Z \hookrightarrow B$ of codimension k has a cohomological pushforward $j_! : H^r(Z) \to H^{r+k}(B)$ constructed by placing the orientation Thom class first in the cup product with the pulled-back cohomology class. Compute this in a tubular neighborhood, excise to the relative pair in B, and forget the complement. We write $\mathrm{PD}_W[A]$ for the class obtained by applying this operation to 1 for $A \hookrightarrow W$. These definitions give $f^* \mathrm{PD}_W[A] = j_!1$ for the transverse square. On a normal fiber the pullback is an orientation isomorphism, which identifies the corresponding Thom classes.

10.3 The excess bundle in a clean intersection

Let $i : B \hookrightarrow W$ and $A \hookrightarrow W$ be closed embedded submanifolds with a clean intersection. This means that $Z = A \cap B$ is a submanifold and $T_z Z = T_z A \cap T_z B$ for every $z \in Z$. For simplicity assume the normal ranks are constant. Clean intersection makes the normal derivative injective, giving

$$0 \longrightarrow \nu_{Z/B} \longrightarrow \nu_{A/W}|_Z \longrightarrow E_{\mathrm{exc}} \longrightarrow 0.$$

The excess bundle E_{exc} measures the normal directions of A that are absent from the image of $T_z B$. Suppose that $\nu_{A/W}$ and $\nu_{Z/B}$ are oriented and orient E_{exc} by the displayed order of the sequence.

Theorem 10.1 (Clean intersection formula)**.** *Under these hypotheses,*

$$i^* \mathrm{PD}_W[A] = j_! e(E_{\mathrm{exc}}) \quad \text{in } H^d(B;\mathbb{Z}), \qquad j : Z \hookrightarrow B.$$

Proof. We compute the pullback near Z in tubular neighborhoods. A neighborhood of Z in B is represented by normal vectors ξ in $\nu_{Z/B}$. Over the corresponding neighborhood of A in W the inclusion is a map into $\nu_{A/W}$. Write it as $\xi \mapsto (\alpha(\xi), v(\xi))$, where the first entry lies in A and the second in the normal fiber over $\alpha(\xi)$. Here $v(0_z) = 0$ and $\alpha(0_z) = z$. Shrink the domain fiberwise if necessary. For $t > 0$ the rescaled map has value

$$(\alpha(t\xi), t^{-1} v(t\xi)).$$

As $t \to 0$, it extends smoothly to the normal derivative over Z, as is seen by integrating the first derivative in a bundle chart along the line from 0_z to ξ. For $t > 0$ it is nonzero off the zero section since $i^{-1}(A) = Z$. At $t = 0$ it is nonzero there since the normal derivative is injective. Thus the relative Thom pullback is computed by the linear injection of normal bundles.

Choose a splitting so that the target normal bundle over Z becomes $\nu_{Z/B} \oplus E_{\mathrm{exc}}$. The Thom class of the sum is the product of the two Thom classes in the displayed order. This follows by restricting to each fiber, where the product is the orientation generator of the oriented direct sum. Pulling back by the inclusion as the first summand leaves its Thom class and pulls back the second class by its zero section. The latter pullback is $e(E_{\mathrm{exc}})$. Extending the resulting relative class to B and forgetting the complement gives precisely $j_! e(E_{\mathrm{exc}})$. □

An ordinary clean pullback displays the manifold Z, but its tangent sequence also contains the cokernel bundle. Transversality makes that cokernel vanish and recovers the rank-zero Thom formula. A tangent-complex description keeps it in a second degree, as developed later for derived intersections. The formula above identifies a classical contribution of that bundle without replacing the clean intersection by a finite set.

10.4 Self-intersection and tangent zeros

Take i to be the inclusion of A itself. The intersection inside $B = A$ is all of A, and the excess bundle is $\nu_{A/W}$. Thus restricting $\mathrm{PD}_W[A]$ to A gives the Euler class of the normal bundle. In $\mathbb{CP}^2$, let A be the projective line $z_0 = 0$. The line $z_1 = 0$ is isotopic to A, through the rotations of the first two homogeneous coordinates. The two lines meet in one point $[0 : 0 : 1]$. In the affine chart at that point they are complementary complex coordinate lines. Complex linear identifications preserve the underlying real orientation, so the intersection sign is positive. The normal Euler number of A is therefore 1. The self-intersection can be computed by a transverse motion even though the original pullback is a two-dimensional manifold.

When $E = TB$ for a closed oriented manifold, a vector field with nondegenerate zeros computes the tangent Euler number by the same formula. Take a Morse function and a metric, using its gradient as the section. At a critical point of index λ, the derivative of the gradient is the inverse metric composed with the Hessian, so its determinant sign is $(-1)^\lambda$. Each corresponding Morse handle contributes that same sign to the Euler characteristic by Chapter 11. Thus $\langle e(TB), [B] \rangle = \chi(B)$, the Poincaré–Hopf equality in the transverse setting [15, Ch. 5]. For example, a height function of S^2 has one minimum and one maximum. Both indices are even, and the two contributions give 2.

CHAPTER 11

Morse functions and handle filtrations

The sublevel sets of a smooth real-valued function form an increasing family of spaces. A regular band contributes a product, while an isolated nondegenerate critical point adds a handle. This turns the real-valued function into a description of the topology of its source. We use the regular and critical level arguments of differential topology [15, Ch. 6] to describe the filtration and its homotopy pushouts explicitly.

11.1 The Hessian and the local form

Let $f : M \to \mathbb{R}$ be smooth and let p satisfy $df_p = 0$. Second partial derivatives in a chart define a symmetric bilinear form on T_pM. Under a change of coordinates the Hessian is conjugated by the Jacobian, since the term involving the second derivatives of the coordinate change is multiplied by $df_p = 0$. A critical point is *nondegenerate* if this form is nonsingular, and its *index* is the dimension of a maximal negative definite subspace. A Morse function has only nondegenerate critical points. They are isolated by the inverse function theorem applied to the coordinate components of df. For a compact source the critical set is closed and therefore finite.

Theorem 11.1 (Morse local form)**.** *If p has index λ in an n-manifold, there are local coordinates (u, v) centered at p, of dimensions λ and $n - \lambda$, such that*

$$f(u, v) = f(p) - \|u\|^2 + \|v\|^2.$$

Proof. In coordinates x near $p = 0$, Taylor's integral formula writes the difference as $x^{\mathsf{T}} A(x) x$ with

$$A(x) = \int_0^1 (1 - t)\, D^2 f(tx)\, dt.$$

This is a smooth symmetric matrix and $A(0)$ is nonsingular. A fixed linear change of basis makes $A(0)$ diagonal with entries -1 and 1. In a neighborhood of 0 all successive pivots in the symmetric elimination of $A(x)$ are nonzero. Eliminating one row and column at a time is the completion of a square, using only division by the current nonzero pivot. The signs of the pivots are locally constant. Taking square roots of their absolute values gives a smooth factorization $A(x) = B(x)^{\mathsf{T}} D B(x)$, where D is the same diagonal sign matrix. The map

$x \mapsto B(x)x$ has invertible derivative $B(0)$ at 0. Its components are the required local coordinates by the inverse function theorem. □

The index is stable under small perturbations of the Hessian through nonsingular forms. If f_s is a smooth family and f_{s_0} has a nondegenerate critical point p, the equation $df_s(x) = 0$ near (s_0, p) has a unique smooth local solution $p(s)$ by the implicit function theorem. The index is constant on this branch. This describes how critical points move through regular parameter values. When the Hessian becomes singular, the equation no longer has to be locally the graph of a smooth section.

11.2 Regular bands

Write $M^a = \{x : f(x) \leq a\}$. Suppose a closed band $f^{-1}([a, b])$ is compact and contains no critical point. Fix a Riemannian metric and use the vector field

$$V = -\nabla f / \|\nabla f\|^2$$

on a neighborhood of the band. Along its flow f decreases at unit rate. A flow line starting at level s therefore reaches level a in time $s - a$. Existence for this duration follows by compactness, since the flow stays in the band until it reaches the designated end. Flowing backward gives the diffeomorphism of the band with $f^{-1}(a) \times [a, b]$. This proves that both boundaries are regular copies of the same level. Flowing each point of $M^b \setminus M^a$ for time $t(f(x) - a)$, for $0 \leq t \leq 1$, and fixing M^a, gives a deformation retraction $M^b \to M^a$. The formula is continuous at level a because the travel time tends to zero.

This argument requires compactness of the band, or another condition giving the same flow existence. A nonproper function on a noncompact manifold can change level behavior at infinity with no critical point. For example, $f : \mathbb{R} \to \mathbb{R}$ given by $f(x) = e^x$ has empty levels at nonpositive values and single-point levels at all positive values. Its derivative is nowhere zero, and a band straddling zero is noncompact.

11.3 Crossing one critical level

Assume M is closed and p is the only critical point in a band around $c = f(p)$. Set $a = c - \varepsilon$, $b = c + \varepsilon$ with $\varepsilon > 0$ small. Write D^j for the closed unit disk in $\mathbb{R}^j$, including $D^0 = \{0\}$ and $\partial D^0 = \varnothing$. The topological handle of index λ is $D^\lambda \times D^{n-\lambda}$ attached along $\partial D^\lambda \times D^{n-\lambda}$.

Theorem 11.2 (Sublevel attachment)**.** *The upper sublevel M^b has the homotopy type, relative to M^a, of M^a with one index-λ handle attached. Consequently there*

is a homotopy pushout of spaces

$$\begin{array}{ccc} \partial D^\lambda & \longrightarrow & M^a \\ \downarrow & & \downarrow \\ D^\lambda & \longrightarrow & M^b. \end{array}$$

Proof. Work in the chart of Theorem 11.1. Choose a metric equal to the Euclidean metric there near p. Choose $r > 0$ with $r^2 < \varepsilon$ and both small enough that

$$H = \{(u, v) : \|v\| \le r, \quad \|u\| \le \sqrt{\varepsilon + \|v\|^2}\}$$

is contained in this metric neighborhood. Dividing u by $\sqrt{\varepsilon + \|v\|^2}$ and rescaling v gives the disk product. The inequality $f \ge a$ on H is an equality precisely on the attaching face. The inequality $f \le c + r^2 < b$ puts the handle inside the upper sublevel. Hence $M^a \cap H$ is the attaching face.

Follow negative gradient flow outside $M^a \cup H$. There are no critical points on the compact closure of the region still to be added, since p is interior to H relative to M^b. The decrease rate of f there is bounded below in absolute value by a positive constant. The flow reaches $M^a \cup H$ in finite time. At its level-a boundary the flow crosses strictly toward the lower sublevel. At the other face of H, $\|v\| = r$, the local gradient equation is $v' = -2v$, so the flow crosses strictly inward there as well. These inequalities give continuous first-entry time, using the first of the two transverse crossings where the faces meet. Defining time zero inside the union yields a deformation retraction of M^b onto $M^a \cup H$ by running a fixed fraction of this time. If a disk factor has dimension zero, its empty boundary requires no crossing argument.

The attaching face and its inclusions admit collars in their respective spaces when the face is nonempty. The resulting pushout computes the homotopy pushout. Projection from the disk product to D^λ is an equivalence of pairs with the corresponding boundary. Using the center $v = 0$ for the attaching map therefore gives the displayed square. □

The smooth handle attachment statement uses a smoothing of the corners of such a model. It describes the boundary surgery as well as the homotopy type [25, §3]. For the homotopy filtration, each interval between critical values can be replaced by a single arrow. If the critical values coincide, a small bump constant near each critical point separates them. The transition regions of these bumps lie in a compact set where $\|df\|$ has a positive lower bound, so sufficiently small bumps create no new critical points. The filtration then starts at the empty set and successively attaches a cell of each critical index. The relative homology of

each step is free of rank one in the critical index and zero in the other degrees. Additivity in the long exact sequences of the pairs gives

$$\chi(M) = \sum_{p \in \mathrm{Crit}(f)} (-1)^{\mathrm{ind}(p)}.$$

On the torus use $f(s,t) = \cos s + 2\cos t$. The four critical points occur where s, t are 0 or π. The Hessian is diagonal with entries $-\cos s$ and $-2\cos t$. The critical values in order are $-3, -1, 1, 3$, with indices $0, 1, 1, 2$. The filtration first creates a component, then adds two one-dimensional cells, and finally fills a two-dimensional cell. It gives $\chi(T^2) = 1 - 2 + 1 = 0$.

Indices in a family do not determine handle transport. For $\theta \in \mathbb{R}/(2\pi\mathbb{Z})$, let $R(\phi)$ be the rotation of $\mathbb{R}^2$ by ϕ and define the quadratic form

$$x^{\mathsf{T}} R(\theta/2) \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix} R(\theta/2)^{\mathsf{T}} x.$$

It is well defined on the circle because adding 2π changes both rotation matrices by a sign. The origin is a nondegenerate critical point of index one in each fiber. Its negative eigenspace is spanned locally by $R(\theta/2)(1,0)$, which reverses sign on a full turn. The negative spaces form a Möbius line, so the local orientations of the index-one disks reverse under transport. The critical point and its index are constant throughout, while the relative homology generator of the local handle has sign monodromy.

CHAPTER 12

Trajectory spaces and Morse homology

A handle filtration determines the ranks of the Morse chain groups. To recover the boundary maps one needs information about the attaching maps. Negative gradient trajectories provide that information through compact spaces whose boundary strata consist of successive trajectories. Counting zero-dimensional spaces produces a chain complex. Retaining all the spaces instead produces a category enriched in spaces.

12.1 Flow lines and the boundary operator

Let M be closed, $f : M \to \mathbb{R}$ Morse, and g a metric. The negative gradient is complete by compactness. For a critical point p, let $W^u(p)$ and $W^s(p)$ be the unstable and stable manifolds. They consist of points whose flow converges to p at time $-\infty$ and $+\infty$, respectively. The stable manifold theorem gives their dimensions as $|p| = \operatorname{ind}(p)$ and $\dim M - |p|$. These dimensions come from the positive and negative eigenspaces of the linearized negative gradient. Assume the pair is Morse–Smale, meaning all $W^u(p)$ and $W^s(q)$ intersect transversely. For $p \neq q$, the intersection, when nonempty, consists of nonstationary trajectories and has dimension $|p| - |q|$. Quotienting by time translation gives the moduli space $\mathcal{M}(p, q)$. When the space is nonempty, $f(p) > f(q)$. Every trajectory crosses a fixed regular level between the two values exactly once, so one can realize the quotient as a smooth level slice. Its dimension is $|p| - |q| - 1$.

The Morse compactification theorem adds broken trajectories to these spaces. The stratum with a sequence of breaks through $p_1, \dots, p_k$ is

$$\mathcal{M}(p, p_1) \times \mathcal{M}(p_1, p_2) \times \cdots \times \mathcal{M}(p_k, q).$$

In particular, under a metric Euclidean in Morse coordinates near the critical points, the spaces have smooth corner models with these strata in codimension k [32, Thms. 2.3 and 2.6]. To see the source of the boundary, cut trajectories at regular levels between the critical values. The compact slices control where segments can travel. If the time between two slices tends to infinity, a limiting segment can approach an intervening critical point. The compactness theorem shows that only finitely many such breaks occur and gives the resulting limiting

tuple. The gluing theorem supplies a nonnegative collar parameter for each break, with zero recording the broken tuple. The proofs below use both compactness and this description of the boundary. A set of broken trajectories alone does not establish the existence of the collars.

Let $C_j(f)$ be the $\mathbb{F}_2$-vector space on critical points of index j. For a difference of indices equal to one, each nonconstant segment drops the index by at least one, so the compactification has no breaks. Thus $\mathcal{M}(p,q)$ is compact and zero-dimensional, hence finite. Set

$$\partial p = \sum_{|q|=|p|-1} \#_2\mathcal{M}(p,q)\, q,$$

where $\#_2$ denotes parity. For $|p|-|r|=2$, $\overline{\mathcal{M}}(p,r)$ is a compact one-manifold whose boundary is the disjoint union of $\mathcal{M}(p,q)\times\mathcal{M}(q,r)$ over $|q|=|p|-1$. The coefficient of r in $\partial^2 p$ is the parity of that boundary. It is zero since each interval has two endpoints. Over the integers, compatible orientations of unstable manifolds give the trajectory signs, and the oriented-boundary formula yields the signed version [28, Chs. 3–4]. The Morse homology theorem identifies the resulting homology with singular homology. Its comparison uses the generator of a handle pair from Chapter 11 and identifies boundary coefficients by flowing an attaching sphere to a transverse normal slice at the next lower index [28, Ch. 4].

12.2 Continuation and homotopies of continuations

Let (f_s, g_s), $s\in\mathbb{R}$, interpolate between two Morse–Smale pairs and be constant for $|s|$ sufficiently large. A continuation trajectory satisfies

$$\dot{x}(s) = -\nabla_{g_s} f_s(x(s)), \qquad \lim_{s\to-\infty} x(s) = p, \quad \lim_{s\to+\infty} x(s) = q.$$

The fixed transition interval distinguishes a time coordinate. In particular one does not divide out an $\mathbb{R}$-translation action for this equation. For regular interpolation data, the dimension is $|p|-|q|$. Indeed, transfer $W^u(p)$ by the time-dependent flow through the transition interval and intersect it transversely with the terminal $W^s(q)$. Counting the zero-dimensional continuation spaces gives a degree-zero map $\Phi : C_*(f_-)\to C_*(f_+)$. Compactified one-dimensional spaces have breaking at the two stationary ends. The two types of boundary points give the two terms in $\partial_+\Phi + \Phi\partial_- = 0$ over $\mathbb{F}_2$.

A regular one-parameter comparison of interpolations has the additional parameter direction. For an index difference of -1, counting it gives a degree-one operator K. In index difference zero, its compactified one-manifold has boundary from the two parameter ends and from breaking at either stationary end. The boundary parity therefore gives

$$\Phi_1 + \Phi_0 = \partial_+ K + K\partial_-.$$

Gluing along a long intermediate interval compares composite continuation with a single interpolation. Reversing the interpolation and comparing the concatenation to the constant interpolation gives inverse maps on homology [28, Ch. 4]. These statements take regularity and compactified gluing for the continuation equations as inputs. The identities explain which boundaries are needed for invariance and why a comparable chain map is specified only up to chain homotopy.

12.3 The enriched flow category

The objects of the flow category $\mathcal{C}_f$ are critical points. Its nonidentity morphism spaces are $\overline{\mathcal{M}}(p, q)$. Concatenating broken tuples is continuous and associative. In particular a codimension-two corner for a three-stage trajectory gives the same ordered tuple under either grouping. The following diagram records the two face inclusions. We abbreviate a compactified space by C_{pq}.

$$\begin{array}{ccc} C_{pq} \times C_{qr} \times C_{rs} & \longrightarrow & C_{pr} \times C_{rs} \\ \downarrow & & \downarrow \\ C_{pq} \times C_{qs} & \longrightarrow & C_{ps}. \end{array}$$

The nerve is a simplicial space, with level n given by composable strings of n morphisms. Faces compose or forget arrows and degeneracies insert identities. Applying the singular complex to each mapping space gives a simplicial category, whose homotopy-coherent nerve is an ∞-category. A path within one trajectory space now gives a higher comparison of arrows. The classifying-space result of Cohen, Jones, and Segal recovers the homotopy type of M from the realization of the topological nerve in this Morse–Smale case [8, Thms. 2.6 and 3.1]. Composition here uses recorded break data. The continuation calculation above explains the one-dimensional coherence needed at the chain level. An analogous construction of continuation on all the enriched data would require compatible multiparameter gluing.

12.4 Two explicit flow categories

On the circle take $f(e^{it}) = \cos t$ and the standard metric. The flow equation is $\dot{t} = \sin t$. The maximum at 0 and the minimum at π are connected by two unparametrized trajectories, one in each open semicircle. There are no intermediate critical points at which to break. The mod-two differential vanishes because the trajectory count is two. The realization of the flow category has two vertices and two edges joining them, which gives a circle directly.

On the round sphere take the height function. The maximum has index two, the minimum has index zero, and the nonstationary trajectories are the

meridians, parametrized by their intersections with the equator. There are no broken trajectories. The morphism space from maximum to minimum is therefore S^1. The nerve realization is obtained from $S^1 \times [0,1]$ by collapsing its two end circles separately to the vertices. This is the suspension S^2. If we retain only the path components of the morphism spaces, the ordinary category has one nonidentity arrow and its realization is an interval. Thus taking components of the mapping spaces can lose the topology of the manifold, even when all the objects and their indices are retained.

The product torus gives an example with actual breaking. Use the flat metric and $f(s,t) = \cos s + 2\cos t$. The negative gradient preserves the two coordinate equations, and the stable and unstable manifolds are products of those on the two circles. Transversality holds because it holds in each factor. For index difference one, exactly one coordinate moves, and the two semicircles in that factor give parity zero. Thus the mod-two complex has one generator in each of degrees zero and two, and two in degree one, with zero differential.

For a trajectory from the maximum to the minimum both coordinates move. Choosing one of the two semicircles in each factor leaves a time shift for each of the two one-dimensional motions. Modulo common translation the difference of shifts is a real coordinate. Each of the four choices therefore gives a real line in the unbroken moduli space. In the two limits of that coordinate, one coordinate transition occurs before the other, giving the two orders of a broken trajectory. The four closed intervals in the compactification exhibit eight boundary points. The two possible intermediate critical points account for four each.

CHAPTER 13

Geometric bordism and the Thom construction

A closed manifold is a possible boundary, and a bordism asks whether two closed manifolds can occur as the ends of a single manifold. The question keeps track of tangent or normal structures when these are part of the geometry. Its stable answer is a homotopy group. The passage from a manifold to that group uses the same two constructions that govern intersections: a tubular neighborhood and a transverse inverse image.

13.1 Bordisms and their elementary groups

A closed d-manifold is compact and without boundary. An unoriented bordism from M_0 to M_1 is a compact $(d+1)$-manifold W with designated ends $\partial W \simeq M_0 \amalg M_1$. In the oriented case the boundary orientation identifies ∂W with $(-M_0) \amalg M_1$. The two conventions use an outward normal first in the orientation of a boundary. One can choose collars, as in Chapter 8, to glue successive bordisms smoothly. The cylinder proves reflexivity, and reversing the roles of the ends proves symmetry, using $-W$ in the oriented setting. Disjoint union gives an abelian group of equivalence classes. The cylinder also gives $[M] + [-M] = 0$ in the oriented group, and $2[M] = 0$ in the unoriented group. Denote the respective groups by Ω_d^{SO} and Ω_d^{O}.

Example 13.1 (Dimension zero)**.** An oriented closed zero-manifold is a finite list of positive and negative points, with multiplicities. Each oriented interval has one endpoint of each sign, and a circle has no endpoints. The signed count consequently gives an invariant in $\mathbb{Z}$. A finite collection with count zero bounds by pairing the positive and negative points and taking one oriented interval for each pair. Thus $\Omega_0^{SO} \simeq \mathbb{Z}$. Dropping the orientations retains the parity of the number of points, and the same interval construction proves $\Omega_0^{O} \simeq \mathbb{Z}/2$.

Products endow the direct sum of the degree groups with graded multiplication. A bordism W for M_0, M_1 produces the bordism $W \times N$ when N is closed. In the oriented case the factor order contributes the sign $(-1)^{de}$ on swapping oriented

manifolds of dimensions d and e. A related construction equips a manifold with a continuous map to a space X, with both this map and its structure required to extend across a bordism. This is a geometric way for a theory of manifolds to run over variable spaces. We first consider the coefficient groups where X is a point.

Example 13.2 (An unoriented surface that does not bound)**.** For a compact smooth 3-manifold W, the double of W is a closed 3-manifold. Its Euler characteristic is zero, by Poincaré duality over $\mathbb{F}_2$. Additivity of Euler characteristic gives $\chi(\partial W) = 2\chi(W)$, so no closed surface of odd Euler characteristic can be its whole boundary. In particular $\mathbb{RP}^2$ represents a nonzero element of Ω_2^O. This computation uses duality only as a vanishing statement on the double. It applies regardless of orientability.

13.2 Collapsing a normal neighborhood

Embed M^d in $\mathbb{R}^{d+k}$ and let ν be its normal bundle of rank k. For the Thom space write $\mathrm{Th}(\nu) = D(\nu)/S(\nu)$, using a metric on ν. A tubular identification of $D(\nu)$ with a closed neighborhood of M in the one-point compactification S^{d+k} gives a pointed collapse. The classifying map of ν induces a bundle map to the universal rank-k bundle γ_k on $BO(k)$. With compatible metrics it induces the second arrow of

$$S^{d+k} \xrightarrow{\text{collapse}} \mathrm{Th}(\nu) \xrightarrow{\text{classifying map}} \mathrm{Th}(\gamma_k).$$

An oriented M induces a normal orientation using the ambient orientation and the order $TM \oplus \nu$. In that case use the universal oriented bundle over $BSO(k)$.

The collapse sends the complement of the open normal disk bundle to the base point. For a smaller radius the annulus between the two boundaries can be absorbed by a radial change of coordinates. An isotopy of tubular neighborhoods gives a homotopy of collapse maps by applying this construction over the interval. Increasing the ambient dimension by one replaces the normal bundle by $\nu \oplus \mathbb{R}$. Under this operation the collapse suspends:

$$\begin{array}{ccc} \Sigma S^{d+k} & \xrightarrow{\Sigma c_\nu} & \Sigma\,\mathrm{Th}(\gamma_k) \\ \downarrow \simeq & & \downarrow \\ S^{d+k+1} & \xrightarrow{c_{\nu\oplus\mathbb{R}}} & \mathrm{Th}(\gamma_{k+1}). \end{array}$$

The right arrow comes from adjoining a trivial line to each normal plane. Choices of embeddings disappear after stabilization. Indeed, for embeddings i_0, i_1 placed in independent coordinate summands, the family $x \mapsto ((1-t)i_0(x), t i_1(x))$ stays embedded for $0 \le t \le 1$. Initial embeddings may therefore be compared using a

smooth cylinder in a common larger ambient space, with orientation-preserving rotations of ambient coordinates if an ordering of the normal directions has been specified.

13.3 Transversality gives the inverse

Theorem 13.3 (Stable Pontryagin–Thom correspondence)**.** *Collapse induces isomorphisms*

$$\Omega_d^O \simeq \pi_d MO, \qquad \Omega_d^{SO} \simeq \pi_d MSO.$$

The kth spaces in these Thom spectra are the Thom spaces of the universal rank-k bundles over $BO(k)$ and $BSO(k)$. The stable group on the right is the colimit of the groups in degrees $d+k$ under suspension and the Thom structure maps.

Proof. We give the normal-bundle argument [15, Ch. 7][9, Lectures 2–5]. Use finite Grassmannians for the bundle classifying maps, enlarging their ambient dimension as needed. Representatives with compact source, and their homotopies, can be taken at finite stages by cellular approximation. This permits the usual smooth transversality argument near the zero section. The disk bundle away from the sphere identification is a smooth total space, with the zero section a closed submanifold of codimension k.

For a pointed map $a : S^{d+k} \to \mathrm{Th}(\gamma_k)$ represented at such a stage, approximate it smoothly near the inverse image of the zero section and make it transverse there. One may leave the map fixed near the base point of the source. To see why local smoothing is enough, first choose a smaller disk bundle about the zero section. The complement of its inverse image maps away from zero and transversality needs to be arranged only on a neighborhood of the smaller inverse image. An arbitrarily small perturbation there, kept equal to the original outside the larger inverse image, introduces no zeros in the intervening closed annulus.

The transverse inverse image M is a closed d-manifold in the finite chart $\mathbb{R}^{d+k}$ of the source. The differential identifies its normal bundle with the pullback of γ_k along $a|_M$. This is just the isomorphism from the normal space to the quotient by the tangent space of the zero section. In the oriented version this gives its orientation by the stated convention. The same argument applied to a homotopy, made transverse relative to its endpoints, yields a bordism in $S^{d+k} \times [0,1]$. It has the required collars after the homotopy is made constant in the parameter near its ends. Suspension appends a trivial direction to the normal bundle of the inverse image.

Conversely, collapse along an embedded manifold is transverse near the zero section and recovers that manifold with its normal data. For a transverse a, shrink to a tubular neighborhood of its inverse image by radially pushing the complement of a smaller bundle neighborhood to the base point. The resulting

map is homotopic to normal collapse: on a smaller tube one compares the normal-coordinate map with its derivative along M. In local bundle coordinates this uses the homotopy given by the rescaled normal component $t^{-1}F(x, tv)$, which extends to the derivative at $t = 0$. Parallel transport in the target bundle identifies the nearby fibers with the fibers over $a|_M$ for this calculation. The derivative is a normal isomorphism. Compactness gives a common tube on which the comparison is defined, and a radial collapse sends its exterior to the base point. This gives inverse operations on stable classes.

An abstract bordism has a neat embedding into a large $\mathbb{R}^{d+k} \times [0, 1]$ with product form near its ends. This can be obtained by prescribing the collar coordinate as the interval coordinate near the boundary and combining it with a relative Euclidean embedding in the remaining coordinates. After stabilization the end embeddings can be adjusted by the embedding isotopies above. Collapsing its normal neighborhood then gives the required homotopy between end collapses. Finally, disjoint union corresponds to addition by placing the two tubes in disjoint balls and pinching the complement to the base point. □

13.4 What stabilization remembers

A normal framing is a specified ordered trivialization of the normal bundle, stabilized by adding the standard last coordinate. It gives a collapse to S^k, since the classifying map of the trivialized bundle factors through the sphere model of the framed theory. The same preimage operation recovers a framing from a transverse inverse image of a regular point of S^k. Stable framed bordism therefore has groups $\pi_d^S = \varinjlim_k \pi_{d+k}(S^k)$. Forgetting a framing and retaining its induced orientation gives the map from the sphere spectrum to MSO on coefficient groups. Distinct framings on the same underlying manifold can give distinct classes. In dimension zero, the relative orientation of an ordered normal frame and the ambient frame contributes the sign of the point, recovering the stable degree.

The theorem classifies manifolds after passing to bordism classes. A space of embeddings and collars also retains paths, homotopies of paths, and automorphisms of individual geometric representatives. The collapse construction runs in families as well as on points. The next chapter retains composable geometric bordisms before taking a classifying space and distinguishes the Thom spectrum in that setting from the normal Thom spectrum used here.

CHAPTER 14

Bordism as a higher category

Collapsing a manifold records its bordism class. Bordisms also form a system of composable arrows, and many of their geometric constructions compare arrows rather than objects. At the next level, a family of diffeomorphisms compares such comparisons. A categorical model must specify which differences are retained at each level.

14.1 Gluing and its space of choices

Fix d. At a real level a, an object is a closed $(d-1)$-manifold embedded in $\{a\} \times \mathbb{R}^N$. An arrow to an object at level $b > a$ is a compact d-submanifold $W \subset [a,b] \times \mathbb{R}^N$ whose only boundary consists of the designated objects and whose embedding is a product in a neighborhood of each end. At an interface level b, two consecutive arrows agree on the object and their product neighborhoods glue. At the level of these embedded representatives the composition is union, so equality of the two composites of three consecutive arrows follows from associativity of union. This defines the positive-length part of an embedded cobordism category C_d. One includes identities when forming its category or uses the corresponding model with nonidentity arrows and a semisimplicial classifying space.

Without fixed ambient levels one can glue two abstract bordisms by identifying boundary collars. The overlap chart on the union uses a signed coordinate on the two inward half collars. Its smooth structure is independent up to a diffeomorphism of collars, and repeated gluing keeps the same signed coordinate along each seam. Retaining the space of such gluings is useful for comparison with embedded families.

For a fixed abstract bordism W with its collared ends, the embedding model has a space of embeddings of W modulo its diffeomorphisms. A diffeomorphism acts freely on parametrized embeddings by precomposition. On a closed manifold, stable embeddings in increasing Euclidean dimensions form a weakly contractible space: a compact family at a finite stage can be deformed to one fixed embedding using an independent block, as in Chapter 13. For a sphere of parameters this deformation fills the sphere by contracting it. In these models one obtains classifying spaces of diffeomorphism groups by taking the homotopy quotient of the embedding space under reparametrization. The use of homotopy quotients retains families with nontrivial monodromy.

Example 14.1 (Parity and the classifying space)**.** Taking the nerve and geometric realization of C_d gives a space BC_d. An arrow becomes a path in this space and can be traversed in either direction. On connected components this identifies objects exactly up to unoriented bordism. To verify the claim in one direction, any abstract bordism can be embedded in a sufficiently large Euclidean slab with collared ends. For the other direction, paths in the space of object embeddings preserve their diffeomorphism type, and one-simplices preserve the bordism class. A finite zigzag with backward arrows gives a bordism after reversing the corresponding slabs. In particular $\pi_0 BC_1 \simeq \mathbb{Z}/2$ via parity of points. This calculation alone leaves the higher homotopy groups of BC_1 undetermined.

14.2 The spectrum of an embedded cobordism category

The Pontryagin–Thom spectrum of the preceding chapter stabilizes the normal rank of individual closed manifolds. For the classifying space of a category with fixed d-dimensional arrows one instead fixes the tangent rank and varies the complementary normal directions. Over $\mathrm{Gr}_d(\mathbb{R}^{d+n})$ let $\gamma_{d,n}$ be the tautological d-plane bundle and $\gamma_{d,n}^{\perp}$ its rank-n orthogonal complement. The spectrum $MTO(d)$ is the Thom spectrum of the virtual bundle $-\gamma_d$ on $BO(d)$. Its $(d+n)$th space can be modeled by $\mathrm{Th}(\gamma_{d,n}^{\perp})$. Adding an ambient coordinate adds a trivial line to the complements along the inclusion of Grassmannians and gives the structure maps.

Theorem 14.2 (Galatius–Madsen–Tillmann–Weiss)**.** *The classifying space of the stable embedded unoriented d-cobordism category is weakly equivalent to $\Omega^{\infty-1} MTO(d)$. For the category with tangent structures classified by a fibration $\theta : X \to BO(d)$, the corresponding spectrum is the Thom spectrum of $-\theta^{*}\gamma_d$.*

Here the shifted infinite loop space is represented by the stabilization of the spaces $\Omega^{d+n-1}\,\mathrm{Th}(\gamma_{d,n}^{\perp})$. The theorem of [10], whose proof uses parametrized geometry, starts with collapse near embedded manifolds over their level coordinate. The tangent d-plane specifies the Grassmannian coordinate of the collapse. The remaining directions supply a point in the complementary normal disk. Compatibility with a cut recovers categorical composition at the level of the classifying space. The theorem's converse compares families of such local data to geometric bordisms, a stronger statement than the inverse-image computation on individual maps in Theorem 13.3. The difference between the spectra is visible in their bundle ranks. In MO the universal normal bundle's rank increases in the stabilization. In $MTO(d)$ the positive tangent plane has fixed rank d, and its negative virtual class is stabilized.

An oriented version uses $BSO(d) \to BO(d)$. The structure on a $(d-1)$-dimensional object is a lift of its tangent bundle with a distinguished added interval direction. This recovers the boundary orientation convention when gluing two oriented arrows. In dimension zero the positive-length category has only the empty manifold as an object type, and its arrows are finite point configurations in open slabs. The emptiness of the boundary in this case is consistent with the shift: $MTO(0)$ is the sphere spectrum and its shifted infinite loop space is connected.

14.3 Corners, duals, and field theories

A fully extended bordism category starts at dimension zero. Its j-morphisms for $0 \le j \le d$ are j-dimensional bordisms with the relevant corners and collars. Geometrically a corner has several independent boundary coordinates, so that composition is gluing along a selected face. As in the ordinary bordism category, one retains diffeomorphisms and their homotopies above the geometric dimensions. With disjoint union this gives a symmetric monoidal (∞, d)-category. In the d-framed version $\mathrm{Bord}_d^{\mathrm{fr}}$, a j-manifold is equipped with a compatible trivialization of its tangent bundle with $d-j$ additional ordered directions. Along collars these added directions identify the boundary framing with the restriction of the framing in the next dimension.

A dual for an object A consists of an object $A^\vee$ and coevaluation and evaluation arrows. They have composites that are equivalent to the respective identities. One of the two triangle identities has diagram

$$\begin{array}{ccc} & \mathrm{coev}\otimes 1 \swarrow & \otimes\, A^\vee \otimes A \\ & & \downarrow 1\otimes \mathrm{ev} \\ & 1_A \searrow & \downarrow \\ & & A. \end{array}$$

In dimension one the intervals with both ends outgoing or both ends incoming supply these maps for the two framings of a point. The triangle identities come from gluing intervals and rectifying the result as a collared interval. In higher dimensions the evaluation and coevaluation morphisms have adjoints and the units and counits have coherent adjointability data through the permitted levels. This is the geometric structure behind full dualizability. Full dualizability is defined in the target symmetric monoidal (∞, d)-category using these successive levels of adjointability.

The framed cobordism hypothesis identifies the space of symmetric monoidal functors out of $\mathrm{Bord}_d^{\mathrm{fr}}$ with the maximal ∞-groupoid of fully dualizable objects of the target, by evaluation at the framed point [21, §2.4]. The functor space here uses equivalences of functors. Other tangential structures require compatible

symmetry data on a fully dualizable object. They do not result from taking the same target space and merely relabeling the point. This universal property refers to functors from the extended higher category. The classifying-space theorem refers to a groupoid completion of the embedded category with adjacent dimensions, so it answers a different question about the geometry of bordisms.

Example 14.3 (A trace from gluing)**.** In dimension one take the target to be finite-dimensional vector spaces over a field k with the tensor product. If the positively framed point is sent to V, its dual point is sent to V^*. The evaluation and coevaluation are the pairing and the tensor $\sum_i e_i \otimes e_i^*$. The tensor is independent of basis because it corresponds to the identity in $\mathrm{End}(V)$. Gluing the outgoing ends of the coevaluation to the incoming ends of the evaluation, with the symmetry exchanging the middle factors, makes a circle. Its image is multiplication by $\sum_i e_i^*(e_i) = \dim_k V$ viewed in k. More generally, inserting an endomorphism on the strand gives its usual trace in the target. The circle evaluation is therefore a consequence of coherent gluing and the choice of dual.

Example 14.4 (A surface cut into pairs of pants)**.** In dimension two the ordinary oriented bordism category has circles and their disjoint unions as objects. In a symmetric monoidal functor to finite-dimensional vector spaces, put A for the vector space assigned to one circle. The twice-incoming pair of pants supplies a multiplication $m : A \otimes A \to A$, and the outgoing disk is its unit. There are two ways to glue two pairs of pants to get three incoming circles and one outgoing circle. The resulting labeled surfaces are diffeomorphic relative to their boundary collars, so their images give associativity. A diffeomorphism interchanging the two incoming circles gives commutativity. Reversing the pants and the disk gives a comultiplication and a counit ϵ.

A cylinder drawn as a cup can be cut into an outgoing disk and a one-incoming pair of pants. Its reverse gives the pairing $(u, v) \mapsto \epsilon(m(u, v))$. The cup and cap satisfy the cylinder identities, so this pairing is nondegenerate. The two ways to cut a surface from two circles to two circles into oppositely directed pairs of pants give the Frobenius equality between comultiplication after multiplication and multiplication on one side of a comultiplication. These statements are consequences of labeled gluing and the elementary classification of the relevant surfaces (Chapter 16). They concern theories defined on circles and surfaces. A fully extended theory also has values on points and compatible additional gluing.

A geometric family also has a different source of higher data: paths in an embedding space act through ambient diffeomorphisms, with a stabilizer of the embedding. The next chapter constructs these lifts, then describes the distinction between a path of images and a chosen lift through the ambient transformations.

CHAPTER 15

Isotopies and ambient transformations

An isotopy is a family of embeddings. An ambient isotopy is a family of diffeomorphisms of the target starting at its identity. The extension theorem turns the first type of family into the second by extending its time-dependent velocity. The resulting factorization has a whole space of choices, which can be understood as a stabilizer in a space of transformations.

15.1 Extending velocities

Theorem 15.1 (Isotopy extension with smooth parameters)**.** *Let A be a closed smooth manifold, and let N be a smooth manifold without boundary. Let B be a compact smooth parameter manifold. Suppose $e : B \times [0,1] \times A \to N$ is smooth and every $e_{b,t}$ is an embedding. There exists a smooth family of compactly supported ambient isotopies $\Phi_{b,t}$ of N with*

$$\Phi_{b,0} = 1_N, \qquad \Phi_{b,t} \circ e_{b,0} = e_{b,t}.$$

The support can be taken in one compact subset of N independent of b, t. If $e_{b,t}$ is constant in t for designated parameters b, then $\Phi_{b,t}$ can be chosen equal to the identity for those parameters.

Proof. The velocity along $e_{b,t}$ is $v_{b,t}(e_{b,t}(a)) = \partial_t e(b,t,a)$. This is well-defined because $e_{b,t}$ is injective. Choose a metric on N. The normal spaces to the embedded slices form a bundle over $B \times [0,1] \times A$. The normal exponential map with b, t unchanged is a tubular coordinate map for a uniform sufficiently small radius. In more detail, it has invertible derivative along the zero section in the manifold with parameters. Local inverse maps give neighborhoods near that section. If there were intersections between distinct normal disks at arbitrarily small radii, compactness would give a limiting pair of points over the same b, t. The injectivity of $e_{b,t}$ forces the pair to coincide, and a local inverse then rules out the intersections for the subsequence. This gives the required uniform radius.

For a point in the tube at parameters b, t, transport the slice velocity from the foot point along the normal geodesic. Multiply by a cutoff in the normal radius that is one near the zero section and zero near the boundary of the tube.

Extend by zero on the complement. The resulting vector field $V_{b,t}$ on N depends smoothly on the parameters and has uniform compact support. The extension uses the velocity at the same b, t and sends a zero velocity on a whole slice to the zero field. The time-dependent flow exists for $0 \leq t \leq 1$ by compact support, and smooth dependence on initial conditions and parameters gives $\Phi_{b,t}$. Solving the equation backward shows each time map is a diffeomorphism. The two curves $t \mapsto e_{b,t}(a)$ and $t \mapsto \Phi_{b,t}(e_{b,0}(a))$ have the same initial point and satisfy the same differential equation. Uniqueness proves the result. □

The theorem gives a local extension principle as well. An open subset of N containing the track of the isotopy contains the tubes of sufficiently small radius. The vector field may therefore be supported there. When the images miss a closed obstacle, this produces an ambient isotopy fixed close to the obstacle. Compactness is used to obtain a common tube and support. Without it a pointwise prescription of velocities may produce a field that escapes to infinity in finite time. For sources with boundary or a target with boundary, prescribed tangencies and compatible collar data require a relative version. The classical extension and gluing results for those settings are developed in [15, Ch. 8].

15.2 Restriction and the stabilizer

For this section take N compact without boundary. Use the C^∞ topologies on spaces of smooth maps and their subspaces. Write $i : A \hookrightarrow N$ for one embedding, $\mathcal{D}_0(N)$ for the identity path component of the diffeomorphism group, and $\mathcal{E}_i$ for the path component of i in the embedding space. The restriction map

$$q : \mathcal{D}_0(N) \longrightarrow \mathcal{E}_i, \qquad q(F) = F \circ i$$

has fiber $K_i = \{F \in \mathcal{D}_0(N) \mid F \circ i = i\}$ at i. In particular K_i fixes the specified parametrized embedding pointwise and may have many path components.

Proposition 15.2 (The restriction bundle)**.** *The map q is a locally trivial principal K_i-bundle. It yields a homotopy fiber sequence*

$$K_i \longrightarrow \mathcal{D}_0(N) \longrightarrow \mathcal{E}_i.$$

Proof. Here is the local-section argument behind the classical restriction theorem [27]. When a smooth embedding j is sufficiently C^1-close to i, the vector $v_j(a) = \exp^{-1}_{i(a)}(j(a))$ is defined and varies continuously with j in the C^∞ topology. Extend this section along i to a vector field V_j on N by normal projection, transport, and a fixed cutoff. For j sufficiently close, the formula $F_{j,t}(x) = \exp_x(tV_j(x))$ defines a local diffeomorphism for each t, by C^1 closeness to the identity. Compactness makes it a covering map on each connected component. Being homotopic to the

identity, it induces an isomorphism on the fundamental group of each component, so the covering has one sheet. Thus it is a diffeomorphism, and $j \mapsto F_{j,1}$ is a continuous section near i with values in $\mathcal{D}_0(N)$.

The same construction based at every embedding shows that all $\mathcal{D}_0(N)$-orbits in the full embedding space are open. An orbit is path connected, and each orbit's complement is open because it is a union of other open orbits. A path cannot meet both an orbit and its complement. Consequently the orbit of i is precisely $\mathcal{E}_i$. Translation of the local section by elements of $\mathcal{D}_0(N)$ gives sections throughout that orbit. If s is a section over U, the map $(j, k) \mapsto s(j)k$ identifies $U \times K_i$ with $q^{-1}(U)$, since any two diffeomorphisms sending i to j differ by a unique element of K_i on the right. The inverse is the restriction together with $F \mapsto s(q(F))^{-1}F$. Local triviality gives the stated homotopy fiber sequence. □

For example, an ambient lift of an embedding loop, started at 1_N, ends at an element of K_i. Its component is independent of the chosen lift. This follows because two lifts F_t, F'_t starting at the identity satisfy $F_t^{-1}F'_t \in K_i$, giving a path from the identity that relates their endpoints. With conventional base-point choices this is the connecting map from $\pi_1(\mathcal{E}_i, i)$ to $\pi_0(K_i)$. It detects when a loop cannot be lifted to a loop of ambient diffeomorphisms.

Example 15.3 (Point pushing on a surface)**.** Take A to be a point and N the closed oriented surface of genus two. Then $\mathcal{E}_i = N$ and q evaluates a diffeomorphism at the designated point x_0. Every loop F_t in $\mathcal{D}_0(N)$ based at the identity evaluates to a loop representing a central element of $\pi_1(N, x_0)$. Indeed, for any based loop γ_s the map $(t, s) \mapsto F_t(\gamma_s)$ defines a map from the torus to N. The two generators of the torus group map to the evaluated loop and γ, so their classes commute.

There exists a loop α in N whose class is noncentral. In the genus-two polygon presentation with generators a, b, c, d and relation $[a, b][c, d] = 1$, sending b, d to 1 realizes a surjection onto the free group on a, c. The image of a does not commute with c, so one may take α to represent a. Its ambient lift, supplied by the restriction bundle, cannot end in the identity component of K_i. If it did, a path in the stabilizer to 1_N would close it into a loop of ambient diffeomorphisms without changing the evaluated class. This contradicts centrality. The returned point is the same, while the endpoint ambient transformation has acquired a nontrivial component in the stabilizer.

The action groupoid of $\mathcal{D}_0(N)$ on $\mathcal{E}_i$ has classifying space equivalent to BK_i. Indeed local triviality identifies $\mathcal{E}_i$ with $\mathcal{D}_0(N)/K_i$. Using any contractible free model $E\mathcal{D}_0(N)$, its homotopy quotient is the space in the middle of

$$E\mathcal{D}_0(N) \times_{\mathcal{D}_0(N)} \mathcal{E}_i \xrightarrow{\simeq} E\mathcal{D}_0(N)/K_i \xrightarrow{\simeq} BK_i.$$

The transitive action therefore gives a single orbit, while the groupoid retains the stabilizer. A problem about continuous choices of ambient extensions can be formulated as a lifting problem for q.

15.3 Changing a seam by isotopy

Let M, N have boundary components $S \subset \partial M$ and $T \subset \partial N$ with collars, and let $\varphi : S \to T$ be a diffeomorphism. If u_t is an isotopy of diffeomorphisms of S beginning at the identity, put $\psi = \varphi \circ u_1$. The manifolds obtained by gluing along ψ and along φ are diffeomorphic. To construct the map, choose a collar $c : S \times [0, \epsilon) \to M$ and a smooth β equal to 1 near zero and to 0 near the far end. On the collar use the diffeomorphism

$$c(x, s) \longmapsto c(u_{\beta(s)}(x), s).$$

Extend it as the identity on the rest of M and use the identity on N. In the source gluing, x and $\psi(x)$ are identified. Their images are $u_1(x)$ and $\varphi(u_1(x))$, which agree in the target. Because the transformation is a product near the seam, the descended map is smooth in the signed collar coordinate and has a smooth inverse. The same construction works with disjoint seams by choosing disjoint collars, so their changes commute on the resulting manifold.

A basic isotopy of disks illustrates how local differential data can fix a model. For an embedding $e : D^r \to \mathbb{R}^n$ smooth up to the boundary, the maps

$$x \longmapsto e(0) + \frac{e(tx) - e(0)}{t}, \qquad 0 < t \leq 1,$$

are embeddings, and the integral formula for the difference quotient gives a smooth limit $x \mapsto e(0) + D_0 e(x)$ at $t = 0$. The limit is still embedded because the differential at zero is injective. This isotopy varies the image and need not preserve boundary values. Relative isotopy classes with fixed boundary can contain further information.

CHAPTER 16

Surfaces and families of models

Surfaces illustrate two different classification problems. The first identifies an individual surface up to existence of a diffeomorphism. The second classifies families of these surfaces, including the diffeomorphisms used when parameters return around a loop. Handles answer the first problem, and the classifying space of the diffeomorphism group specifies the second.

16.1 Disk and band presentations

Let S be a closed connected smooth surface. Start with a finite triangulation. Thicken vertices to disks, edges to bands, and faces to the complementary disks. This is the smooth 0, 1, and 2-handle description after rounding corners. It is also the type of decomposition produced by Morse theory. It can be simplified to a single starting disk and a single final disk as follows.

First merge along the bands of a spanning tree T of the 1-skeleton. A pair of disks joined by a band between their distinct boundary circles forms a disk after rounding. As a result the union over the tree is one disk. Next take the dual graph on faces, using only edges which do not cross T. This graph is connected. Otherwise a nonempty proper union of its components would give a mod-2 sum of faces whose boundary uses only edges of T. That boundary is nonzero because the full dual graph is connected. It would be a nonzero 1-cycle in a tree. Choose a dual spanning tree and merge the final face disks across each crossed band in that tree. This leaves one final disk and some number r of bands. Consequently $\chi(S) = 2 - r$. In particular, if $\chi(S) = 2$ there are no bands and S is a sphere. Two disks glued along their boundaries give the usual sphere because any diffeomorphism of their boundary circles extends to a disk. Indeed an orientation-preserving circle diffeomorphism is isotopic to the identity by linear interpolation of a monotone lift with the identity lift. The collar construction of Chapter 15 extends this isotopy. A reflection accounts for the other sign.

16.2 Cutting along simple curves

A mod-2 1-cycle on a surface can be represented by a disjoint collection of embedded smooth circles. For a simplicial representative there is an even number

of incident selected edges at each vertex. In a small disk around the vertex, pair neighboring endpoints and round their arcs. The modification has the same mod-2 relative class in this disk because the set of endpoints is unchanged and the disk has zero first homology. The modified edges give the required collection.

Suppose first that S is oriented and has $\chi(S) < 2$. Poincaré duality over $\mathbb{F}_2$ gives nonzero $H_1(S;\mathbb{F}_2)$. In a representative collection for a nonzero class at least one circle is itself nonzero in homology, hence nonseparating. Otherwise it would bound one side modulo 2. An oriented surface gives each embedded circle an orientable normal line bundle, so such a circle has an annular neighborhood. Cutting along it leaves a connected surface with two additional boundary circles. Cap them with disks to obtain S', with $\chi(S') = \chi(S) + 2$.

Recover S from S' by adding a torus summand. To see the summand, choose an arc joining the capped disks in S'. A neighborhood of the arc together with the two disk caps is itself a disk. Removing the two caps from this neighborhood gives a disk with two holes, and the original annulus joins the two inner circles. This is a punctured torus, using the usual square model of the torus cut along one of its basic curves. Replacing the disk by the punctured torus is the connected sum. Thus iteration, terminating when $\chi = 2$, gives a connected sum of tori, with the empty sum denoting the sphere.

When S is nonorientable, its orientation character $w_1(TS) : H_1(S;\mathbb{F}_2) \to \mathbb{F}_2$ is nonzero. Apply the circle representation above to a class detected by this character. Some component is one-sided, since the tangent to a circle is orientable and the normal line has nontrivial orientation. Its neighborhood is a Möbius band. The complement is connected: every component must meet the new boundary, which is just one circle, since S is connected. Removing the band and capping the new boundary gives S' with $\chi(S') = \chi(S) + 1$. Recovering S replaces the cap by a Möbius band, which is the punctured projective plane. Iteration leaves an orientable surface together with at least one projective-plane summand.

Lemma 16.1 (Exchanging a handle and a crosscap)**.** *The surface $T^2\#\mathbb{RP}^2$ is diffeomorphic to the connected sum of three projective planes.*

Proof. Use mod-2 homology and let a, b be the two basic torus classes and c the core of the projective-plane band. Their only nonzero pairings under intersection are $a \cdot b = b \cdot a = 1$ and $c \cdot c = 1$. Here the self-intersection of the one-sided core is one, as can be computed by a perturbation in the nontrivial normal line. The orientation character is zero on a, b and one on c.

Join disjoint representatives of a and c by an embedded band. Replacing short subarcs by the opposite sides of the band gives an embedded circle of class $x = a + c$. It is one-sided. Let C denote the complement of the interior of its

Möbius neighborhood. The Mayer–Vietoris sequence over $\mathbb{F}_2$ gives

$$H_1(S;\mathbb{F}_2) \simeq H_1(C;\mathbb{F}_2) \oplus \langle x \rangle.$$

Indeed the sole gluing circle is zero in the homology of both sides, where it bounds the mod-2 fundamental class relative to that circle, and both sides are connected. The image of $H_1(C;\mathbb{F}_2)$ is therefore the orthogonal complement of x: it pairs trivially with x, and $x \cdot x = 1$. The class $b + c$ belongs to this complement and has nonzero orientation character. Hence C, and also its capped version, is nonorientable.

The capped version has Euler characteristic zero. Removing another one-sided band and capping leaves Euler characteristic one, which cannot occur for a closed orientable surface by the preceding torus reduction. One more such removal leaves Euler characteristic two and therefore a sphere. Reattaching the three bands proves the claim by the description of capping as connected sum. □

Theorem 16.2 (Classification of compact connected surfaces)**.** *Let S have b boundary components. Up to diffeomorphism, S is a surface obtained by deleting b open disks from one of the following closed models:*

$$S_g = \#^g T^2 \quad (g \geq 0), \qquad N_h = \#^h \mathbb{RP}^2 \quad (h \geq 1).$$

The type and its index are determined by orientability, b, and $\chi(S)$. The formulas are $\chi(S_g) - b = 2 - 2g - b$ and $\chi(N_h) - b = 2 - h - b$.

Proof. For closed surfaces, the reductions and Lemma 16.1 give the models. Connected sum subtracts two from the sum of Euler characteristics, since each puncture removes a disk and the gluing circle has Euler characteristic zero. This yields the displayed formulas, which distinguish different indices within each orientability type.

For a surface with boundary, cap all the circles. Orientability of the capped model is equivalent to orientability of S, since an orientation on S extends over each disk with the matching boundary convention. The positions of the disjoint cap disks carry no extra invariant for an unmarked surface. Shrink each disk inside a collar neighborhood, move the resulting disks along paths of centers in the complement of the other disks, and use diffeomorphisms in the coordinate neighborhoods of the disk images to match the final disks. This is the disk isotopy construction of [15, Chs. 8–9]. The set of boundary circles may be permuted because no boundary labels are part of the classification. □

16.3 Models in a family

For any fixed smooth surface S, the groupoid with its diffeomorphisms as arrows has classifying space $B\,\mathrm{Diff}(S)$ when the arrows have the C^∞ topology. A locally

trivial S-bundle with smooth structure on the fibers has continuous transition functions valued in this C^∞ group. On a paracompact base where the cover is numerable, such bundles are classified up to isomorphism by homotopy classes of maps to $B\,\mathrm{Diff}(S)$. The classifying map of the principal frame bundle sends the fiber model to a copy of S:

$$\begin{array}{ccc} E & \longrightarrow & E\,\mathrm{Diff}(S) \times_{\mathrm{Diff}(S)} S \\ \downarrow & & \downarrow \\ B & \text{classifying map} & B\,\mathrm{Diff}(S). \end{array}$$

The square describes the pullback of the universal bundle in a model for which the classifying map is chosen. A proper smooth submersion of boundaryless manifolds with compact surface fibers provides such local trivializations. Choose a horizontal complement to the vertical tangent bundle and lift straight paths in a small coordinate ball in the base. Properness gives existence for the bounded time interval, and reversal of the paths gives the inverse maps. This identifies all nearby fibers smoothly with the chosen reference fiber.

For example, form a torus bundle over a circle by gluing the ends of a product interval with the torus diffeomorphism induced by $\left(\begin{smallmatrix} 1 & 1 \\ 0 & 1 \end{smallmatrix}\right)$ on $\mathbb{R}^2/\mathbb{Z}^2$. The homology groups of the fibers form a local system with this monodromy on $H_1(T^2;\mathbb{Z})$. A product bundle would have identity monodromy. Conjugation cannot identify the displayed matrix with the identity. Thus the genus specifies the fiber of this family, while the classifying map records its nontrivial gluing over the parameter circle.

CHAPTER 17

Smooth spaces and quotient stacks

A smooth manifold can be studied through the objects that map into it. A chart is a particularly useful test map, and a parametrized family is another. By allowing all smooth test manifolds, one retains families that pass through singular orbits and through changes of local trivialization. The resulting category also admits quotients with automorphisms.

17.1 Smooth test manifolds

Use the category Man of finite-dimensional smooth manifolds with the open-cover topology. In this part the test manifolds have no boundary, and the underlying site is made small in a chosen universe. One can enlarge it to a site of manifolds with corners when boundary charts are required. A smooth space in this chapter is a sheaf of spaces on the site. Write $\mathcal{E}$ for the resulting ∞-topos and $\underline{M}$ for the sheaf $U \mapsto C^\infty(U, M)$ represented by a manifold M. Its values are sets regarded as discrete spaces. This differs from sending M to its singular complex: the latter has paths and forgets the smooth structure, while $\underline{M}$ records smooth families.

The presheaf represented by M is a sheaf. Given maps from an open cover of U to M that agree on pairwise intersections, their glued map is continuous and locally smooth, hence smooth. By Yoneda every morphism $\underline{M} \to \underline{N}$ comes from a unique smooth map $M \to N$. Products are preserved: a map $U \to M \times N$ is exactly a pair of smooth maps. A fiber product in smooth spaces always exists. If the two maps to a representable target are transverse, the ordinary smooth fiber product represents it, because the ordinary submanifold cut out by the transverse equations has the universal lifting property of smooth maps. For singular equations the sheaf fiber product still makes sense. Derived smooth geometry carries additional information about those equations, as developed in Chapter 21.

Chapter 1 gives a smooth map that is a homeomorphism of its underlying spaces without being a diffeomorphism. The failure of invertibility is detected by the identity test family on its target. Thus recovering a space from all of its smooth probes also recovers the differential structure of its maps.

17.2 Locally represented geometry

Write $\mathcal{X}(U)$ for the space of U-families in a smooth space $\mathcal{X}$. Paths in this space compare families, and higher simplices compare such paths. In a stack valued in groupoids these paths are isomorphisms. Descent replaces literal agreement on overlaps by compatible identifications. For an open cover U_i of U, the comparison is

$$\mathcal{X}(U) \simeq \lim_{[n]\in\mathbf{\Delta}} \prod_{i_0,\dots,i_n} \mathcal{X}(U_{i_0} \cap \cdots \cap U_{i_n}).$$

For a groupoid-valued example, local objects and isomorphisms on double overlaps must satisfy a cocycle condition on triple overlaps. Morphisms between global objects are computed by maps of descent data. A map from a manifold into a stack can therefore be locally represented by maps to an atlas with transitions between the local lifts.

An atlas for a differentiable stack is a representable surjective submersion from a manifold. Such a map admits local sections after pullback to a test manifold. It is therefore an effective epimorphism in $\mathcal{E}$. Its Čech groupoid consists of the atlas, the space of arrows between its two lifts, and the higher fiber products encoding composition. The stack is the sheaf colimit of this groupoid [20, §6.1]. The manifold structure of the arrow space distinguishes a smooth groupoid from its classifying homotopy type.

17.3 An action groupoid as a quotient

Let a Lie group G act smoothly on M from the left. Its action groupoid has object space M and arrow space $G \times M$, with source $(g, x) \mapsto x$ and target $(g, x) \mapsto gx$. The quotient $[M/G]$ is the colimit of the simplicial nerve in smooth spaces. It classifies local representatives of points modulo gauge transformations.

Proposition 17.1 (Families in an action quotient)**.** *A map from a test manifold U to $[M/G]$ can be described by a smooth principal right G-bundle $\pi : P \to U$ and a smooth map $a : P \to M$ satisfying $a(pg) = g^{-1}a(p)$. Arrows between such data are bundle isomorphisms preserving these maps.*

Proof. Choose a local section s_i of the torsor over each open U_i and write $s_j = s_i g_{ij}$ on overlaps. The transition maps satisfy $g_{ij}g_{jk} = g_{ik}$. The points $x_i = a(s_i)$ obey $x_j = g_{ij}^{-1}x_i$, so the action arrows from x_j to x_i carry exactly the transition data of the torsor. Conversely, descent data (x_i, g_{ij}) glue $U_i \times G$ by $(u, h)_j \sim (u, g_{ij}(u)h)_i$. The functions $(u, h)_i \mapsto h^{-1}x_i(u)$ agree with this identification and give a. Changing local sections changes the data by a groupoid isomorphism. Every bundle isomorphism preserving the map to M is locally described by such changes. The constructions thus identify both objects and arrows in the quotient stack. □

The map $[M/G] \to [*/G]$ forgets the equivariant map and keeps its torsor. Pulling it back along the trivial torsor with a chosen trivialization gives back M.

$$\begin{array}{ccc} \underline{M} & \longrightarrow & [M/G] \\ \downarrow & & \downarrow \\ * \text{ trivial torsor} & & [*/G]. \end{array}$$

For a free proper action, $M \to M/G$ is a principal bundle. Writing a map to the manifold quotient as the pullback of this bundle with its tautological map to M gives an inverse to the map $[M/G] \to \underline{M/G}$. In general this equivalence fails at any point with a nontrivial stabilizer. The stabilizer is precisely the automorphism group of the corresponding family over a point.

17.4 The reflection quotient

For $G = C_2$ acting on $\mathbb{R}$ by reflection, the topological orbit space is the closed half-line, with coordinate t^2. The origin of the quotient stack has automorphism group C_2 while a nonzero orbit has none. A smooth family supported at the origin on a connected base is a principal C_2-bundle equipped with the zero map. Over S^1 the trivial cover and the connected double cover give distinct families, even though the same coarse function $t^2 = 0$ represents both.

The connected cover with total space $\mathbb{R}/(4\pi\mathbb{Z})$ admits a weight $w([\theta]) = \cos(\theta/2)$ that is odd under the deck transformation. Its square descends to the smooth function $(1 + \cos\theta)/2$. There is no smooth real function on the base with that square. In a local coordinate centered at the unique zero, the square vanishes quadratically with positive quadratic coefficient. A smooth square root must have nonzero derivative and thus change sign there. The complement of the unique zero is connected, so a real nonzero function there could not change sign. The quotient stack permits the pair of local roots together with a sign transition. Smooth differential data see an obstruction absent from the coarse observable.

The next chapter keeps such transition data while forming spaces of maps and composing families. Chapter 19 uses a related descent construction for orientations of bundles.

CHAPTER 18

Mapping objects and families over a base

The space of maps between manifolds has several useful topologies, as in Chapter 3. Smooth families can also be specified without choosing a topology on the set of maps. This uses all smooth test objects and an evaluation morphism. The same construction makes sense for families over a base that is itself a stack.

18.1 Smooth mapping objects

For smooth spaces X, Y in the ∞-topos $\mathcal{E}$, write $[X, Y]$ for their exponential. Thus

$$\mathrm{Map}_{\mathcal{E}}(T, [X, Y]) \simeq \mathrm{Map}_{\mathcal{E}}(T \times X, Y)$$

naturally in T. When $X = \underline{M}$ and $Y = \underline{N}$, Yoneda computes its value on a test manifold U:

$$[\underline{M}, \underline{N}](U) = C^{\infty}(U \times M, N).$$

This specifies the regularity of derivatives in both the test and input variables. It has no requirement that the exponential be a finite-dimensional manifold. Even the smooth maps from a circle to a line have infinite-dimensional space of variations. The sheaf model also distinguishes changes of parameter that are smooth from continuous paths for a chosen function-space topology. The comparison with convenient calculus can be made in settings in which both sides' smooth test plots agree [2, 17].

As a basic example, a family of real functions on M is a map $P \to [\underline{M}, \underline{\mathbb{R}}]$. Taking its f at a point p gives a function f_p. Restricting this family along $a : P' \to P$ gives $f_{a(p')}$. Differentiating with respect to the input gives the vertical one-form $d_M f$ on $P \times M$, a section of the pullback of T^*M. The restriction of this section along a is d_M of the restricted function, since the maps in the input directions are unchanged by varying a. Differentiation with respect to the parameter instead pulls its covectors back along Da. The two directions have different transformation laws even for a trivial product.

18.2 Evaluation over a base

Let B be any object of $\mathcal{E}$. Products and exponentials in its slice are written $\times_B$ and $[X, Y]_B$. This is the category of families varying over B. The defining evaluation and its adjunction are

$$\mathrm{ev}_B : [X, Y]_B \times_B X \longrightarrow Y,$$
$$\mathrm{Map}_B(P \times_B X, Y) \simeq \mathrm{Map}_B(P, [X, Y]_B).$$

So a map $f : P \times_B X \to Y$ gives its classifying morphism $r_f : P \to [X, Y]_B$. A reparametrization includes a map between parameter objects over B commuting with the classifying maps up to a specified homotopy. These reparametrizations and their higher comparisons form the slice $\mathcal{E}_{/[X,Y]_B}$. The identity family has parameter B. A family from X to Y with parameter P, followed by one from Y to Z with parameter Q, has parameter $P \times_B Q$. Its evaluation is the same as applying the two evaluations in order. The associativity comparison comes from the product comparison on three parameters.

Proposition 18.1 (Base change for relative evaluation)**.** *For $a : C \to B$, pullback sends $[X, Y]_B$ to $[a^*X, a^*Y]_C$ by an equivalence preserving evaluations. It preserves composition and reparametrization of families.*

Proof. The functor $a_!$ left adjoint to pullback sends an object over C to its structure map composed with a. For $T \to C$,

$$a_!(T \times_C a^*X) \simeq (a_!T) \times_B X$$

because each side is the same pullback equipped with the map to B. Therefore

$$\begin{aligned}\mathrm{Map}_C(T, a^*[X, Y]_B) &\simeq \mathrm{Map}_B(a_!T, [X, Y]_B)\\ &\simeq \mathrm{Map}_B((a_!T) \times_B X, Y)\\ &\simeq \mathrm{Map}_C(T \times_C a^*X, a^*Y).\end{aligned}$$

These comparisons are natural in T. Yoneda gives the equivalence, and the construction evaluates the same map under the adjunctions. The product that defines composition is preserved by pullback. Naturality with respect to P gives the last assertion. □

Proposition 18.2 (Descent of families)**.** *Let $U \to B$ be an effective epimorphism with Čech nerve $U_\bullet$. Set $X_n = X \times_B U_n$ and $Y_n = Y \times_B U_n$. Then the category of parametrized maps from X to Y is equivalent to*

$$\mathcal{E}_{/[X,Y]_B} \simeq \lim_{[n]\in\mathbf{\Delta}} \mathcal{E}_{/[X_n,Y_n]_{U_n}}.$$

Proof. Descent for slices of an ∞-topos identifies $\mathcal{E}_{/B}$ with the limit of $\mathcal{E}_{/U_n}$ [20, §6.1.3]. Slicing at a compatible target commutes with taking that limit. Indeed an arrow into the target is a compatible family of local arrows into its components, with the same statement for all mapping spaces. Under the first equivalence the image of the target exponential is the compatible family of local exponentials by Proposition 18.1. □

In ordinary charts with transition equivalences chosen, this result glues local evaluations when the comparisons commute. If an input and output are bundles with changes of frames a_{ij} and b_{ij}, the local maps transform as $f_i = b_{ij} f_j a_{ij}^{-1}$ when the parameter coordinates have already been transported. If the same parameter supplies two stages of a composite, the pullback along the diagonal $P \to P \times_B P$ implements that sharing. In sheaves with higher morphisms, local comparisons themselves need coherence on multiple overlaps. The limit in the proposition records exactly those data.

18.3 Transport over a homotopy type

In the ∞-category $\mathcal{S}$ of spaces, straightening identifies $\mathcal{S}_{/B}$ with $\mathrm{Fun}(B, \mathcal{S})$, where B is an ∞-groupoid [20, §3.2]. A path $\gamma : b \to c$ induces equivalences $t_X : X_b \simeq X_c$ and $t_Y : Y_b \simeq Y_c$. A map on the fiber at b is transported by $h \mapsto t_Y h t_X^{-1}$. The higher comparisons required for this formula come from the path composition of the original functors. The evaluation square

$$\begin{array}{ccc} [X_b, Y_b] \times X_b & \xrightarrow{\mathrm{ev}} & Y_b \\ \downarrow{\scriptstyle \mathrm{transport}} & & \downarrow{\scriptstyle t_Y} \\ [X_c, Y_c] \times X_c & \xrightarrow{\mathrm{ev}} & Y_c \end{array}$$

commutes with this transport. To check the formula for the exponential, first evaluate it at each b and choose these conjugation maps along paths. A coherent map from P to this family is a collection of maps $P_b \times X_b \to Y_b$ commuting with all transports. Thus the family satisfies the exponential adjunction.

For an ordinary discrete group action this specializes to $(g \cdot h)(x) = g\, h(g^{-1}x)$, and a family $f(p, x)$ obeys $f(gp, gx) = g f(p, x)$. A particular fiber of the family can have less symmetry than the total family. It is equivariant under the stabilizer of p. For smooth actions one may work over the smooth classifying stack of Chapter 17 to retain the dependence on smooth generalized points. As an example, $SO(2)$ acts on $\mathbb{R}^2$ and its line directions. Orthogonal projection onto a line transforms by conjugation under rotation, so the whole family is equivariant with the moving line as parameter. Projection onto one fixed line commutes only with rotations preserving that line.

CHAPTER 19

Orientation and characteristic classes in families

An oriented manifold has locally chosen positive frames whose transitions preserve orientation. For a family of manifolds there can instead be a local orientation on every fiber and a change of sign around a parameter loop. The orientation is then a double cover of the base or of the total space, as appropriate. This cover is itself descent data.

19.1 Determinants and orientation covers

For a real vector bundle E of rank r over M, its determinant $\det E$ is $\bigwedge^r E$. The orientation cover $\mathrm{Or}(E) \to M$ has fiber the two components of nonzero frames of $\det E$. It is a principal C_2-bundle. Local frames of E give transition matrices $g_{ij} : U_i \cap U_j \to GL(r, \mathbb{R})$, with $g_{ij} g_{jk} = g_{ik}$. Their images under $\mathrm{sgn} \circ \det$ are the transitions of the cover. An orientation of E is a section of that cover. This is the same as a choice of frame component on each connected chart such that all transition signs are positive.

The determinant signs form a Čech 1-cocycle with values in $C_2 \simeq \mathbb{Z}/2$. Replacing the local trivializations changes this cocycle by a coboundary. On a good cover of a manifold it represents the first Stiefel–Whitney class $w_1(E)$. In smooth spaces $\mathrm{Or}(E)$ is classified by a map $\underline{M} \to [*/C_2]$ as in Proposition 17.1. For $r \geq 1$, lifting a family of frames from $GL(r, \mathbb{R})$ to the subgroup with positive determinant is equivalent to specifying a smooth trivialization of the sign torsor. The two descriptions give the same obstruction, and both commute with pullback.

For real line bundles the sign class is also complete. Choose a fiberwise positive inner product and take unit local frames. All transition functions now have values in $\{1, -1\}$. A sign cocycle determines a line by gluing, and two cocycles that differ by a coboundary give isomorphic bundles by changing unit frames. Thus orientation and triviality coincide for a real line. This reflects the contractibility of the subgroup of positive real scalars. On S^1, a line with sign monodromy -1 has no nowhere-zero section: a section on the interval parametrizing the circle has endpoint values of opposite sign unless both vanish. The intermediate value theorem then forces a zero. For the trivial bundle the constant unit section avoids this obstruction.

19.2 An exact sequence of relative tangent spaces

An exact sequence of vector bundles

$$0 \longrightarrow K \longrightarrow E \longrightarrow F \longrightarrow 0$$

gives a natural isomorphism $\det E \simeq \det K \otimes \det F$. To construct it, start with a basis of K and append lifts of a basis of F. A change of lifts adds linear combinations of the earlier basis vectors and leaves the top exterior product unchanged. Changes of the two bases scale the wedge by the corresponding determinants. This constructs the isomorphism in local frames, where it is visibly natural under pullback. The orientation cover of E is accordingly the tensor product of the two covers, and $w_1(E) = w_1(K) + w_1(F)$.

Let $f : M \to N$ and $g : Q \to N$ be transverse smooth maps. At a point of their fiber product Z the derivative $df - dg$ is surjective. The preimage normal form of Chapter 5 identifies TZ with its kernel. If $p : Z \to M$, $q : Z \to Q$, and $h = fp = gq$, this yields

$$\det TZ \simeq p^* \det TM \otimes q^* \det TQ \otimes (h^* \det TN)^\vee.$$

The convention uses the ordered sum $TM \oplus TQ$ and the displayed derivative. Choosing oriented frames of the three ambient bundles orients the intersection via this convention. When maps are replaced by parameter-dependent families transverse in the vertical directions, the same construction uses the vertical bundles. In a model for derived smooth intersections the two-term tangent complex continues to record this determinant when the derivative fails to be surjective.

$$\begin{array}{ccc} Z & \xrightarrow{\ q\ } & Q \\ \downarrow{\scriptstyle p} & & \downarrow{\scriptstyle g} \\ M & \xrightarrow{\ f\ } & N \end{array}$$

This determinant relation is an orientation formula for the pullback, rather than a choice of a splitting of its normal sequence.

19.3 An oriented count over a parameter space

Suppose $\pi_M : M \to B$ and $\pi_N : N \to B$ are smooth submersions of boundaryless manifolds of equal relative dimension d, with $f : M \to N$ a smooth map over B. Let $s : B \to N$ be a smooth section, and suppose that for each b, $s(b)$ is a regular value of f_b. Put $Z = M \times_N B$. If $Z \to B$ is proper, it is a finite covering over each connected component of B, allowing the empty covering. In local submersion coordinates the defining equations of Z have an invertible vertical Jacobian and

can be solved as a smooth graph over B. Thus $Z \to B$ is a local diffeomorphism. Properness makes each fiber compact and discrete, hence finite, and prevents additional inverse images from entering a sufficiently small neighborhood. These local graphs form the sheets.

The derivative along each sheet is an isomorphism between the two vertical tangent bundles. If both are oriented globally, this determines a sign that is locally constant on the sheet. Monodromy permutes the sheets and preserves their signs. Consequently the sum of signs on a fiber is constant on each connected component of B. For proper fiber maps of oriented manifolds this agrees with the regular-value formula for degree in Chapter 9, when the conventions for degree there apply. The unsigned count modulo two remains defined if orientations are absent. More generally one can keep the orientation line of each sheet. Locally changing a frame reverses its corresponding sign, and the local counts use the transitions of those lines. This description explains which orientation data must be fixed before adding signs as ordinary integers.

For example, let ξ be the Möbius real line on S^1 and consider the fiber map from ξ to itself given by multiplication by any fixed negative scalar. Each fiber has degree -1 when the same frame is used for both source and target. Neither of the vertical lines is orientable, but the tensor product of their orientation lines is trivial. The relative orientation is supplied by the identification of the two lines. By contrast, identifying ξ with a trivial line by a fiberwise linear isomorphism would contradict its nonzero sign class.

19.4 Euler classes and pullback

For an oriented rank-r bundle ξ, the Thom class $u_\xi \in H^r(E, E \setminus M; \mathbb{Z})$ restricts to the chosen positive generator on each fiber pair. It is natural for pullback bundles [26, 14]. The Euler class is $e(\xi) = 0_\xi^* u_\xi$, where the zero section maps M to the total space. A nowhere-zero section has zero pullback of the relative class. Homotopy through sections to the zero section then gives $e(\xi) = 0$. When a smooth section is transverse to zero, the local intersection computation identifies the Euler class with the cohomology class dual to its zero locus under the usual orientation and compactness hypotheses of Chapter 10. At a regular isolated zero this follows from identifying the local pullback of the Thom class with the local degree of the section in a bundle chart.

One can package these classes using a bundle classifying space and cohomology classes there. The smooth classifying stack retains smooth gauge data, whereas a topological classifying space classifies the bundle’s underlying topological class in the settings of Chapter 7. The resulting characteristic classes commute with pullback and therefore apply to families described by either local frames or their topological classification. The Euler class measures one aspect of the bundle and

can coincide for distinct smooth families. Classifying maps retain the gluing data from which these invariants are extracted.

CHAPTER 20

Discriminants and the topology of families

A smooth family can move through maps with singular fibers. The discriminant marks parameter values where the implicit function theorem fails in the input direction. Outside it, under a properness hypothesis, fibers form an ordinary bundle even though a drawing of two neighboring fibers may make their equivalence hard to see.

20.1 A regular total space with singular fibers

Let P, X, Y be smooth finite-dimensional manifolds without boundary, and let $F : P \times X \to Y$ be smooth. Fix $y \in Y$ and suppose that F is transverse to the point y. Thus $Z = F^{-1}(y)$ is a manifold. Write $\pi : Z \to P$ for the projection. For $(p, x) \in Z$ the tangent equation is

$$D_p F\, \dot{p} + D_x F\, \dot{x} = 0.$$

The derivative of π sends $(\dot{p}, \dot{x})$ to $\dot{p}$. It is surjective if and only if $D_x F$ is onto. In one direction this follows by solving the tangent equation for every $\dot{p}$. In the other direction, surjectivity of $D\pi$ implies that the image of $D_p F$ is contained in the image of $D_x F$. The sum of the two images is all of $T_y Y$ by transversality of F.

Let $C \subset Z$ consist of the points where $D\pi$ is not onto. This is closed by the rank criterion. Define the discriminant of this family at y to be $D = \pi(C)$. The corresponding map on the complementary locus is

$$\begin{array}{ccc} \pi^{-1}(P \setminus D) & \hookrightarrow & Z \\ \downarrow{\scriptstyle \pi} & & \downarrow{\scriptstyle \pi} \\ P \setminus D & \hookrightarrow & P. \end{array}$$

The square is a pullback, and on the left the projection is a submersion. By Sard's theorem the discriminant has measure zero in parameter charts. This conclusion requires the regular total space assumption; it makes no assertion about a family with an identically singular evaluation.

Proposition 20.1 (Bundles away from the discriminant)**.** *If $\pi : Z \to P$ is proper, then D is closed and over each connected component of $P \setminus D$ in the image, the restricted projection is a smooth fiber bundle. On any component outside the image the fibers are all empty.*

Proof. A proper map between manifolds is closed, so it sends C to a closed set. Restriction to the open complement remains proper by base change. On that open set the image is both open (submersions have local sections near image points) and closed (properness). It is therefore a union of connected components.

Here is the local triviality argument for a proper submersion. Choose a Riemannian metric on its domain and define horizontal vectors to be orthogonal to $\ker D\pi$. The derivative identifies this horizontal bundle with the pullback of the tangent bundle of the base. Take a coordinate ball centered at p with compact closure contained in the base of the submersion. Lift each radial path in the ball horizontally starting at an arbitrary point of the fiber over p. The lifted path stays in the inverse image of the closed ball, which is compact. The smooth ordinary differential equation therefore has a solution over the entire radial segment. Its dependence on endpoints and initial points is smooth. The reverse radial path gives the inverse identification with the fiber over p. This gives a smooth trivialization over the ball. □

A different metric may give different transports, and a loop can give nontrivial bundle monodromy. The change of trivialization is part of the family, as in Chapter 18. Properness holds automatically for the projection of this closed preimage when X is compact. With noncompact X, roots can escape to infinity without ever becoming singular. For example $F(t, x) = tx - 1$ has one root when $t \neq 0$ and none at 0, with $D_x F = t$ nonzero at every root. The entire zero locus is regular, but its projection is not proper, since the roots for $t \to 0$ are unbounded.

20.2 Fold and cusp calculations

For the family of maps $F(t, x) = x^2 - t$, the total zero set is a parabola, parametrized by x. The projection to t has one critical point at $x = 0$ and is proper. For $t < 0$ the fiber is empty, and for $t > 0$ its two points are $x = \pm\sqrt{t}$. The derivatives $2x$ give opposite local degree signs and their sum is zero. At $t = 0$ the vertical derivative vanishes but the parameter derivative is -1, so the total space remains smooth. This is the local fold mechanism by which a pair of solutions is created or lost. For a given one-dimensional family, the fold form follows locally if the parameter derivative and the second input derivative are nonzero at a critical root. Indeed one can solve the equation for t as a function $\phi(x)$ by the implicit function theorem. Then $\phi'(0) = 0$ and $\phi''(0) = -F_{xx}/F_t \neq 0$.

The one-dimensional Morse lemma gives the asserted normal form up to local coordinates and a possible reversal of parameter (Chapter 11).

Now take parameters (a, b) and equation

$$F(a, b, x) = x^3 + ax + b = 0.$$

The total locus is smooth since $D_b F = 1$. It projects properly to the plane of coefficients: a root with $|x| > 1$ satisfies $|x|^3 \le |a||x| + |b|$, which bounds its size when (a, b) lies in a compact set. A critical root satisfies $3x^2 + a = 0$. In terms of that root s the discriminant is parametrized by

$$a = -3s^2, \qquad b = 2s^3, \qquad 4a^3 + 27b^2 = 0.$$

For $s \neq 0$ the critical point above (a, b) is a fold of the projection. The root $s = 0$ is a triple root at the cusp. If $a \ge 0$ a polynomial off the discriminant has one real root. If $a < 0$, evaluate it at the local maximum and minimum $x = \mp\sqrt{-a/3}$. Their values are $b + 2(-a/3)^{3/2}$ and $b - 2(-a/3)^{3/2}$. There are three simple real roots when these have opposite signs, exactly when $4a^3 + 27b^2 < 0$, and otherwise one off the discriminant. In the three-root region the ordered derivative signs are $+, -, +$. They have the same sum as the sign of the unique root elsewhere. Topology of the fiber distinguishes these regions although degree does not.

20.3 Transverse jet conditions in a family

The vertical r-jet of a family is the map

$$(p, x) \longmapsto j^r_x F_p$$

from $P \times X$ to $J^r(X, Y)$. A submanifold of the jet bundle can impose a rank condition, a root condition, or a constraint on higher derivatives. If this total jet map is transverse to a smooth submanifold S, then its inverse image is smooth and the parametric transversality theorem of Chapter 6 applies. The derivative test for the projection shows that a regular parameter has $j^r F_p$ transverse to S. If $\operatorname{codim} S > \dim X$, such a parameter has empty pullback of S. For a finite or countable collection of smooth strata, one can discard a countable union of measure-zero exceptional sets, provided total transversality holds to each. This test addresses the actual set of allowed parameters. It can fail for a constrained family even if a different unconstrained family contains all the needed perturbation directions.

Classical bundles describe the complementary locus pointwise and along paths. At a discriminant one can retain the original equation by a derived pullback, as in Chapter 22. When the parameter space carries an appropriate stratification, one can also study transport adapted to the strata (Chapter 23). In the cubic

example the cusp stratum, fold stratum, and open regions record which regular configurations approach a given critical configuration. They preserve information absent from a single fiber count.

CHAPTER 21

Derived smooth spaces

A smooth inverse image of a regular value has a manifold structure because its equations can be placed among the coordinates. The same equations at a critical value can define a singular set, a smooth set with excess dimension, or a point at which infinitesimal motions remain. Derived smooth geometry keeps the equations and their relations in the structure sheaf. Its pullback represents the compatibility of two maps even when their differentials fail to be transverse.

21.1 The local model of a zero set

Use derived manifolds in smooth geometry, with ordinary manifolds included by their smooth function rings. We work with finite local models of the form

$$Z(F) = U \times^{h}_{\mathbb{R}^m} \{0\}, \qquad F : U \longrightarrow \mathbb{R}^m,$$

where U is open in $\mathbb{R}^n$ and the superscript h denotes a homotopy pullback in the category of derived manifolds. The underlying topological space is $F^{-1}(0)$. Its functions form a sheaf of derived C^∞-rings. Spivak constructs derived manifolds by gluing such models [30]. Carchedi and Steffens characterize a finite-limit, idempotent-complete extension of the category of manifolds in which transverse pullbacks are preserved [5]. We use the finite local zero-locus case. More general iterated limits may have cotangent complexes of wider amplitude.

The ordered list of equations gives a concrete model. Write $\mathcal{A}$ for the sheaf of smooth germs from U restricted to the topological zero set, and let $F_1, \dots, F_m$ be the components of F. Adjoin generators η_i of cohomological degree -1:

$$\mathcal{K}_F = \mathcal{A} \otimes_{\mathbb{R}} \bigwedge\nolimits_{\mathbb{R}}(\eta_1, \dots, \eta_m), \qquad d\eta_i = F_i.$$

The differential extends by the signed Leibniz rule. The degree-zero smooth structure is that of $\mathcal{A}$. This is the local Koszul model inside the differential graded realization of derived smooth geometry [4]. One can see the origin of the formula by resolving the point in the coordinate target. In the smooth germs at the origin of $\mathbb{R}^m$, the coordinate functions $t_1, \dots, t_m$ form a regular sequence. Division by a coordinate for functions vanishing on its zero hyperplane follows

by integrating the partial derivative along that coordinate. After quotienting by the first $i-1$ coordinates, the same argument applies on their common zero plane. Multiplication by t_i there is injective, since a smooth function annihilated by t_i vanishes on the dense locus where $t_i \neq 0$. The Koszul resolution for (t_i) thus resolves the origin. In the differential graded model the resolution is formed by adjoining the degree -1 generators with differentials t_i. Pullback transports these differentials to the corresponding components of F. In particular,

$$H^0(\mathcal{K}_F) = \mathcal{A}/(F_1, \ldots, F_m).$$

Negative cohomology describes relations of the equations modulo the boundaries of the Koszul complex.

If F is submersive near a zero, the submersion chart writes it as the projection onto the first m coordinates. The regular sequence argument then gives $H^i(\mathcal{K}_F) = 0$ for $i < 0$ and identifies H^0 with smooth functions on the usual level submanifold. This proves the agreement with the classical construction in a transverse coordinate chart. Another useful check is stabilization. Replacing $F(u)$ by $(F(u), s)$ on $U \times \mathbb{R}^r$ changes the local model by the independent regular-coordinate equations $s_1 = \cdots = s_r = 0$. Successively resolving those equations reduces the model to $\mathcal{K}_F$. Equivalently, the two derived spaces agree by the composition and product laws for pullbacks.

21.2 The tangent complex

For a point x of $Z(F)$ write $T_x Z(F)$ for the tangent complex, dual to the cotangent complex evaluated at x. Its local model uses the cohomological degrees $0, 1$.

Proposition 21.1 (Tangent complex of equations)**.** *At a zero x of F there is a natural identification*

$$T_x Z(F) \simeq [\, T_x U \xrightarrow{D_x F} \mathbb{R}^m \,],$$

where $T_x U$ has degree zero. In particular the cohomology spaces are $\ker D_x F$ *in degree zero and* $\operatorname{coker} D_x F$ *in degree one.*

Proof. In the differential graded chart, the cotangent module is generated over the function algebra by the differentials of the smooth coordinates in degree zero and the differentials of the η_i in degree -1. The smooth chain rule gives $d(\mathrm{d}\eta_i) = \mathrm{d}F_i$. This is the derived cotangent model because the coordinates and the Koszul generators form the smooth free chart with successively adjoined generators, with an open localization on U. Evaluating at x gives the two-term cotangent complex

$$[\, (\mathbb{R}^m)^* \xrightarrow{(D_x F)^*} T_x^* U \,]$$

in degrees $-1, 0$. Dualizing this finite-dimensional complex gives the asserted two terms. Equivalently, the cotangent pushout formula at the point computes the same complex from the pullback of the two maps to $\mathbb{R}^m$ [30, §7]. □

The degree-zero cohomology is the space of compatible tangent vectors. Given a first-order right-hand side b for $F(u) = b$, the degree-one cohomology records its class modulo the changes attainable by moving u. The index

$$\dim H^0(T_x Z(F)) - \dim H^1(T_x Z(F)) = n - m$$

is the *virtual dimension* at x. The equality follows by subtracting the same rank from the two terms. Stabilization adds the acyclic complex $[\mathbb{R}^r \xrightarrow{1} \mathbb{R}^r]$, which leaves the index unchanged.

Example 21.2 (Tangent information and higher functions)**.** Let R be the ring of real smooth germs at 0. For $F(t) = t^2$, the Koszul differential is multiplication by t^2. It is injective, since a smooth germ vanishing off the origin vanishes on a neighborhood. Taylor's formula with smooth remainder gives $R/(t^2) \cong \mathbb{R}[\epsilon]/(\epsilon^2)$: every germ has a unique constant and linear term modulo t^2. Thus this example has a nonreduced ordinary function ring. Replacing $F(t)$ by t^3 gives $\mathbb{R}[\epsilon]/(\epsilon^3)$ with the same tangent complex $[\mathbb{R} \xrightarrow{0} \mathbb{R}]$. If F is identically zero on $\mathbb{R}$, the same complex appears at every point, and the degree -1 Koszul generator now survives in the function algebra. These descriptions distinguish the three models through functions, even though their pointwise tangent complexes agree at the selected zeros.

Example 21.3 (Redundant constraints)**.** The map $t \mapsto (t, 0)$ to $\mathbb{R}^2$ has a single zero. Its Koszul algebra has one regular constraint and one generator of zero differential, so the germ is equivalent to $\bigwedge_{\mathbb{R}}(\eta)$ with $|\eta| = -1$. The tangent differential $v \mapsto (v, 0)$ is injective with one-dimensional cokernel. The virtual dimension is -1, accounting for the redundant target direction. The pair of constant maps $\{0\} \rightrightarrows \mathbb{R}$ also has a derived pullback whose function algebra is this exterior algebra. Its homotopy pullback in spaces is contractible, since the based loop space of the underlying real line is contractible. Here the extra information comes from the smooth structure on the target and is detected by the derived function algebra.

21.3 Compatibility as a mapping property

If $X \to W \leftarrow Y$ are ordinary smooth maps, the derived pullback has the universal property

$$\mathrm{Map}(Q, X \times^h_W Y) \simeq \mathrm{Map}(Q, X) \times^h_{\mathrm{Map}(Q,W)} \mathrm{Map}(Q, Y)$$

for derived test objects Q. A point of the right-hand space consists of the two maps and a comparison between their composites to W. The comparison is part of the universal property at all higher simplicial degrees.

$$\begin{array}{ccc} X \times^{h}_{W} Y & \longrightarrow & Y \\ \downarrow & & \downarrow{\scriptstyle g} \\ X & \xrightarrow{\ f\ } & W \end{array}$$

Around a pair (x, y) with $f(x) = g(y)$, choose a single chart on W and express compatibility by $F(u, v) = \phi(f(u)) - \phi(g(v))$. The diagonal of the coordinate space is its ordinary transverse zero set for the difference map. Pullback associativity therefore identifies this local derived model with $Z(F)$. The tangent complex in Proposition 21.1 becomes

$$[\, T_x X \oplus T_y Y \xrightarrow{Df - Dg} T_{f(x)} W \,].$$

Changing target coordinates identifies the last term by the derivative of the chart transition. More intrinsically, the complex is the homotopy fiber of the difference of the tangent maps, a description that also makes sense after the local equations have been glued.

A further base change can be performed one arrow at a time or in a single pullback. Both comparisons represent the same space of maps from every Q, by the repeated universal property and associativity of limits. This observation is useful for families of constraints: imposing incidence on a family and then restricting its parameter gives the same derived incidence space as restricting the family first. Surjectivity of a derivative is instead a chartwise test that identifies a particular derived intersection with an ordinary one. The next chapter uses this difference to calculate excess intersection classes.

CHAPTER 22

Derived intersections and excess

Let A and B be embedded submanifolds of W, with the manifolds taken without boundary. Their derived intersection represents the pullback of their inclusions. Its underlying points are precisely the points of $A \cap B$. The linearization at c is the two-term complex

$$[\, T_cA \oplus T_cB \xrightarrow{(u,v)\mapsto u-v} T_cW \,],$$

using the result of Chapter 21. If A is cut out locally by a submersion s, the same derived intersection is the zero locus of $s|_B$. The two presentations are equivalent by pasting pullback squares. In particular the normal equation model at c has source T_cB, target $N_{A/W,c}$, and differential induced by the inclusion into T_cW.

22.1 Clean intersections

Definition 22.1. An intersection of embedded submanifolds is *clean* if $C = A \cap B$ is an embedded submanifold and $T_cC = T_cA \cap T_cB$ for all $c \in C$.

This condition allows the tangent sum to be smaller than T_cW. Since its dimension is locally constant on C, the quotient

$$E = TW|_C \,/\, (\, TA|_C + TB|_C \,)$$

is a smooth vector bundle. It is the *excess bundle*. The degree-zero tangent cohomology is TC by identifying a kernel pair with its common vector, and degree-one cohomology is E. The natural map on normal bundles gives the exact sequence

$$0 \longrightarrow N_{C/B} \longrightarrow N_{A/W}\Big|_C \longrightarrow E \longrightarrow 0. \tag{22.1}$$

Indeed, $TB|_C \to TW|_C / TA|_C$ has kernel TC, and its cokernel is the displayed quotient. Thus an excess direction can be read either from the derived tangent or from the normal bundles of the inclusion square.

$$\begin{array}{ccc} C & \hookrightarrow & B \\ \downarrow & & \downarrow{\scriptstyle i_B} \\ A & \hookrightarrow & W \end{array}$$

For a transverse intersection the excess bundle has rank zero. If $A = B$, the inclusion on the left of (22.1) has source zero and $E = N_{A/W}$. Both are clean cases. Even when TC and E exhaust the tangent cohomology, the entire local function algebra of the derived intersection belongs to its geometric structure.

22.2 Local functions in the clean case

A clean zero set has a local simplification that explains the appearance of the excess bundle in the linearization. Let $F : U \to \mathbb{R}^m$ have a smooth zero set C with $\ker D_c F = T_c C$ along C. Near any chosen c, select $r = \operatorname{codim}_U C$ components with independent differentials. Use them as normal coordinates $u_1, \dots, u_r$. Their common zero set locally agrees with C: the latter is an embedded submanifold of it of the same dimension, and hence open in it near c. Shrinking to a product coordinate box, write each remaining function as

$$F_{r+j}(a, u) = \sum_{i=1}^{r} h_{ji}(a, u)\, u_i.$$

The coefficients can be obtained by integrating $\partial_{u_i} F_{r+j}(a, tu)$ over $0 \le t \le 1$. In the Koszul model the invertible change of generators $\eta_{r+j} \mapsto \eta_{r+j} - \sum_i h_{ji}\eta_i$ makes the new last $m - r$ generators closed. The first r generators still resolve the coordinate plane. Taking their cohomology leaves the exterior algebra over the local functions of C on $m - r$ generators of degree -1.

For the intersection of A and B, the degree -1 generators on C span the dual of the cokernel of the normal derivative, namely E^*. For example, a vector in the target dual annihilates the image of the derivative exactly when its linear combination of the differentials of the equations is zero on C. The coordinate calculation then identifies the negative cohomology sheaves locally with exterior powers of E^*. The calculation uses a splitting of the equations on the coordinate box. A splitting on each box does not by itself specify compatible choices on overlaps, so it supplies local models for gluing.

22.3 The Euler factor

Write u_V for the Thom class of an oriented real vector bundle V, considered in the cohomology of its total space relative to the complement of the zero section. Its Euler class is $e(V) = 0_V^* u_V$ under the zero section 0_V. Suppose the submanifolds A and B are closed as subsets of W, and choose orientations of $N_{C/B}$ and $N_{A/W}$. Orient E by the exact sequence (22.1), taking the subbundle first. Let $u_A \in H^r(W, W \setminus A; \mathbb{Z})$ be the class obtained by excision from the oriented normal bundle of A, where $r = \operatorname{rank}(N_{A/W})$. Assume the ranks are constant on the components in question.

Theorem 22.2 (Clean excess intersection)**.** *Under excision to the normal bundle of C in B, the image of $i_B^* u_A$ is*

$$u_{N_{C/B}} \smile p^* e(E), \qquad p : N_{C/B} \to C.$$

Proof. Choose tubular coordinates at A in W and at C in B as in Chapter 8. Shrink the latter tube to fit inside the preimage of the former under i_B. In these coordinates the inclusion is a map to $N_{A/W}$ of the form

$$(c, v) \longmapsto (a(c, v),\, w(c, v)).$$

It sends $v = 0$ to the zero vector over c, and $w(c, v) \neq 0$ for nonzero v in the tube. Choose a radius disk bundle for the source tube, so that (c, tv) remains in the tube for $0 \leq t \leq 1$. For $t > 0$ the maps $(c, v) \mapsto (a(c, tv), t^{-1} w(c, tv))$ take values in the full normal bundle $N_{A/W}$. Taylor's formula in local bundle frames extends these maps smoothly to $t = 0$, where the map is $(c, v) \mapsto (c, J_c v)$ for the inclusion J in (22.1). Each map sends the complement of the zero section into the complement. For $t > 0$ this uses the description of the preimage of A. For $t = 0$ it uses the injectivity of J. Consequently the original pullback of the Thom class and its pullback along J agree as relative classes.

A bundle metric splits the exact sequence so that J is inclusion into $N_{C/B} \oplus E$ as the first factor. The Thom class of a direct sum is the product of the Thom classes pulled back along the projections, with the given orientation order [3]. On the subbundle in question, the second projection is the zero section of E after p. Its pullback contributes $p^* e(E)$, giving the formula. □

The corresponding absolute class on B is the Gysin image of $e(E)$ under $C \hookrightarrow B$. The transverse formula is recovered since the Euler class of the zero bundle is the unit. Over $\mathbb{F}_2$, the same argument uses the canonical modulo-two Thom classes and dispenses with the orientation choices. The exact sequence still distinguishes a failure of transversality from the dimensions of the ordinary intersection.

Example 22.3 (Self-intersection of a diagonal)**.** Let $W = S^2 \times S^2$ and let $A = B$ be the diagonal oriented by S^2. Its normal bundle is identified with TS^2 by sending (v, w) to $w - v$. The orientation of this normal agrees with the coorientation induced by the product: in the ordered basis formed by (v, v) followed by $(0, w)$, the change-of-basis matrix to the two factor bases has determinant one. The excess theorem gives the self-intersection number $\langle e(TS^2), [S^2] \rangle$. For an explicit Euler calculation, project the constant unit vertical vector e_3 onto each tangent plane. The result is the section $x \mapsto e_3 - (e_3 \cdot x)x$. Its two zeros are the poles. The tangent derivatives are $-\mathrm{id}$ at the north pole and id at the south pole. Both preserve orientation in dimension two, giving total index 2. The ordinary set of self-intersection points is the whole diagonal. The normal excess determines its numerical contribution.

22.4 Tangencies and families

Consider the x-axis and the graph of x^2 in $\mathbb{R}^2$. Their intersection consists of the origin, while the intersection of tangent spaces there is a line. The clean condition fails. Pulling the defining equation $y = 0$ back to the graph nevertheless gives the derived zero locus of t^2 in Example 21.2. The function ring has a nonzero nilpotent linear coordinate, and the tangent cohomology has one dimension in each of degrees zero and one. The virtual dimension is zero.

The nearby family of graphs $y = t^2 - \lambda$ makes the variation visible. For $\lambda > 0$ there are two transverse intersections, with opposite local signs in the normal coordinate y, since the derivatives are $2\sqrt{\lambda}$ and $-2\sqrt{\lambda}$. For $\lambda < 0$ there are none. On a small interval in t and for sufficiently small $|\lambda|$, the total zero set meets neither vertical endpoint of the t-interval. The equation on the parameter product is submersive, with derivative -1 in the λ direction. Its zero locus is the parabola $\lambda = t^2$, and restriction to an interval about the critical parameter yields a compact oriented cobordism between the signed zero sets at the two regular endpoints. Hence the oriented number is zero on both sides while the cardinality changes. This is a direct local instance of the conservation of a signed intersection under a regular total-space deformation.

Here the total incidence manifold is smooth, but its map to the parameter has a critical point. The derived fiber at that point is still computed by the single equation after base change. This is a useful separation: transversality of the total incidence equation gives a smooth total space, and regularity of the projection governs which ordinary fibers vary smoothly. Derived base change retains a fiber at every parameter. The next chapter describes how the ordinary nearby topology can be recorded along different parameter strata.

CHAPTER 23

Stratifications and exit paths

A parameter may cross a discriminant at which the regular fibers of an incidence map change. On each regular region one can transport nearby fibers along paths. At the discriminant one also needs comparisons from its strata into the nearby regions. Exit paths assemble both kinds of comparisons into a single category. The construction uses a stratified space as its source and can take values in spaces, retaining homotopies of the propagated data.

23.1 Conical neighborhoods and stratified simplices

Let I be a partially ordered set, with its open upward subsets, and let $X \to I$ be continuous. The fiber X_a over a is the stratum with that label. We order the labels so that a point of a lower stratum may be approached from strata with higher labels. In the conical models used here, a neighborhood of a point of X_a is a product of a neighborhood in that stratum with the open cone on a link. The cone vertex has label a, and the positive radial coordinates keep the labels of the link in $I_{>a}$. For a compact link L this cone is $(L \times [0,\infty))/(L \times \{0\})$. Its sections at fixed positive radius are copies of the link.

The exit-path simplicial set has simplices given by maps of a standard simplex to X that follow a nondecreasing chain of labels $a_0 \le \cdots \le a_n$. In barycentric coordinates $(t_0, \dots, t_n)$ the point with largest nonzero index i must map into X_{a_i}. A one-simplex starts on its initial stratum and can enter a higher stratum continuously. A two-simplex supplies a homotopy compatible with such changes of stratum. Within one stratum this is the usual singular complex, and all its edges are invertible as paths. Between strata the allowed arrows are directional. We write $\mathrm{Exit}(X)$ for the resulting ∞-category under conical hypotheses [19, Appendix A].

The cone on finitely many points is a useful first example. Give the vertex its own stratum and each positive open ray its own stratum. Between different rays there is no exit arrow because a path between them must return to the vertex. From the vertex to a fixed point of a ray there is a single exit arrow up to a contractible space of choices. Radial interpolation compares each such path to the straight one in that ray, preserving the initial vertex and the exit condition. In the stratified simplex model, interpolating radial coordinates preserves positivity on the prescribed positive faces. Thus the same comparison contracts choices

through higher simplices. Each ray itself is contractible. This gives the category with an initial vertex and one object for each ray, and only the arrows from the vertex to those objects apart from identities. This is the finite discrete link case of the cone computation in [1, Lemma 3.3.5].

23.2 Constructible sheaves and their transport

A sheaf of spaces on X is stratified constructible if its restriction to each stratum is locally constant. For clarity, one version of the exit-path theorem assumes that X is paracompact, locally of singular shape, and conically stratified with respect to a poset satisfying the ascending chain condition. Locally of singular shape means that the shape of the sheaf topos on each open set is given by its singular homotopy type. In particular the usual manifold and finite conical examples below satisfy the hypotheses. In this setting the theorem gives an equivalence

$$\mathrm{Shv}_{\mathrm{cstr}}(X;\mathcal{S}) \simeq \mathrm{Fun}\big(\mathrm{Exit}(X),\mathcal{S}\big)$$

between constructible sheaves and functors to spaces [19, Theorem A.9.3]. One direction sends a constructible sheaf to its stalks, with transport along the exit arrows. The other assigns to an open set the space of coherent sections of the corresponding left fibration over its exit category. Ayala, Francis, and Rozenblyum develop a related theory using conically smooth stratifications and describe the descent of the exit construction itself [1]. The stated assumptions are part of the comparison.

For a single stratum, the exit category is the fundamental ∞-groupoid, so the equivalence recovers local systems of spaces. For the cone on finitely many points, it gives a value V at the vertex, a value W_j at each ray, and a map $V \to W_j$. In the ordinary sheaf case this follows directly using a basis of small stars at the origin and intervals in the rays. Restrictions between nested sufficiently small stars induce the same vertex stalk. Restrictions to each ray are the respective specialization maps and restrictions inside the ray are locally constant. Conversely, the indicated diagram defines the restrictions on a basis. On an overlap inside a ray sections agree precisely when their images in that ray agree, which is the ordinary sheaf gluing condition. The topological cone is contractible, but direction and specialization persist in its stratification.

23.3 A real fold and a complex branch

Let $\pi : \mathbb{R} \to \mathbb{R}$ be $u \mapsto u^2$, with the target stratified as the origin and the two complementary rays. For a field k, push the constant sheaf of k-valued locally constant functions forward along π. The resulting sheaf on the parameter line has stalks 0 on the negative ray and k^2 on the positive ray. This follows

by counting components of the preimages of small intervals. A small interval containing the origin has connected preimage, so its stalk there is k. Restricting from that interval to the positive side sends one constant to the same value on both components. Its exit diagram is

$$0 \longleftarrow k \xrightarrow{a \mapsto (a,a)} k^2.$$

There is another constructible sheaf with the same three stalks and zero specialization on the positive side. The star calculation constructs it from that diagram. These two sheaves are inequivalent because the ranks of their specialization maps differ. Even matching the dimension of every nearby stalk leaves the gluing undetermined.

Over the complex line consider $z \mapsto z^2$. On the punctured plane this is a double covering: a choice of one square root continues uniquely along a path and changes sign along one positive loop around the origin. The pushforward has stalk k^2 with the two factors exchanged by this monodromy. At the origin its value is k, since the preimage of a small disk is again a disk. An exit along a radial segment again induces the diagonal map to k^2. Its image is fixed by the exchange, which gives compatibility of specialization with the loop. For example, the sections on the punctured disk have the same constant value on its two nearby sheets after transport around a full turn. The regular stalk dimension alone does not detect this condition.

In the real case the total incidence is the ordinary parabola $\lambda = u^2$ from Chapter 22. The exceptional derived fiber also remembers the relation $u^2 = 0$ in its local functions. The exit diagram of the pushforward above is instead computed from locally constant functions on sets of preimages. Both structures can be associated to the same family, with the exit diagram recording transport of nearby observations and the derived fiber retaining its infinitesimal equations.

23.4 Ordinary transport on a regular region

A condition for obtaining local systems from an incidence space is properness over its regular region. Let $q : E \to V$ be a proper surjective submersion of smooth manifolds without boundary. For any Riemannian metric on E, the orthogonal complement of $\ker Dq$ is a horizontal bundle. Its projection to TV is a vector bundle isomorphism over E. Take a small coordinate ball about $v \in V$ with compact closure in a larger coordinate chart. To transport a point of $q^{-1}(v)$ to another fiber over the ball, lift the straight line from v horizontally. This is an ordinary differential equation on E using the horizontal lift of the coordinate velocity. Its projected path solves the chosen equation on the base. As long as its projected path remains in the closed ball, its lift stays in the compact inverse image of the ball's closure. Compactness supplies existence for the whole time

interval. Lifting the reversed straight line gives the inverse map to the original fiber. Smooth dependence on initial points and velocities identifies the restriction of q to the ball with a product. This is the local trivialization in the proper submersion theorem.

Local trivializations supply transport of the homotopy types and cohomology of the fibers along regular paths. It is locally constant because each trivialization uses the same fiber throughout its ball. The fold restricts on the positive ray to such a locally trivial two-point family. The origin has no ordinary submersion chart and its comparison with that family uses the separate specialization map.

Properness is significant here. The disjoint union $\mathbb{R} \amalg (\mathbb{R} \setminus \{0\})$ maps to $\mathbb{R}$ by the identity on each summand. The map is a surjective submersion, and its fiber cardinalities still change at zero. The missing point in the second summand prevents the horizontal transport argument from extending to that fiber. This example separates local regularity of individual points of the incidence space from control of a whole family of fibers.

23.5 Specialization between local systems

On $S^1 \times [0, \infty)$, stratify by the lower circle and its open complement. In the stratified simplex construction, the label of a point is determined by its second coordinate. Consequently exit simplices in the product are pairs of ordinary singular simplices on S^1 and exit simplices of the half-line. The half-line has the exit category of the one-ray cone, hence is equivalent to the arrow category $[1]$. We obtain

$$\operatorname{Exit}\bigl(S^1 \times [0, \infty)\bigr) \simeq \Pi_\infty(S^1) \times [1].$$

A constructible sheaf of spaces on this cylinder specifies two local systems and a natural transformation from the lower one to the upper one. For vector spaces, writing their loop transport as A and B, a specialization s satisfies $Bs = sA$. Indeed the two ways to move once around the circle and pass to the upper stratum form the corresponding commuting square. For sheaves of spaces the comparison is coherent along all paths and homotopies in the circle. This illustrates why even simple discriminant surfaces can require comparison maps between representations of fundamental groups, rather than separate lists of regular and singular fiber types.

CHAPTER 24

Spaces of immersions and embeddings

The immersion condition depends on the first derivative at one point. The embedding condition also compares pairs of points. This distinction persists in families: rank conditions give a space of formal immersion data, while a theory of embeddings must also track configurations. We begin with the parametric form of immersion theory for manifolds without boundary.

24.1 Formal immersions and regular homotopy

For M^m, N^n, a *formal immersion* is a pair (f, F) consisting of a smooth map $f : M \to N$ and a smooth vector bundle monomorphism $F : TM \to f^*TN$. It is holonomic when $F = Df$. Write $\mathrm{FImm}(M, N)$ for its space with the compact-open smooth topology on the maps and their bundle data, and $\mathrm{Imm}(M, N)$ for the corresponding space of immersions. Equivalently the formal space is the space of sections over M of the bundle whose fiber at x consists of (y, L) with $L : T_xM \hookrightarrow T_yN$. The differential gives the section corresponding to an immersion.

Theorem 24.1 (Smale–Hirsch)**.** *For smooth manifolds without boundary and $n > m$, the map*

$$\mathrm{Imm}(M, N) \longrightarrow \mathrm{FImm}(M, N), \qquad f \longmapsto (f, Df),$$

is a weak homotopy equivalence.

This is the parametric immersion theorem [16, 33]. It identifies the homotopy types of families of immersions and formal monomorphisms in the stated dimension range. In particular it gives a regular homotopy classification on path components. A relative consequence for parameters follows from factoring the differential map through its mapping path fibration. The fibration has weakly contractible fibers. A family of genuine immersions on a subcomplex of a finite CW parameter space, with matching formal data on the full parameter space, lifts to that fibration. The lift extends from the subcomplex by successive disk lifts. Projection back gives genuine immersions on the full parameter space, together with a homotopy of their differentials to the prescribed formal family, fixed on the subcomplex. This conclusion is about continuous families for the topology in the theorem.

Example 24.2 (Plane curves)**.** Trivializations of TS^1 and $T\mathbb{R}^2$ identify a formal derivative of a parametrized plane curve with a map $S^1 \to \mathbb{R}^2 \setminus \{0\}$. The underlying map to $\mathbb{R}^2$ can be contracted independently. Radial normalization and Theorem 24.1 classify path components of the immersion space by degrees of maps $S^1 \to S^1$, or tangent winding numbers. For instance $t \mapsto (\sin t, \sin 2t)$ is an immersion: $\cos t = 0$ forces $\cos 2t = -1$. Its normalized tangent at $t + \pi$ is the reflection of its tangent at t in the vertical axis. Translation of the parameter is orientation preserving, while reflection of the tangent circle reverses orientation. Thus the winding number equals its own negative and is zero. A smooth plane embedding of the circle has winding number 1 or -1 by the turning tangent theorem [29, Lecture 9]. The regular homotopy component of this figure eight therefore contains no embeddings.

24.2 Tests on pairs of points

Put $M^{[2]} = (M \times M) \setminus \Delta_M$. For a map $f : M \to N$, injectivity means that $(x, y) \mapsto (f(x), f(y))$ avoids Δ_N on $M^{[2]}$. If M is compact, an injective immersion is an embedding: it is a homeomorphism onto its image by compactness and the Hausdorff property, and its local immersion charts then give the smooth embedded structure. The deleted product expresses the pairwise test independently of the first-jet condition.

For compact M the embedding locus is open in the C^1 topology. To see this, let f be an embedding. At each point choose coordinates on a small source ball and on the target such that projection of the coordinate image of f to its first m coordinates has invertible derivative at the center. Shrink the source ball until this derivative is uniformly close to its center value. For every sufficiently C^1-close map g, the projected coordinate derivative stays close to the same invertible linear map. Integrating along line segments in the ball then bounds the distance between the projected images of distinct points below by a positive multiple of their coordinate distance. This gives injectivity and full rank on that ball. Finitely many such balls cover M. Pairs sharing a ball constitute a neighborhood of the diagonal in $M \times M$. On the compact remaining set of pairs, the distances between their f-images have a positive minimum. A small uniform perturbation keeps those images distinct as well. This proves openness and explains how local derivative margins and global pair separation give different stability conditions.

Proposition 24.3 (Generic embedding slices of a family)**.** *Let M be compact of dimension m and let $f : P \times M \to N^n$ be a smooth family, with P, M, N finite-dimensional, second-countable manifolds without boundary. Assume $n > 2m$. Suppose the vertical first jet $(p, x) \mapsto j^1(f_p)(x)$ is transverse to each rank-deficient*

stratum in $J^1(M,N)$, and the map

$$(p,x,y)\longmapsto (f_p(x),f_p(y))\colon P\times M^{[2]}\longrightarrow N\times N$$

is transverse to Δ_N. Then f_p is an embedding outside a set of parameters of measure zero in local charts.

Proof. For each total map take the inverse image of the specified stratum or diagonal, a smooth submanifold by transversality. Sard's theorem applied to its projection to P gives a full-measure set of regular parameters. Regularity of this projection is equivalent to transversality in the slice at that parameter. To check the equivalence, write A, B for the parameter and source parts of the derivative into the normal quotient to the specified submanifold. Total transversality means that $\operatorname{im} A + \operatorname{im} B$ is that quotient. Projection of the kernel onto the parameter tangent space is onto exactly when every A-image lies in $\operatorname{im} B$, which under total transversality is precisely surjectivity of B.

There are finitely many rank strata. The stratum of rank $k < m$ in the space of linear maps $\mathbb{R}^m \to \mathbb{R}^n$ has codimension $(m-k)(n-k)$: in a chart where a k-minor is invertible, the remaining block is prescribed by its Schur complement. This also gives the corresponding codimension in jet coordinates. It is at least $n-m+1$, which exceeds m under the stated inequality. Therefore a transverse slice avoids every deficient stratum. The diagonal has codimension $n > 2m$, so its transverse preimage in the $2m$-dimensional deleted product is also empty. The slice is an injective immersion, and compactness gives the embedding. □

One can verify the diagonal hypothesis by showing that the parameter derivatives can move the difference of the two outputs in all normal directions at each collision. This requires independent control at the pair under consideration. A transversality theorem for the unrestricted smooth mapping space gives no such control for a prescribed family of finite dimension. For example the family

$$f_p(t) = \big(p\cos(2t), p\sin(2t)\big), \qquad p \in \mathbb{R}\setminus\{0\},$$

immerses the circle for every p, and it identifies t with $t+\pi$ for every p. Embedding the targets in a higher-dimensional Euclidean space by zero padding leaves this obstruction intact. Adding two output coordinates $a\cos t, a\sin t$ with $a \neq 0$ yields an embedding, since the added coordinates separate points and already form an immersion. In an architecture using trigonometric features, the choice of features therefore changes the class of topological realizations before the weights are selected.

24.3 The configuration stages

Let $\mathcal{O}(U)$ be the poset of open subsets of an open manifold domain $U \subset M$, and let $\mathcal{O}_k(U)$ be its subposet of opens diffeomorphic to unions of at most k disjoint open balls. Embedding calculus studies the restriction functor $V \mapsto \mathrm{Emb}(V, N)$. Its approximations can be expressed by the homotopy right Kan extensions

$$T_k \,\mathrm{Emb}(U, N) = \operatorname*{holim}_{V \in \mathcal{O}_k(U)^{\mathrm{op}}} \mathrm{Emb}(V, N).$$

Restriction from U gives the comparison to each stage. Increasing k adds simultaneous compatibility on configurations spread across more balls. In the range $n > m$, the first stage is weakly equivalent to the immersion space [33]. One way to understand the first stage is that one-ball restrictions of embeddings and immersions have the same local homotopy information, while immersion spaces satisfy homotopy descent for the ordinary open covers used at that stage. Later stages retain compatibility beyond this one-point condition.

$$\begin{array}{ccc} \mathrm{Emb}(U, N) & \longrightarrow & T_2 \,\mathrm{Emb}(U, N) \\ & \searrow & \downarrow \\ & & T_1 \,\mathrm{Emb}(U, N) \xrightarrow{\simeq} \mathrm{Imm}(U, N) \end{array}$$

The convergence theorem in codimension at least three identifies the embedding space for a compact source without boundary with the homotopy inverse limit of its stages [12]. Both the codimension condition and the passage to the limit matter. The rank-and-pair test gives a criterion on an actual smooth map in a suitable dimension range. The tower compares entire spaces of embeddings and keeps their homotopies.

For families, an actual map from a parameter space into $\mathrm{Emb}(M, N)$ induces maps to all of these stages and to the formal immersion space. This expresses the comparison using functors between spaces of models. Applying the same construction to restrictions of the parameter commutes with the maps in the tower. Imposing an architecture supplies a particular subset or family of realizations inside these mapping spaces. The topology of its immersion and embedding loci depends on how that realization map meets the topology detected by differentials and configurations.

CHAPTER 25

Compact group actions and orbit types

The stabilizer of a point determines which symmetries act on its tangent and normal spaces. This linear information also governs a neighborhood of the orbit. We use it to describe what remains when orbits are treated as objects with automorphisms, and to locate the strata on which ordinary smooth quotient charts suffice.

Throughout this chapter, G is a compact Lie group acting smoothly on a Hausdorff second-countable manifold M without boundary. We write $G_x = \{g \in G : gx = x\}$ and $Gx = \{gx : g \in G\}$. An equivariant map $f : M \to N$ satisfies $f(gx) = gf(x)$. Differentiation at x gives

$$D_{gx}f \circ D_x g = D_{f(x)}g \circ D_x f.$$

In particular, $G_x \subseteq G_{f(x)}$, and $D_x f$ is an intertwiner for G_x. This condition persists for smooth families, with the stabilizer taken in the total source when parameters themselves move.

25.1 Invariant metrics and fixed points

Choose a Riemannian metric h on M. With probability Haar measure dg, set

$$\bar{h}_x(v, w) = \int_G h_{gx}(D_x g(v), D_x g(w))\, dg.$$

Compactness permits differentiation in x in local charts. The average of the positive functions obtained by setting $w = v \neq 0$ is positive, and right invariance of Haar measure makes $\bar{h}$ invariant. Every element of G therefore acts by isometries. Its action commutes with the exponential map wherever that map is defined.

Proposition 25.1 (Tangent spaces to fixed loci)**.** *Let K be a closed subgroup of G. The set M^K of points fixed by every element of K is a closed smooth submanifold, allowing different dimensions on different connected components. At $x \in M^K$ its tangent space is $(T_x M)^K$.*

Proof. An exponential chart at x identifies an invariant small ball in T_xM with a neighborhood of x. The chart is K-equivariant because K fixes x and acts isometrically. If a point of its image is fixed, injectivity implies that its tangent coordinate is fixed by K. Conversely an exponential of a fixed vector is fixed. The fixed vectors form a linear subspace, so these charts prove both the local submanifold assertion and the tangent formula. Closedness follows by intersecting the closed fixed sets of individual elements. □

For example, S^1 acts on S^2 by rotation about its vertical axis. The two poles form the fixed locus and their entire tangent planes are rotating representations. The nonpolar points have trivial stabilizer. A tangent bundle on the orbit space alone therefore cannot account for the two polar tangent representations.

25.2 A slice from the normal bundle

Theorem 25.2 (Slice at a compact orbit)**.** *For $x \in M$, put $H = G_x$ and $V_x = (T_xGx)^\perp \subset T_xM$ for an invariant metric. There is a ball $D \subset V_x$ about 0 and an equivariant diffeomorphism*

$$G \times_H D \xrightarrow{\sim} U, \qquad [g, v] \longmapsto g \exp_x(v),$$

where $U \subset M$ is an invariant open neighborhood of the orbit. The balanced product identifies (gh, v) with (g, hv).

Proof. The subgroup H is closed. The orbit map induces an injective map $G/H \to M$. The differential of $G \to M$ at the identity sends a Lie algebra element ξ to the velocity of $t \mapsto \exp(t\xi)x$. A zero velocity means that the resulting fundamental vector field vanishes at x. Uniqueness for the flow makes the entire one-parameter subgroup fix x, so the kernel is the Lie algebra of H. Translating establishes that $G/H \to M$ is an immersion, and compactness of G/H makes it an embedding.

The normal spaces of the orbit at gx are the translates of V_x by D_xg. Thus its normal bundle is $G \times_H V_x$, via $[g, v] \mapsto D_xg(v)$. The tubular-neighborhood construction of Chapter 8 applies to the normal exponential. Compactness of Gx allows a single small radius, and the metric makes the resulting diffeomorphism equivariant. Its fiber over x is the ball D. □

This proof is the compact-group case of the slice theorem. Proper Lie group actions admit a more general version [23]. In the slice above, the stabilizer of $[g, v]$ is gH_vg^{-1} with $H_v = \{h \in H : hv = v\}$. Every nearby stabilizer is therefore conjugate to a subgroup of H. Moreover, the points of the same orbit type as x, meaning that their stabilizer is conjugate to H, correspond locally to $G \times_H D^H$. Indeed, if $H_v \subseteq H$ is conjugate in G to H, it has the same dimension and number

of components as H. The inclusion then identifies their identity components and reaches each component of H, hence $H_v = H$. The local orbit-type stratum is $(G/H) \times D^H$, with orbit quotient D^H. The entire neighborhood has orbit quotient D/H. In particular a free action has charts with $H = 1$, giving a smooth quotient and a principal bundle over it.

There is also a reduction of the symmetry group along any fixed orbit type. Fix a closed subgroup H which occurs as a stabilizer and put $M_H = \{x : G_x = H\}$. In M^H, this set is open. To see this near $x \in M_H$, a nearby stabilizer is conjugate to a subgroup of H by the slice theorem. If it contains H, equality follows by comparing dimensions and component counts as above. The normalizer $N_G(H) = \{g : gHg^{-1} = H\}$ preserves M_H, and $N_G(H)/H$ acts freely on it. The G-orbit of any point of type (H) meets M_H, and two points of M_H in a common G-orbit differ by an element of $N_G(H)$. It follows that $M_{(H)}/G$ identifies with $M_H/(N_G(H)/H)$. In the slice at a point of M_H, the exact-stabilizer locus is $N_G(H) \times_H D^H$: a point $[g, v]$ has stabilizer H precisely when $v \in D^H$ and $g \in N_G(H)$. Thus both quotient charts identify with D^H, proving that the identification is smooth for their quotient manifold structures. This describes the global quotient within one type using an action with free orbits.

25.3 The quotient stack in slice coordinates

Write $[M/G]$ for the smooth quotient stack. Its objects over a test manifold T are principal right G-bundles $P \to T$ with a smooth map $F : P \to M$ satisfying $F(pg) = g^{-1}F(p)$. Morphisms are bundle isomorphisms commuting with these maps. Restricting to the neighborhood of Theorem 25.2 gives

$$\begin{array}{ccc} D & \hookrightarrow & G \times_H D \\ \downarrow & & \downarrow \\ [D/H] & \xrightarrow{\simeq} & [U/G]. \end{array}$$

To prove the lower equivalence, use the equivariant projection $U \to G/H$. The inverse image of the coset H under $P \xrightarrow{F} U \to G/H$ is a principal right H-bundle $Q \subset P$. Locally, a section of P and a smooth lift of the resulting map into G/H produce a section of Q. The restriction $F|_Q$ lands in the fiber D. In the other direction a bundle Q with equivariant $u : Q \to D$ extends to $Q \times_H G$ with map $[q, g] \mapsto g^{-1}[1, u(q)]$. The two constructions recover each other and their bundle isomorphisms. In particular they commute with pullback in T.

The infinitesimal action at x is the linear map

$$\mathfrak{g} \longrightarrow T_xM, \qquad \xi \longmapsto \left.\frac{d}{dt}\right|_{t=0} \exp(t\xi)x.$$

Its kernel is $\mathfrak{h} = \mathrm{Lie}(H)$ and its cokernel is the normal representation V_x. They are the cohomology of the tangent complex of the action groupoid in degrees -1 and 0, respectively. The reason for the complex is that an infinitesimal group element is an arrow from a displacement v to $v + \xi_M(x)$. Quotienting the displacement gives the cokernel, and arrows from a displacement to itself give the kernel. A finite stabilizer has zero Lie algebra, so the tangent complex can miss finite automorphisms even when the slice stack retains them.

Example 25.3 (An exceptional orbit with a disk quotient)**.** Let S^1 act on $S^3 \subset \mathbb{C}^2$ by $\lambda(z_1, z_2) = (\lambda z_1, \lambda^2 z_2)$. A point with $z_1 \neq 0$ has trivial stabilizer. At $x = (0, 1)$ the stabilizer is $\{1, -1\}$ and the orbit tangent is the imaginary z_2-axis. Its orthogonal slice is the complex z_1-line, with -1 acting by multiplication by -1. The coarse chart $D/\{1, -1\}$ is topologically a disk, by the map $z \mapsto z^2$, whereas $[D/\{1, -1\}]$ assigns the order-two automorphism group to the central point. Both the Lie algebra of the stabilizer and the local coarse topology would leave that finite group undetected.

The slice converts a question about changing orbit types into one about linear representations. The next chapter uses this reduction for a local obstruction: the normal derivative of an equivariant map has to be surjective within an appropriate category of representations.

CHAPTER 26

Equivariant transversality and fixed-locus constraints

The ordinary transversality argument perturbs a map in directions spanning the target normal space. Equivariance restricts those directions to intertwiners of stabilizer representations. At a fixed point even a tiny invariant target may require a normal representation absent from the source. We identify the obstruction at a single orbit, then show how the ordinary density argument returns for free source actions.

26.1 The normal map at a stabilizer

Let G be a compact Lie group acting smoothly on manifolds M, N without boundary. Let $A \subset N$ be a G-invariant embedded submanifold and $f : M \to N$ an equivariant smooth map. Fix $x \in M$ with $y = f(x) \in A$, set $H = G_x$, and let $N_y A = T_y N / T_y A$ denote the fiber of the normal bundle. The classical transversality condition at x is surjectivity of $T_x M \xrightarrow{D_x f} T_y N \to N_y A$. This map is H-equivariant. Its restriction to $T_x Gx$ vanishes, since $f(Gx) \subseteq A$. It therefore descends to an H-map

$$\overline{D_x f} : T_x M / T_x Gx \longrightarrow N_y A.$$

Proposition 26.1 (Transverse germs at an orbit)**.** *Let $x \in M$, $y \in A$, and assume $H = G_x$ fixes y. There exists an equivariant map on an invariant neighborhood of Gx taking x to y and transverse to A along Gx if and only if $T_x M / T_x Gx$ admits a surjective H-linear map to $N_y A$.*

Proof. Any such map gives the stated normal derivative. Conversely choose invariant metrics and identify the two quotient spaces with orthogonal complements. Let $L : V_x \to (T_y A)^{\perp}$ be a surjective H-linear map. In the slice chart of Theorem 25.2, use

$$[g, v] \longmapsto g \exp_y(Lv).$$

For a small enough ball this exponential is defined. Because y is fixed by H, the chart at y commutes with the H-action, making the formula well defined. Its normal derivative on the slice at $v = 0$ is L, and translation carries surjectivity to each point of the orbit. □

There is a local perturbation version of the same criterion. Fix an equivariant f with $f(x) = y$. On a smaller slice ball the expression $u(v) = \exp_y^{-1}(f(\exp_x v))$ is defined and is H-equivariant. Choose an invariant bump χ on this ball, equal to one near zero. Replacing $u(v)$ by $u(v) + t\chi(v)Lv$, where L is a surjection as in the proof, gives a smooth perturbation through y with support near Gx. Extend it by equivariance on the tube and by the old f off the support. Its normal derivative at x is $\overline{D_x f} + tL$. A full-rank minor of L provides a polynomial minor of this derivative with nonzero leading coefficient. Hence for arbitrarily small t the new germ is transverse along Gx. When the normal fiber is zero the original germ is already transverse. The construction concerns a single orbit and leaves global density as a separate question.

The fixed-point loci M^H, N^H, A^H are submanifolds by Chapter 25. Invariant metrics split the normal sequence, yielding

$$T_y N^H / T_y A^H \;\cong\; (N_y A)^H.$$

If f is transverse at x, its normal derivative is surjective on fixed vectors. In fact, averaging any preimage of an invariant vector gives an invariant preimage. Thus the restricted map $f|_{M^H} : M^H \to N^H$ is transverse to A^H at x. Testing only the restriction for the full stabilizer H tests only the trivial part of the normal representation.

Example 26.2 (A persistent singular zero)**.** Let $M = \mathbb{C}$ have the S^1-action λz, and $N = \mathbb{C}$ the action $\lambda^2 w$. Put $A = \{0\}$. Every equivariant map sends 0 to 0, since the only point of N fixed by the full group is 0. The element -1 acts by -1 on $T_0 M$ and by $+1$ on $T_0 N$. Hence $D_0 f(-v) = D_0 f(v)$ and real linearity gives $D_0 f = 0$. No equivariant perturbation achieves ordinary transversality at the origin. The fixed-locus restriction is a map from a point to a point and is transverse there. For the concrete map $z \mapsto z^2$, the zero set is a single point, even though the linear normal test fails in both real directions.

Example 26.3 (Moving a parameter changes the constraint)**.** Let $P = \mathbb{R}$ and $N = \mathbb{R}$ with the sign action of C_2 on each. The evaluation $P \to N,\ p \mapsto p$ at a one-point input is an equivariant submersion. At the fixed parameter 0 the constant evaluation at the sole input hits $\{0\}$ and fails transversality. For $p \neq 0$, the evaluation misses $\{0\}$. Holding $p \neq 0$ fixed leaves no sign symmetry of the parameter. That is why ordinary parameterized transversality can hold for the family while the restriction to fixed parameters has no transverse member.

26.2 Recovery for free source actions

Theorem 26.4 (Equivariant density for a free compact source)**.** *Let M be compact without boundary with a free smooth G-action, and let $A \subset N$ be invariant, closed,*

and embedded. In the smooth equivariant mapping space from M to N, the maps transverse to A are dense and open in the C^∞ topology.

Proof. Write $B = M/G$. The slice theorem supplies the smooth principal bundle $M \to B$. Equivariant maps $M \to N$ correspond to sections of the associated bundle $E = (M \times N)/G \to B$, where G acts diagonally. In a principal bundle chart an equivariant map is $(g, b) \mapsto g\, s(b)$, with $s(b)$ the fiber coordinate of the section. Its normal derivative to A vanishes in the group directions and is the normal derivative of $s(b)$ in the base directions, up to translation by g. Thus transversality of the map to A is equivalent to transversality of its section to $E_A = (M \times A)/G \subset E$.

Fix a section s. Its image is compact. Cover the image by finitely many local bundle and fiber coordinate charts, with smaller charts still covering it and bump functions supported in the larger charts. Multiplying fiber coordinate vector fields by these bumps gives globally defined compactly supported vertical vector fields $V_1, \dots, V_\ell$ spanning the vertical space at every point of $s(B)$. Their flows preserve the projection and can be composed to give sections s_t for $t \in \mathbb{R}^\ell$. The evaluation $(b, t) \mapsto s_t(b)$ has at $t = 0$ a differential whose parameter directions span the vertical tangent, while its base directions project onto T_bB. It is a submersion on $B \times U$ for a sufficiently small open neighborhood U of zero, by continuity and compactness. Parametric transversality from Chapter 6 supplies transverse sections s_t for parameters arbitrarily close to zero. The flow construction converges to s in the smooth topology as $t \to 0$. Transversality is open in the C^1 topology on the compact source because A is closed. Any sequence of failed normal-surjectivity tests would have base points with a convergent subsequence and would give a failed test for the limiting map. These statements transfer to equivariant maps by the local bundle charts. □

The proof uses perturbations of sections of a bundle on a manifold. When stabilizers occur, admissible perturbations have to respect their linear actions and the same submersion argument may become impossible. For related equivariant general position questions see [31]. The classical parametric argument used in the proof is developed in [15, Ch. 3].

26.3 Pullbacks of quotient stacks

For an everywhere transverse equivariant map, $Z = f^{-1}(A)$ is an invariant smooth submanifold. Its quotient stack fits into

$$\begin{array}{ccc} [Z/G] & \longrightarrow & [M/G] \\ \downarrow & & \downarrow{\scriptstyle [f]} \\ [A/G] & \longrightarrow & [N/G]. \end{array}$$

This is a homotopy pullback of smooth stacks. Over a test manifold an object of the pullback consists of two principal bundles with equivariant maps to M and A, together with an isomorphism between the images in $[N/G]$. Use the isomorphism to identify the bundles. Their two maps then have the same composite to N. The common bundle map lands in the ordinary transverse fiber product $M \times_N A$, identified with Z, and the construction also matches bundle isomorphisms. Taking quotient stacks hence retains the stabilizer at each point of the transverse solution set. The derived pullbacks of Chapter 22 give another question to ask when the normal map fails to be surjective: which tangent and obstruction directions belong to the intersection with its equations retained?

CHAPTER 27

Factorization through an intermediate manifold

A smooth map $f : X \to Y$ may be specified by two maps through an intermediate manifold L. The factorization retains how the map is constructed. Rank bounds give local restrictions, and homotopy invariants can rule out a factorization even when every local rank test passes. We will also distinguish a pointwise choice of factorization from a continuous family of choices.

27.1 Differential constraints

Suppose $f = d \circ e$ with $e : X \to L$ and $d : L \to Y$. The chain rule says that $D_x f$ factors as

$$T_x X \xrightarrow{D_x e} T_{e(x)} L \xrightarrow{D_{e(x)} d} T_{f(x)} Y.$$

Consequently $\operatorname{rank} D_x f \leq \dim L$. The constant-rank normal form gives a local converse at points of an open region where the rank is some fixed r: in suitable coordinates the map is the projection to r coordinates followed by the inclusion of those coordinates in the target chart. This describes a local factorization through an open subset of $\mathbb{R}^r$. The coordinate projections on two such neighborhoods need not provide one global intermediary.

Proposition 27.1 (An exact reconstruction embeds the source)**.** *If X is compact and $d \circ e = \mathrm{id}_X$ with smooth $e : X \to L$ and $d : L \to X$, then e is a smooth embedding. If X and L have the same dimension, both have no boundary, and L is connected, then e is a diffeomorphism for nonempty X.*

Proof. The left-inverse identity makes e injective, and differentiating it makes De injective at every point. Compactness of X and the Hausdorff property of L give a homeomorphism onto the image. Together with immersion, this is the embedding condition. In equal dimensions the inverse function theorem makes $e(X)$ open. It is also compact, hence closed, and connectedness of L implies $e(X) = L$. The local smooth inverses then give the global inverse. □

The derivative already imposes a representation constraint when a group acts on the spaces and both maps are equivariant. At a fixed input x, the identity factors through $T_{e(x)} L$ by two intertwiners. Thus $T_x X$ must occur as a split

subrepresentation for the stabilizer of x. This may exclude an intermediate manifold of sufficient ordinary dimension.

Conversely, every smooth embedding of compact X in L admits an exact reconstruction defined on a neighborhood of the encoded image. By the tubular-neighborhood theorem, a small neighborhood U of $e(X)$ is identified with a neighborhood of the zero section in its normal bundle. Projection from the bundle to $e(X)$, followed by e^{-1}, gives a smooth decoder $U \to X$. The existence of a decoder on all of L asks for an extension of a left inverse. A local normal form alone only supplies the neighborhood decoder, even when the embedded image is closed. For example, the usual inclusion of S^n in $\mathbb{R}^{n+1}$ has the radial decoder on an annulus, yet a global sphere-valued decoder would induce an impossible retraction on top cohomology.

27.2 Obstructions preserved by homotopy

Suppose now that $d \circ e$ is homotopic to id_X. For any commutative coefficient ring R, functoriality in singular cohomology gives unital graded ring maps

$$H^*(X;R) \xrightarrow{d^*} H^*(L;R) \xrightarrow{e^*} H^*(X;R), \qquad e^*d^* = \mathrm{id}.$$

A nonzero cup product in $H^*(X;R)$ must remain nonzero after applying d^*. Rank tests on the groups H^k alone can miss this restriction. Similarly, on homotopy groups a path supplied by the homotopy allows the base point of $d \circ e$ to be compared with the starting point in X. With that comparison in place, the map on each homotopy group induced by e has a left inverse induced by d. See [13] for the cohomology and base-point conventions.

For a closed connected oriented n-manifold X, with $n > 0$, a degree-one reconstruction of X induces a nonzero map on $H^n(X;\mathbb{Z})$. It cannot factor through a manifold of dimension less than n, since that manifold has no cohomology in degree n. Here one can use a triangulation of the intermediate smooth manifold and its dimension bound. This is a homotopy obstruction, so it also applies when the reconstruction has critical points. A factorization through a contractible L makes the composite nullhomotopic by applying d to a contraction after e. Positive dimension allows degree to distinguish this case from the identity.

Example 27.2 (The cup product on a torus)**.** Let $X = T^2$ and take L to be the connected sum of two copies of $S^1 \times S^2$. The first cohomology groups of X and L both have rank two over $\mathbb{Z}$. The product of the two coordinate classes in $H^1(T^2;\mathbb{Z})$ generates $H^2(T^2;\mathbb{Z})$. On L, all products of degree-one classes are zero. Indeed, collapsing the separating sphere of the connected sum gives a map into the wedge of the two closed summands, inducing an isomorphism in H^1 by the Mayer–Vietoris sequence. In the wedge, classes on different summands have

zero product. Each summand has one degree-one generator pulled back from S^1, so its square vanishes too. Thus the ring retraction displayed above cannot exist, even though comparing H^1 alone would permit it.

27.3 A sharp uniform obstruction on the sphere

Proposition 27.3 (Reconstruction with a Euclidean intermediary)**.** *Let $n > 0$, and view S^n as the unit sphere in $\mathbb{R}^{n+1}$. The following statements use the Euclidean norm on that ambient space.*

1. *For every m, if $e : S^n \to \mathbb{R}^m$ and $d : \mathbb{R}^m \to S^n$ are continuous, then $\sup_{x \in S^n} \|d(e(x)) - x\| = 2$.*
2. *If d is allowed to land in $\mathbb{R}^{n+1}$ and $m \leq n$, the minimum possible uniform reconstruction error is 1. For $m \geq n + 1$, there is an exact ambient reconstruction.*

Proof. In the first case $d \circ e$ has degree zero by contractibility. If it had no antipodal coincidence $d(e(x)) = -x$, then the straight segment between x and $d(e(x))$, normalized to the unit sphere, would be a homotopy to the identity. Such a segment passes through the origin only for antipodal endpoints. The degrees therefore force an antipodal coincidence and the claimed error.

For the second part, the Borsuk–Ulam theorem applied to e, after including $\mathbb{R}^m$ into $\mathbb{R}^n$, gives a pair $x, -x$ with the same encoded value [22]. Both reconstructions are a common vector u. The triangle inequality gives $2 \leq \|u - x\| + \|u + x\|$, so at least one error is at least 1. The constant reconstruction 0 attains that bound. For $m \geq n + 1$, include the sphere in the first $n + 1$ coordinates and decode by their projection. □

The distinction between the codomains matters. If a reconstruction lands close to the sphere, one can normalize along its realized image. This does not imply that the decoder can be normalized on the entire intermediate $\mathbb{R}^m$, where it may take value zero.

27.4 Spaces of factorizations

In the category of spaces the factorizations of f through L, together with their chosen homotopies, form the homotopy pullback

$$\begin{array}{ccc} \mathrm{Fact}_L(f) & \longrightarrow & \mathrm{Map}(X, L) \times \mathrm{Map}(L, Y) \\ \downarrow & & \downarrow{\scriptstyle \circ} \\ \{f\} & \longrightarrow & \mathrm{Map}(X, Y). \end{array}$$

A point is a triple (e, d, γ), with γ a path from $d \circ e$ to f. Paths between such points track changes of both maps and homotopies. This is a homotopy problem, while an exact smooth factorization requires equality $d \circ e = f$ and smooth maps. For example, the smooth map $e^{it} \mapsto \exp(i \sin t)$ on S^1 is nullhomotopic and therefore factors through a point in the homotopy sense. It is nonconstant, so it does not admit an exact factorization through a point.

Families can have an obstruction even when every individual target value has a lift. With decoder $d(z) = z^2$ on $L = S^1$, every single $b \in S^1$ has a preimage. The family $f(b) = b$ has no continuous choice of encoded value $e(b)$ satisfying $d(e(b)) = b$. Taking degrees would yield $2 \deg(e) = 1$, which also rules out a choice accompanied by a homotopy to f. Local choices are related on chart overlaps by the sign action. The obstruction is in the family of choices, invisible to the existence of each separate factorization.

CHAPTER 28

Persistence of sublevel families

Morse theory studies a smooth function by following its sublevel sets across regular intervals and critical values. Keeping the sublevel sets together with their inclusion maps gives a diagram of spaces. Applying homology turns this diagram into a representation of the ordered set of thresholds. This chapter gives a direct stability estimate on the diagram and compares it with the stronger local triviality available for smooth parameter families at regular levels.

28.1 The filtered diagram

For a continuous function $f : X \to \mathbb{R}$, let

$$X_f(a) = \{x \in X : f(x) \leq a\}.$$

For $a \leq b$ there is an inclusion $X_f(a) \hookrightarrow X_f(b)$. These data give a functor from the real line with its order arrows to topological spaces, and hence a diagram in the category of homotopy types. In particular, using singular chains before passing to homology keeps a chain-level diagram. For any coefficient field k, the homology groups $a \mapsto H_j(X_f(a); k)$ form a persistence module. Its structure map is the homology map induced by inclusion.

Given $\varepsilon \geq 0$, translation takes a diagram F to $a \mapsto F(a + \varepsilon)$. A strict ε-interleaving of diagrams of spaces consists of natural maps

$$\alpha_a : F(a) \longrightarrow E(a + \varepsilon), \qquad \beta_a : E(a) \longrightarrow F(a + \varepsilon).$$

The two shifted composites must equal the respective structure maps for a step of length 2ε. Passing from actual spaces to homotopy types sends these identities to coherent comparisons. A strict interleaving is enough for the applications below.

Proposition 28.1 (Uniform control of filtered spaces)**.** *If $f, g : X \to \mathbb{R}$ are continuous and $\sup_{x \in X} |f(x) - g(x)| \leq \varepsilon$, their sublevel diagrams have a strict ε-interleaving. It is natural under pullback of functions along continuous maps. When f, g are invariant under a common group action, the maps of the interleaving are equivariant.*

Proof. The inequalities $g(x) \le f(x) + \varepsilon$ and $f(x) \le g(x) + \varepsilon$ give inclusions in the indicated directions. All arrows are the restrictions of the identity map on X. In particular the composition

$$X_f(a) \hookrightarrow^{\alpha_a} X_g(a+\varepsilon) \hookrightarrow^{\beta_{a+\varepsilon}} X_f(a+2\varepsilon)$$

is the structure inclusion, and the reverse composite is the analogous inclusion on g-sublevels. For a continuous map $u : Y \to X$, use u on sublevel sets of the pulled-back functions. It commutes with the interleaving inclusions. Invariance makes each sublevel set a G-space and makes the inclusions equivariant. □

This estimate survives functorial operations on the diagram. For instance one can take fixed spaces under a subgroup H, followed by homology. The fixed sublevel $X_f(a)^H$ is the sublevel of $f|_{X^H}$ inside X^H. Alternatively, one can apply homology to $EG \times_G X_f(a)$, the homotopy quotient at threshold a. The resulting invariants answer different questions about symmetry, and the same uniform bound at the diagram level applies to both. A bottleneck bound for their homology still requires a tameness hypothesis for the resulting modules.

28.2 When persistence diagrams are available

A persistence module is called q-tame if every structure map over a strictly positive step has finite rank. This does not require finite-dimensional homology exactly at each threshold. When X is the realization of a finite simplicial complex, the homology persistence modules of every continuous real-valued f on X are q-tame. In fact, fix $a < b$ and subdivide so that the oscillation of f on every simplex is less than $b - a$, using compactness and uniform continuity. The union of simplices that meet $X_f(a)$, together with their faces, is a finite subcomplex K with

$$X_f(a) \subseteq |K| \subseteq X_f(b).$$

The corresponding homology map factors through a finite-dimensional vector space $H_j(|K|; k)$. This also applies to compact smooth manifolds by smooth triangulability.

The stability theorem for persistence diagrams bounds their bottleneck distance by the interleaving distance for q-tame modules [6]. Consequently the bound in Proposition 28.1 gives a bottleneck bound ε in the finite-complex setting. One way to read the estimate is that a feature with lifetime exceeding 2ε cannot simply be matched to the diagonal under a bottleneck matching of cost at most ε. See [7] for the classical stability result for sublevel-set persistence. The inclusions proved here are assertions at the level of spaces, independent of the tameness used to assign persistence diagrams. Applying homology forgets information,

such as cup products, that can distinguish otherwise identical homology groups or homology persistence modules.

Example 28.2 (Equal homology modules with different filtered spaces)**.** Equip $S^2 \vee S^4$ and $\mathbb{CP}^2$ with the constant function zero. In both cases the sublevel is empty before zero, and the entire space thereafter. The two spaces have one homology generator in each of degrees $0, 2, 4$, with field coefficients, so the ordinary homology persistence modules agree in every degree. The degree-two generator on $\mathbb{CP}^2$ has a nonzero square, whereas the degree-two class on the wedge is pulled back along the projection to S^2 and has square zero [13]. An equivalence of their filtered homotopy types would give a homotopy equivalence at every nonnegative threshold and would preserve these products. Thus the diagrams of spaces can distinguish filtrations with identical homology persistence.

28.3 Transport at finitely many regular thresholds

Let X be a compact smooth manifold without boundary, P a smooth parameter manifold, and $F : P \times X \to \mathbb{R}$ smooth. Write $f_p(x) = F(p, x)$. The map $p \mapsto f_p$ is continuous for the uniform norm by compactness of X. More quantitatively, along a smooth path c of parameters of length $\ell(c)$, a bound L on the norm of the partial derivative D_pF over its image times X gives $\|f_{c(1)} - f_{c(0)}\|_\infty \leq L\ell(c)$. This follows by differentiating $F(c(t), x)$, integrating, and taking the supremum over x. It gives an interleaving estimate even along paths where the topology at a chosen threshold changes.

Proposition 28.3 (Local trivialization of regular sublevel flags)**.** *If $a_1 < \cdots < a_r$ are regular values of f_{p_0}, there is a neighborhood of p_0 on which the family of finite flags $X_{f_p}(a_1) \subseteq \cdots \subseteq X_{f_p}(a_r)$ is smoothly trivialized by diffeomorphisms of X. An empty level set counts as regular.*

Proof. Compactness allows the parameter neighborhood and small disjoint strips around the levels a_i to be chosen so that $\nabla_x F$ is nonzero throughout each strip. Otherwise there would be sequences converging to a critical point of f_{p_0} at one of those levels. Choose coordinates on a small ball in P around p_0 and a metric on X. To lift a constant coordinate vector v on the ball, start with $(v, 0)$ on $P \times X$. Near each level modify its X-component to

$$-\frac{D_pF(v)}{\|\nabla_x F\|^2}\nabla_x F.$$

Use a cutoff in the strip equal to one at the level to extend this correction by zero. The resulting vector is tangent to each hypersurface $F = a_i$. Integrating these lifts along radial coordinate paths from p_0 to p gives diffeomorphisms from the starting fiber to the ending fiber. Compactness of X gives existence along every

such path, and reversing the path gives the inverse. The flow preserves each level and its chosen sublevel side, since a trajectory cannot cross a hypersurface to which the field is tangent. Smooth dependence on the path endpoint gives the required local trivialization. □

Example 28.4 (A regular family with monodromy)**.** Take $P = X = S^1$ and $F(\phi, \theta) = \cos(2\theta - \phi)$, with both angles defined modulo 2π. Each level a strictly between -1 and 1 is regular for every f_ϕ, and the sublevel is a pair of closed arcs. In a local lift of the angle ϕ, translation $\theta \mapsto \theta + \phi/2$ identifies the sublevels with those at $\phi = 0$. A loop in ϕ changes this translation by π, interchanging the arcs. Thus the local system of connected components has nontrivial monodromy, although every sublevel in the family has the same homotopy type. Degree-zero homology records the same permutation on the two component generators.

Regularity controls a finite chosen flag. A moving critical value can cross a threshold and change its sublevel topology even under uniformly small perturbations. The interleaving estimate then compares shifted thresholds. A stratified description of the parameter and threshold space requires control of the critical locus and the hypotheses for stratified transport, as in Chapter 23. A smooth parameter space by itself does not supply such a stratification. The local trivialization and the quantitative interleaving therefore give complementary information about a parametrized filtration.

APPENDIX A

Notation

TM, T^*M, Df
: The tangent bundle, cotangent bundle, and derivative of a smooth map. See Chapters 2 and 11.

$\mathrm{Emb}(M, N)$, $\mathrm{Imm}(M, N)$
: Spaces of smooth embeddings and immersions. The topology is specified when using continuity or homotopy of families. See Chapters 3, 15, and 24.

$J^r(M, N)$, $j^r_x f$
: The bundle of order-r jets of maps and the jet of f at x. The map $j^r f$ is its holonomic section over M. See Chapter 4.

$f \pitchfork S$, $M \times_N Q$
: Transversality to a submanifold and the fiber product. The ordinary smooth fiber product represents the transverse pullback. A superscript $\mathbf{R}$ or an explicit label distinguishes a derived pullback. See Chapters 6 and 22.

$\det E$, $\mathrm{Or}(E)$, $w_1(E)$
: The determinant line, orientation cover, and first Stiefel–Whitney class of a real vector bundle. See Chapter 19.

$\deg f$, $e(E)$, $\chi(M)$
: Degree, Euler class, and Euler characteristic. An integral degree requires the appropriate orientation data. See Chapters 9 and 10.

$\mathrm{Th}(E)$, $MT\theta$
: Thom space of a bundle and the Madsen–Tillmann spectrum for a tangential structure. See Chapters 13 and 14.

$\underline{M}$, $[M/G]$, $\mathcal{E}_{/B}$
: A manifold represented as a smooth sheaf, a smooth quotient stack, and a slice over a base. See Chapters 17 and 18.

$[X, Y]_B$, $\mathrm{Map}_B(X, Y)$
: The internal function object and the mapping space in the slice. The former represents all parametrized maps over B. See Chapter 18.

$M_{\leq a}$, $H_k(M_{\leq a})$
: A sublevel set and its degree-k homology, with coefficients specified in context. See Chapters 11 and 28.

APPENDIX B

Subject index

References point to the chapters where each term is developed.

Approximation 3.
Atlas Smooth, 1. Stack, 17.
Bordism 13, 14.
Classifying space 7, 14.
Collar 8, 14.
Connected sum 16.
Degree 9, 19.
Derived intersection 21, 22.
Descent 17, 18, 19.
Discriminant 20, 23.
Embedding 2, 15, 24.
Equivariance 25, 26.
Euler class 10, 19.
Exit paths 23.
Factorization 18, 27.
Fixed points 25, 26.
Folds 20.
Handle attachment 11, 16.
Immersion 2, 24.
Intersection number 10.
Isotopy extension 15.
Jet 4, 20.
Mapping space 3, 18.
Monodromy 17, 19.
Morse theory 11, 12.
Orbit type 25.
Orientation 7, 19.
Persistence 28.
Pontryagin–Thom construction 13.
Quotient stack 17, 25.
Regular value 5, 20.
Sard's theorem 5, 6.
Slice theorem 25.
Sublevel sets 11, 28.
Tangent complex 21, 22.
Transversality 6, 26.
Tubular neighborhood 8.
Vector bundle 7, 19.

Bibliography

[1] David Ayala, John Francis, and Nick Rozenblyum. A stratified homotopy hypothesis. *Journal of the European Mathematical Society* 21 (2019), 1071–1178. DOI.

[2] John C. Baez and Alexander E. Hoffnung. Convenient categories of smooth spaces. *Transactions of the American Mathematical Society* 363 (2011), 5789–5825. arXiv:0807.1704.

[3] Raoul Bott and Loring W. Tu. *Differential Forms in Algebraic Topology*. Graduate Texts in Mathematics 82, Springer, 1982. DOI.

[4] David Carchedi. Derived manifolds as differential graded manifolds. Preprint, 2023. arXiv:2303.11140.

[5] David Carchedi and Pelle Steffens. On the universal property of derived manifolds. arXiv:1905.06195.

[6] Frédéric Chazal, Vin de Silva, Marc Glisse, and Steve Oudot. *The Structure and Stability of Persistence Modules*. SpringerBriefs in Mathematics. Springer, 2016. DOI.

[7] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. *Discrete & Computational Geometry* 37 (2007), 103–120. DOI.

[8] Ralph L. Cohen, John D. S. Jones, and Graeme B. Segal. *Morse Theory and Classifying Spaces*. Preprint, 1995. Available manuscript.

[9] Daniel S. Freed. *Bordism: Old and New*. Lecture notes, 2013. Author's notes.

[10] Søren Galatius, Ib Madsen, Ulrike Tillmann, and Michael Weiss. The homotopy type of the cobordism category. *Acta Mathematica* 202 (2009), 195–239. DOI.

[11] Martin Golubitsky and Victor Guillemin. *Stable Mappings and Their Singularities*. Graduate Texts in Mathematics 14. Springer, 1973. DOI.

[12] Thomas G. Goodwillie and Michael Weiss. Embeddings from the point of view of immersion theory II. *Geometry & Topology* 3 (1999), 103–118. DOI.

[13] Allen Hatcher. *Algebraic Topology*. Cambridge University Press, 2002. Author's copy.

[14] Allen Hatcher. *Vector Bundles and K-Theory*. Version 2.2, 2017. Author's notes.

[15] Morris W. Hirsch. *Differential Topology*. Graduate Texts in Mathematics 33. Springer, 1976. Publisher's edition.

[16] Morris W. Hirsch. Immersions of manifolds. *Transactions of the American Mathematical Society* 93 (1959), 242–276. DOI.

[17] Andreas Kriegl and Peter W. Michor. *The Convenient Setting of Global Analysis*. Mathematical Surveys and Monographs 53. American Mathematical Society, 1997. DOI.

[18] John M. Lee. *Introduction to Smooth Manifolds*. Second edition. Graduate Texts in Mathematics 218. Springer, 2013. DOI.

[19] Jacob Lurie. *Higher Algebra*. Manuscript, 2017. Author's manuscript.

[20] Jacob Lurie. *Higher Topos Theory*. Annals of Mathematics Studies 170. Princeton University Press, 2009. Author's copy.

[21] Jacob Lurie. On the classification of topological field theories. *Current Developments in Mathematics 2008*, 129–280, 2009. arXiv:0905.0465.

[22] Jiří Matoušek. *Using the Borsuk–Ulam Theorem*. Universitext. Springer, 2003. DOI.

[23] Eckhard Meinrenken. *Group Actions on Manifolds*. Lecture notes. Author's notes.

[24] John W. Milnor. *Topology from the Differentiable Viewpoint*. Princeton University Press, 1997 reissue. Publisher's edition.

[25] John W. Milnor. *Morse Theory*. Annals of Mathematics Studies 51. Princeton University Press, 1963. Publisher's edition.

[26] John W. Milnor and James D. Stasheff. *Characteristic Classes*. Annals of Mathematics Studies 76. Princeton University Press, 1974. Publisher's edition.

[27] Richard S. Palais. Local triviality of the restriction map for embeddings. *Commentarii Mathematici Helvetici* 34 (1960), 305–312. Journal scan.

[28] Matthias Schwarz. *Morse Homology*. Progress in Mathematics 111. Birkhäuser, 1993. DOI.

[29] Paul Seidel. *Differential Geometry*. MIT OpenCourseWare 18.950, Fall 2008. Chapter 1, Lecture 9. Lecture notes.

[30] David I. Spivak. Derived smooth manifolds. *Duke Mathematical Journal* 153 (2010), 55–128. arXiv:0810.5174.

[31] Arthur G. Wasserman. Equivariant differential topology. *Topology* 8 (1969), 127–150. DOI.

[32] Katrin Wehrheim. Smooth structures on Morse trajectory spaces, featuring finite ends and associative gluing. *Geometry & Topology Monographs* 18 (2012), 369–450. DOI.

[33] Michael Weiss. Embeddings from the point of view of immersion theory I. *Geometry & Topology* 3 (1999), 67–101. DOI.