Pseudo-holomorphic curves and virtual fundamental cycles

John Pardon

Spring 2021

Note: This text currently takes the form of informal (and very incomplete) lecture notes. Its release in such form is rather unconventional and probably ill-advised, but is made due to popular request. I do not warrant that it is free of errors, let alone useful for anything, and it is subject to change without notice. *Use at your own risk!* Comments and corrections are always welcome.

Contents

Contents **i**

Foreword **iii**

1 Topology **1**
1.1 Categories . . . 1
1.2 2-categories . . . 2
1.3 Properties of morphisms . . . 4
1.4 Sheaves and stacks . . . 9
1.5 Representable maps . . . 11
1.6 Topological stacks . . . 12
1.7 Mapping stacks . . . 13
1.8 Stability . . . 16
1.9 Grothendieck topologies . . . 18

2 Algebraic topology **20**
2.1 Simplicial sets . . . 20
2.2 ∞-categories . . . 26
2.3 Spaces . . . 38
2.4 Sheaves . . . 44
2.5 Shape theory . . . 47
2.6 Spanier–Whitehead duality . . . 50
2.7 Proper pushforward . . . 50
2.8 Fundamental cycles . . . 51
2.9 Verdier duality . . . 53
2.10 Derived category of vector bundles . . . 55
2.11 Global spaces and orbispaces . . . 56
2.12 Global ∞-categories . . . 58
2.13 Equivariant site . . . 58

3 Analysis **60**
3.1 Function spaces . . . 60
3.2 Differential operators . . . 64

3.3 Elliptic operators 66
3.4 Ellipticity in cylindrical ends 73
3.5 Families of elliptic operators 74

4 Riemann surfaces **77**
4.1 Riemann surfaces 77
4.2 Families of Riemann surfaces 77
4.3 Cauchy–Riemann operators 78
4.4 Families of Cauchy–Riemann operators 78

5 Pseudo-holomorphic curves **79**
5.1 Pseudo-holomorphic maps 79
5.2 A priori estimates 79
5.3 Families of pseudo-holomorphic maps 85
5.4 Moduli stacks of pseudo-holomorphic maps 86
5.5 Stability 86
5.6 Linear deformation theory 87
5.7 Non-linear deformation theory 88
5.8 Inverse exponential maps 89
5.9 Local structure of moduli spaces 92

Bibliography **95**

Foreword

Pseudo-holomorphic curves were introduced to symplectic topology in a seminal paper of Gromov [6], and their study has by now evolved into a mature subject. The aim of this text is to provide a foundational treatment of the enumerative theory of pseudo-holomorphic curves.

A map $u : C \to X$ from a Riemann surface C to an almost complex manifold X (i.e. a manifold equipped with an endomorphism $J : TX \to TX$ squaring to $-\mathbf{1}$) is called *pseudo-holomorphic* when its differential $du : TC \to TX$ is $\mathbb{C}$-linear. The equation asserting $\mathbb{C}$-linearity of du is called the *pseudo-holomorphic curve equation*; it is a non-linear elliptic partial differential equation.

Our central objects of study in this text are the moduli spaces of solutions to the pseudo-holomorphic curve equation and its variants. Each such moduli space carries a canonical cycle, known as its *virtual fundamental cycle*, which encapsulates its enumerative content. The ultimate goal for us is to construct this cycle.

Acknowledgements

The presentation in this text has benefitted in innumerable ways from many conversations over many years with many different people. With advance apologies for any omissions, let me acknowledge in particular communication with Mohammed Abouzaid, Dustin Clausen, Yasha Eliashberg, Tobias Ekholm, Helmut Hofer, Daniel Litt, Stefan Schwede, Paul Seidel, Mohan Swaminathan.

Chapter 1

Topology

1.1 Categories

1.1.1 Definition (Category)**.** A *category* $\mathcal{C}$ consists of the following data:

(1.1.1.1) A set $\mathcal{C}$, whose elements are called the *objects* of $\mathcal{C}$.

(1.1.1.2) For every pair of objects $X, Y \in \mathcal{C}$, a set $\operatorname{Hom}(X, Y)$, whose elements are called the *morphisms* $X \to Y$ in $\mathcal{C}$.

(1.1.1.3) For every triple of objects $X, Y, Z \in \mathcal{C}$, a map

$$\operatorname{Hom}(X, Y) \times \operatorname{Hom}(Y, Z) \to \operatorname{Hom}(X, Z)$$

called *composition*, such that for every composable triple of morphisms a, b, c, the two compositions $(ab)c$ and $a(bc)$ are equal.

(1.1.1.4) For every object $X \in \mathcal{C}$, an element $\mathbf{1}_X \in \operatorname{Hom}(X, X)$ called the *identity morphism* such that composition with $\mathbf{1}_X$ defines the identity map $\operatorname{Hom}(X, Y) \to \operatorname{Hom}(X, Y)$ and $\operatorname{Hom}(Z, X) \to \operatorname{Hom}(Z, X)$.

1.1.2 Example. Sets and functions form a category Set. That is, an object of Set is a set X, and a morphism $X \to Y$ in Set is a map of sets. Topological spaces and continuous maps form a category Top. Groups and group homomorphisms form a category Grp.

1.1.3 Exercise. Show that in a category, the identity morphisms (1.1.1.4) are uniquely determined by the rest of the data (1.1.1.1)–(1.1.1.3) provided they exist.

1.1.4 Exercise. A morphism $X \to Y$ in a category is called an *isomorphism* iff there exists a morphism $Y \to X$ such that the compositions $X \to Y \to X$ and $Y \to X \to Y$ are the identity morphisms $\mathbf{1}_X$ and $\mathbf{1}_Y$. Show that a given morphism $X \to Y$ as at most one such 'inverse' morphism $Y \to X$.

1.1.5 Definition (Functor)**.** A *functor* $F : \mathcal{C} \to \mathcal{D}$ between categories consists of the following data:

(1.1.5.1) For every object $X \in \mathcal{C}$, an object $F(X) \in \mathcal{D}$.

(1.1.5.2) For every pair of objects $X, Y \in \mathcal{C}$, a map $F : \operatorname{Hom}(X, Y) \to \operatorname{Hom}(F(X), F(Y))$ such that $F(\mathbf{1}_X) = \mathbf{1}_{F(X)}$ and such that composing and applying F in either order define the same map

$$\operatorname{Hom}(X, Y) \times \operatorname{Hom}(Y, Z) \to \operatorname{Hom}(F(X), F(Z)).$$

Given categories $\mathcal{C}$ and $\mathcal{D}$, there is a category $\mathsf{Fun}(\mathcal{C}, \mathcal{D})$ whose objects are functors $F : \mathcal{C} \to \mathcal{D}$ and whose morphisms $F \to G$ (called *natural transformations*) consist of, for every object $X \in \mathcal{C}$, a morphism $F(X) \to G(X)$ such that for every morphism $X \to Y$, the two compositions $F(X) \to G(X) \to G(Y)$ and $F(X) \to F(Y) \to G(Y)$ agree.

1.1.6 Definition (Equivalence of categories)**.** A functor $\mathcal{C} \to \mathcal{D}$ is called an *equivalence* iff there exists a functor $\mathcal{D} \to \mathcal{C}$ such that the compositions $\mathcal{C} \to \mathcal{D} \to \mathcal{C}$ and $\mathcal{D} \to \mathcal{C} \to \mathcal{D}$ are naturally isomorphic (in $\mathsf{Fun}(\mathcal{C}, \mathcal{C})$ and $\mathsf{Fun}(\mathcal{D}, \mathcal{D})$, respectively) to the identity functors $\mathbf{1}_{\mathcal{C}}$ and $\mathbf{1}_{\mathcal{D}}$.

1.1.7 Exercise. Let $\mathcal{C}$ be any category, and let $\mathcal{C}_0 \subseteq \mathcal{C}$ be the full subcategory spanned by a subset of the objects of $\mathcal{C}$ which contains at least one object in every isomorphism class. Show that the tautological functor $\mathcal{C}_0 \to \mathcal{C}$ is an equivalence.

1.2 2-categories

1.2.1 Definition (2-category)**.** A *2-category* $\mathcal{C}$ consists of the following data:

(1.2.1.1) A set $\mathcal{C}$, whose elements are called the *objects* of $\mathcal{C}$.

(1.2.1.2) For every pair of objects $X, Y \in \mathcal{C}$, a groupoid $\operatorname{Hom}(X, Y)$, whose objects are called the *morphisms* $X \to Y$ in $\mathcal{C}$.

(1.2.1.3) For every triple of objects $X, Y, Z \in \mathcal{C}$, a functor

$$\operatorname{Hom}(X, Y) \times \operatorname{Hom}(Y, Z) \to \operatorname{Hom}(X, Z)$$

called *composition*.

(1.2.1.4) For every quadruple of objects $X, Y, Z, W \in \mathcal{C}$, a natural isomorphism between the two ways of composing twice to obtain a functor

$$\operatorname{Hom}(X, Y) \times \operatorname{Hom}(Y, Z) \times \operatorname{Hom}(Z, W) \to \operatorname{Hom}(X, W)$$

such that for every quadruple of morphisms a, b, c, d, the cyclic composition

$$\begin{array}{ccccc} & & (ab)(cd) & & \\ & \nearrow\!\!\swarrow & & \nwarrow\!\!\searrow & \\ ((ab)c)d & & & & a(b(cd)) \\ \updownarrow & & & & \updownarrow \\ (a(bc))d & & \longleftrightarrow & & a((bc)d) \end{array}$$

is the identity map.

(1.2.1.5) For every object $X \in \mathcal{C}$, an object $\mathbf{1}_X \in \operatorname{Hom}(X, X)$ together with natural isomorphisms between pairing with $\mathbf{1}_X$ and the identity functors $\operatorname{Hom}(X, Y) \to \operatorname{Hom}(X, Y)$ and $\operatorname{Hom}(Z, X) \to \operatorname{Hom}(Z, X)$, such that for every pair of morphisms a, b, the cyclic compositions

$$ab \leftrightarrow (\mathbf{1}a)b \leftrightarrow \mathbf{1}(ab) \leftrightarrow ab$$
$$ab \leftrightarrow (a\mathbf{1})b \leftrightarrow a(\mathbf{1}b) \leftrightarrow ab$$
$$ab \leftrightarrow a(b\mathbf{1}) \leftrightarrow (ab)\mathbf{1} \leftrightarrow ab$$

are the identity maps.

1.2.2 Exercise. Show that in the identity axiom (1.2.1.5), the cyclic composition $ab \leftrightarrow (a\mathbf{1})b \leftrightarrow a(\mathbf{1}b) \leftrightarrow ab$ being the identity map implies the same for the other two (use the pentagon identity (1.2.1.4)).

1.2.3 Definition (2-functor)**.** A functor between 2-categories $F : \mathcal{C} \to \mathcal{D}$ consists of the following data:

(1.2.3.1) For every object $X \in \mathcal{C}$, an object $F(X) \in \mathcal{D}$.

(1.2.3.2) For every pair of objects $X, Y \in \mathcal{C}$, a functor

$$\operatorname{Hom}(X, Y) \to \operatorname{Hom}(F(X), F(Y)).$$

(1.2.3.3) For every triple of objects $X, Y, Z \in \mathcal{C}$, a natural isomorphism between the two compositions

$$\begin{array}{ccc} \operatorname{Hom}(X,Y) \times \operatorname{Hom}(Y,Z) & \longrightarrow & \operatorname{Hom}(X,Z) \\ \downarrow & & \downarrow \\ \operatorname{Hom}(F(X),F(Y)) \times \operatorname{Hom}(F(Y),F(Z)) & \longrightarrow & \operatorname{Hom}(F(X),F(Z)) \end{array}$$

which is compatible with composition in the sense that for any triple of morphisms a, b, c, the cyclic composition

$$\begin{array}{ccc} F(abc) & \longleftrightarrow & F(a)F(bc) \\ \updownarrow & & \updownarrow \\ F(ab)F(c) & \longleftrightarrow & F(a)F(b)F(c) \end{array}$$

is the identity.

(1.2.3.4) For every object $X \in \mathcal{C}$, an isomorphism between $F(\mathbf{1}_X)$ and $\mathbf{1}_{F(X)}$ such that for every $a \in \operatorname{Hom}(Y, X)$, the following cyclic composition is the identity map

$$\begin{array}{ccc} F(a) & \longleftrightarrow & F(a\mathbf{1}_X) \\ \updownarrow & & \updownarrow \\ F(a)\mathbf{1}_{F(X)} & \longleftrightarrow & F(a)F(\mathbf{1}_X) \end{array}$$

and similarly for objects of $\operatorname{Hom}(X, Z)$.

Given 2-categories $\mathcal{C}$ and $\mathcal{D}$, there is a 2-category $\mathsf{Fun}(\mathcal{C}, \mathcal{D})$ whose objects are functors $F : \mathcal{C} \to \mathcal{D}$ and whose morphisms $\phi : F \to G$ consist of the following data:

(1.2.3.5) For every object $X \in \mathcal{C}$, a morphism $\phi_X : F(X) \to G(X)$.

(1.2.3.6) For every pair of objects $X, Y \in \mathcal{C}$, a natural isomorphism between the two compositions

$$\begin{array}{ccc} \mathrm{Hom}(X,Y) & \longrightarrow & \mathrm{Hom}(F(X),F(Y))) \\ \downarrow & & \downarrow \\ \mathrm{Hom}(G(X),G(Y)) & \longrightarrow & \mathrm{Hom}(F(X),G(Y)) \end{array}$$

which is compatible with composition in the sense that for any pair of morphisms a, b, the cyclic composition

$$\begin{array}{ccccc} F(a)F(b)\phi_Z & \longleftrightarrow & F(a)\phi_Y G(b) & \longleftrightarrow & \phi_X G(a)G(b) \\ \updownarrow & & & & \updownarrow \\ F(ab)\phi_Z & & \longleftrightarrow & & \phi_X G(ab) \end{array}$$

is the identity and for every object $X \in \mathcal{C}$ the cyclic composition

$$\begin{array}{ccccc} F(\mathbf{1}_X)\phi_X & & \longleftrightarrow & & \phi_X G(\mathbf{1}_X) \\ \updownarrow & & & & \updownarrow \\ \mathbf{1}_{F(X)}\phi_X & \longleftrightarrow & \phi_X & \longleftrightarrow & \phi_X \mathbf{1}_{G(X)} \end{array}$$

is the identity.

An isomorphism of morphisms $\phi \to \psi$ consists of:

(1.2.3.7) For every $X \in \mathcal{C}$, an isomorphism $\phi_X \to \psi_X$ such that the cyclic composition

$$\begin{array}{ccc} \phi_X F(a) & \longleftrightarrow & G(a)\phi_X \\ \updownarrow & & \updownarrow \\ \psi_X F(a) & \longleftrightarrow & G(a)\psi_X \end{array}$$

is the identity.

1.3 Properties of morphisms

1.3.1 Definition (Property of morphisms)**.** A *property of morphisms* $\mathcal{P}$ in a category $\mathcal{C}$ is a set $\mathcal{P}$ of isomorphism classes in $\mathsf{Fun}(\Delta^1, \mathcal{C})$ (an object of $\mathsf{Fun}(\Delta^1, \mathcal{C})$ is a morphism $X \to Y$ in $\mathcal{C}$, and a morphism in $\mathsf{Fun}(\Delta^1, \mathcal{C})$ is a commutative square). We will say "$X \to Y$ has/is $\mathcal{P}$" to mean $(X \to Y) \in \mathcal{P}$.

1.3.2 Example. In any category $\mathcal{C}$, properties of morphisms include being an isomorphism, being a monomorphism, or being an epimorphism.

1.3.3 Example. In the category Top of topological spaces, properties of morphisms include being injective, being surjective, having dense image, being open (image of every open set is open), being closed (image of every closed set is closed).

1.3.4 Definition (Property preserved under pullback)**.** A property of morphisms $\mathcal{P}$ in a category with fiber products is said to be *preserved under pullback* iff $X \to Y$ having $\mathcal{P}$ implies $X \times_Y Z \to Z$ has $\mathcal{P}$ for every morphism $Z \to Y$.

1.3.5 Definition (Property preserved under composition)**.** A property of morphisms $\mathcal{P}$ is said to be *preserved under composition* iff the composition of any two morphisms with $\mathcal{P}$ has $\mathcal{P}$.

1.3.6 Exercise. Suppose $\mathcal{P}$ is a property of morphisms which is preserved under pullback and composition. Show that $\mathcal{P}$ is preserved under fiber product, meaning that if $X \to Y$ and $X' \to Y'$ have $\mathcal{P}$ and $Y \to Z \leftarrow Y'$ are arbitrary, then $X \times_Z X' \to Y \times_Z Y'$ has $\mathcal{P}$.

1.3.7 Exercise. Show that the following properties of morphisms of sets are preserved under pullback and composition: injective, surjective, bijective. Conclude the same for morphisms of topological spaces by noting that the forgetful functor $\mathsf{Top} \to \mathsf{Set}$ is continuous.

1.3.8 Exercise. The following properties of morphisms of topological spaces $f : X \to Y$ are preserved under pullback and composition:

(1.3.8.1) f is open.

(1.3.8.2) f is an embedding (i.e. a homeomorphism onto its image).

(1.3.8.3) f is a closed embedding.

(1.3.8.4) f has local sections (i.e. there is an open cover $Y = \bigcup_i U_i$ such that each inclusion $U_i \to Y$ factors through $X \to Y$).

1.3.9 Definition. A *swarm* is a pair (S, s) consisting of a topological space S and a closed point $s \in S$ such that $S \setminus s$ has the discrete topology and $\overline{S \setminus s} = S$. A map of swarms $f : (S, s) \to (T, t)$ is a continuous map with $f^{-1}(t) = s$. (I would be grateful to anyone who shows me an existing name for this notion.)

1.3.10 Exercise. Show that a subset A of a topological space X is closed iff for every swarm (S, s) and every continuous map $f : S \to X$, we have $f(S \setminus s) \subseteq A$ implies $f(s) \in A$.

1.3.11 Exercise. Show that there is a unique topology on $\mathbb{Z}_{\geq 0} \cup \{\infty\}$ with the property that a map $f : \mathbb{Z}_{\geq 0} \cup \{\infty\} \to X$ is continuous iff the sequence $f(0), f(1), f(2), \ldots$ converges to $f(\infty)$. Show that $(\mathbb{Z}_{\geq 0} \cup \{\infty\}, \{\infty\})$ with this topology is a swarm.

1.3.12 Exercise. Show that being closed (for maps of topological spaces) is not preserved under pullback.

1.3.13 Proposition. *For a map of topological spaces $f : X \to Y$, the following are equivalent:*

(1.3.13.1) (Universally closed) *For every map $Z \to Y$, the pullback $X \times_Y Z \to Z$ is closed.*

(1.3.13.2) (Finite subcover property) *For every $\{U_i \subseteq X\}_i$ covering $f^{-1}(y)$, there exists a finite subcollection which cover $f^{-1}(V)$ for some open neighborhood $y \in V \subseteq Y$.*

(1.3.13.3) (Swarm lifting property) *For every swarm S and every diagram of solid arrows*

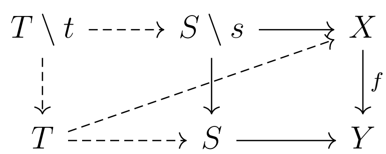

there exists a map of swarms $T \to S$ and dotted arrows making the diagram commute.

These conditions are a relative version of compactness: a topological space X is compact iff the map $X \to *$ is universally closed.

Proof. We show the swarm lifting property (1.3.13.3) implies universal closedness (1.3.13.1). Suppose $X \to Y$ satisfies the swarm lifting property (1.3.13.3). The swarm lifting property is evidently preserved under pullback, so it suffices to show that it implies that $X \to Y$ is closed. Let $A \subseteq X$ be closed. To show its image is closed, we just need to show that any limit of any swarm in $f(A)$ also lies in $f(A)$. In other words, we have a diagram of solid arrows as in (1.3.13.3), so the swarm lifting property means the limit also lies in $f(A)$.

We show that universal closedness (1.3.13.1) implies the swarm lifting property (1.3.13.3). Let $X \to Y$ be universally closed (1.3.13.1). Consider the closure of the image of $S \setminus s \to X \times_Y S$. The image of this closure in S certainly contains $S \setminus s$, and it is closed since $X \to Y$ is universally closed, so it hence contains s. In view of (1.3.10), there is a swarm T and a map $T \to X \times_Y S$ whose projection down to S is a map of swarms, which is exactly the desired swarm lifting diagram.

We show the swarm lifting property (1.3.13.3) implies the finite subcover property (1.3.13.2). Suppose $X \to Y$ satisfies the swarm lifting property (1.3.13.3). Let U_α be a covering of $f^{-1}(y)$, and suppose for sake of contradiction that no finite subcollection covers $f^{-1}(V)$ for any neighborhood $y \in V \subseteq Y$. We consider the directed set indexed by pairs consisting of finite subsets of the U_α and neighborhoods $y \in V \subseteq Y$. By assumption, for each element α of this directed set, we may find an $x_\alpha \in f^{-1}(V_\alpha) \setminus \bigcup_{i \in I_\alpha} U_i$. By construction, $f(x_\alpha) \to y$, so by (1.3.13.3), there exists a subnet of x_α converging to some x with $f(x) = y$. On the other hand, since x is in some U_i, the net x_α is eventually outside U_i, hence can have no subnet converging to x, contradiction.

We show that the finite subcover property (1.3.13.2) implies universal closedness (1.3.13.1). The finite subcover property specialized to the case of a single open set is precisely the property of being closed (it says that if $U \subseteq X$ is open and contains $f^{-1}(y)$, then it contains $f^{-1}(V)$ for some open neighborhood $y \in V \subseteq Y$, which is precisely the assertion that $f(X \setminus U) \subseteq Y$ is closed). It thus suffices to note that the finite subcover property is preserved under pullback. □

1.3.14 Exercise. Show that a map of topological spaces is universally closed iff it is closed and has compact fibers.

1.3.15 Exercise (Separated). Show that for a morphism of topological spaces $f : X \to Y$, the following are equivalent:

(1.3.15.1) Every pair of distinct points $x_1, x_2 \in X$ in the same fiber $f(x_1) = f(x_2)$ have disjoint open neighborhoods $U_1 \cap U_2 = \varnothing$, $x_i \in U_i \subseteq X$.

(1.3.15.2) The relative diagonal $X \to X \times_Y X$ is a closed embedding.

(1.3.15.3) For every swarm S and every diagram of solid arrows

$$\begin{array}{ccc} S \setminus s & \longrightarrow & X \\ \downarrow & \nearrow & \downarrow \\ S & \longrightarrow & Y \end{array}$$

there exists at most one dotted arrow making the diagram commute.

A morphism satisfying these conditions is called *separated*; this is a relative version of the Hausdorff property (X is Hausdorff iff $X \to *$ is separated). Show that being separated is preserved under pullback and composition.

1.3.16 Definition (Properties of the diagonal)**.** For any property $\mathcal{P}$ of morphisms in a category in which all self fiber products $X \times_Y X$ exist, we may define a new property $\mathcal{P}_\Delta$ to mean that the relative diagonal has $\mathcal{P}$. That is, $X \to Y$ has $\mathcal{P}_\Delta$ iff its relative diagonal $X \to X \times_Y X$ has $\mathcal{P}$.

1.3.17 Exercise. Show that if $\mathcal{P}$ is 'is an isomorphism' then $\mathcal{P}_\Delta$ is 'is a monomorphism'. Show that if $\mathcal{P}$ is 'is a split monomorphism' then every morphism has $\mathcal{P}_\Delta$.

1.3.18 Exercise. For maps of topological spaces, show that if $\mathcal{P}$ is 'is an embedding' then $\mathcal{P}_\Delta$ is everything.

1.3.19 Exercise (The diagonal of a pullback is a pullback of the diagonal). Show that if the left square is a fiber square, then so are the right two squares:

$$\begin{array}{ccc} X' & \longrightarrow & Y' \\ \downarrow & & \downarrow \\ X & \longrightarrow & Y \end{array} \quad \Longrightarrow \quad \begin{array}{ccccc} X' & \longrightarrow & X' \times_{Y'} X' & \longrightarrow & Y' \\ \downarrow & & \downarrow & & \downarrow \\ X & \longrightarrow & X \times_Y X & \longrightarrow & Y \end{array} \tag{1.3.19.1}$$

Conclude that if $\mathcal{P}$ is preserved under pullback then so is $\mathcal{P}_\Delta$.

1.3.20 Exercise (The diagonal of a composition is a composition of pullbacks of the diagonals). Show that if $X \to Y \to Z$ are morphisms, then $X \times_Y X \to X \times_Z X$ is a pullback of $Y \to Y \times_Z Y$. Conclude that if $\mathcal{P}$ is preserved under pullback and composition then so is $\mathcal{P}_\Delta$.

1.3.21 Lemma (Cancellation)**.** *Let $\mathcal{P}$ be a property of morphisms preserved under pullback and under composition. If the composition $X \to Y \to Z$ has $\mathcal{P}$ and $Y \to Z$ has $\mathcal{P}_\Delta$, then $X \to Y$ has $\mathcal{P}$.*

Proof. We have $X \to Y$ factors as $X \to X \times_Z Y \to Y$. The first map $X \to X \times_Z Y$ is a pullback of $Y \to Y \times_Z Y$ so has $\mathcal{P}$. The second map is a pullback of $X \to Z$ so has $\mathcal{P}$. □

1.3.22 Exercise. If $X \to Y$ and $Y \to Z$ are maps of sets whose composition $X \to Z$ is injective, then the first map $X \to Y$ is also injective. Prove this using the abstract cancellation property (1.3.21).

1.3.23 Exercise. Prove both directly and using cancellation that if $X \to Y \to Z$ are maps of topological spaces whose composition $X \to Z$ is separated, then the first map $X \to Y$ is separated.

1.3.24 Exercise. Prove both directly and using cancellation that if $X \to Y \to Z$ are maps of topological spaces whose composition $X \to Z$ is an embedding, then the first map $X \to Y$ is an embedding.

1.3.25 Definition (Proper)**.** A map of topological spaces is called *proper* iff it is separated and universally closed. This is a relative version of being compact Hausdorff.

1.3.26 Exercise. Show that for an embedding of topological spaces, the three conditions closed, universally closed, and proper are all equivalent. Conclude that a proper map has proper diagonal.

1.3.27 Remark (Discussion of properness). The diagonal of any map of topological spaces is an embedding, so being separated is the same as having universally closed diagonal. Thus we may state the definition of properness more cleanly as $X \to Y$ is proper iff $X \to Y$ and $X \to X \times_Y X$ are both universally closed. The diagonal being a monomorphism (since it has a retraction) implies that the diagonal of the diagonal is an isomorphism. Thus properness can be stated even more cleanly as the condition that $X \to Y$ and all of its iterated diagonals are universally closed. The significance of properness is also explained by phrasing it in terms of the lifting properties (1.3.13.3) and (1.3.15.3), i.e. it is an existence and uniqueness condition for swarm lifting (allowing for precomposition by an arbitrary map of swarms). An elementary analogy is the statement that map of sets is an isomorphism iff it and its diagonal are both surjective.

1.3.28 Definition (Local properties)**.** A property $\mathcal{P}$ of morphisms of topological spaces which is preserved under pullback is called *local on the target* iff for every morphism $X \to Y$, if there exists an open cover $Y = \bigcup_i U_i$ such that each pullback $X \times_Y U_i \to U_i$ has $\mathcal{P}$, then $X \to Y$ has $\mathcal{P}$.

1.3.29 Exercise (Descent for local properties). Let $\mathcal{P}$ be a property of morphisms of topological spaces which is preserved under pullback. Show that $\mathcal{P}$ is local on the target iff for every map $X \to Y$ and every map $Z \to Y$ admitting local sections, we have $(X \to Y) \in \mathcal{P}$ iff $(X \times_Y Z \to Z) \in \mathcal{P}$.

1.3.30 Exercise. Show that the following properties of morphisms of topological spaces preserved under pullback are local on the target:

(1.3.30.1) open

(1.3.30.2) embedding
(1.3.30.3) closed embedding
(1.3.30.4) separated
(1.3.30.5) universally closed

1.4 Sheaves and stacks

1.4.1 Definition (Sheaf)**.** A functor $F : \mathsf{Top}^{\mathsf{op}} \to \mathsf{Set}$ is called a *sheaf* iff for every topological space U and every open cover $U = \bigcup_i U_i$, the natural map

$$F(U) \xrightarrow{\sim} \left\{ \{\alpha_i \in F(U_i)\}_{i \in I} \,\middle|\, \alpha_j|_{U_i \cap U_j} = \alpha_i|_{U_i \cap U_j} \right\} \tag{1.4.1.1}$$

is a bijection. A morphism of sheaves is a natural transformation of functors. That is, the category of sheaves is, by definition, the full subcategory $\mathsf{Shv}(\mathsf{Top}) \subseteq \mathsf{Fun}(\mathsf{Top}^{\mathsf{op}}, \mathsf{Set})$ spanned by sheaves.

1.4.2 Exercise. Show that $\operatorname{Hom}(-, X) : \mathsf{Top}^{\mathsf{op}} \to \mathsf{Set}$ is a sheaf for any topological space X.

Recall that a presheaf $\mathcal{C}^{\mathsf{op}} \to \mathsf{Set}$ is called representable iff it is isomorphic to $\operatorname{Hom}(-, X)$ for some $X \in \mathcal{C}$. Thus (1.4.2) says that every representable presheaf is a sheaf. That is, the Yoneda embedding $\mathsf{Top} \hookrightarrow \mathsf{PShv}(\mathsf{Top})$ lands inside the full subcategory $\mathsf{Shv}(\mathsf{Top}) \subseteq \mathsf{PShv}(\mathsf{Top})$.

1.4.3 Remark. We wish to view $\mathsf{Shv}(\mathsf{Top})$ as a category of 'generalized topological spaces' containing the usual category of topological spaces as a full subcategory $\mathsf{Top} \subseteq \mathsf{Shv}(\mathsf{Top})$. In particular, we will see how to manipulate objects of $\mathsf{Shv}(\mathsf{Top})$ in similar ways to how we manipulate topological spaces.

1.4.4 Remark. It is often the case that one can produce an object of $\mathsf{Shv}(\mathsf{Top})$ quite easily, and then through some hard work show it in fact lies in the full subcategory Top, hence is a genuine space. Often this is the case with moduli spaces: it is easy to define a 'moduli functor' sending X to 'families over X', and it is usually self-evident that it is a sheaf. Proving representability is more difficult, but is made easier by having first shown the moduli functor to be a sheaf (e.g. representability can be proven 'locally').

1.4.5 Exercise. Consider the functor $\mathsf{Top}^{\mathsf{op}} \to \mathsf{Set}$ which sends a topological space Z to the collection of open subsets of Z is a sheaf. Show that this functor is a sheaf; then, show that it is representable.

1.4.6 Example. Consider the functor $\mathsf{Top}^{\mathsf{op}} \to \mathsf{Set}$ which sends a space X to the set of double covers $\tilde{X} \to X$ up to isomorphism ($\tilde{X} \to X$ is a double cover iff there exists an open cover $X = \bigcup_i U_i$ such that each pullback $\tilde{X} \times_X U_i \to U_i$ is isomorphic to $U_i \sqcup U_i \to U_i$). Show that this functor is not a sheaf.

The reason the presheaf of isomorphism classes of double covers (1.4.6) fails to be a sheaf is that two double covers may be isomorphic in multiple ways, and we must remember the specific isomorphisms in order to glue. A set does not know 'isomorphisms' between its elements; a groupoid does. Thus in order to have descent, we should regard double covers of X not as a set, but rather as a groupoid, i.e. double covers should be regarded as a functor $\mathsf{Top}^{\mathsf{op}} \to \mathsf{Grpd}$ where Grpd is the 2-category of groupoids. This is a problem any time we are trying to parameterize objects which have automorphisms (in the case of (1.4.6), two-element sets).

1.4.7 Definition (Stack). A functor $F : \mathsf{Top}^{\mathsf{op}} \to \mathsf{Grpd}$ is called a stack iff for every open cover $U = \bigcup_i U_i$, the natural map

$$F(U) \xrightarrow{\sim} \left\{ \begin{matrix} \{\alpha_i \in F(U_i)\}_{i\in I} \\ \{\beta_{ij} : \alpha_j|_{U_i\cap U_j} \xrightarrow{\sim} \alpha_i|_{U_i\cap U_j}\}_{i,j\in I} \end{matrix} \middle| \begin{matrix} \beta_{ii} = \mathbf{1}_{\alpha_i} \\ \beta_{ij}\beta_{jk}\beta_{ki} = \mathbf{1}_{\alpha_i}|_{U_i\cap U_j\cap U_k} \end{matrix} \right\} \tag{1.4.7.1}$$

is an equivalence of groupoids, where on the right hand side, an isomorphism

$$(\{\alpha_i\}_{i\in I}, \{\beta_{ij}\}_{i,j\in I}) \to (\{\alpha'_{i'}\}_{i'\in I'}, \{\beta'_{i'j'}\}_{i',j'\in I'}) \tag{1.4.7.2}$$

consists of a collection $\{\gamma_{ii'} : \alpha_i|_{U_i\cap U'_{i'}} \xrightarrow{\sim} \alpha'_{i'}|_{U_i\cap U'_{i'}}\}_{i\in I, i'\in I'}$ satisfying $\alpha_{ii'}\beta_{i'j'} = \alpha_{ij'}$ and $\beta_{ij}\alpha_{jj'} = \alpha_{ij'}$. The category of stacks is denoted $\mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})$ or $\mathsf{Stk}(\mathsf{Top})$.

1.4.8 Exercise. Continuing (1.4.6), consider now the functor $\mathsf{Top}^{\mathsf{op}} \to \mathsf{Grpd}$ which sends a topological space X to the groupoid whose objects are double covers $\tilde{X} \to X$ and whose morphisms are isomorphisms over X. Show that this functor is a stack.

Every sheaf is a stack, by composing with the fully faithful inclusion $\mathsf{Set} \subseteq \mathsf{Grpd}$. Thus $\mathsf{Shv}(\mathsf{Top}, \mathsf{Set}) \subseteq \mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})$. Hence we have full faithul inclusions

$$\mathsf{Top} \subseteq \mathsf{Shv}(\mathsf{Top}) \subseteq \mathsf{Stk}(\mathsf{Top}), \tag{1.4.8.1}$$

and we will also regard stacks as 'generalized topological spaces'.

Stacks which are not sheaves (such as any moduli functor where the objects being parameterized have automorphisms, e.g. (1.4.8)) are certainly not representable. Nevertheless, there are other weaker properties that such stacks may satisfy which make them 'close' to representable, hence essentially geometric objects.

Recall that the 2-Yoneda lemma implies that the natural map $\mathrm{Hom}(\mathrm{Hom}(-, X), F) \to F(X)$ given by evaluation at $\mathbf{1}_X$ is an equivalence.

1.4.9 Definition (Points of stacks). A *point* x of a stack X is a map $x : * \to X$, i.e. it is an object $x \in X(*)$ (also simply written $x \in X$). Such a point has an *automorphism group* denoted $\mathrm{Aut}(x)$ (its automorphism group in the groupoid $X(*)$).

1.4.10 Exercise (Automorphism stack $\underline{\mathrm{Aut}}$). Given a point $x : * \to X$, let $\underline{\mathrm{Aut}}(x) = * \times_X *$. Show that the points of $\underline{\mathrm{Aut}}(x)$ all have trivial automorphism group and are in natural bijection with $\mathrm{Aut}(x)$. Show that $\underline{\mathrm{Aut}}(x) = * \times_{X\times X} X$, and conclude that if $X \to X \times X$ satisfies $\mathcal{P}$ then so does $\underline{\mathrm{Aut}}(x) \to *$, for any property $\mathcal{P}$ preserved under pullback. More generally, show that for any map $f : X \to Y$, the induced map $\underline{\mathrm{Aut}}(x) \to \underline{\mathrm{Aut}}(f(x))$ is a pullback of the relative diagonal of $X \to Y$, hence satisfies any property preserved under pullback which holds for $X \to X \times_Y X$.

1.5 Representable maps

1.5.1 Definition (Representable maps)**.** A map of stacks $X \to Y$ is called *representable* iff $X \times_Y U$ is representable for every map $U \to Y$ from a topological space U.

1.5.2 Exercise. Show that representability is preserved under pullback and composition.

1.5.3 Definition (Properties of representable maps)**.** Let $\mathcal{P}$ be a property of morphisms of topological spaces which is preserved under pullback. A representable map of stacks $X \to Y$ is said to have $\mathcal{P}$ iff for every map $U \to Y$ from a topological space U, the pullback map $X \times_Y U \to U$ has $\mathcal{P}$.

1.5.4 Exercise. Show that if $\mathcal{P}$ is a property of morphisms of topological spaces which is preserved under pullback, then the induced property for maps of stacks is also preserved under pullback. Show that if $\mathcal{P}$ preserved under composition, then so is the induced property for maps of stacks.

1.5.5 Definition (Maps admitting local sections)**.** A map of stacks $X \to Y$ is said to *admit local sections* iff for every map $U \to Y$ from a topological space U, there exists an open cover $U = \bigcup_i U_i$ so that each restriction $U_i \to Y$ lifts to X (note that this agrees with (1.5.3) when $X \to Y$ is representable).

1.5.6 Exercise. Show that admitting local sections is preserved under pullback and composition.

1.5.7 Definition (Local properties)**.** A property $\mathcal{P}$ of morphisms of stacks which is preserved under pullback is called *local on the target* iff for every map $X \to Y$ and every map $U \to Y$ admitting local sections, the pullback $X \times_Y U \to U$ has $\mathcal{P}$ iff $X \to Y$ has $\mathcal{P}$.

1.5.8 Lemma (Representability is a local property)**.** *If $G \to Y$ admits local sections, then $X \to Y$ is representable iff $X \times_Y G \to G$ is representable.*

Proof. By pulling back everything under a map from a topological space to Y, we may assume without loss of generality that Y is representable. Since Y is representable and $G \to Y$ admits local sections, there exists an open cover $Y = \bigcup_i U_i$ such that each $X \times_Y U_i \to U_i$ is representable, hence $X \times_Y U_i$ is representable. Gluing these topological spaces together over their common intersections $X \times_Y (U_i \cap U_j)$ yields a topological space representing X. □

1.5.9 Exercise (Descent for local properties)*.* Let $\mathcal{P}$ be a property of morphisms of topological spaces which is preserved under pullback and local on the target. Show that the induced property of morphisms of stacks (1.5.3) is local on the target.

1.5.10 Exercise. Let X be a stack. A subset $E \subseteq |X(*)|$ ($|\cdot|$ denotes isomorphism classes) determines an assignment to every map $f : Z \to X$ of a subset $Z_{E,f} \subseteq Z$ which is compatible with pullback in the sense that $Z'_{E,f\circ g} = g^{-1}(Z_{E,f})$. Show that this defines a bijection between subsets of $|X|$ and pullback compatible assignments of subsets of Z to maps $Z \to X$.

1.5.11 Exercise (Classification of embedded substacks). Let X be a stack, let $E \subseteq |X(*)|$ be any subset, and let X_E denote the stack for which a map $Z \to X_E$ is a map $Z \to X$ whose specialization to every point of Z lies in E. Show that for $f : Z \to X$, the natural diagram

$$\begin{array}{ccc} Z_{E,f} & \longrightarrow & X_E \\ \downarrow & & \downarrow \\ Z & \xrightarrow{f} & X \end{array} \tag{1.5.11.1}$$

is a pullback square. Conclude that $X_E \to X$ is an embedding (1.3.8.2) (in particular, it is representable) and that $X_E \to X$ satisfies a property $\mathcal{P}$ preserved under pullback iff every $Z_{E,f} \subseteq Z$ satisfies $\mathcal{P}$. Moreover, show that every embedding $X' \to X$ is uniquely isomorphic to an $X_E \to X$.

1.5.12 Exercise (Coarse space). Given a stack X, define its *coarse space* $|X|$ to be the topological space whose points are the isomorphism classes of points of X, and for which $U \subseteq |X|$ is open iff the associated embedded substack of X is an open substack (i.e. for every map $Z \to X$ from a topological space Z, the locus of points $z \in Z$ mapping to U is an open subset). Show that $X \mapsto |X|$ is a functor $\mathsf{Stk}(\mathsf{Top}) \to \mathsf{Top}$ which is left adjoint to the inclusion $\mathsf{Top} \hookrightarrow \mathsf{Stk}(\mathsf{Top})$.

1.5.13 Exercise. Show that if $X \to Y$ is open then $|X| \to |Y|$ is open. Show that if $X \to Y$ is universally closed then $|X| \to |Y|$ is closed.

1.6 Topological stacks

1.6.1 Definition (Atlas)**.** An *atlas* on a stack X is a topological space U and a representable map $U \to X$ admitting local sections.

1.6.2 Definition (Topological stack)**.** A stack is called *topological* iff it admits an atlas.

We shall see soon give another characterization of topological stacks: they are precisely those stacks which admit a presentation by a topological groupoid. In fact, topological groupoid presentations correspond precisely to atlases.

1.6.3 Lemma. *If X is topological, then every map $Z \to X$ from a topological space Z is representable.*

Proof. Let $U \to X$ be an atlas. Since $U \to X$ is representable, the fiber product $Z \times_X U$ is representable, so since U is a topological space, the map $Z \times_X U \to U$ is representable. Since $U \to X$ admits local sections, it follows that $Z \to X$ is representable by descent. □

1.6.4 Definition (Topological groupoid)**.** A *topological groupoid* $M \rightrightarrows O$ consists of two topological spaces O ('objects') and M ('morphisms') together with continuous maps $s, t : M \to O$ ('source' and 'target'), $e : O \to M$ ('identity'), $i : M \to M$ ('inverse'), and $c : M \times_O M \to M$ ('composition') satisfying the axioms of a groupoid: $c(c \times \mathbf{1}) = c(\mathbf{1} \times c)$ ('associativity of composition'), $se = te = \mathbf{1}_O$

1.6.5 Exercise (Stack associated to a topological groupoid). Given a topological groupoid $M \rightrightarrows O$, we define a stack $[M \rightrightarrows O] : \mathsf{Top}^{\mathsf{op}} \to \mathsf{Grpd}$ as follows. A map $U \to [M \rightrightarrows O]$ consists of the data of an open cover $\{U_i \subseteq U\}_i$ and collections of maps $\{\alpha_i : U_i \to O\}_i$ and $\{\beta_{ij} : U_i \cap U_j \to M\}_{ij}$ such that $s\beta_{ij} = \alpha_i|_{U_i \cap U_j}$, $t\beta_{ij} = \alpha_j|_{U_i \cap U_j}$, $\beta_{ii} = \mathbf{1}_{\alpha_i}$ and $\beta_{ij}\beta_{jk}\beta_{ki} = \mathbf{1}_{\alpha_i}|_{U_i \cap U_j \cap U_k}$. An isomorphism between two such maps $(U_i, \{\alpha_i\}_i, \{\beta_{ij}\}_{ij}) \xrightarrow{\sim} (U_{i'}, \{\alpha_{i'}\}_{i'}, \{\beta_{i'j'}\}_{i'j'})$ consists of maps $\gamma_{ii'} : U_i \cap U_{i'} \to M$ such that $s\gamma_{ii'} = \alpha_i|_{U_i \cap U_{i'}}$, $t\gamma_{ii'} = \alpha_{i'}|_{U_i \cap U_{i'}}$, and $\beta_{ij}\gamma_{jj'} = \gamma_{ii'}\beta_{i'j'}$ over $U_i \cap U_j \cap U_{i'} \cap U_{j'}$. Show that $[M \rightrightarrows O]$ is a stack.

1.6.6 Example (Stack quotient X/G). For a topological space X equipped with an action of a topological group G, let X/G denote the stack associated to the groupoid $G \times X \rightrightarrows X$ (forget and action, respectively).

1.6.7 Definition. A $\mathcal{P}$ *atlas* is an atlas $U \to X$ with property $\mathcal{P}$.

1.7 Mapping stacks

1.7.1 Definition (Mapping sheaf $\underline{\mathrm{Hom}}(X, Y)$)**.** For topological spaces X and Y, the sheaf $\underline{\mathrm{Hom}}(X, Y)$ is defined by declaring a map $Z \to \underline{\mathrm{Hom}}(X, Y)$ to be a continuous map $Z \times X \to Y$.

1.7.2 Example. The set of maps $* \to \underline{\mathrm{Hom}}(X, Y)$ is the set $\mathrm{Hom}(X, Y)$ of continuous maps $X \to Y$.

1.7.3 Exercise. Show that $\underline{\mathrm{Hom}}(X, Y)$ is a sheaf on Top.

1.7.4 Exercise. Show that the natural map $\underline{\mathrm{Hom}}(X, Y \times Y') \to \underline{\mathrm{Hom}}(X, Y) \times \underline{\mathrm{Hom}}(X, Y')$ is an isomorphism.

1.7.5 Exercise. Show that there is a tautological 'evaluation' map $X \times \underline{\mathrm{Hom}}(X, Y) \to Y$.

1.7.6 Definition. A *condition* $\mathcal{C}$ on morphisms $X \to Y$ is a subset $\mathrm{Hom}_{\mathcal{C}}(X, Y) \subseteq \mathrm{Hom}(X, Y)$. A condition $\mathcal{C}$ determines a sheaf $\underline{\mathrm{Hom}}_{\mathcal{C}}(X, Y)$ sending Z to the continuous maps $Z \times X \to Y$ whose specialization to every $z \in Z$ lies in $\mathrm{Hom}_{\mathcal{C}}(X, Y)$.

Recall (1.5.10) that a condition on morphisms $X \to Y$ is the same as an assignment to every map $f : Z \to \underline{\mathrm{Hom}}(X, Y)$ of a subset $Z_{\mathcal{C},f} \subseteq Z$ which is compatible with pullback in the sense that $Z_{\mathcal{C},f \circ g} = g^{-1}(Z'_{\mathcal{C},f})$. Recall also (1.5.11) that $\underline{\mathrm{Hom}}_{\mathcal{C}}(X, Y) \to \underline{\mathrm{Hom}}(X, Y)$ is an embedding and satisfies a property $\mathcal{P}$ preserved under pullback iff every $Z_{\mathcal{C},f} \subseteq Z$ satisfies $\mathcal{P}$. We will say this condition as "$\mathcal{C}$ is $\mathcal{P}$".

1.7.7 Lemma. *For $x \in X$ and $V \subseteq Y$ closed, the condition $f(x) \in V$ is closed.*

Proof. For $F : Z \times X \to Y$, the relevant subset of Z is $F|_{Z \times \{x\}}^{-1}(V)$. □

1.7.8 Exercise. Show that for any subset $A \subseteq X$ and any closed set $V \subseteq Y$, the condition $f(A) \subseteq V$ is closed.

1.7.9 Exercise. Show that if X is Hausdorff and $A \subseteq X$ is any subset, then the condition $f|_A = \mathbf{1}_A$ is closed.

1.7.10 Lemma. *For $K \subseteq X$ compact and $U \subseteq Y$ open, the condition $f(K) \subseteq U$ is open.*

Proof. Equivalently, we show that the condition $f(K) \cap V \neq \varnothing$ is closed for $V \subseteq Y$ closed. This condition may be alternatively stated as $f^{-1}(V) \cap K \neq \varnothing$. Given a map $F : Z \times X \to Y$, the subset $S_F \subseteq Z$ of maps satisfying this condition is the projection of $F|_{Z\times K}^{-1}(V)$ to Z. The inverse image $F|_{Z\times K}^{-1}(V)$ is closed, so its projection to Z is closed since $K \to *$ is universally closed (1.3.13.1). □

1.7.11 Lemma. *The diagonal of* $\underline{\mathrm{Hom}}(X, Y)$ *is an embedding* (1.3.8.2).

Proof. The diagonal of $\underline{\mathrm{Hom}}(X, Y)$ is the map $\underline{\mathrm{Hom}}(X, Y) \to \underline{\mathrm{Hom}}(X, Y \times Y)$ (1.7.4). Since $Y \to Y \times Y$ is an embedding, $\underline{\mathrm{Hom}}(X, Y) = \underline{\mathrm{Hom}}_{\mathcal{C}}(X, Y \times Y)$ where $\mathcal{C}$ is the condition of having image contained in the diagonal. The inclusion of the subsheaf of maps satisfying any condition $\mathcal{C}$ is an embedding (1.5.11). □

1.7.12 Exercise. Show that if Y is Hausdorff, then $\underline{\mathrm{Hom}}(X, Y)$ is separated (i.e. its diagonal is a closed embedding) (use (1.7.8) and the proof of (1.7.11)).

We now turn to representability of $\underline{\mathrm{Hom}}(X, Y)$ itself. As remarked earlier, the set of maps $* \to \underline{\mathrm{Hom}}(X, Y)$ is simply the set of maps $X \to Y$. It follows that $\underline{\mathrm{Hom}}(X, Y)$ is representable iff there is a topology $\mathcal{T}$ on the set $\mathrm{Hom}(X, Y)$ such that a map $Z \times X \to Y$ is continuous iff the induced map $Z \to \mathrm{Hom}(X, Y)_{\mathcal{T}}$ is continuous.

1.7.13 Definition (Compact-open topology)**.** The compact-open topology on the set $\mathrm{Hom}(X, Y)$ is the topology generated by declaring to be open the locus of maps $f : X \to Y$ satisfying $f(K) \subseteq U$ for all compact sets $K \subseteq X$ and open sets $U \subseteq Y$. The resulting topological space is denoted $\mathrm{Hom}(X, Y)_{\mathrm{cptopen}}$.

For any map $Z \to \underline{\mathrm{Hom}}(X, Y)$, the induced map $Z \to \mathrm{Hom}(X, Y)_{\mathrm{cptopen}}$ is continuous (1.7.10). There is thus a tautological map $\underline{\mathrm{Hom}}(X, Y) \to \mathrm{Hom}(X, Y)_{\mathrm{cptopen}}$. Hence if $\underline{\mathrm{Hom}}(X, Y)$ is representable, necessarily by $\mathrm{Hom}(X, Y)_{\mathcal{T}}$ for some topology $\mathcal{T}$, then $\mathcal{T}$ is at least as fine as the compact-open topology.

1.7.14 Definition. A topological space is called *locally compact* iff every point has a fundamental system of compact neighborhoods.

1.7.15 Proposition. *If X is locally compact, then the map* $\underline{\mathrm{Hom}}(X, Y) \to \mathrm{Hom}(X, Y)_{\mathrm{cptopen}}$ *is an isomorphism. In particular,* $\underline{\mathrm{Hom}}(X, Y)$ *is representable.*

Proof. We are to show that if $Z \to \mathrm{Hom}(X, Y)_{\mathrm{cptopen}}$ is continuous, then the resulting map $Z \times X \to Y$ is also continuous. What we should show is that if (z, x) is sent inside an open set $U \subseteq Y$, then a neighborhood of (z, x) is as well. Since X is locally compact, there is a compact neighborhood $K \subseteq X$ such that $z \times K$ is sent inside U. Since $Z \to \underline{\mathrm{Hom}}(X, Y)$ is continuous in the compact-open topology, there is an open set $V \subseteq Z$ such that $V \times K$ is sent inside U. Thus we are done, noting that K is a neighborhood of x. □

1.7.16 Lemma. *If X is locally compact and second countable and Y is second countable, then $\underline{\mathrm{Hom}}(X,Y)$ is second countable.*

Proof. This follows immediately from the definition of the compact-open topology, once we observe that since X is locally compact (has a compact basis) and second countable (has a countable open basis), it has a compact countable basis (proof: for every U, V in the countable open basis, choose a compact K satisfying $U \subseteq K \subseteq V$ if such exists; the collection of all such K is countable and is a basis). □

1.7.17 Lemma. *If X is σ-locally compact and Y is metrizable, then $\underline{\mathrm{Hom}}(X,Y)$ is metrizable.*

Proof. Let d_Y be a metric on Y, and let $V_1 \subseteq V_2 \subseteq \cdots X$ be an exhaustion by compact subsets (meaning every compact subset of X is contained in some V_i). Then

$$d(f,g) := \sum_{i=1}^{\infty} \min(2^{-i}, \sup_{x \in V_i} d_Y(f(x), g(x))) \tag{1.7.17.1}$$

is a metric on $\mathrm{Hom}(X,Y)$. It suffices to check that the topology it induces coincides with the compact-open topology.

Fix a map $f : X \to Y$. A neighborhood of $f \in \mathrm{Hom}(X,Y)_{\mathrm{cptopen}}$ is given by a finite collection of compact sets $K_1, \ldots, K_n \subseteq X$ and open sets $U_1, \ldots, U_n \subseteq Y$ and consists of those maps satisfying $g(K_i) \subseteq U_i$. Since $f(K_i) \subseteq U_i$ and $f(K_i)$ is compact, it follows that $\inf_{y \notin U_i} d_Y(f(K_i), y) > 0$. This implies that there exists $\delta > 0$ such that $d(f,g) < \delta$ implies $g(K_i) \subseteq U_i$.

Conversely, let us show that any neighborhood of $f : X \to Y$ with respect to d contains some neighborhood of f in the compact-open topology. A neighborhood of f with respect to d is the locus of maps g satisfying $\sup_{x \in V} d_Y(f(x), g(x)) < \delta$ for some $\delta > 0$ and some compact set $V \subseteq X$. Each $x \in V$ has a compact neighborhood $W_x \subseteq V$ over which f lands inside the open ball of radius $\frac{1}{3}\delta$ around $f(x)$. Cover V by finitely many of these neighborhoods $W_{x_1}, \ldots, W_{x_n}$, and consider the locus of maps $g : X \to Y$ for which $g(W_{x_i})$ is sent inside the open ball of radius $\delta/2$ around $f(x_i)$. For any such g and any $x \in V$, we have $x \in W_{x_i}$ for some i, so $d_Y(f(x), g(x)) \leq d_Y(f(x), f(x_i)) + d_Y(f(x_i), g(x)) < \frac{1}{3}\delta + \frac{1}{3}\delta = \frac{2}{3}\delta$. □

The basic mapping sheaf $\underline{\mathrm{Hom}}(-,-)$ admits several important generalizations, such as the sheaf parameterizing sections of a fixed map $E \to X$ or maps between fibers $X_b \to Y_b$ of maps $X, Y \to B$. Here is the most general notion we will consider.

1.7.18 Definition. Let B be a sheaf, and let $E \to X \to B$ be representable maps. We let $\underline{\mathrm{Sec}}(E/X/B)$ denote the sheaf in which a map $Z \to \underline{\mathrm{Sec}}(E/X/B)$ is a map $Z \to B$ together with a section of $E \times_B Z \to X \times_B Z$ over Z. That is, a map $Z \to \underline{\mathrm{Sec}}(E/X/B)$ is a diagram

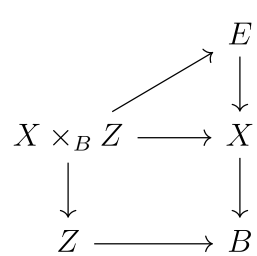

(1.7.18.1)

The universal section $\mathrm{ev} : \underline{\mathrm{Sec}}(E/X/B) \times_B X \to E$ is called the evaluation map.

1.7.19 Example. A point of $\underline{\mathrm{Sec}}(E/X/B)$ is a point $b \in B(*)$ together with a section of $E_b \to X_b$.

It is evident that a diagram

$$\begin{array}{ccc} E & \longrightarrow & E' \\ \downarrow & & \downarrow \\ X & \longrightarrow & X' \\ \downarrow & & \downarrow \\ B & \longrightarrow & B' \end{array} \tag{1.7.19.1}$$

induces a map $\underline{\mathrm{Sec}}(E/X/B) \to \underline{\mathrm{Sec}}(E'/X'/B')$. The following induced maps are evidently isomorphisms

$$\underline{\mathrm{Sec}}((E \times_B B')/(X \times_B B')/B') \xrightarrow{\sim} \underline{\mathrm{Sec}}(E/X/B) \times_B B' \tag{1.7.19.2}$$

$$\underline{\mathrm{Sec}}((E \times_X E')/X/B) \xrightarrow{\sim} \underline{\mathrm{Sec}}(E/X/B) \times_B \underline{\mathrm{Sec}}(E'/X/B) \tag{1.7.19.3}$$

It follows that the diagonal of $\underline{\mathrm{Sec}}(E/X/B) \to B$ is an embedding (since $E \to E \times_X E$ is an embedding), hence in particular representable.

1.7.20 Exercise. Let $s : B \to X$ be a section, and let $F \subseteq s^*E := E \times_X B$ be a closed substack. Show that the condition on $\underline{\mathrm{Sec}}(E/X/B)$ of sending s to F is a closed condition.

1.7.21 Lemma. *Fix $E \to X \to B$. Let $X_0 \to X$ be any map such that $X_0 \to B$ is open, and let $E_0 \subseteq E \times_X X_0$ be a closed substack. The condition on $\underline{\mathrm{Sec}}(E/X/B)$ of mapping X_0 to E_0 is closed.*

Proof. In (1.7.18.1), we are to pull back to X_0, take inverse image of E_0, which is a closed substack of the pullback $X_0 \times_B Z$, and take the locus of points of z whose fibers are contained in this subset. In other words, we take complement (which is open), pushforward (preserves open sets since $X_0 \to B$ is open), and take complement again (giving us a closed subset of Z). □

1.7.22 Example. If $E' \subseteq E$ is a closed substack and $X \to B$ is open, then $\underline{\mathrm{Sec}}(E'/X/B) \to \underline{\mathrm{Sec}}(E/X/B)$ is a closed embedding.

1.7.23 Exercise. Show that if $E \to X$ is separated and $X \to B$ is open, then the diagonal of $\underline{\mathrm{Sec}}(E/X/B) \to B$ is a closed embedding.

1.7.24 Exercise. Find an example where $E \to X$ is separated but $\underline{\mathrm{Sec}}(E/X/B) \to B$ is not.

1.8 Stability

1.8.1 Definition (Stable diagonal)**.** A *stable diagonal* of a stack X is a factorization of its diagonal $X \to X \times X$ into an open embedding $X \to \widehat{X}$ and a proper map $\widehat{X} \to X \times X$.

1.8.2 Exercise. Let $X \to \widehat{X} \to X \times X$ be a stable diagonal of X. Show that for any embedded substack $A \subseteq X$, the factorization

$$A \to A \times_X \widehat{X} \times_X A \to A \times A \tag{1.8.2.1}$$

is a stable diagonal of A. More generally, show that for any map $A \to X$ whose diagonal is a closed embedding, a stable diagonal of A is given by taking

$$\widehat{A} = \big(A \times_X \widehat{X} \times_X A\big) \setminus \big((A \times_X A) \setminus A\big) \tag{1.8.2.2}$$

where the difference operations $\setminus$ indicate taking the complement of an open or closed substack.

1.8.3 Definition (Stable locus)**.** A point x of a stack X is called *stable* iff $\underline{\mathrm{Aut}}(x)$ is proper. The embedded substack of X spanned by its stable points is denoted $X^s \subseteq X$.

1.8.4 Example. Consider the quotient stack $X = \widehat{\mathbb{C}}/\mathbb{C}^\times$. Its coarse space is the topological quotient, which has three points, namely the orbits $\{0\}$, $\{\infty\}$, and $\mathbb{C}^\times$. The diagonal of X is not proper (it pulls back under an atlas to $\widehat{\mathbb{C}} \times \mathbb{C}^\times \to \widehat{\mathbb{C}} \times \widehat{\mathbb{C}}$). The point of X corresponding to the orbit $\mathbb{C}^\times$ has trivial automorphism group, so is stable. The points of X corresponding to the orbits $\{0\}$ and $\{\infty\}$ have automorphism group $\mathbb{C}^\times$, so are not stable. Thus $X^s \subseteq X$ is an open substack, and it is isomorphic to the topological space $*$. In particular, it is separated, in contrast to X.

On its own, this notion of stability (1.8.3) is not so useful or well behaved. It becomes so only under additional assumptions, such as the existence of a stable diagonal with certain properties.

1.8.5 Proposition. *Let X be a topological stack with stable diagonal $X \to \widehat{X} \to X \times X$. Suppose that every point of $\widehat{X} \setminus X$ has non-proper automorphism group. Then $X^s \subseteq X$ is an open substack and has proper diagonal.*

Proof. Consider the proper map $\widehat{X} \setminus X \to X \times X$ and pull back under the diagonal $X \to X \times X$ to obtain a proper map

$$X \times_{X \times X} (\widehat{X} \setminus X) \to X. \tag{1.8.5.1}$$

The image of this map is thus a closed substack of X (1.5.13). If $x \in X$ is not in im (1.8.5.1), then $\underline{\mathrm{Aut}}(x) = * \times_{X \times X} X = * \times_{X \times X} \widehat{X}$ is proper, so $x \in X^s$. Conversely, we claim that X^s is disjoint from im (1.8.5.1). In fact, we claim that $X^s \times X^s \subseteq X \times X$ is disjoint from the image of $\widehat{X} \setminus X \to X \times X$. Since this map is proper, it has proper diagonal, so if $\hat{x} \in \widehat{X}$ maps to $(x, x') \in X \times X$, the map $\underline{\mathrm{Aut}}(\hat{x}) \to \underline{\mathrm{Aut}}(x) \times \underline{\mathrm{Aut}}(x')$ is proper. If $(x, x') \in X^s \times X^s$, the target is proper, so $\underline{\mathrm{Aut}}(\hat{x})$ is also proper, hence $\hat{x} \in X$ as we have assumed that otherwise $\underline{\mathrm{Aut}}(\hat{x})$ is non-proper.

We have thus shown that X^s is the complement of the closed substack im (1.8.5.1), so is an open substack. We have also show that $\widehat{X} \times_{X \times X} (X^s \times X^s) = X \times_{X \times X} (X^s \times X^s)$. Thus $X^s = X^s \times_X \widehat{X} \times_X X^s$, which means the diagonal of X^s is a pullback of $\widehat{X} \to X \times X$, hence is proper. □

1.9 Grothendieck topologies

A Grothendieck topology is a very general structure which provides a 'notion of locality' which is sufficient to make sense of the definition of a sheaf.

1.9.1 Definition (Sieve)**.** Let $\mathcal{C}$ be a category. A *sieve* S on an object $c \in \mathcal{C}$ is a collection of morphisms $x \to c$ which is closed under pre-composition (meaning if $x \to c$ is in S then so is every composition $y \to x \to c$).

A sieve on $c \in \mathcal{C}$ may be regarded as specifying which maps to c have 'small image'.

1.9.2 Example. The following conditions on morphisms to a set X all define sieves: (0) all maps, (1) image consists of at most one point, (2) image consists of at most finitely many points, (3) not surjective, (4) image is contained in a given subset $X' \subseteq X$, (5) image is contained in at least one of a given collection of subsets $X_i \subseteq X$, (6) no maps.

1.9.3 Exercise. Show that the following conditions on morphisms to a set X do not define sieves: (1) injective, (2) surjective, (3) bijective.

1.9.4 Definition (Grothendieck topology)**.** Let $\mathcal{C}$ be a category. A *Grothendieck topology* J on $\mathcal{C}$ specifices for every object $c \in \mathcal{C}$ a set $J(c)$ of sieves on c (called *covering sieves*) satisfying the following properties:

(1.9.4.1) (Identity) The sieve of all maps to X is a covering sieve.

(1.9.4.2) (Stability under pullback) For any morphism $X \to Y$ and any covering sieve S on Y, its pullback to X (consisting of all those maps $Z \to X$ whose composition with $X \to Y$ lies in S) is a covering sieve.

(1.9.4.3) (Locality) For any covering sieve S on X, a sieve T on X is a covering sieve iff its pullback to every element of S is a covering sieve.

A *site* is a category with a Grothendieck topology $(\mathcal{C}, J)$.

1.9.5 Exercise. Show that in any Grothendieck topology, if $S \subseteq S'$ and S is a covering sieve then so is S'. Show that in any Grothendieck topology, the intersection of two covering sieves (on the same object) is a covering sieve.

1.9.6 Exercise (Surjective topology on Set). Declare a sieve S on $X \in \mathsf{Set}$ to be a covering sieve iff every map $* \to X$ is in S. Show that this is a Grothendieck topology.

1.9.7 Exercise (Usual topology on $\mathsf{Op}(X)$). Let X be a topological space. The category $\mathsf{Op}(X)$ consists of open subsets $U \subseteq X$ and inclusions. Declare a sieve S on U to be a covering sieve iff $\bigcup_{V \in S} V = U$. Show that this is a Grothendieck topology.

1.9.8 Definition (Partition of unity)**.** Let X be a topological space. A *partition of unity* on X is a collection of functions $\varphi_i : X \to [0,1]$ which is locally finite (every point of X has a neighborhood over which all but finitely many of the φ_i are identically zero) and satisfies $\sum_i \varphi_i \equiv 1$. A partition of unity *subordinate* to an open cover $X = \bigcup_i U_i$ is one indexed by the same indexing set such that $\operatorname{supp} \varphi_i \subseteq U_i$, where $\operatorname{supp} \varphi$ (the *support*) is the complement of the largest open subset of X over which φ is identically zero.

1.9.9 Exercise (Numerable topology on $\mathsf{Op}(X)$ [3]). Let X be a topological space. Declare a sieve S on U to be a covering sieve iff $\bigcup_{V \in S} V = U$ and there exists a partition of unity subordinate to S. Show that this is a Grothendieck topology.

Chapter 2

Algebraic topology

2.1 Simplicial sets

2.1.1 Definition (Simplex category $\mathbf{\Delta}$)**.** Let $\mathbf{\Delta}$ denote the category whose objects are non-empty finite ordered sets and whose morphisms are weakly order preserving maps. In other words, every object of $\mathbf{\Delta}$ is isomorphic to $[n] = \{0, \ldots, n\}$ for some integer $n \geq 0$, and a morphism $f : [n] \to [m]$ is a map of sets satisfying $f(i) \leq f(j)$ for $i \leq j$.

2.1.2 Example. There is a functor $\mathbf{\Delta} \to \mathsf{Top}$ sending a finite ordered set S to the complete simplex on S and a map $S \to T$ to the induced affine linear map of simplices.

2.1.3 Example. There is a functor $\mathbf{\Delta} \to \mathsf{Cat}$ which regards a finite ordered set S as a category whose objects are elements $s \in S$ and whose morphisms are pairs $s \leq s'$. This functor is fully faithful.

2.1.4 Definition (Simplicial objects)**.** For any category $\mathcal{C}$, a simplicial object of $\mathcal{C}$ is a functor $\Delta^{\mathsf{op}} \to \mathcal{C}$. Simplicial objects of $\mathcal{C}$ form a category denoted $\mathsf{s}\mathcal{C} = \mathsf{Fun}(\Delta^{\mathsf{op}}, \mathcal{C})$. (A simplicial object of $\mathcal{C}$ is, despite the terminology, evidently not an object of $\mathcal{C}$.)

2.1.5 Exercise. Show that sending $X \in \mathcal{C}$ to the constant functor $\Delta^{\mathsf{op}} \to * \xrightarrow{X} \mathcal{C}$ defines a fully faithful functor $\mathcal{C} \hookrightarrow \mathsf{s}\mathcal{C}$.

2.1.6 Definition (Simplicial sets)**.** The category of simplicial sets is $\mathsf{sSet} = \mathsf{Fun}(\Delta^{\mathsf{op}}, \mathsf{Set})$.

2.1.7 Example. The Yoneda functor of $\mathbf{\Delta}$ is an embedding $\mathbf{\Delta} \to \mathsf{sSet}$, and the image of $[n]$ under this embedding is denoted $\Delta^n_\bullet$. For any simplicial set X, the Yoneda embedding identifies elements of $X([n])$ with maps $\Delta^n_\bullet \to X$; these are called the 'n-simplices of X', also denoted X_n. One should view a simplicial set as a combinatorial/categorical specification of a way of 'gluing' together simplices along simplicial maps (in fact, this assertion can be made precise categorically: $\mathsf{sSet} = \mathsf{Fun}(\mathbf{\Delta}^{\mathsf{op}}, \mathsf{Set})$ is the free cocompletion of $\mathbf{\Delta}$).

2.1.8 Exercise. What are the k-simplices of $\Delta^1_\bullet$? (There are precisely $k+2$ of them.)

2.1.9 Exercise. Show that a map of simplicial sets is a monomorphism iff it induces an injection on k-simplices for every $k \geq 0$. Show that a map of simplicial sets is an epimorphism iff it induces a surjection on k-simplices for every $k \geq 0$.

2.1.10 Definition (Geometric realization)**.** The simplex functor $\mathbf{\Delta} \to \mathsf{Top}$ (2.1.2) admits a unique cocontinuous extension $|\cdot| : \mathsf{sSet} \to \mathsf{Top}$ called geometric realization. Concretely, the geometric realization is the colimit of $\Delta^n \in \mathsf{Top}$ over all maps $\Delta^n_\bullet \to X$ (i.e. the colimit of the composition $\mathbf{\Delta}_{/X} \to \mathbf{\Delta} \to \mathsf{Top}$).

2.1.11 Definition (Singular simplicial set)**.** The right adjoint to geometric realization is called the singular simplicial set functor $S_* : \mathsf{Top} \to \mathsf{sSet}$; it is defined by the property that $\operatorname{Hom}([n], S_*X) = \operatorname{Hom}(\Delta^n, X)$ (continuous maps of topological spaces).

Since geometric realization is a right adjoint, it preserves colimits. Singular simplicial set, being a left adjoint, preserves limits.

2.1.12 Definition (Non-degenerate simplices [4, (8.3)])**.** Let X be a simplicial set. Every simplex $\Delta^n \to X$ admits a unique factorization $\Delta^n \to \Delta^r \to X$ with minimal r. When $r = n$ we say $\Delta^n \to X$ is *non-degenerate*.

Proof. Suppose we have factorizations

$$\begin{array}{ccc} \Delta^n & \longrightarrow & \Delta^r \\ \downarrow & & \downarrow \\ \Delta^s & \longrightarrow & X \end{array} \tag{2.1.12.1}$$

We may assume without loss of generality that the two maps $\Delta^n \to \Delta^r$ and $\Delta^n \to \Delta^s$ are surjective. Since they are surjective, these maps just ‘collapse’ certain arrows in the quiver $0 \to \cdots \to n$. Collapsing the union of these sets of arrows defines a surjection $\Delta^n \to \Delta^q$ fitting into a square

$$\begin{array}{ccc} \Delta^n & \longrightarrow & \Delta^r \\ \downarrow & & \downarrow \\ \Delta^s & \longrightarrow & \Delta^q \end{array} \tag{2.1.12.2}$$

and it suffices to show that this square is a pushout of simplicial sets (the universal property then gives the desired factorization of the map to X). Since simplicial sets is a presheaf category, a diagram is a pushout iff $\operatorname{Hom}(\Delta^k, \cdot)$ sends it to a pushout of sets for every k. Thus we should show that

$$\begin{array}{ccc} \operatorname{Hom}([k],[n]) & \longrightarrow & \operatorname{Hom}([k],[r]) \\ \downarrow & & \downarrow \\ \operatorname{Hom}([k],[s]) & \longrightarrow & \operatorname{Hom}([k],[q]) \end{array} \tag{2.1.12.3}$$

is a pushout of sets. Each of the maps in (2.1.12.2) has a section, which means the same for (2.1.12.3), so each of these maps is also surjective. Fix sections of the maps in (2.1.12.2) which send a given object to the smallest element in its inverse image. Composing $\mathrm{Hom}([k],[n]) \to \mathrm{Hom}([k],[r])$ with this section gives an endomorphism A of $\mathrm{Hom}([k],[n])$, and similarly for $\mathrm{Hom}([k],[n]) \to \mathrm{Hom}([k],[s])$ we get an endomorphism B. A sufficiently high power $(AB)^N$ coincides with the endomorphism associated to the section of $\mathrm{Hom}([k],[n]) \to \mathrm{Hom}([k],[q])$. This shows that (2.1.12.3) is a pushout. □

2.1.13 Exercise (Cardinality of a simplicial set). The cardinality of a simplicial set is the cardinality of the set of its non-degenerate simplices. Show that for a simplicial set of cardinality $\kappa > 0$, the set of all its simplices has cardinality $\min(\kappa, \aleph_0)$.

2.1.14 Exercise. Show that there are exactly $\frac{(n+m)!}{n!m!}$ non-degenerate $(n+m)$-simplices in $\Delta^n \times \Delta^m$. Identify these simplices with paths from $(0,0)$ to (n,m) in the $n \times m$ unit grid.

2.1.15 Exercise (k-skeleton). For $k \geq 0$, let $\mathbf{\Delta}_{\leq k} \subseteq \mathbf{\Delta}$ denote the full subcateogry spanned by simplices of dimension $\leq k$. Show that the restriction functor $\mathsf{Fun}(\mathbf{\Delta}^{\mathsf{op}}, \mathsf{Set}) \to \mathsf{Fun}(\mathbf{\Delta}^{\mathsf{op}}_{\leq k}, \mathsf{Set})$ has a left adjoint. The composition of restriction and this left adjoint is called the k-skeleton functor $\mathrm{sk}_{\leq k} : \mathsf{sSet} \to \mathsf{sSet}$. Show that the counit map $\mathrm{sk}_{\leq k} X \to X$ is injective (use (2.1.12)).

2.1.16 Lemma. *The functor $S_* \circ |\cdot|$ preserves products.*

Proof. We should show that given continuous maps $\Delta^n \to |X|$ and $\Delta^n \to |Y|$, the induced map $\Delta^n \to |X \times Y|$ is continuous. This is a nontrivial statement since the continuous bijection $|X \times Y| \to |X| \times |Y|$ is not always a homeomorphism. It is, however, a homeomorphism, when both X and Y are finite. It thus suffices to show that every continuous map $\Delta^n \to |X|$ factors through (the realization of) a finite subcomplex $X_0 \subseteq X$. In fact, this holds for any compact space in place of Δ^n (2.1.17). □

2.1.17 Lemma. *Every map $Z \to |X|$ from a compact topological space Z factors through a finite subcomplex of X.*

Proof. Suppose not, namely suppose there is a map $Z \to |X|$ and a sequence of points $z_1, z_2, \ldots \in Z$ which are sent to the interiors of distinct simplices of $|X|$. A map $|X| \to \mathbb{R}$ is continuous iff its restriction to every simplex is continuous. Every map $\partial\Delta^n \to \mathbb{R}$ extends to Δ^n, and this extension can be chosen to take any given prescribed value at a given point in its interior. We may thus construct, by induction on skeleta, a continuous function $|X| \to \mathbb{R}$ which takes the value i at the image of z_i. The composition $Z \to |X| \to \mathbb{R}$ is thus an unbounded continuous function on a compact space, contradiction. We conclude that every continuous map $Z \to |X|$ factors through a finite subcomplex of X, as was to be shown. □

2.1.18 Definition (Horns). The ith horn $\Lambda^n_i \subseteq \Delta^n$ $(0 \leq i \leq n)$ consists of those k-simplices of Δ^n which omit at least one vertex of Δ^n other than vertex i.

2.1.19 Exercise. Draw $\Lambda^n_i \subseteq \Delta^n$ for all $n \leq 3$.

2.1.20 Definition (Kan fibration)**.** A map of simplicial sets $X \to Y$ is called a *Kan fibration* iff it satisfies the *right lifting property* for every horn (Δ^n, Λ^n_i), meaning that for every commuting diagram of solid arrows

$$\begin{array}{ccc} \Lambda^n_i & \longrightarrow & X \\ \downarrow & \nearrow & \downarrow \\ \Delta^n & \longrightarrow & Y \end{array} \tag{2.1.20.1}$$

there exists a dotted arrow making the diagram commute. A simplicial set X is called a *Kan complex* iff the map $X \to *$ is a Kan fibration, namely iff X satisfies the *extension property* for maps from (Δ^n, Λ^n_i) (alternatively, the extension property holds for maps $(\Delta^n, \Lambda^n_i) \to X$), meaning that every map $\Lambda^n_i \to X$ extends to a map $\Delta^n \to X$.

2.1.21 Exercise. Show that the property of being a Kan fibration is preserved under pullback and composition.

2.1.22 Exercise. Show that the singular simplicial set (2.1.11) of any topological space is a Kan complex.

2.1.23 Exercise. Let X be a simplicial set. Its set of components $\pi_0 X$ is its set of vertices modulo the equivalence relation closure of the relation given by the edges. Show that if X is a Kan complex, then the relation of edges ($x \sim x'$ iff there exists an edge $x \to x'$) is an equivalence relation.

2.1.24 Definition (Simplicial set pair)**.** A *simplicial set pair* (X, A) is a simplicial set X and a subcomplex $A \subseteq X$.

2.1.25 Definition. Let $\mathcal{M}$ be a set of simplicial set pairs. A simplicial set pair (X, A) is said to be *filtered by pushouts of* $\mathcal{M}$ iff there exists a filtration $A = R_0 \subseteq R_1 \subseteq \cdots$ with $X = \bigcup_{i=0}^{\infty} R_i$ such that each morphism $R_i \to R_{i+1}$ is a pushout of a coproduct of morphisms in $\mathcal{M}$, namely there is a pushout diagram

$$\begin{array}{ccc} R_i & \longrightarrow & R_{i+1} \\ \downarrow & & \downarrow \\ \bigsqcup_\alpha Y_\alpha & \longrightarrow & \bigsqcup_\alpha B_\alpha \end{array} \tag{2.1.25.1}$$

for some set of pairs $(Y_\alpha, B_\alpha) \in \mathcal{M}$.

2.1.26 Exercise. Show that if $X \to Y$ satisfies the right lifting property with respect to every pair in $\mathcal{M}$, then it satisfies the right lifting property with respect to any pair which is filtered by pushouts of $\mathcal{M}$.

2.1.27 Exercise. Use (2.1.15) to show that every simplicial set pair (X, A) is filtered by pushouts of pairs $(\Delta^k, \partial\Delta^k)$.

2.1.28 Exercise (Small object argument [22]). Let $\mathcal{M}$ be a set of *finite* (2.1.13) simplicial set pairs. Let $X \to Y$ be any map of simplicial sets. We define a sequence of inclusions $X_0 \hookrightarrow X_1 \hookrightarrow \cdots$ of simplicial sets over Y beginning with $X_0 = X$. Given $X_i \to Y$, we consider the set of all diagrams

$$\begin{array}{ccc} A & \longrightarrow & X_i \\ \downarrow & & \downarrow \\ B & \longrightarrow & Y \end{array} \tag{2.1.28.1}$$

for pairs $(B, A) \in \mathcal{M}$. We define X_{i+1} to be the pushout of $X_i \leftarrow \coprod A \to \coprod B$ where the disjoint union is over all diagrams (2.1.28.1). Let $X_\infty = \bigcup_i X_i$ be the ascending union, so $X \to X_\infty$ is, by construction, filtered by pushouts of $\mathcal{M}$. Show that $X_\infty \to Y$ satisfies the right lifting property with respect to $\mathcal{M}$. To treat the case of arbitrary $\mathcal{M}$ (its constituent simplicial set pairs not assumed to be finite), it suffices to induct over any ordinal whose cofinality is greater than the cardinality of any simplicial set in $\mathcal{M}$.

2.1.29 Exercise. Let $\mathcal{M}$ be a set of finite simplicial set pairs. Using the factorization from (2.1.28), show that if $A \to B$ satisfies the left lifting property with respect to every morphism satisfying the right lifting property with respect to $\mathcal{M}$, then $A \to B$ is a retract of a morphism filtered by pushouts of $\mathcal{M}$.

2.1.30 Definition (Product of simplicial set pairs)**.** For simplicial set pairs (X, A) and (Y, B), their product is defined to be

$$(X, A) \times (Y, B) = (X \times Y, (X \times B) \cup_{A \times B} (A \times Y)). \tag{2.1.30.1}$$

2.1.31 Lemma. *Every product $(\Delta^n, \Lambda^n_i) \times (\Delta^k, \partial\Delta^k)$ is filtered by pushouts of horns.*

2.1.32 Exercise (Simplicial mapping space). Let X and Y be simplicial sets. Define a simplicial set $\underline{\mathrm{Hom}}(X, Y)$ by the universal property that a map $Z \to \underline{\mathrm{Hom}}(X, Y)$ is a map $Z \times X \to Y$. Use (2.1.27) and (2.1.31) to show that if Y is a Kan complex then so is $\underline{\mathrm{Hom}}(X, Y)$.

2.1.33 Exercise (Homotopy category of Kan complexes $h\mathcal{S}$). For a Kan complex X and a simplicial set K, call maps $f, g : K \to X$ *homotopic* iff there exists a map $K \times \Delta^1 \to X$ whose restrictions to $K \times 0$ and $K \times 1$ coincide with f and g, respectively. Use (2.1.31) to show that homotopy is an equivalence relation on the set of maps $K \to X$. The homotopy category of Kan complexes $h\mathcal{S}$ has objects Kan complexes and morphisms homotopy classes of maps. A map of Kan complexes is called a homotopy equivalence iff it is an isomorphism in $h\mathcal{S}$.

2.1.34 Exercise. Let $X \to Y$ be a Kan fibration. Show that an edge $y \to y'$ in Y induces a canonical map $X_y \to X_{y'}$ in the homotopy category of Kan complexes $h\mathcal{S}$ (lift the pair $X_y \times (\Delta^1, 0)$). Show that for any 2-simplex in Y with vertices y, y', y'', the induced triangle in $h\mathcal{S}$ commutes. Conclude that this defines a functor $Y \to h\mathcal{S}$.

2.1.35 Definition (Trivial Kan fibration)**.** A map of simplicial sets is called a *trivial Kan fibration* iff it satisfies the right lifting property for every pair $(\Delta^n, \partial\Delta^n)$. A simplicial set is called a *trivial Kan complex* iff the map $X \to *$ is a trivial Kan fibration.

2.1.36 Exercise. Show that a trivial Kan fibration is a Kan fibration. In fact, show that a trivial Kan fibration satisfies the right lifting property for every pair (X, A) where $A \subseteq X$ is a subcomplex.

2.1.37 Exercise. Show that a trivial Kan fibration is a homotopy equivalence.

2.1.38 Exercise. Show that a Kan complex is trivial iff it is contractible (homotopy equivalent to $*$).

2.1.39 Exercise (Functor $\mathsf{sSet} \to h\mathcal{S}$). Show that if $X \hookrightarrow Y$ is filtered by pushouts of horns and Z is a Kan complex, then the map $\underline{\mathrm{Hom}}(Y, Z) \to \underline{\mathrm{Hom}}(X, Z)$ is a trivial Kan fibration. A *Kanification* is an inclusion $X \hookrightarrow \hat{X}$ which is filtered by pushouts of horns with $\hat{X}$ a Kan complex (this always exists by the small object argument (2.1.28)). Show that given any map $X \to Y$ and any choice of Kanifications $X \hookrightarrow \hat{X}$ and $Y \hookrightarrow \hat{Y}$, there exists a dotted arrow making the following diagram commute:

$$\begin{array}{ccc} X & \longrightarrow & \hat{X} \\ \downarrow & & \vdots \\ Y & \longrightarrow & \hat{Y} \end{array} \tag{2.1.39.1}$$

Show that sending X to (any choice of) $\hat{X}$ and sending a map $X \to Y$ to (any choice of) extension $\hat{X} \to \hat{Y}$ defines a functor from the category of simplicial sets sSet to the homotopy category of spaces $h\mathcal{S}$. Finally, show that this functor sends any inclusion of simplicial sets which is filtered by pushouts of horns to an isomorphism.

2.1.40 Lemma. *A Kan fibration is trivial iff its fibers are trivial.*

Proof. The set of diagrams of solid arrows

$$\begin{array}{ccc} \partial\Delta^k & \longrightarrow & X \\ \downarrow & \nearrow & \downarrow \\ \Delta^k & \longrightarrow & Y \end{array} \tag{2.1.40.1}$$

is the set of vertices of $\underline{\mathrm{Hom}}(\partial\Delta^k, X) \times_{\underline{\mathrm{Hom}}(\partial\Delta^k, Y)} \underline{\mathrm{Hom}}(\Delta^k, Y)$, and the set of such diagrams equipped with a lift is the set of vertices of $\underline{\mathrm{Hom}}(\Delta^k, X)$. The forgetful map

$$\underline{\mathrm{Hom}}(\Delta^k, X) \to \underline{\mathrm{Hom}}(\partial\Delta^k, X) \times_{\underline{\mathrm{Hom}}(\partial\Delta^k, Y)} \underline{\mathrm{Hom}}(\Delta^k, Y) \tag{2.1.40.2}$$

is a Kan fibration: lifting a pair (Δ^n, Λ^n_i) against this map is the same as lifting the product $(\Delta^n, \Lambda^n_i) \times (\Delta^k, \partial\Delta^k)$ against $X \to Y$. Since (2.1.40.2) is a Kan fibration, its image is a union of connected components of the target.

The forgetful map $\underline{\mathrm{Hom}}(\partial\Delta^k, X) \times_{\underline{\mathrm{Hom}}(\partial\Delta^k, Y)} \underline{\mathrm{Hom}}(\Delta^k, Y) \to \underline{\mathrm{Hom}}(\Delta^k, Y)$ is a pullback of $\underline{\mathrm{Hom}}(\partial\Delta^k, X) \to \underline{\mathrm{Hom}}(\partial\Delta^k, Y)$ which is a Kan fibration since $X \to Y$ is. Every map $\Delta^k \to Y$ is homotopic to a constant map, so every diagram (2.1.40.1) is homotopic to one in which the map $\Delta^k \to Y$ is constant. Lifting for such diagrams is precisely the assertion that the fibers are trivial. □

2.1.41 Lemma. *A Kan fibration is trivial iff it is a homotopy equivalence.*

Proof. Let $f : X \to Y$ be a Kan fibration, and let $g : Y \to X$ be a homotopy inverse. The composition $fg : Y \to Y$ is homotopic to $\mathbf{1}_Y$ via a homotopy $H : Y \times \Delta^1 \to Y$. This homotopy gives us a lifting problem

$$\begin{array}{ccc} Y & \xrightarrow{g} & X \\ \downarrow{\scriptstyle \times 0} & \nearrow & \downarrow{\scriptstyle f} \\ Y \times \Delta^1 & \xrightarrow{H} & Y \end{array} \tag{2.1.41.1}$$

whose solution is a homotopy from g to a section of f. We may thus assume that our homotopy inverse $g : Y \to X$ is a section, namely $fg = \mathbf{1}_Y$.

We now consider homotopies between gf and $\mathbf{1}_X$, i.e. maps $H : X \times \Delta^1 \to X$ with $H|_{X\times 0} = gf$ and $H|_{X\times 1} = \mathbf{1}_X$. We consider the equivalence relation of homotopy rel $X \times \partial\Delta^1$ on such H. The equivalence classes are thus the connected components of the fiber of $\underline{\mathrm{Hom}}(X \times \Delta^1, X) \to \underline{\mathrm{Hom}}(X \times \partial\Delta^1, X)$ over $gf \sqcup \mathbf{1}_X$. This map is a Kan fibration, so homotopy of maps $X \times \partial\Delta^1 \to X$ induces a bijection between homotopy classes of extensions to $X \times \Delta^1$. In particular, post-composition with the homotopy equivalence f induces a bijection on homotopies H modulo homotopy rel boundary. Now $fgf = f\mathbf{1}_X$, so there is a canonical constant homotopy between them, and we consider any homotopy H whose post-composition fH is homotopic rel boundary to this constant homotopy. Now the homotopy between fH and the constant homotopy can be lifted to a homotopy of H itself using the fact that $X \to Y$ is a Kan fibration, so we obtain a homotopy H such that fH is itself constant. Specializing H to the fibers of f, we concude that the fibers are contractible, hence trivial (2.1.38), so f is a trivial Kan fibration (2.1.40). □

2.2 ∞-categories

The definition of ∞-categories is due to Bordman–Vogt, and their development into a working theory is due to Joyal [14] and Lurie [17]. The subject of ∞-categories belongs to algebraic topology, specifically abstract/categorical homotopy theory as founded by Quillen [22].

2.2.1 Definition (∞-category)**.** An ∞-*category* is a simplicial set which satisfies the extension property for all *inner horns* $\Lambda^n_i \subseteq \Delta^n$ (inner means $0 < i < n$).

2.2.2 Exercise (Nerve of a category). For any category $\mathcal{C}$, its *nerve* is the simplicial set whose n-simplices are functors from the quiver $0 \to \cdots \to n$ to $\mathcal{C}$ (note that $\mathsf{Fun}([n], \mathcal{C})$ is indeed a contravariant functor of $[n]$). Show that a simplicial set is the nerve of a category iff every inner horn has a unique filling (so, in particular, the nerve of a category is a ∞-category).

2.2.3 Discussion (Objects, morphisms, and composition in an ∞-category). An *object* x of an ∞-category $\mathcal{C}$ is a vertex of $\mathcal{C}$, and a *morphism* $f : x \to y$ is an edge in $\mathcal{C}$. For an object $x \in \mathcal{C}$, its identity morphism $\mathbf{1}_x$ is the degenerate edge over x. A 2-simplex in $\mathcal{C}$ with boundary

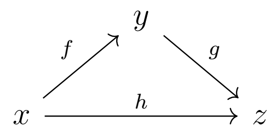

(2.2.3.1)

should be thought of as a *homotopy* between the 'composition of f and g' (which is not itself a morphism in $\mathcal{C}$, as it is not an edge) and h. A given horn Λ^2_1 typically has many different fillings to Δ^2, so we cannot call h 'the' composition of f and g (merely 'a' composition). The higher horn filling conditions do imply, however, that extending a given map $\Lambda^2_1 \to \mathcal{C}$ to Δ^2 is a contractible choice (precisely, $\mathrm{Hom}(\Delta^2, \mathcal{C}) \to \mathrm{Hom}(\Lambda^2_1, \mathcal{C})$ is a trivial Kan fibration). They also encode the data to guarantee that composition, in this sense, is associative up to coherent homotopy.

2.2.4 Definition. Given an ∞-category $\mathcal{C}$, its opposite $\mathcal{C}^{\mathsf{op}}$ is the opposite simplicial set (i.e. its precomposition with $\mathsf{op} : \boldsymbol{\Delta} \to \boldsymbol{\Delta}$).

2.2.5 Definition (Functor and diagram categories)**.** For any simplicial set (often an ∞-category) K and any ∞-category $\mathcal{C}$, denote by $\mathsf{Fun}(K, \mathcal{C})$ the simplicial mapping space; that is, an n-simplex $\Delta^n \to \mathsf{Fun}(K, \mathcal{C})$ is a map of simplicial sets $K \times \Delta^n \to \mathcal{C}$. There is an evident composition map $\mathsf{Fun}(\mathcal{A}, \mathcal{B}) \times \mathsf{Fun}(\mathcal{B}, \mathcal{C}) \to \mathsf{Fun}(\mathcal{A}, \mathcal{C})$.

2.2.6 Exercise. For any simplicial set K, and any category $\mathcal{C}$, show that $\mathsf{Fun}(K, \mathcal{C})$ is (the nerve of) the category of K-shaped diagrams in $\mathcal{C}$. Here a K-shaped diagram in $\mathcal{C}$ consists of an object of $\mathcal{C}$ for every vertex of K and for every edge of K a morphism in $\mathcal{C}$ between the objects associated to its endpoints (required to the identity when the edge is degenerate) such that the boundary of every 2-simplex commutes. A morphism of K-shaped diagrams consists of morphisms between the objects at every vertex such that the square associated to each edge commutes. In particular, when K is the nerve of a category $\mathcal{D}$, conclude that $\mathsf{Fun}(\mathcal{D}, \mathcal{C})$ is (the nerve of) the category of functors $\mathcal{D} \to \mathcal{C}$.

2.2.7 Proposition. $\mathsf{Fun}(K, \mathcal{C})$ *is an* ∞*-category for any* ∞*-category* $\mathcal{C}$*.*

Proof. It suffices to show that the ∞-category $\mathcal{C}$ satisfies the extension property for pairs $(\Delta^n, \Lambda^n_i) \times K$ with $0 < i < n$. By induction on the skeleta of K (2.1.27), it suffies to show that $\mathcal{C}$ satisfies the extension property for pairs $(\Delta^n, \Lambda^n_i) \times (\Delta^k, \partial\Delta^k)$ with $0 < i < n$ and $k \geq 0$. It thus suffices to show that for $0 < i < n$, the product $(\Delta^n, \Lambda^n_i) \times (\Delta^k, \partial\Delta^k)$ is filtered by inner horn fillings. We verify this property next (2.2.8) (stated separately for later use). □

2.2.8 Lemma. *The product* $(\Delta^n, \Lambda^n_i) \times (\Delta^k, \partial\Delta^k)$ *with* $0 < i < n$ *is filtered by pushouts of inner horn pairs.*

Proof. The product $\Delta^n \times \Delta^k$ is the nerve of the category $[n] \times [k]$. Its non-degenerate $(n + k)$-simplices are thus in bijection with lattice paths from $(0, 0)$ to (n, k). There is a partial order on such lattice paths by requiring one to lie always above the other. We filter the pair $(\Delta^n, \Lambda^n_i) \times (\Delta^k, \partial\Delta^k)$ by adding these $(n+k)$-simplices one at a time, always adding one which is maximal (among the ones not already added) with respect to this partial order. Thus we just have to show that each individual such simplex addition can be realized by filling some inner horns.

Let Q denote $(\Lambda^n_i \times \Delta^k) \cup (\Delta^n \cup \partial\Delta^k) \cup$ (some upward closed set S of $(n+k)$-simplices). What we should do is, for any maximal $(n+k)$-simplex σ not in S, identify a vertex v of σ such that (1) for every simplex $\tau \subseteq \sigma$ lying in Q, the simplex spanned by τ and v inside σ also lies in Q, and (2) every simplex $\tau \subseteq \sigma$ consisting only of vertices $\leq v$ or only vertices $\geq v$ is in Q. The simplex σ is a lattice path $(0,0) \to (n,k)$. A simplex $\tau \subseteq \sigma$ corresponds to a subset of the vertices in this path, and it lies in Q iff it satisfies

$$\text{(image in } [n] \text{ does not contain } [n]-\{i\}) \vee \text{(image in } [k] \text{ is not everything)} \vee \text{(misses some cliffbottom corner)} \quad (2.2.8.1)$$

Take v to be the vertex of σ with first coordinate i and largest second coordinate. Property (2) is apparent: only vertices $\leq v$ or only vertices $\geq v$ would not surject onto $[n]-\{i\}$ since $0 < i < n$. For property (1), image in $[n]$ not containing $[n]-\{i\}$ is certainly preserved by adding v, as is missing some cliffbottom corner (v cannot be a cliffbottom corner). If adding v destroys the property of image in $[k]$ not being everything, then v must be the only vertex of σ with its second coordinate; if this second coordinate is $< k$, then this means σ misses a cliffbottom corner; if this second coordinate is $= k$, then it means σ cannot surject to $[n]-\{i\}$ as it will miss everything $> i$. □

2.2.9 Definition (Homotopy category of an ∞-category)**.** Let $\mathcal{C}$ be an ∞-category. For objects $x, y \in \mathcal{C}_0$, the relation of *right homotopy* on the set of morphisms $x \to y$ is defined by $e \sim e'$ iff there exists a 2-simplex

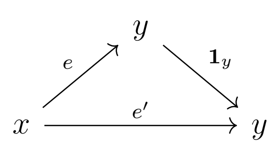

(2.2.9.1)

Right homotopy is an equivalence relation: reflexivity holds by taking a degenerate 2-simplex over e, and symmetry and transitivity follow from the following two inner horn fillings (the boxed vertex is the cone point of the horn):

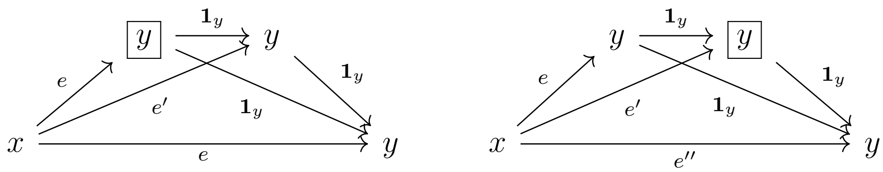

(2.2.9.2)

There is a corresponding equivalence relation *left homotopy*. Right homotopy implies left homotopy by filling the following inner horn (so by symmetry the converse is true as well)

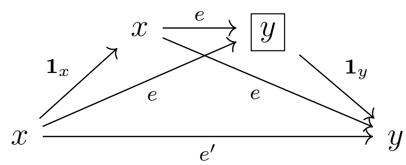

(2.2.9.3)

Since right homotopy and left homotopy are the same, we may simply call this relation *homotopy* of morphisms $x \to y$. Filling the following inner horn

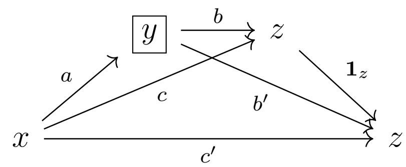

(2.2.9.4)

shows that if b and b' are homotopic, then any two fillings of $x \xrightarrow{a} y \xrightarrow{b} z$ and $x \xrightarrow{a} y \xrightarrow{b'} z$ give homotopic morphisms $x \to z$. By symmetry (using left homotopy = right homotopy) we conclude that composition is well-defined on homotopy classes. The *homotopy category* $h\mathcal{C}$ has the same objects as $\mathcal{C}$ (i.e. the vertices of $\mathcal{C}$) and has morphisms the homotopy classes of morphisms in $\mathcal{C}$. There is a tautological functor $\mathcal{C} \to h\mathcal{C}$.

2.2.10 Exercise. Show that the homotopy category of a category is itself.

2.2.11 Lemma. *The map* $\mathcal{C} \to h\mathcal{C}$ *satisfies the lifting property for* $(\Delta^2, \partial\Delta^2)$.

Proof. Fill the inner horn

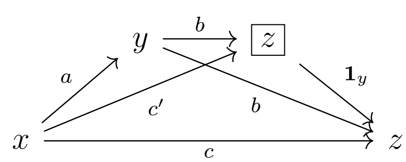

(2.2.11.1)

to show the equivalent assertion that a 2-simplex with edges a, b, c exists in $\mathcal{C}$ iff the boundary commutes in $h\mathcal{C}$. □

2.2.12 Definition (Isomorphisms in an ∞-category)**.** A morphism in an ∞-category $\mathcal{C}$ is called an isomorphism iff its image in $h\mathcal{C}$ is an isomorphism.

2.2.13 Definition (Equivalence of ∞-categories)**.** A functor of small ∞-categories $F : \mathcal{C} \to \mathcal{D}$ is called an *equivalence* iff there exists a functor $G : \mathcal{D} \to \mathcal{C}$ such that $G \circ F \simeq \mathbf{1}_\mathcal{C}$ and $F \circ G \simeq \mathbf{1}_\mathcal{D}$ (isomorphisms in the functor categories $\mathsf{Fun}(\mathcal{C}, \mathcal{C})$ and $\mathsf{Fun}(\mathcal{D}, \mathcal{D})$, respectively).

2.2.14 Remark. Consider the category $h\mathsf{Cat}_\infty$ whose objects are ∞-categories and whose morphisms are isomorphism classes of functors. A functor of ∞-categories is an equivalence iff it is an isomorphism in $h\mathsf{Cat}_\infty$. It follows that equivalences of ∞-categories satisfy the 2-out-of-3 property and that a retract of an equivalence is an equivalence.

2.2.15 Exercise. Show that for an equivalence of ∞-categories $\mathcal{C} \to \mathcal{D}$, the induced map $h\mathcal{C} \to h\mathcal{D}$ is an equivalence of categories.

2.2.16 Exercise. Show that a trivial Kan fibration $\mathcal{C} \to \mathcal{D}$ between ∞-categories is an equivalence.

2.2.17 Exercise. Let $F : \mathcal{C} \to \mathcal{D}$ be an equivalence of ∞-categories. Show that the induced maps $\mathsf{Fun}(K, \mathcal{C}) \to \mathsf{Fun}(K, \mathcal{D})$ and $\mathsf{Fun}(\mathcal{D}, \mathcal{E}) \to \mathsf{Fun}(\mathcal{C}, \mathcal{E})$ are equivalences.

2.2.18 Exercise. Show that if pre-composition with $\mathcal{A} \to \mathcal{B}$ induces an equivalence $\mathsf{Fun}(\mathcal{B}, \mathcal{C}) \to \mathsf{Fun}(\mathcal{A}, \mathcal{C})$ for every ∞-category $\mathcal{C}$, then $\mathcal{A} \to \mathcal{B}$ is an equivalence (in fact, it suffices to consider just the two cases $\mathcal{C} = \mathcal{A}$ and $\mathcal{C} = \mathcal{B}$).

2.2.19 Definition (Join)**.** For simplicial sets X and Y, their join $X \star Y$ is defined by the universal property that map $Z \to X \star Y$ is a map $p : Z \to \Delta^1$ and a pair of maps $p^{-1}(0) \to X$ and $p^{-1}(1) \to Y$.

2.2.20 Exercise. Show that $(X \star Y)^{\mathsf{op}} = Y^{\mathsf{op}} \star X^{\mathsf{op}}$.

2.2.21 Exercise. Show that $\Delta^n \star \Delta^m = \Delta^{n+m+1}$ for $n, m \geq -1$ (where $\Delta^{-1} = \varnothing$).

2.2.22 Exercise. For any simplicial set K, let $K^{\triangleright} = K \star \Delta^0$ and $K^{\triangleleft} = \Delta^0 \star K$. Prove that the geometric realization of $K^{\triangleright}$ is contractible for every simplicial set K.

2.2.23 Definition (Slice categories)**.** Given a diagram $K \to \mathcal{C}$, the over-category $\mathcal{C}_{/K}$ is defined by the universal property that a map $X \to \mathcal{C}_{/K}$ is a map $X \star K \to \mathcal{C}$ extending the given map $K \to \mathcal{C}$. The under-category $\mathcal{C}_{K/}$ is defined analogously using maps $K \star X \to \mathcal{C}$.

2.2.24 Exercise. For simplicial set pairs (X, A) and (Y, B) (meaning $A \to X$ and $B \to Y$ are level-wise injections), define $(X, A) \star (Y, B) := (X \star Y, (X \star B) \cup_{A \star B} (A \star Y))$. Show that $(\Delta^n, \Lambda^n_i) \star (\Delta^m, \partial\Delta^m) = (\Delta^{n+m+1}, \Lambda^{n+m+1}_i)$. Conclude that $(\Delta^n, \Lambda^n_i) \star (K, \varnothing)$ is filtered by inner horn fillings. Conclude that $\mathcal{C}_{/K}$ is an ∞-category.

2.2.25 Exercise. Suppose that a functor of ∞-categories $\mathcal{C} \to \mathcal{D}$ has the right lifting property for pairs (Δ^2, Λ^2_0). Show that $\mathcal{C} \to \mathcal{D}$ reflects split monomorphisms (let us call a morphism in an ∞-category a split monomorphism iff its image in the homotopy category is a split monomorphism i.e. has a retraction). Show that any functor which reflects split monomorphisms also reflects isomorphisms.

2.2.26 Exercise. Combine (2.2.24) and (2.2.25) to show that every forgetful functor $\mathcal{C}_{/K} \to \mathcal{C}$ reflects isomorphisms.

2.2.27 Lemma (Isomorphism as an extension property)**.** *A morphism e in an ∞-category is an isomorphism iff every left horn $\Lambda^n_0 \subseteq \Delta^n$ with initial (i.e. 01) edge e can be filled.*

Proof. This is due to Joyal; we follow the proof from [17, 1.2.4.3]. If $e : x \to y$ satisfies the left horn filling condition for $n = 2, 3$, then filling the following two left horns produces an inverse f to e in $h\mathcal{C}$:

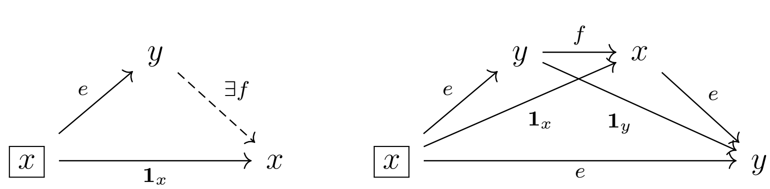

(2.2.27.1)

Conversely, suppose that e is an isomorphism in $h\mathcal{C}$, and fix a left horn $\Lambda^n_0 \to \mathcal{C}$ with initial edge e. We have $(\Delta^n, \Lambda^n_0) = (\Delta^1, \Lambda^1_0) \star (\Delta^{n-2}, \partial\Delta^{n-2})$, so we have, in particular, a map $\bar{e} : \Delta^1 \to \mathcal{C}_{/\partial\Delta^{n-2}}$ whose composition with the forgetful map $\mathcal{C}_{/\partial\Delta^{n-2}} \to \mathcal{C}$ is e. Since the functor $\mathcal{C}_{/\partial\Delta^{n-2}} \to \mathcal{C}$ reflects isomorphisms (2.2.26), the edge $\bar{e}$ is an isomorphism. There is thus a map $\Delta^2 \to \mathcal{C}_{/\partial\Delta^{n-2}}$ whose 01 edge is $\bar{e}$ and whose 02 edge is degenerate. This extends our given extension problem $(\Delta^1, \Lambda^1_0) \star (\Delta^{n-2}, \partial\Delta^{n-2}) \to \mathcal{C}$ to an extension problem $(\Delta^2, 02) \star (\Delta^{n-2}, \partial\Delta^{n-2}) \to \mathcal{C}$, which is solvable since $(\Delta^2, 02)$ is filtered by right horns, which upon join with $(\Delta^{n-2}, \partial\Delta^{n-2})$ become inner since $n \geq 2$ (2.2.24). □

2.2.28 Exercise. Show that given a morphism $e : x \to y$ in $\mathcal{C}$, the set of inverses on one side, i.e. 2-simplices $\Delta^2 \to \mathcal{C}$ whose edge 01 is e and whose edge 02 is $\mathbf{1}_x$, form a trivial Kan complex. More precisely, show to be a trivial Kan complex the simplicial set whose n-simplices are maps $\Delta^n \times \Delta^2 \to \mathcal{C}$ whose restriction to $\Delta^n \times \Lambda^2_0$ is the composition $\Delta^n \times \Lambda^2_0 \to \Lambda^2_0 \to \mathcal{C}$ where the map $\Lambda^2_0 \to \mathcal{C}$ is composed of e on (01) and $\mathbf{1}_x$ on (02).

2.2.29 Exercise. Let $m : p_0 \to p_1$ be a morphism in $\mathsf{Fun}(K, \mathcal{C})$; that is $m : K \times \Delta^1 \to \mathcal{C}$ and $p_i = m|_{K\times i}$ for $i = 0, 1$. Show that $\mathcal{C}_{/m} \to \mathcal{C}_{/p_0}$ is a trivial Kan fibration. Show that if m is an isomorphism, then $\mathcal{C}_{/m} \to \mathcal{C}_{/p_1}$ is a trivial Kan fibration. Conclude that there is a canonical contractible family of maps $\mathcal{C}_{/p_0} \to \mathcal{C}_{/p_1}$, which are equivalences if m is an isomorphism.

2.2.30 Definition (∞-groupoid)**.** An *∞-groupoid* is an ∞-category in which every morphism is an isomorphism (by (2.2.27), this is equivalent to being a Kan complex).

2.2.31 Definition (Core of an ∞-category)**.** For an ∞-category $\mathcal{C}$, its core is the subcomplex $\mathcal{C}_\simeq \subseteq \mathcal{C}$ defined as those simplices all of whose edges are isomorphisms. A functor $\mathcal{C} \to \mathcal{D}$ evidently restricts to a functor $\mathcal{C}_\simeq \to \mathcal{D}_\simeq$.

2.2.32 Exercise. Show that $\mathcal{C}_\simeq$ is an ∞-groupoid.

2.2.33 Exercise. Show that a map of Kan complexes is a homotopy equivalence iff it is an equivalence of ∞-groupoids.

2.2.34 Proposition (Isomorphisms in diagram categories)**.** *The functor* $\mathcal{F}\mathrm{un}(K, \mathcal{D}) \to \mathcal{F}\mathrm{un}(K_0, \mathcal{D}) = \prod_{c\in K} \mathcal{D}$ *reflects isomorphisms.*

Proof. Let an edge of $\mathcal{F}\mathrm{un}(K, \mathcal{D})$ be given, and let us study when $\mathcal{F}\mathrm{un}(K, \mathcal{D})$ has the extension property for left horns (Δ^n, Λ^n_0) with initial edge e. Equivalently, this means we study the extension property for $(\Delta^n, \Lambda^n_0) \times K$ for $\mathcal{D}$. We can filter K by simplices, so it suffices to understand the extension property for $(\Delta^n, \Lambda^n_0) \times (\Delta^k, \partial\Delta^k)$ for $\mathcal{D}$. It suffices to show that $(\Delta^n, \Lambda^n_0) \times (\Delta^k, \partial\Delta^k)$ admits a filtration by inner horns and left horns with initial edge $e \times \{0\}$. We use the same filtration from (2.2.7), and the same choice of vertices v. Property (1) continues to hold, with the same proof (which used $i < n$ but not $i > 0$). Let us investigate property (2). If the vertical coordinate of v is > 0, then a simplex all of whose vertices are before v will be in Q (it will miss the vertical coordinate of v in $[k]$), as will a simplex all of whose vertices are after v (it will miss 0). If the vertical coordinate of v is 0, then $v = (0, 0)$, so all faces of this Δ^{n+m} not in Q will be left horns to fill. These faces all contain $(1, 0)$, since not containing $(1, 0)$ either means not surjecting onto $[n] - \{0\}$ (not hitting 1) or missing a cliffbottom corner $(1, 0)$. Thus their initial edge is $(0, 0) \to (1, 0)$, which is $e \times \{0\}$ as desired. □

2.2.35 Definition (Essential surjectivity)**.** A functor of ∞-categories is called *essentially surjective* iff its action on homotopy categories is essentially surjective.

2.2.36 Definition (Full subcategories)**.** The *full subcategory* of $\mathcal{C}$ spanned by a set of its objects $S \subseteq \mathcal{C}_0$ is the subcomplex of $\mathcal{C}$ consisting of those simplices all of whose vertices are in S.

2.2.37 Definition (Final objects)**.** An object $c \in \mathcal{C}$ is called a *final object* iff the extension property holds for maps $(\Delta^n, \partial\Delta^n) \to \mathcal{C}$ sending the final vertex of Δ^n to c (for $n \geq 1$). Dually, an initial object in $\mathcal{C}$ is a final object in $\mathcal{C}^{\mathsf{op}}$.

2.2.38 Exercise. Show that if $c \in \mathcal{C}$ is final, then any map $K \to \mathcal{C}$ extends to a map $K^{\triangleright} \to \mathcal{C}$ sending the cone point to c.

2.2.39 Exercise. Show that the full subcategory of $\mathcal{C}$ spanned by final objects is either a trivial Kan complex or empty.

2.2.40 Exercise. Show that if $c \in \mathcal{C}$ is a final object, then $\mathcal{C}_{/c} \to \mathcal{C}$ is a trivial Kan fibration (hence an equivalence (2.2.16)).

2.2.41 Exercise. Show that if $c \in \mathcal{C}$ is final, then its image in $h\mathcal{C}$ is final. Conclude that every morphism between final objects in $\mathcal{C}$ is an isomorphism.

2.2.42 Exercise. Show that $(c \xrightarrow{1_c} c) \in \mathcal{C}_{/c}$ is a final object.

2.2.43 Exercise. Show that if $\mathcal{C}$ is (the nerve of) a category, then an object is final iff it is final in the ∞-categorical sense.

2.2.44 Exercise. Use (2.2.27) to show that the property of being a final object is preserved under isomorphism (show that if $x \to y$ is an isomorphism and y is final then x is also final).

2.2.45 Exercise. Show that an equivalence of ∞-categories sends final objects to final objects. One can proceed as follows: given $x \in \mathcal{C}$, suppose $F(x) \in \mathcal{D}$ is final, fix a map $\partial\Delta^n \to \mathcal{C}$ with final vertex x, fill its image in $\mathcal{D}$, consider the image of this filling under an inverse to F, and solve the remaining extension problem $(\Delta^n, \partial\Delta^n) \times (\Delta^1, 0) \to \mathcal{C}$.

2.2.46 Definition (Reflection and lifting of final objects)**.** A functor $F : \mathcal{C} \to \mathcal{D}$ is said to *reflect final objects* iff $F(c) \in \mathcal{D}$ final implies $c \in \mathcal{C}$ is final. It is said to *lift final objects* iff every final object of $\mathcal{D}$ lifts (up to isomorphism) to an object of $\mathcal{C}$.

2.2.47 Proposition. *The functor* $\mathsf{Fun}(K, \mathcal{D}) \to \mathsf{Fun}(K_0, \mathcal{D})$ *reflects and lifts final objects.*

Proof. An object of $\mathsf{Fun}(K_0, \mathcal{D}) = \prod_{x \in K} \mathcal{D}$ is final iff its value at every $x \in K$ is a final object of $\mathcal{D}$. To lift a final object of $\mathsf{Fun}(K_0, \mathcal{D})$ to $\mathsf{Fun}(K, \mathcal{D})$ consists of solving a series of extension problems $(\Delta^k, \partial\Delta^k) \to \mathcal{D}$ where all vertices map to final objects of $\mathcal{D}$; these extension problems are solvable by definition of final objects (2.2.37). To show that an object of $\mathsf{Fun}(K, \mathcal{D})$ is final if its image in $\prod_{x \in K} \mathcal{D}$ is final, we should solve an extension problem $(\Delta^k, \partial\Delta^k) \times K \to \mathcal{D}$ when the final vertex of Δ^k times any vertex of K is sent to a final object of $\mathcal{D}$. By filtering K by simplices, we reduce to maps from $(\Delta^k, \partial\Delta^k) \times (\Delta^m, \partial\Delta^m)$. This pair can further be filtered by simplices $(\Delta^n, \partial\Delta^n)$, and in each case the final vertex of Δ^n will lie over the final vertex of Δ^k (otherwise Δ^n would lie entirely over $\partial\Delta^k$). We are thus reduced to the extension problem whose solution is the definition of final objects (2.2.37). □

2.2.48 Definition (Limits and colimits in an ∞-category)**.** A *limit diagram* is a diagram $K^{\triangleleft} \to \mathcal{C}$ which is a final object in $\mathcal{C}_{/K}$. The *limit* of a diagram $K \to \mathcal{C}$ is the image $\lim(K) \in \mathcal{C}$ of a final object in $\mathcal{C}_{/K}$ (if one exists; otherwise the limit is not defined). Colimits and colimit diagrams are defined analogously in terms of initial objects of $\mathcal{C}_{K/}$.

2.2.49 Remark. The prefix 'homotopy' can be added in front of 'limit' or 'colimit' to indicate the ∞-categorical context as opposed to the context of ordinary categories.

2.2.50 Exercise. Show that maps $K' \to K \to \mathcal{C}$ induce a map $\mathcal{C}_{/K} \to \mathcal{C}_{/K'}$. Show that this induces, up to contractible choice, a morphism $\lim(K) \to \lim(K')$ in $\mathcal{C}$.

2.2.51 Definition (Final and initial functors)**.** A map $K \to K'$ is called *final* iff the induced map $\mathcal{C}_{K'/} \to \mathcal{C}_{K/}$ is an equivalence for every ∞-category $\mathcal{C}$ and every diagram $K' \to \mathcal{C}$. Similarly, a map $K \to K'$ is called initial when $\mathcal{C}_{K'/} \to \mathcal{C}_{K/}$ is always an equivalence. (This definition follows [17, 4.1.1.8].)

2.2.52 Exercise (Final functors induce isomorphisms on colimits). Show that if $K \to K'$ is final and $K' \to \mathcal{C}$ is any diagram, then the induced map $\operatorname{colim}(K) \to \operatorname{colim}(K')$ (2.2.50) is an isomorphism.

2.2.53 Exercise. Show that if $K \subseteq K'$ is filtered by pushouts of right horns, then $\mathcal{C}_{K'/} \to \mathcal{C}_{K/}$ is a trivial Kan fibration (use the fact that $(\Delta^n, \Lambda^n_n) \star (\Delta^k, \partial\Delta^k)$ is an inner horn for $k \geq 0$), so in particular the inclusion $K \subseteq K'$ is final.

2.2.54 Exercise. Show that the inclusion of the final vertex $* \in K^{\triangleright}$ is filtered by pushouts of right horns, hence is final.

2.2.55 Exercise (Inclusion of a final object is a final functor). Let $\mathcal{C}$ be an ∞-category with a final object $x \in \mathcal{C}$. Consider $\mathcal{C}^{\triangleright} = \mathcal{C} \star \Delta^0$ and its subcomplex $\{x\}^{\triangleright} = \{x\} \star \Delta^0 = \Delta^1$. Using the fact that x is a final object, show that the inclusion of pairs $(\mathcal{C}, \{x\}) \hookrightarrow (\mathcal{C}^{\triangleright}, \{x\}^{\triangleright})$ admits a retraction (proceed by induction on simplices). Observe that $\{x\}^{\triangleright} \to \mathcal{C}^{\triangleright}$ is final by noting that $\varnothing^{\triangleright} \subseteq \{x\}^{\triangleright}$ is the first step of the filtration by right horns from (2.2.54). Finally, show that a retract of a final map is final, so $\{x\} \to \mathcal{C}$ is final.

2.2.56 Definition (Cocomplete ∞-categories)**.** Let κ be any infinite cardinal. An ∞-category $\mathcal{C}$ is called *κ-cocomplete* iff every diagram $K \to \mathcal{C}$ for K of cardinality $< \kappa$ has a colimit in $\mathcal{C}$.

2.2.57 Lemma ([17, 4.4.2.6, 4.4.2.7])**.** *An ∞-category $\mathcal{C}$ is κ-cocomplete iff it admits coproducts over index sets of cardinality $< \kappa$ and admits pushouts. A functor $\mathcal{C} \to \mathcal{D}$ is κ-cocontinuous iff it preserves coproducts over index sets of cardinality $< \kappa$ and preserves pushouts.*

2.2.58 Definition (Mapping space $\operatorname{Hom}_{\mathcal{C}}$)**.** Given objects $x, y \in \mathcal{C}$, the mapping space $\operatorname{Hom}_{\mathcal{C}}(x, y) \in h\mathcal{S}$ is the object given by the simplicial set $\operatorname{Hom}_{\mathcal{C}}^{\mathrm{cyl}}(x, y)$ whose k-simplices are maps $\Delta^k \times \Delta^1 \to \mathcal{C}$ sending $\Delta^k \times 0$ to x and sending $\Delta^k \times 1$ to y.

2.2.59 Exercise. Prove that $\operatorname{Hom}_{\mathcal{C}}^{\mathrm{cyl}}(x, y)$ is a Kan complex (filter $(\Delta^n, \Lambda^n_i) \times (\Delta^1, \partial\Delta^1)$ by horns and appeal to (2.2.27)).

2.2.60 Exercise (Enrichment of $h\mathcal{C}$ over $h\mathcal{S}$)*.* Consider the simplicial set whose k-simplices are maps $\Delta^k \times \Delta^2 \to \mathcal{C}$ which map $\Delta^k \times \{0, 1, 2\}$ to objects $x, y, z \in \mathcal{C}$, respectively. Show that the forgetful map from this simplicial set to $\operatorname{Hom}_{\mathcal{C}}(x, y) \times \operatorname{Hom}_{\mathcal{C}}(y, z)$ (whose k-simplices are maps $\Delta^k \times \Lambda^2_1 \to \mathcal{C}$ sending $\Delta^k \times \{0, 1, 2\}$ to x, y, z) is a trivial Kan fibration. Choosing a section of this trivial Kan fibration determines a 'composition' morphism

$$\operatorname{Hom}_{\mathcal{C}}(x, y) \times \operatorname{Hom}_{\mathcal{C}}(y, z) \to \operatorname{Hom}_{\mathcal{C}}(x, z) \tag{2.2.60.1}$$

in $h\mathcal{S}$ which is independent of the choice of section. Show that composition is associative. Show that composition with $\mathbf{1}_x$ and $\mathbf{1}_y$ on either side both define the identity map $\operatorname{Hom}_{\mathcal{C}}(x, y) \to \operatorname{Hom}_{\mathcal{C}}(x, y)$. Conclude that this defines an enrichment of $h\mathcal{C}$ over $h\mathcal{S}$, equipped with the monoidal structure $\times$ and the functor $\pi_0 : h\mathcal{S} \to \mathsf{Set}$.

2.2.61 Exercise (Enrichment of functors over $h\mathcal{S}$)*.* Show that a functor $F : \mathcal{C} \to \mathcal{D}$ induces maps $\operatorname{Hom}_{\mathcal{C}}(x, y) \to \operatorname{Hom}_{\mathcal{D}}(F(x), F(y))$ which are compatible with composition. Show that for functors $F, G : \mathcal{C} \to \mathcal{D}$ and a natural transformation $F \Rightarrow G$, the following diagram commutes

$$\begin{array}{ccc} \operatorname{Hom}_{\mathcal{C}}(x, y) & \longrightarrow & \operatorname{Hom}_{\mathcal{D}}(F(x), F(y)) \\ \downarrow & & \downarrow \\ \operatorname{Hom}_{\mathcal{D}}(G(x), G(y)) & \longrightarrow & \operatorname{Hom}_{\mathcal{D}}(F(x), G(y)) \end{array} \tag{2.2.61.1}$$

(one can reduce to the case $\mathcal{D} = \mathcal{C} \times \Delta^1$). Conclude that this defines a lift of the map $\mathsf{Fun}(\mathcal{C}, \mathcal{D}) \to \mathsf{Fun}(h\mathcal{C}, h\mathcal{D})$ to the category of enriched functors $h\mathcal{C} \to h\mathcal{D}$. Conclude that if $F : \mathcal{C} \to \mathcal{D}$ is an equivalence, then the induced maps $\operatorname{Hom}_{\mathcal{C}}(x, y) \to \operatorname{Hom}_{\mathcal{D}}(F(x), F(y))$ are isomorphisms in $h\mathcal{S}$ for every $x, y \in \mathcal{C}$.

2.2.62 Definition (Full faithfulness)**.** A functor $F : \mathcal{C} \to \mathcal{D}$ is called *fully faithful* iff the induced map $\operatorname{Hom}_{\mathcal{C}}(x, y) \to \operatorname{Hom}_{\mathcal{D}}(F(x), F(y))$ is an isomorphism in $h\mathcal{S}$ for every $x, y \in \mathcal{C}$ (in other words, F induces an equivalence on homotopy categories enriched over $h\mathcal{S}$).

2.2.63 Definition (Left and right fibrations)**.** A map of simplicial sets $X \to Y$ is called a *left fibration* iff it satisfies the right lifting property with respect to all *left horns* (Δ^n, Λ^n_i) (meaning $0 \leq i < n$). Right fibrations are defined analogously via right horns ($0 < i \leq n$).

2.2.64 Lemma. *An object $x \in \mathcal{C}$ is final iff $\operatorname{Hom}_{\mathcal{C}}(a, x) = *$ for every $a \in \mathcal{C}$.*

Proof. The assertion that $\operatorname{Hom}_{\mathcal{C}}(a, x) = *$ for every $a \in \mathcal{C}$ is the statement that the fibers of $\mathcal{C}_{/x} \to \mathcal{C}$ are contractible. The assertion that x is final the statement that $\mathcal{C}_{/x} \to \mathcal{C}$ is a trivial Kan fibration (2.2.40). These are equivalent for any right fibration in place of $\mathcal{C}_{/x} \to \mathcal{C}$ (2.2.65). $\square$

2.2.65 Lemma. *A left fibration is a trivial Kan fibration iff each of its fibers is contractible.*

Proof. We follow the proof from [17, 2.1.3.4]. Let $Q \to X$ be a left fibration. Fix a simplex $\Delta^k \to X$ and a section of $Q \times_X \partial\Delta^k \to \partial\Delta^k$. Consider the map $\Delta^k \times \Delta^1 \to \Delta^k$ which is the identity on $\Delta^k \times 0$ and on $\Delta^k \times 1$ sends everything to the final vertex of Δ^k. The pair $\partial\Delta^k \times (\Delta^1, 0)$ is filtered by pushouts of left horns, so our section can be extended to $\partial\Delta^k \times \Delta^1$. Now we can extend it over $\Delta^k \times 1$ since this takes place within a single fiber of $Q \to X$. Finally, we have a lifting problem $(\Delta^k, \partial\Delta^k) \times (\Delta^1, 1)$ against $Q \to X$. This pair admits a filtration by k inner horns and one final right horn. To lift along the right horn, note that it suffices to find any extension of the right horn mapping to $Q \times_X \Delta^k$ since its projection to Δ^k is automatically compatible. Now an extension exists by (2.2.27) since $Q \times_X \Delta^k$ is an ∞-category and the boundary edge of the right horn lies over the final vertex of Δ^k, thus entirely in a fiber, which is a Kan complex, hence an ∞-groupoid (all morphisms are isomorphisms). $\square$

2.2.66 Lemma. *A left fibration over a Kan complex is a Kan fibration.*

Proof. We follow the proof from [17, 2.1.3.3]. Let $X \to Y$ be a left fibration with Y a Kan complex. We should show that $\underline{\mathrm{Hom}}(\Delta^k, X) \to \underline{\mathrm{Hom}}(\Lambda^k_j, X) \times_{\underline{\mathrm{Hom}}(\Lambda^k_j, Y)} \underline{\mathrm{Hom}}(\Delta^k, Y)$ is surjective on vertices. The map $\underline{\mathrm{Hom}}(\Delta^k, Y) \to \underline{\mathrm{Hom}}(\Lambda^k_j, Y)$ is a trivial Kan fibration since Y is Kan, so the map from the target above to $\underline{\mathrm{Hom}}(\Lambda^k_j, X)$ is a trivial Kan fibration. The map $\underline{\mathrm{Hom}}(\Delta^k, X) \to \underline{\mathrm{Hom}}(\Lambda^k_j, X)$ is a trivial Kan fibration, so the key map above is a homotopy equivalence. It is also a left fibration (a left horn times any pair is filtered by pushouts of left horns). Any left fibration of Kan complexes which is a homotopy equivalence is surjective on vertices (for any vertex in the target, being a homotopy equivalence means there is an edge terminating at it whose source lifts, and now this is a left horn to lift). $\square$

2.2.67 Example (Representability of right fibrations). Let $\mathcal{C}$ be an ∞-category. The right fibration $\mathcal{C}_{/c} \to \mathcal{C}$ has the particular property that $\mathcal{C}_{/c}$ has a final object (2.2.42).

Now let $p : \mathcal{M} \to \mathcal{C}$ be any right fibration over $\mathcal{C}$ (so $\mathcal{M}$ is also an ∞-category). For any object $m \in \mathcal{M}$, the map

$$\mathcal{M}_{/m} \xrightarrow{\sim} \mathcal{C}_{/p(m)} \tag{2.2.67.1}$$

is a trivial Kan fibration (this amounts to the right lifting property for p with respect to right horns with final vertex $m \mapsto p(m)$). The forgetful map $\mathcal{M}_{/m} \to \mathcal{M}$ thus determines canonical maps

$$\mathrm{Hom}_{\mathcal{C}}(x, p(m)) \to \mathcal{M}_x \tag{2.2.67.2}$$

in $h\mathcal{S}$ for every object $x \in \mathcal{C}$. If $m \in \mathcal{M}$ is a final object, then $\mathcal{M}_{/m} \to \mathcal{M}$ is a trivial Kan fibration (2.2.40), so the induced maps $\mathrm{Hom}_{\mathcal{C}}(x, p(m)) \to \mathcal{M}_x$ are isomorphisms (even better we have produced a zig-zag of trivial Kan fibrations connecting the right fibrations $\mathcal{M} \to \mathcal{C}$ and $\mathcal{C}_{/p(m)} \to \mathcal{M}$).

In fact, the converse holds: $m \in \mathcal{M}$ is final iff the induced maps $\mathrm{Hom}_{\mathcal{C}}(x, p(m)) \to \mathcal{M}_x$ are isomorphisms in $h\mathcal{S}$. Indeed, $\mathrm{Hom}_{\mathcal{C}}(x, p(m)) \to \mathcal{M}_x$ being a homotopy equivalence means that $(\mathcal{M}_{/m})_x \to \mathcal{M}_x$ is a homotopy equivalence. This map $(\mathcal{M}_{/m})_x \to \mathcal{M}_x$ is a right fibration over a Kan complex, hence is a Kan fibration (2.2.66). Hence it being a homotopy equivalence implies it is a trivial Kan fibration (2.1.41). Its fibers are thus trivial Kan complexes, whence the same holds for $\mathcal{M}_{/m} \to \mathcal{M}$, so m is a final object (2.2.64).

Motivated by these facts, we will say that a right fibration $\mathcal{M} \to \mathcal{C}$ is *representable* iff $\mathcal{M}$ has a final object, and we will call the image in $\mathcal{C}$ of this final object the *representing object* for $\mathcal{M} \to \mathcal{C}$. For example, the limit of a diagram $K \to \mathcal{C}$ exists iff the right fibration $\mathcal{C}_{/K} \to \mathcal{C}$ is representable, and in this case $\lim(K)$ is the representing object.

2.2.68 Definition. Let $F : \mathcal{C} \to \mathcal{D}$ be a functor, and let $d \in \mathcal{D}$. The category $\mathcal{C}_{F(\cdot)/d}$ is the fiber product

$$\begin{array}{ccc} \mathcal{C}_{F(\cdot)/d} & \longrightarrow & \mathcal{D}_{/d} \\ \downarrow & & \downarrow \\ \mathcal{C} & \xrightarrow{F} & \mathcal{D} \end{array} \tag{2.2.68.1}$$

Thus an object of $\mathcal{C}_{F(\cdot)/d}$ is an object $c \in \mathcal{C}$ together with a morphism $F(c) \to d$.

2.2.69 Definition (Adjoint functors)**.** Let $F : \mathcal{C} \to \mathcal{D}$ and $G : \mathcal{C} \to \mathcal{D}$ be functors of ∞-categories. An *adjunction* (F, G) (F termed the left adjoint and G the right adjoint) is a specific piece of data which, in particular, encodes equivalences $\operatorname{Hom}_{\mathcal{D}}(F(c), d) = \operatorname{Hom}_{\mathcal{C}}(c, G(d))$ for all $c \in \mathcal{C}$ and $d \in \mathcal{D}$. We give two notions of adjunction here, and we will see later that they are equivalent in a homotopical sense.

A *counit adjunction* $(F, G, F \circ G \Rightarrow \mathbf{1}_{\mathcal{D}})$ is a triple such that for every $c \in \mathcal{C}$ and $d \in \mathcal{D}$, the induced map

$$\operatorname{Hom}_{\mathcal{C}}(c, G(d)) \to \operatorname{Hom}_{\mathcal{D}}(F(c), F(G(d))) \to \operatorname{Hom}_{\mathcal{D}}(F(c), d) \tag{2.2.69.1}$$

is an isomorphism (equivalently, $(F(G(d)) \to d) \in \mathcal{C}_{F(\cdot)/d}$ is a final object for every $d \in \mathcal{D}$). A *unit adjunction* $(F, G, \mathbf{1}_{\mathcal{C}} \Rightarrow G \circ F)$ is defined analogously (so $(G^{\mathsf{op}}, F^{\mathsf{op}}, G^{\mathsf{op}} \circ F^{\mathsf{op}} \Rightarrow \mathbf{1}_{\mathcal{C}^{\mathsf{op}}})$ is a counit adjunction).

2.2.70 Exercise (Adjoints and full faithfulness)*.* Let $(F, G, \mathbf{1}_{\mathcal{C}} \Rightarrow G \circ F)$ be a unit adjunction. Show that the composition

$$\operatorname{Hom}_{\mathcal{C}}(x, y) \to \operatorname{Hom}_{\mathcal{D}}(F(x), F(y)) = \operatorname{Hom}_{\mathcal{C}}(x, G(F(y))) \tag{2.2.70.1}$$

coincides with composition with the unit map $y \to G(F(y))$ (use the fact that $\mathbf{1}_{\mathcal{C}} \Rightarrow G \circ F$ is a natural transformation). Conclude that F is fully faithful iff the unit map is a natural isomorphism of functors.

2.2.71 Definition (Inner fibrations)**.** A map of simplicial sets $X \to Y$ is called an *inner fibration* iff it satisfies the right lifting property with respect to inner horns.

2.2.72 Exercise. Follow the argument of (2.2.7) to show that for any inner fibration $Q \to X$, the simplicial set of sections $\underline{\operatorname{Sec}}(Q/X)$ (whose k-simplices are sections of $Q \times \Delta^k \to X \times \Delta^k$) is an ∞-category.

2.2.73 Exercise. Follow the argument of (2.2.34) to show that for any inner fibration $Q \to X$, the functor $\underline{\mathrm{Sec}}(Q/X) \to \prod_{x\in X} Q_x$ reflects isomorphisms.

2.2.74 Definition (Cartesian and cocartesian fibrations)**.** Let $X \to Y$ be a map of simplicial sets. An edge e in X is called *cocartesian* iff $X \to Y$ satisfies the right lifting property with respect to left horns whose 01 edge maps to e. The map $X \to Y$ is called a *cocartesian fibration* iff it is an inner fibration and every left horn (Δ^1, Λ^1_0) lifts to a cocartesian edge in X. A map $X \to Y$ is called a cartesian fibration iff $X^{\mathsf{op}} \to Y^{\mathsf{op}}$ is a cocartesian fibration.

2.2.75 Exercise. Show that being a cocartesian fibration is preserved under pullback and composition.

2.2.76 Exercise. Show that a left fibration is a cocartesian fibration.

2.2.77 Exercise. Show that $\mathcal{C} \to *$ is a cocartesian fibration iff $\mathcal{C}$ is an ∞-category. Show moreover that the cocartesian edges are precisely the isomorphisms in $\mathcal{C}$.

2.2.78 Proposition. *For any cocartesian fibration $Q \to X$, the functor* $\mathrm{Sec}(Q/X) \to \prod_{x\in X} Q_x$ *reflects and lifts final objects.*

Proof. The proof of (2.2.47) applies provided we show that for any cocartesian fibration $Q \to \Delta^k$, any section of Q over $\partial\Delta^k$ taking the value a final object over the final vertex of Δ^k extends to all of Δ^k. To do this, we follow the strategy from (2.2.65). We consider the map $r : \Delta^k \times \Delta^1 \to \Delta^k$ which is the identity on $\Delta^k \times 0$ and which sends $\Delta^k \times 1$ to the final vertex. Our section over $\partial\Delta^k$ extends to a section of r^*Q over $\partial\Delta^k \times \Delta^1$ which sends each edge $v \times \Delta^1$ to a cocartesian edge. This section now extends over $\Delta^k \times 1$ by the assumption of the value at the final vertex being a final object. What remains is an extension problem $(\Delta^k, \partial\Delta^k) \times (\Delta^1, 1) \to Q$ (map to Δ^k is determined by action on vertices, so it is really now just an extension problem, not a lifting problem). The usual filtration of $(\Delta^k, \partial\Delta^k) \times (\Delta^1, 1)$ by pushouts of horns begins with k inner horns (evidently fillable) and one final right horn. To fill the right horn, we want to know that a certain edge in the total space Q is an isomorphism (note Q is an ∞-category) but it is since it is contained in a single fiber (the fiber over the final vertex of Δ^k) and is assumed cocartesian there. ☐

2.2.79 Proposition (Localization of ∞-categories)**.** *Let $\mathcal{C}$ be an ∞-category and let W be a set of morphisms in $\mathcal{C}$. There exists an ∞-category $\mathcal{C}[W^{-1}]$ and a functor $\mathcal{C} \to \mathcal{C}[W^{-1}]$ such that the induced functor $\mathcal{F}\mathrm{un}(\mathcal{C}[W^{-1}], \mathcal{D}) \to \mathcal{F}\mathrm{un}_W(\mathcal{C}, \mathcal{D})$ is an equivalence, where $\mathcal{F}\mathrm{un}_W(\mathcal{C}, \mathcal{D}) \subseteq \mathcal{F}\mathrm{un}(\mathcal{C}, \mathcal{D})$ denotes the full subcategory spanned by the functors $\mathcal{C} \to \mathcal{D}$ which send W to isomorphisms.*

Proof. We use the small object argument (2.1.28). Let $\mathcal{C}_0 := \mathcal{C}$, and for each $i \geq 1$ define $\mathcal{C}_i$ from $\mathcal{C}_{i-1}$ by filling all horns $\Lambda^n_k \subseteq \Delta^n$ which are either inner or whose boundary morphism lies in W. Define $\mathcal{C}[W^{-1}] := \bigcup_i \mathcal{C}_i$ to be the ascending union. It is immediate that $\mathcal{C}[W^{-1}]$ is an ∞-category, and the morphisms W are isomorphisms in $\mathcal{C}[W^{-1}]$ by the extension criterion for isomorphisms (2.2.27). Now it suffices to note that each map $\mathcal{F}\mathrm{un}(\mathcal{C}_i, \mathcal{D}) \to \mathcal{F}\mathrm{un}(\mathcal{C}_{i-1}, \mathcal{D})$ is a trivial Kan fibration over the locus of functors which send W to isomorphisms (to see this, apply to $\mathcal{D}$ the fact that inner horn fillings are unique up to contractible choice, as are outer horn fillings with an isomorphism as a boundary morphism). ☐

2.3 Spaces

Topological spaces form an ∞-category Top^∞. Homotopy theory is more interested in the full subcategory $\mathcal{S} \subseteq \mathsf{Top}^\infty$ spanned by CW-complexes (equivalently, geometric realizations of simplicial sets). This is the ∞-category of spaces, and it plays a central role in the theory of ∞-categories (analogous to the role of the category of sets Set in the theory of categories and the role of the 2-category of groupoids in the theory of 2-categories). The ∞-category $\mathcal{S}$ may also be described purely combinatorially via Kan complexes, their mapping spaces, and the simplicial nerve construction. It is complete and cocomplete; in fact, it is the universal cocomplete ∞-category with a marked object.

2.3.1 Definition. A *fibrant simplicial category* is a category C enriched in Kan complexes.

2.3.2 Definition. The *simplicial nerve* $\mathcal{C}$ of a fibrant simplicial category C has n-simplices the tuples of objects $X_0, \ldots, X_n \in \mathsf{C}$ together with maps $f_{ij} : (\Delta^1)^{\{i+1,\ldots,j-1\}} \to \mathsf{C}(X_i, X_j)$ such that $f_{ik}|_{\{t_j=1\}} = f_{ij} \times f_{jk}$. Given a map $s : \Delta^m \to \Delta^n$, such data pulls back as follows. We take $Y_i := X_{s(i)}$, and we take $g_{ij} = f_{s(i)s(j)}$ precomposed with the map $(\Delta^1)^{\{i+1,\ldots,j-1\}} \to (\Delta^1)^{\{s(i)+1,\ldots,s(j)-1\}}$ given on vertices by the formula $t_k = \max_{s(a)=k} t_a$ (interpreted to be 0 when $s^{-1}(k)$ is empty).

2.3.3 Exercise. Describe explicitly the 0-simplices (objects), 1-simplices (morphisms), and 2-simplices of the simplicial nerve $\mathcal{C}$ of a fibrant simplicial category C. Consider the subcategory of $\mathcal{C}$ consisting of those simplices in which every f_{ij} is constant; how is this related to C?

2.3.4 Lemma. *The simplicial nerve of a fibrant simplicial category is an ∞-category.*

Proof. The data of an inner horn $\Lambda^n_i \to \mathcal{C}$ determines the data of a simplex $\Delta^n \to \mathcal{C}$ except for the map $f_{0n} : (\Delta^1)^{\{1,\ldots,n-1\}} \to \mathsf{C}(X_0, X_n)$ being defined only when the i-coordinate is 1 or one of the other coordinates is 0 or 1. Thus to fill an inner horn, we should solve the extension problem for maps $(\Delta^1, \{1\}) \times (\Delta^1, \partial\Delta^1)^{n-2} \to \mathsf{C}(X_0, X_n)$. This extension problem is solvable since $\mathsf{C}(X_0, X_n)$ is Kan and the domain pair is filtered by horn inclusions. □

2.3.5 Definition (∞-category of spaces $\mathcal{S}$)**.** The ∞-category $\mathcal{S}$ is the simplicial nerve of the fibrant simplicial category whose objects are Kan complexes and whose complex of morphisms from $X \to Y$ is the simplicial mapping space $\underline{\mathrm{Hom}}(X, Y)$ (2.1.32). Concretely, an n-simplex in $\mathcal{S}$ consists of the data of Kan complexes $X_0, \ldots, X_n$ and maps $f_{ij} : X_i \times (\Delta^1)^{\{i+1,\ldots,j-1\}} \to X_j$ for $i < j$ such that for every triple $i < j < k$, the following diagram commutes:

$$\begin{array}{ccc} X_i \times (\Delta^1)^{\{i+1,\ldots,j-1\}} \times (\Delta^1)^{\{j+1,\ldots,k-1\}} & \xrightarrow{f_{ij} \times \mathbf{1}} & X_j \times (\Delta^1)^{\{j+1,\ldots,k-1\}} \\ \downarrow{\scriptstyle \times\{1\}_j} & & \downarrow{\scriptstyle f_{jk}} \\ X_i \times (\Delta^1)^{\{i+1,\ldots,k-1\}} & \xrightarrow{f_{ik}} & X_k \end{array} \tag{2.3.5.1}$$

Given a map $s : \Delta^m \to \Delta^n$, such data pulls back by taking $Y_i = X_{s(i)}$ and $g_{ij} = f_{s(i)s(j)}$ precomposed with the map $(\Delta^1)^{\{i+1,\ldots,j-1\}} \to (\Delta^1)^{\{s(i)+1,\ldots,s(j)-1\}}$ given by $t_k = \max_{s(a)=k} t_a$.

2.3.6 Exercise. Describe explicitly the 0-simplices (objects), 1-simplices (morphisms), and 2-simplices of $\mathcal{S}$. Identify a simplicial subset of $\mathcal{S}$ which is the full subcategory of sSet spanned by Kan complexes. Show that the homotopy category of $\mathcal{S}$ is the category $h\mathcal{S}$ introduced earlier (2.1.33).

2.3.7 Lemma (Coproducts in $\mathcal{S}$)**.** *Disjoint unions of Kan complexes are coproducts in* $\mathcal{S}$*.*

Proof. Fix a family of Kan complexes Z_α indexed by $\alpha \in A$ for some set A. The coproduct of these objects is, by definition (2.2.48), an initial object of the under-category $\mathcal{S}_{A/}$ (regarding A as a disjoint union of 0-simplices). The family of all maps $Z_\alpha \to \bigsqcup_\alpha Z_\alpha$ defines a lift of $\bigsqcup_\alpha Z_\alpha \in \mathcal{S}$ to $\mathcal{S}_{A/}$, and our task is to show that this object of $\mathcal{S}_{A/}$ is initial. This means, by definition (2.2.37), that any map $\partial\Delta^n \to \mathcal{S}_{A/}$ whose initial vertex maps to (this distinguished lift of) $\bigsqcup_\alpha Z_\alpha$ should extend to Δ^n. Concretely, this amounts to showing the extension property for diagrams

$$\begin{array}{ccc} \coprod_\alpha Z_\alpha \times (\Delta^1,\{0\})^{\{0\}} \times (\Delta^1,\partial\Delta^1)^{\{1,\ldots,n-1\}} & & \\ {\scriptstyle \times\{1\}^{\{0\}}}\uparrow & \searrow & \\ \coprod_\alpha Z_\alpha \times (\Delta^1,\partial\Delta^1)^{\{1,\ldots,n-1\}} & \longrightarrow & X_n \end{array} \tag{2.3.7.1}$$

where we have implicitly appealed to the fact that products distribute over coproducts in sSet (because they do so in Set). Now the bottom map is determined uniquely by the diagonal map, and the extension problem for the diagonal map has a solution since $(\Delta^1,\{0\})$ is a horn, so its product with anything is filtered by pushouts of horns (2.1.31). □

2.3.8 Exercise (Products in $\mathcal{S}$). Show that products of Kan complexes are products in $\mathcal{S}$ by reducing to the extension property for diagrams

$$\begin{array}{ccc} X_0 \times (\Delta^1,\partial\Delta^1)^{\{1,\ldots,n-1\}} & & \\ {\scriptstyle \times\{1\}^{\{n\}}}\downarrow & \searrow & \\ X_0 \times (\Delta^1,\partial\Delta^1)^{\{1,\ldots,n-1\}} \times (\Delta^1,0)^{\{n\}} & \longrightarrow & \prod_\alpha Z_\alpha \end{array} \tag{2.3.8.1}$$

and appealing to (2.1.31).

2.3.9 Lemma (Pullbacks in $\mathcal{S}$)**.** *Let A be a Kan complex and let $B_\pm \to A$ be Kan fibrations. The fiber product $C = B_+ \times_A B_-$ in simplicial sets is the fiber product in* $\mathcal{S}$*.*

Proof. By definition (see (2.2.37) and (2.2.48)), the diagram

$$\begin{array}{ccc} C & \longrightarrow & B_+ \\ \downarrow & & \downarrow \\ B_- & \longrightarrow & A \end{array} \tag{2.3.9.1}$$

is a pullback in $\mathcal{S}$ iff the extension property holds for maps $(\Delta^n, \partial\Delta^n) \star (\bullet \leftarrow \bullet \rightarrow \bullet) \to \mathcal{S}$ which send $n \star (\bullet \leftarrow \bullet \rightarrow \bullet) = \Delta^1 \times \Delta^1$ (where $n \in \Delta^n$ is the final vertex) to (2.3.9.1). Let us label the vertices of the second term $\beta_- \to \alpha \leftarrow \beta_-$ and the vertices of the first term $0, \ldots, n-1, \gamma$ (thus $\alpha, \beta_\pm, \gamma$ correspond to $A, B_\pm, C$). Now the extension property in question amounts to the extension property for maps to (2.3.9.1) from $X_0 \times (\Delta^1, \partial\Delta^1)^{\{1,\ldots,n-1\}}$ times the diagram

$$\begin{array}{ccc} * & \xrightarrow{\times\{1\}^\gamma} & (\Delta^1, 0)^\gamma \\ & \nwarrow & \\ \downarrow{\scriptstyle \times\{1\}^\gamma} & ((-) \leftarrow 0 \rightarrow (+)) & \downarrow{\scriptstyle \times(+)} \\ & \searrow{\scriptstyle \times\{1\}^\gamma} & \\ (\Delta^1, 0)^\gamma & \xrightarrow{\times(-)} & (\Delta^1, 0)^\gamma \times ((-) \leftarrow 0 \rightarrow (+)) \end{array} \tag{2.3.9.2}$$

Since (2.3.9.1) is a pullback in sSet, the map from the upper left corner is determined uniquely by the other maps, so we can just consider maps to $B_- \to A \leftarrow B_+$ from

$$X_0 \times (\Delta^1, \partial\Delta^1)^{\{1,\ldots,n-1\}} \times \left(\begin{array}{ccc} & & (\Delta^1, 0)^\gamma \\ & & \downarrow{\scriptstyle \times(+)} \\ (\Delta^1, 0)^\gamma & \xrightarrow{\times(-)} & (\Delta^1, 0)^\gamma \times ((-) \leftarrow 0 \rightarrow (+)) \end{array} \right) \tag{2.3.9.3}$$

where the restriction of the map from the lower right corner to $(\cdots \times \{1\}^\gamma \times \cdots)$ must be independent of the $((-) \leftarrow 0 \rightarrow (+))$ factor.

We solve this extension problem as follows. First solve the extension problem $X_0 \times (\Delta^1, \partial\Delta^1)^{\{1,\ldots,n-1\}} \times (\Delta^1, 0)^\gamma \times (0) \to A$ from the lower right corner; the map is then forced on $X_0 \times (\Delta^1)^{\{1,\ldots,n-1\}} \times 1^\gamma \times ((-) \leftarrow 0 \rightarrow (+))$. The remaining extension problem for the lower left corner now takes the form $X_0 \times (\Delta^1, \partial\Delta^1)^{\{1,\ldots,n-1\}} \times (\Delta^1, \partial\Delta^1)^\gamma \times ((-) \leftarrow 0 \rightarrow (+), 0) \to A$ which again has a solution since the last factor is filtered by horn pushouts. What remains is now to lift $X_0 \times (\Delta^1, \partial\Delta^1)^{\{1,\ldots,n-1\}} \times (\Delta^1, 0)^\gamma$ against $B_+ \to A$ and against $B_- \to A$, which are both possible since $B_\pm \to A$ are Kan fibrations. □

2.3.10 Exercise. Show that (2.3.9) continues to hold under the assumption that only $B_+ \to A$ (i.e. not necessarily $B_- \to A$) is a Kan fibration, by solving the extension problem in a slightly different order (start with $X_0 \times (\Delta^1, \partial\Delta^1)^{\{1,\ldots,n-1\}} \times (\Delta^1, 0)^\gamma \to B_-$, then do $X_0 \times (\Delta^1, \partial\Delta^1)^{\{1,\ldots,n-1\}} \times (\Delta^1, \partial\Delta^1)^\gamma \times ((-) \leftarrow 0, (-)) \to A$, and then proceed as before).

2.3.11 Lemma (Pushouts in $\mathcal{S}$). *Let $B_- \leftarrow A \to B_+$ be inclusions of Kan complexes, and let C be the pushout (in simplicial sets). Let $C \hookrightarrow \hat{C}$ be an inclusion into a Kan complex filtered by horn inclusions. The square*

$$\begin{array}{ccc} A & \longrightarrow & B_+ \\ \downarrow & & \downarrow \\ B_- & \longrightarrow & \hat{C} \end{array} \tag{2.3.11.1}$$

is a pushout in $\mathcal{S}$*.*

Proof. By definition, the diagram (2.3.11.1) is a pushout iff the extension property holds for maps $(\bullet \leftarrow \bullet \rightarrow \bullet) \star (\Delta^n, \partial\Delta^n) \rightarrow \mathcal{S}$ which send $(\bullet \leftarrow \bullet \rightarrow \bullet) \star 0 = \Delta^1 \times \Delta^1$ (where $0 \subseteq \Delta^n$ is the initial vertex) to (2.3.11.1).

Let us label the vertices of the first term $\beta_- \leftarrow \alpha \rightarrow \beta_+$ and the vertices of the second term $\gamma, 1, \ldots, n$ (thus $\alpha, \beta_\pm, \gamma$ correspond to $A, B_\pm, \hat{C}$). Now the extension property in question amounts to the extension property for a collection of maps

$$f_{\gamma n} : \hat{C} \times (\Delta^1, \partial\Delta^1)^{\{1,\ldots,n-1\}} \rightarrow X_n \tag{2.3.11.2}$$
$$f_{\beta_\pm n} : B_\pm \times (\Delta^1, 0)^{\{\gamma\}} \times (\Delta^1, \partial\Delta^1)^{\{1,\ldots,n-1\}} \rightarrow X_n \tag{2.3.11.3}$$
$$f^\pm_{\alpha n} : A \times (\Delta^1)^{\{\beta_\pm\}} \times (\Delta^1, 0)^{\{\gamma\}} \times (\Delta^1, \partial\Delta^1)^{\{1,\ldots,n-1\}} \rightarrow X_n \tag{2.3.11.4}$$

subject to the compatibility relations

$$f_{\beta_\pm n}|_{\cdots\times\{1\}^{\{\gamma\}}\times\cdots} = f_{\gamma n}|_{B_\pm\times\cdots} \tag{2.3.11.5}$$
$$f^\pm_{\alpha n}|_{\cdots\times\{1\}^{\{\beta_\pm\}}\times\cdots} = f_{\beta_\pm n}|_{A\times\cdots} \tag{2.3.11.6}$$
$$f^\pm_{\alpha n}|_{\cdots\times\{1\}^{\{\gamma\}}\times\cdots} = f_{\gamma n}|_{A\times\cdots} \tag{2.3.11.7}$$
$$f^+_{\alpha n}|_{\cdots\times\{0\}^{\{\beta_+\}}\times\cdots} = f^-_{\alpha n}|_{\cdots\times\{0\}^{\{\beta_-\}}\times\cdots} \tag{2.3.11.8}$$

We solve this extension problem as follows. Consider the restriction $f^\pm_{\alpha n}|_{\cdots\times\{0\}^{\{\beta_+\}}\times\cdots} : A \times (\Delta^1, 0)^{\{\gamma\}} \times (\Delta^1, \partial\Delta^1)^{\{1,\ldots,n-1\}} \rightarrow X_n$. This pair on the left is filtered by horn inclusions, so an extension exists. In accordance with (2.3.11.7) (independence of the β-coordinate when the γ-coordinate equals 1), this reduces $f^\pm_{\alpha n}$ to the extension problem $A \times (\Delta^1, 0)^{\{\beta_\pm\}} \times (\Delta^1, \partial\Delta^1)^{\{\gamma\}} \times (\Delta^1, \partial\Delta^1)^{\{1,\ldots,n-1\}} \rightarrow X_n$; this remains filtered by horns so it has an extension. Now extension of $f_{\beta_\pm n}$ is reduced to extension for $(B_\pm, A) \times (\Delta^1, 0)^{\{\gamma\}} \times (\Delta^1, \partial\Delta^1)^{\{1,\ldots,n-1\}}$; this is filtered by horns so we are done. Now finally $f_{\gamma n}$ is the extension problem $(\hat{C}, C) \times (\Delta^1, \partial\Delta^1)^{\{1,\ldots,n-1\}} \rightarrow X_n$, and again the domain is filtered by horns. □

2.3.12 Proposition. $\mathcal{S}$ *is cocomplete.*

Proof. It suffices to show that $\mathcal{S}$ admits pushouts and arbitrary coproducts (2.2.57). Coproducts are disjoint unions (2.3.7) and so always exist. Now let us show that every pushout $B_- \leftarrow A \rightarrow B_+$ exists. Factor each $A \rightarrow B_\pm$ as an inclusion followed by a trivial Kan fibration (exists by small object argument as usual). Now a trivial Kan fibration is an isomorphism in $\mathcal{S}$, so by replacing $B_\pm$ by the total spaces of these trivial Kan fibrations, we may assume that each $A \rightarrow B_\pm$ is injective. Now the pushout exists by (2.3.11). □

2.3.13 Proposition. $\mathcal{S}$ *is complete.*

Proof. Products of Kan complexes are Kan and are the categorical products (2.3.8). For pullbacks, note that any $B \rightarrow A$ can be turned into a Kan fibration by the small object argument (2.1.28) by adding horn fillings to B, and in this case the pullback is simply the pullback in simplicial sets (2.3.9). We should check that if $B \to B'$ is an inclusion of Kan complexes which is filtered by inner horn fillings, then it is an equivalence in $\mathcal{S}$, but this follows since $\underline{\mathrm{Hom}}(B', X) \to \underline{\mathrm{Hom}}(B, X)$ is then a trivial Kan fibration for any Kan complex X. □

2.3.14 Definition (∞-category of topological spaces Top^∞)**.** Let $I := [0,1]$. An n-simplex in Top^∞ consists of the data of topological spaces $X_0, \ldots, X_n$ and continuous maps $f_{ij} : X_i \times I^{\{i+1,\ldots,j-1\}} \to X_j$ for $i < j$ (which are *collared*, meaning independent of each I-coordinate over some neighborhood of ∂I) such that for every triple $i < j < k$, the following diagram commutes:

$$\begin{array}{ccc} X_i \times I^{\{i+1,\ldots,j-1\}} \times I^{\{j+1,\ldots,k-1\}} & \xrightarrow{f_{ij}\times \mathbf{1}} & X_j \times I^{\{j+1,\ldots,k-1\}} \\ \downarrow{\scriptstyle \times\{1\}_j} & & \downarrow{\scriptstyle f_{jk}} \\ X_i \times I^{\{i+1,\ldots,k-1\}} & \xrightarrow{f_{ik}} & X_k \end{array} \tag{2.3.14.1}$$

Given a map $s : \Delta^m \to \Delta^n$, such data pulls back as follows. We take $Y_i := X_{s(i)}$, and we take $g_{ij} = f_{s(i)s(j)}$ precomposed with the map $I^{\{i+1,\ldots,j-1\}} \to I^{\{s(i)+1,\ldots,s(j)-1\}}$ given by $t_k = \max_{s(a)=k} t_a$.

2.3.15 Exercise. Describe explicitly the 0-simplices (objects), 1-simplices (morphisms), and 2-simplices of Top^∞. Identify a simplicial subset of Top^∞ which is Top. When does a map $\partial\Delta^2 \to \mathsf{Top}^\infty$ extend to Δ^2? What is the homotopy category of Top^∞?

2.3.16 Lemma. Top^∞ *is an* ∞*-category.*

Proof. We should show the extension property for maps $(\Delta^n, \Lambda^n_i) \to \mathsf{Top}^\infty$. Given a map $\Lambda^n_i \to \mathsf{Top}^\infty$, we get a subset of the data needed to specify a map $\Delta^n \to \mathsf{Top}^\infty$: we get topological spaces $X_0, \ldots, X_n$, and we get maps f_{ab} except that the map f_{0n} is defined only on the boundary of $(I, \partial I)^{\{1,\ldots,n\}\setminus\{i\}} \times (I, \{1\})^{\{i\}}$. Here we have used the requirement of collaring to glue various maps. Now to solve the extension problem for maps

$$X_0 \times (I, \partial I)^{\{1,\ldots,n\}\setminus\{i\}} \times (I, \{1\})^{\{i\}} \to X_n \tag{2.3.16.1}$$

we just need to observe that the pair $(I, \partial I)^{\{1,\ldots,n\}\setminus\{i\}} \times (I, \{1\})^{\{i\}}$ admits a retraction (in a way ensuring the collaring continues to hold). □

2.3.17 Lemma. *Disjoint unions of topological spaces are coproducts in* Top^∞*.*

Proof. The coproduct of objects $X_\alpha \in \mathsf{Top}$ is simply their disjoint union $\bigsqcup_\alpha X_\alpha$. To show that the diagram of all maps $X_\alpha \to \bigsqcup_\alpha X_\alpha$ remains a coproduct diagram in Top^∞, we should show a certain extension property about being an initial object (2.2.48)(2.2.37). This amounts to showing the extension property for a collection of maps

$$\Big\{ f_{\alpha n} : X_\alpha \times (I, \{0\})^{\{0\}} \times (I, \partial I)^{\{1,\ldots,n\}} \to X_n \Big\}_{\alpha \in A} \tag{2.3.17.1}$$

$$f_{0n} : \Big(\bigsqcup_\alpha X_\alpha \Big) \times (I, \partial I)^{\{1,\ldots,n\}} \to X_n \tag{2.3.17.2}$$

with the requirement that f_{0n} restricted to a given X_α agrees with $f_{\alpha n}$ restricted to $\{1\}$ in the $(I,\{0\})^{\{0\}}$ factor. Now the maps $f_{\alpha n}$ always have the desired extension by the argument used in (2.3.16), and this induces an extension of f_{0n} by restricting and taking their disjoint union (using the fact that $\bigsqcup_\alpha X_\alpha$ is the coproduct in Top). □

2.3.18 Lemma. *Let $C \in \mathsf{Top}$, let $B_+, B_- \subseteq C$ be open subsets, and let $A := B_+ \cap B_-$. If there exists a continuous map $f : C \to \mathbb{R}$ such that $f^{-1}((-\varepsilon,\infty)) \subseteq B_+$ and $f^{-1}((-\infty,\varepsilon)) \subseteq B_-$, then the diagram*

$$\begin{array}{ccc} A & \longrightarrow & B_+ \\ \downarrow & & \downarrow \\ B_- & \longrightarrow & C \end{array} \tag{2.3.18.1}$$

is a pushout in Top^∞.

Proof. By definition (see (2.2.37) and (2.2.48)), we must show the extension property holds for maps $(\bullet \leftarrow \bullet \rightarrow \bullet) \star (\Delta^n, \partial\Delta^n) \to \mathsf{Top}^\infty$ which send $(\bullet \leftarrow \bullet \rightarrow \bullet) \star \Delta^0 = \Delta^1 \times \Delta^1$ ($\Delta^0 \subseteq \Delta^n$ being the initial vertex) to (2.3.18.1).

Let us label the vertices of the first term $\beta_- \leftarrow \alpha \rightarrow \beta_+$ and the vertices of the second term $\gamma, 1, \ldots, n$ (thus $\alpha, \beta_\pm, \gamma$ correspond to $A, B_\pm, C$). Now the extension property in question amounts to the extension property for a collection of maps

$$f_{\gamma n} : C \times (I,\partial I)^{\{1,\ldots,n-1\}} \to X_n \tag{2.3.18.2}$$
$$f_{\beta_\pm n} : B_\pm \times (I,0)^{\{\gamma\}} \times (I,\partial I)^{\{1,\ldots,n-1\}} \to X_n \tag{2.3.18.3}$$
$$f^\pm_{\alpha n} : A \times I^{\{\beta_\pm\}} \times (I,0)^{\{\gamma\}} \times (I,\partial I)^{\{1,\ldots,n-1\}} \to X_n \tag{2.3.18.4}$$

subject to the compatibility relations

$$f_{\beta_\pm n}|_{\cdots \times \{1\}^{\{\gamma\}} \times \cdots} = f_{\gamma n}|_{B_\pm \times \cdots} \tag{2.3.18.5}$$
$$f^\pm_{\alpha n}|_{\cdots \times \{1\}^{\{\beta_\pm\}} \times \cdots} = f_{\beta_\pm n}|_{A \times \cdots} \tag{2.3.18.6}$$
$$f^\pm_{\alpha n}|_{\cdots \times \{1\}^{\{\gamma\}} \times \cdots} = f_{\gamma n}|_{A \times \cdots} \tag{2.3.18.7}$$
$$f^+_{\alpha n}|_{\cdots \times \{0\}^{\{\beta_+\}} \times \cdots} = f^-_{\alpha n}|_{\cdots \times \{0\}^{\{\beta_-\}} \times \cdots} \tag{2.3.18.8}$$

We solve this extension problem as follows. We begin with $f^\pm_{\alpha n}$, which we prefer to view as a single function $f_{\alpha n}$ with domain $A \times J^\beta \times (I,0)^{\{\gamma\}} \times (I,\partial I)^{\{1,\ldots,n-1\}}$, where J^β denotes the union of I^{β_+} and I^{β_-} glued at 0 in accordance with (2.3.18.8). For any $t \in J^\beta$, the pair $(J^\beta, t) \times (I,0)^{\{\gamma\}} \times (I,\partial I)^{\{1,\ldots,n-1\}}$ admits a retraction, independent of the β-coordinate near $t_\gamma = 1$ in accordance with (2.3.18.7). Moreover, this retraction can be made to depend continuously on t. We take t to be a function of f, so that $t = 1 \in I^{\{\beta_\pm\}} \subseteq J^\beta$ when f is near $\pm\varepsilon$. Now the pair $\{t\} \times (I,0)^{\{\gamma\}} \times (I,\partial I)^{\{1,\ldots,n-1\}}$ also admits a retraction, so by composition we have extended $f^\pm_{\alpha n}$ as desired. This induces an extension of $f_{\beta_\pm n}$ which near $f = \pm\varepsilon$ is induced by a retraction of the pair $(I,0)^{\{\gamma\}} \times (I,\partial I)^{\{1,\ldots,n-1\}}$, hence it may be extended to all of $B_\pm$ by the same. We have now determined $f_{\gamma n}$ uniquely, and are done. □

2.3.19 Exercise. Use (2.3.18) to show that the pushout of $* \leftarrow S^{n-1} \to *$ in Top^∞ is S^n.

2.3.20 Proposition. Top^∞ *is cocomplete.*

Proof. It suffices to show that Top^∞ admits pushouts and arbitrary coproducts (2.2.57). Arbitrary coproducts exist in Top^∞ by (2.3.17). Certain diagrams are shown to be pushouts in (2.3.18), so it suffices to show that every diagram $X \leftarrow A \to Y$ is isomorphic in Top^∞ to one of the form (2.3.18.1). This can be done by replacing X with the mapping cylinder $(A \times I) \cup_A X$ and similarly with Y. □

2.3.21 Proposition. *Geometric realization* $|\cdot| : \mathcal{S} \to \mathsf{Top}^\infty$ *is cocontinuous.*

Proof. Preserving coproducts is immediate since coproducts are just disjoint unions in both $\mathcal{S}$ (2.3.7) and Top^∞ (2.3.17).

Now we should check preserving pushouts! We know from the proof of cocompletness of $\mathcal{S}$ (2.3.12) that every pushout in $\mathcal{S}$ takes, up to isomorphism, the form (2.3.11). Now we can consider the diagram with C in place of $\hat{C}$. In this case, we can find a continuous function $f : C \to \mathbb{R}$ vanishing precisely on A, non-negative on B_+, non-positive on B_-, such that a standard regular neighborhood of A contains $f^{-1}((-1,1))$ (indeed, can just let f take a standard form on each standard simplex, something like bump function near simplices which are part of A). This shows that the pushout in Top^∞ is C, whereas the pushout in $\mathcal{S}$ is $\hat{C}$, so we just need to show that the map $C \to \hat{C}$ is an isomorphism in Top^∞. We will show that $\hat{C}$ deformation retracts to $C \subseteq \hat{C}$. A given horn filling certainly deformation retracts, as does simultaneous separate horn fillings. Now at time 2^{-k} we execute the deformation retraction of the kth horn filling in the construction of $C \hookrightarrow \hat{C}$. Note that to check that a map out of $\hat{C}$ is continuous, we just need to check it is continuous on each simplex, and this certainly is true because each simplex of $\hat{C}$ is contained in some finite stage of the construction, so does not move until time 2^{-k} for some finite $k < \infty$. □

2.4 Sheaves

We discussed sheaves earlier in the context of categories and 2-categories. We now discuss sheaves in the context of ∞-categories, where the discussion takes a more homotopy theoretic and algebrotopological flavor.

2.4.1 Definition (Sheaves)**.** A *presheaf* on a topological space X valued in an ∞-category $\mathcal{C}$ is a functor $\mathsf{Op}(X)^{\mathsf{op}} \to \mathcal{C}$. The category of presheaves is denoted $\mathsf{PShv}(X, \mathcal{C}) = \mathsf{Fun}(\mathsf{Op}(X)^{\mathsf{op}}, \mathcal{C})$. A *sheaf* is a presheaf which satisfies *descent*, meaning that for every open set $U \subseteq X$ and every covering sieve $\mathcal{U}$ on U, the composition

$$\mathcal{U}^{\triangleleft} \to \mathsf{Op}(X)^{\mathsf{op}} \xrightarrow{F} \mathcal{C} \tag{2.4.1.1}$$

is a limit diagram, where $\mathcal{U}^{\triangleleft} \to \mathsf{Op}(X)^{\mathsf{op}}$ is the inclusion of the full subcategory spanned by $\mathcal{U}$ and U. The full subcategory of $\mathsf{PShv}(X, \mathcal{C})$ spanned by sheaves is denoted $\mathsf{Shv}(X, \mathcal{C})$.

2.4.2 Definition (Cosheaves)**.** The categories of cosheaves and precosheaves valued in $\mathcal{C}$ are $\mathsf{PCoshv}(X,\mathcal{C}) = \mathsf{PShv}(X,\mathcal{C}^{\mathsf{op}})^{\mathsf{op}}$ and $\mathsf{Coshv}(X,\mathcal{C}) = \mathsf{Shv}(X,\mathcal{C}^{\mathsf{op}})^{\mathsf{op}}$. In other words, a precosheaf on X is a functor $\mathsf{Op}(X) \to \mathcal{C}$, and cosheaf is a precosheaf which sends covers to colimit diagrams in $\mathcal{C}$.

2.4.3 Lemma. *If a sieve S is generated by open sets U_i, then the natural map*

$$\lim_{U\in S} \mathcal{F}(U) \to \lim_{0<|I|<\infty} \mathcal{F}\biggl(\bigcap_{i\in I} U_i\biggr) \tag{2.4.3.1}$$

is an isomorphism.

Proof. We have a chain of isomorphisms

$$\lim_{U\in S} \mathcal{F}(U) \to \lim_{U\in S} \lim_{\substack{0<|I|<\infty \\ U_i\supseteq U}} \mathcal{F}(U) = \lim_{0<|I|<\infty} \lim_{U\subseteq\bigcap_{i\in I} U_i} \mathcal{F}(U) \leftarrow \lim_{0<|I|<\infty} \mathcal{F}\biggl(\bigcap_{i\in I} U_i\biggr). \tag{2.4.3.2}$$

The first is because every $U \in S$ is contained in some U_i, the second is because limits commute with limits, and the last is because $U = \bigcap_{i\in I} U_i$ is initial. □

2.4.4 Example. The precosheaf $U \mapsto C_*(U)$ (singular chains) is a cosheaf $\mathsf{Top} \to \mathcal{D}(\mathbb{Z})$. This amounts to the assertion that for any topological space X with open covering $X = \bigcup_i U_i$, we have $C_*(X)$ is the colimit of $C_*(U_i)$. Consider the direct limit $C_*^b = \underrightarrow{\mathrm{colim}}(C_* \xrightarrow{b} C_* \xrightarrow{b} \cdots)$ of barycentric subdivision. Since each barycentric subdivision map is a quasi-isomorphism and directed colimits are exact, it follows that $C_* \to C_*^b$ is a quasi-isomorphism (hence an isomorphism in $\mathcal{D}(\mathbb{Z})$). Therefore it is equivalent to show that C_*^b is a cosheaf. For any open covering $X = \bigcup_i U_i$, the total complex of the double complex

$$C_*^b(X) \leftarrow \bigoplus_i C_*^b(U_i) \leftarrow \bigoplus_{i,j} C_*^b(U_i \cap U_j) \leftarrow \cdots \tag{2.4.4.1}$$

is exact. This amounts to the cosheaf property for C_*^b. Applying the duality functor $\underline{\mathrm{Hom}}(-,\mathbb{Z}) : \mathcal{D}(\mathbb{Z})^{\mathsf{op}} \to \mathcal{D}(\mathbb{Z})$ turns C_* into C^*, so we conclude that C^* is a sheaf (duality is continuous).

Similarly, the functor $S_* : \mathsf{Top} \to \mathcal{S}$ is a cosheaf.

2.4.5 Exercise. Show that the identity functor $\mathsf{Top} \to \mathsf{Top}$ is a cosheaf.

2.4.6 Lemma. *The tautological functor* $\mathsf{Top} \to \mathsf{Top}^\infty$ *is a cosheaf in the numerable topology.*

Proof. Let $\mathcal{U}$ be a numerable open cover of a topological space X. We are to show that $\mathcal{U} \to X$ is a colimit diagram in Top^∞. By (2.4.9), we are free to refine the cover $\mathcal{U}$. Thus suppose that $\mathcal{U}$ is locally finite and has a subordinate partition of unity $\varphi_i : X \to [0,1]$.

We claim that the colimit of $\mathcal{U}$ in Top^∞ is the identification space

$$\bigcup_{i_0<\cdots<i_k} \Delta^k \times (U_{i_0} \cap \cdots \cap U_{i_k}) \Big/ \sim \tag{2.4.6.1}$$

(order the cover arbitrarily for the ordering $<$). The proof of this is similar to (2.3.18). The point is that specifying a map out of this identification space is the same as specifying maps out of each U_i, homotopies between them over the overlaps $U_i \cap U_j$, homotopies of homotopies over triple overlaps $U_i \cap U_j \cap U_k$, etc.

Now we claim that the natural map from (2.4.6.1) to X is a homotopy equivalence. The partition of unity defines a section of this map. The identity map and the section are homotopic via the evident linear homotopy. □

2.4.7 Definition (Dagger construction)**.** Suppose $\mathcal{C}$ is complete and admits filtered colimits. The functor $\dagger : \mathsf{PShv}(X, \mathcal{C}) \to \mathsf{PShv}(X, \mathcal{C})$, denoted $\mathcal{F} \mapsto \mathcal{F}^\dagger$, is defined informally by the formula

$$\mathcal{F}^\dagger(X) = \underset{S \in J(X)}{\underrightarrow{\operatorname{colim}}} \lim_{U \in S} \mathcal{F}(U), \tag{2.4.7.1}$$

where $J(X)$ denotes the set of covering sieves on X. A precise definition of $\dagger$ uses Kan extensions, see [17, 6.2.2.9]. There is a tautological natural transformation $\mathbf{1} \Rightarrow \dagger$.

2.4.8 Exercise (Applying dagger to a constant presheaf). Suppose $\mathcal{C}$ is complete and admits filtered colimits, and denote by $\underline{A}$ the constant presheaf sending everything to $A \in \mathcal{C}$. Show that

$$\underline{A}^\dagger(X) = \begin{cases} A & X \neq \varnothing \\ * & X = \varnothing \end{cases} \tag{2.4.8.1}$$

where $* \in \mathcal{C}$ denote the terminal object, and show that

$$\underline{A}^{\dagger\dagger}(X) = \underset{S \in J(X)}{\underrightarrow{\operatorname{colim}}} \lim_{U \in S \setminus \{\varnothing\}} A. \tag{2.4.8.2}$$

2.4.9 Lemma ([17, 6.2.2.15])**.** *$\mathcal{F}$ is a sheaf iff $\mathcal{F} \to \mathcal{F}^\dagger$ is an isomorphism.*

Proof. The map $\mathcal{F} \to \mathcal{F}^\dagger$ being an isomorphism is the statement that

$$\mathcal{F}(X) \to \underset{S \in J(X)}{\underrightarrow{\operatorname{colim}}} \lim_{U \in S} \mathcal{F}(U) \tag{2.4.9.1}$$

is an isomorphism. On the other hand, $\mathcal{F}$ is a sheaf iff $\mathcal{F}(X) \to \lim_{U \in S} \mathcal{F}(U)$ is an isomorphism for every $S \in J(X)$. Thus the nontrivial direction is to show that if $\mathcal{F} \to \mathcal{F}^\dagger$ is an isomorphism then $\mathcal{F}$ is a sheaf.

It suffices to verify the sheaf property at a given object X, and by replacing our site with the slice category over X, we can assume X is the final object. Fix a covering sieve $S \in J(X)$, and let us show that $\mathcal{F}$ satisfies the sheaf property for S. The map $\mathcal{F} \to \mathcal{F}^\dagger$ factors through $\mathcal{F}^S$, where $\mathcal{F}^S(U)$ is the limit of $\mathcal{F}$ over the pullback of S to U. Now we have a diagram

$$\begin{array}{ccccc} \mathcal{F}(X) & \longrightarrow & \mathcal{F}^S(X) & \longrightarrow & \mathcal{F}^\dagger(X) \\ \downarrow & & \downarrow & & \downarrow \\ \lim\limits_{U \in S} \mathcal{F}(U) & \longrightarrow & \lim\limits_{U \in S} \mathcal{F}^S(U) & \longrightarrow & \lim\limits_{S \in U} \mathcal{F}^\dagger(U) \end{array} \tag{2.4.9.2}$$

The middle vertical map is an isomorphism by inspection (since S pulled back to any object of S is the full over-category). Since $\mathcal{F} \to \mathcal{F}^\dagger$ is an isomorphism, this means that the left/right vertical maps (which are equal) are retracts of the middle vertical map (and a retract of an isomorphism in any category is an isomorphism). □

2.4.10 Definition (Sheafification)**.** Suppose $\mathcal{C}$ is complete and has filtered colimits. The inclusion $\mathsf{Shv}(X,\mathcal{C}) \subseteq \mathsf{PShv}(X,\mathcal{C})$ has a left adjoint called the *sheafification* functor. It may be defined by a transfinite iteration of † [17, 6.2.2.7]. We need a regular cardinal κ such that $\lim_{U\in S}$ commutes with κ-filtered colimits for every covering sieve S on any subset of X. Then we define $\dagger^\alpha$ for ordinals $\alpha < \kappa$, and we claim that $\dagger^\kappa$ (identifying κ with the first ordinal of cardinality κ) is sheafification. To see that $\dagger^\kappa$ is a sheaf, we use the factorization from (2.4.9): each map $\dagger^\gamma\mathcal{F} \to \dagger^{\gamma+1}\mathcal{F}$ factors through a presheaf for which the sheaf property holds at X. Thus $\dagger^\kappa\mathcal{F}$ is a κ-filtered colimit of presheaves for which the sheaf property holds at X. Since κ-filtered colimits commute with the limit over any sieve, we conclude that $\dagger^\kappa\mathcal{F}$ satisfies the sheaf property at X. Since † is compatible with restriction, the same holds for any open subset of X. For any sheaf $\mathcal{G}$, the natural map $\operatorname{Hom}(\dagger^\kappa\mathcal{F},\mathcal{G}) \to \operatorname{Hom}(\mathcal{F},\mathcal{G})$ is an isomorphism (the case $\kappa = 1$ is (**??**), and this implies the general case by transfinite induction on κ and the universal property of colimits). It is immediate from its incarnation as $\dagger^\kappa$ that sheafification commutes with any functor $\mathcal{C} \to \mathcal{D}$ which is continuous and preserves filtered colimits.

2.4.11 Definition (Locally constant sheaf)**.** There is a tautological map $\mathcal{C} \to \mathsf{PShv}(-,\mathcal{C})$ which we may compose with sheafification $\mathsf{PShv}(-,\mathcal{C}) \to \mathsf{Shv}(-,\mathcal{C})$. A sheaf in the essential image of this map $\mathcal{C} \to \mathsf{Shv}(-,\mathcal{C})$ is called a constant sheaf. A sheaf is called locally constant iff its restriction to every element of some open cover is constant. We denote by $\mathsf{Shv}_{\mathsf{locconst}} \subseteq \mathsf{Shv}$ the full subcategory spanned by sheaves which are locally constant.

2.4.12 Definition (Pushforward and pullback)**.** Given a map of topological spaces $f : X \to Y$, the pushforward functor $f_* : \mathsf{Shv}(X,\mathcal{C}) \to \mathsf{Shv}(Y,\mathcal{C})$ is defined by $(f_*\mathcal{F})(U) := \mathcal{F}(f^{-1}(U))$ (this defines a functor $\mathsf{PShv}(X,\mathcal{C}) \to \mathsf{PShv}(Y,\mathcal{C})$ which sends sheaves to sheaves (**??**)). Pushforward has a left adjoint pullback $f^* : \mathsf{Shv}(Y,\mathcal{C}) \to \mathsf{Shv}(X,\mathcal{C})$ which is the composition of the presheaf pullback $U \mapsto \varinjlim_{f(U)\subseteq V} \mathcal{F}(V)$ with sheafification.

2.5 Shape theory

We saw in (2.3) that the homotopy theory of Kan complexes $\mathcal{S}$ embeds as a full subcategory of the homotopy theory of arbitrary topological spaces Top^∞ via the geometric realization functor $|\cdot| : \mathcal{S} \to \mathsf{Top}^\infty$, whose essential image consists of those topological spaces with the homotopy type of a CW-complex. This inclusion has a right adjoint $S_* : \mathsf{Top}^\infty \to \mathcal{S}$ (singular simplicial set), which gives one way of going in the other direction, i.e. from arbitrary topological spaces to Kan complexes / CW-complexes. Such a functor lets one apply invariants from algebraic topology (which are defined on $\mathcal{S}$) to arbitrary topological spaces.

There are many pertinent topological phenomena which are not captured by the singular simplicial set functor $S_* : \mathsf{Top}^\infty \to \mathcal{S}$. In particular, it is not what is usually relevant when discussing the homotopy theory of compact Hausdorff spaces. For example, if $W \in \mathsf{Top}^\infty$ denotes the Warsaw circle (2.5.1), then $S_*W = *$ is a single point, whereas we would much prefer a functor $\mathsf{Top}^\infty \to \mathcal{S}$ which sends W to the cricle S^1. In fact, a theory of virtual fundamental cycles essentially requires having such a functor.

2.5.1 Example (Warsaw circle). The *Warsaw circle* W is a well-known 'pathological' topological space.

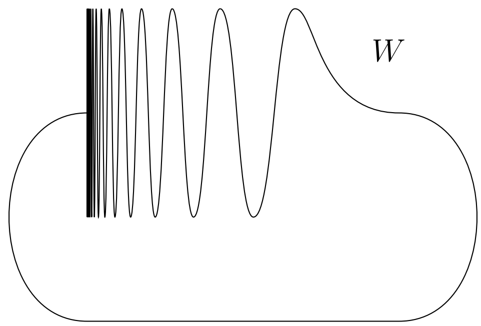

(2.5.1.1)

The fundamental group of W is trivial: maps from $\mathbb{R}$ cannot pass across the singularity, and in fact $S_*W = *$. The shape of W is the circle S^1. Thus the shape captures the cycle of 'going around once' on the Warsaw circle, whereas S_* does not. If we imagine the Warsaw circle cut out of $\mathbb{R}^2$ as the zero set of a function which is negative inside and positive outside, then the induced cycle on W should 'go around once', i.e. it should be the fundamental cycle of S^1, which exists in $\operatorname{shape}(W)$ but not in S_*W. If we take an annular neighborhood A of W, it should certainly be the case that the pushforward of $[W]^{\mathrm{vir}}$ to A is the generator of $H_1(A)$. This is evidently impossible if we view $[W]^{\mathrm{vir}}$ as living on S_*W. We should therefore view it as living on $\operatorname{shape}(W)$.

The *shape functor* is the desired functor from topological spaces to spaces (more precisely, pro-spaces). It arises by asking for the left adjoint to the inclusion $\mathcal{S} \subseteq \mathsf{Top}^\infty$. Concretely speaking, this just means that, whereas the singular simplicial set functor asks for a given topological space how CW-complexes map to it, the shape functor asks how it maps to CW-complexes.

2.5.2 Definition (Shape functor)**.** Asking for the left adjoint of geometric realization $|\cdot| : \mathcal{S} \to \mathsf{Top}^\infty$ is the same as asking for the composition

$$\mathsf{Top}^\infty \xrightarrow{\mathcal{Y}^{\mathsf{op}}_{(\mathsf{Top}^\infty)^{\mathsf{op}}}} \mathcal{P}((\mathsf{Top}^\infty)^{\mathsf{op}})^{\mathsf{op}} \xrightarrow{|\cdot|^*} \mathcal{P}(\mathcal{S}^{\mathsf{op}})^{\mathsf{op}} \tag{2.5.2.1}$$

to land inside representable functors $\mathcal{S} \subseteq \mathcal{P}(\mathcal{S}^{\mathsf{op}})^{\mathsf{op}}$. This is not the case, however it does land inside pro-objects (2.5.3), and this defines the shape functor $\operatorname{shape} : \mathsf{Top}^\infty \to \mathsf{Pro}\,\mathcal{S}$.

2.5.3 Lemma. *The essential image of* (2.5.2.1) *is contained in* $\mathsf{Pro}\,\mathcal{S} \subseteq \mathcal{P}(\mathcal{S}^{\mathsf{op}})^{\mathsf{op}}$.

Proof. We are to show that for $X \in \mathsf{Top}^\infty$, the comma category $\mathcal{S}_{X/|\cdot|}$ is cofiltered. It is equivalent to show that the functor $\operatorname{Hom}_{\mathsf{Top}^\infty}(X, |\cdot|) : \mathcal{S} \to \mathcal{S}$ preserves finite limits. The Yoneda functor $\operatorname{Hom}_{\mathsf{Top}^\infty}(X, \cdot) : \mathsf{Top}^\infty \to \mathcal{S}$ preserves all limits (**??**), and $|\cdot| : \mathcal{S} \to \mathsf{Top}^\infty$ preserves finite limits by (**??**). $\square$

2.5.4 Exercise. Show that the shape functor is cocontinuous.

2.5.5 Exercise. Show that the composition $\mathsf{Top} \to \mathsf{Top}^\infty \to \mathsf{Pro}\,\mathcal{S}$ is a cosheaf in the numerable topology (use (2.4.6)).

Let us derive another nice description of the cosheaf $U \mapsto \operatorname{shape}(U)$. Consider the constant precosheaf $\underline{*}$ defined by $\underline{*}(U) = *$; there is an evident map $\operatorname{shape}(-) \to \underline{*}$. Let us regard $\underline{*}$ as valued in $\mathsf{Pro}\,\mathcal{S}$. Since $\mathcal{S}$ is complete, $\mathsf{Pro}\,\mathcal{S}$ is copresentable, so the dagger functor (2.4.7) is defined, and its transfinite iteration gives the cosheafification functor for precosheaves valued in $\mathsf{Pro}\,\mathcal{S}$. Since shape is a cosheaf in the numerable topology, there is thus an induced map $\operatorname{shape}(-) \to \underline{*}^{\mathrm{cosh}}_{\mathrm{num}}$, the target being the cosheafification with respect to the numerable topology.

2.5.6 Proposition. *The map* $\operatorname{shape}(-) \to \underline{*}^{\dagger\dagger}_{\mathrm{num}}$ *is an isomorphism.*

Proof. The functors $\operatorname{Hom}(-, Z) : (\mathsf{Pro}\,\mathcal{S})^{\mathsf{op}} \to \mathcal{S}$ for $Z \in \mathcal{S}$ are continuous and preserve filtered colimits, and together they reflect isomorphisms. It thus suffices to show that $\underline{Z}^{\dagger\dagger}_{\mathrm{num}} \to \operatorname{Hom}(\operatorname{shape}(-), Z)$ is an isomorphism for every $Z \in \mathcal{S}$. Realizing Z as a topological space, this is the map of presheaves $\underline{S_* Z} \to S_* \underline{\operatorname{Hom}}(-, Z)$ (presheaves on Top valued in $\mathcal{S}$) after applying $\dagger$ twice to the domain.

We know what $\dagger\dagger$ does to a constant presheaf (2.4.8), so we are to show for every topological space X that the following map is an isomorphism in $\mathcal{S}$:

$$\varinjlim_{S \in J(X_{\mathrm{num}})} S_* \underline{\operatorname{Hom}}(N(S), Z) \to S_* \underline{\operatorname{Hom}}(X, Z) \tag{2.5.6.1}$$

(the existence of this map relies on (2.5.5), ultimately on (2.4.6)). This is proven for X paracompact in [17, 7.1.6.8]; we will given an argument for general X.

We may take Z to be a simplicial complex with its metric induced from the embedding into ℓ^1 (such spaces are precisely $\mathcal{S} \subseteq \mathsf{Top}^\infty$). Since Z is metrizable, the mapping space $\underline{\operatorname{Hom}}(\Delta^p, Z)$ is metrizable (1.7.17), hence paracompact Hausdorff. We consider the open cover of $\underline{\operatorname{Hom}}(\Delta^p, Z)$ by the sets $\{f : f(K_i) \subseteq \star(\sigma_i)\}$ for a finite cover by compact sets $\Delta^n = \bigcup_i K_i$ and simplices $\sigma_i \subseteq Z$ (where $\star$ denotes the open star). We will call this the *star cover* of $\underline{\operatorname{Hom}}(\Delta^p, Z)$. Since $\underline{\operatorname{Hom}}(\Delta^p, Z)$ is paracompact Hausdorff, the star cover is numerable (this is in some sense the key fact underlying in this proof).

The target $S_* \underline{\operatorname{Hom}}(X, Z)$ is the simplicial set $[n] \mapsto \operatorname{Hom}(X, \underline{\operatorname{Hom}}(\Delta^n, Z))$. For any numerable open cover S, we can consider the simplicial subset $\operatorname{Hom}^S(X, \underline{\operatorname{Hom}}(\Delta^n, Z))$ consisting of those maps $X \to \underline{\operatorname{Hom}}(\Delta^n, Z)$ for which S refines the pullback of the star cover. Since the star cover is numerable, we have

$$\operatorname{Hom}(X \times \Delta^n, Z) = \operatorname{Hom}(X, \underline{\operatorname{Hom}}(\Delta^n, Z)) = \varinjlim_{S \in J(X_{\mathrm{num}})} \operatorname{Hom}^S(X, \underline{\operatorname{Hom}}(\Delta^n, Z)). \tag{2.5.6.2}$$

Each simplicial set $[n] \mapsto \operatorname{Hom}^S(X, \underline{\operatorname{Hom}}(\Delta^n, Z))$ is a Kan complex by the same argument that $S_* \underline{\operatorname{Hom}}(X, Z)$ is (2.1.22).

Now we contemplate the simplicial set $[n] \to \mathrm{Hom}^S(X, \underline{\mathrm{Hom}}(\Delta^n, Z))$. It parameterizes families of maps $X \to Z$ which send every element of S to a subset of some star of Z. Such a map determines a map $N(S) \to Z$, so this simplicial set admits a map to $S_* \underline{\mathrm{Hom}}(N(S), Z)$. We have thus defined a map in the reverse direction of (2.5.6.1). Some inspection will show these maps are inverses. □

2.5.7 Exercise. Let $\mathcal{C}$ be complete and admit filtered colimits, let $A \in \mathcal{C}$, and consider the constant presheaf $\underline{A}$ on topological spaces. Use (2.4.8) and (2.5.6) to show that $\underline{A}^{\dagger\dagger}_{\mathrm{num}}(X)$ is given by the image of shape$(X) \in \mathsf{Pro}\,\mathcal{S}$ under the functor

$$(\mathsf{Pro}\,\mathcal{S})^{\mathsf{op}} = \mathsf{Ind}\,\mathcal{S}^{\mathsf{op}} \xrightarrow{\mathsf{Ind}(\lim A)} \mathsf{Ind}\,\mathcal{C} \xrightarrow{\mathrm{colim}} \mathcal{C} \tag{2.5.7.1}$$

where $\lim A : \mathcal{S}^{\mathsf{op}} \to \mathcal{C}$ is the unique continuous functor sending $*$ to A.

2.6 Spanier–Whitehead duality

Spanier–Whitehead duality is an operation on spectra which is analogous to the functor sending a vector space to its dual.

2.6.1 Exercise. Show that the functor $\underline{\mathrm{Hom}}(-, k) : \mathrm{Vect}_k \to \mathrm{Vect}_k$ is cocontinuous and sends k to k. Show that these properties characterize it uniquely up to unique natural isomorphism.

2.6.2 Definition (Spanier–Whitehead duality)**.** The functor $D : \mathrm{Sp} \to \mathrm{Sp}^{\mathsf{op}}$ is the unique cocontinuous functor satisfying $D(S^0) = S^0$ (existence and uniqueness of D follows from the universal property of the category of spectra (**??**)).

Since D is cocontinuous, it preserves finite colimits, hence also finite limits. Its square $D^2 = D^{\mathsf{op}} \circ D : \mathrm{Sp} \to \mathrm{Sp}$ thus preserves finite limits and finite colimits and satisfies $D^2(S^0) = S^0$. It follows that D^2 is the identity functor on the category of finite spectra. More generally, the property that $D^2(S^0) = S^0$ determines, by the universal property of Kan extension ($\mathsf{Fun}(\mathrm{Sp}^f, \mathrm{Sp}) \hookrightarrow \mathsf{Fun}(\mathrm{Sp}, \mathrm{Sp})$), a natural transformation $\mathbf{1} \Rightarrow D^2$.

2.6.3 Example. Since D is cocontinuous, there is a natural isomorphism $\Omega D = D\Sigma$ and natural isomorphisms $D(X \vee Y) = DX \vee DY$ (note that $\vee$ is both the product and coproduct in Sp). In particular, $DS^n = D\Sigma^n S^0 = \Omega^n DS^0 = \Omega^n S^0 = S^{-n}$.

2.6.4 Exercise. Describe $D(\bigvee_{i=1}^{\infty} S^1)$.

Once we get to Verdier duality (2.9.1), we will see how it can be used to derive a concrete 'formula' for the Spanier–Whitehead dual of a finite complex (2.9.2).

2.7 Proper pushforward

For a locally proper map $f : X \to Y$, there is a proper pushforward functor $f_!$ on sheaves which can be thought of as sections with proper support. Under favorable hypotheses, it has a right adjoint $f^!$.

2.7.1 Proposition. *Suppose $\mathcal{C}$ is complete, admits filtered colimits, and has an initial object. For any open inclusion $j : U \hookrightarrow X$, pullback j^* has a left adjoint $j_!$.*

Proof. Consider the functor $j_!^{\mathrm{pre}} : \mathsf{Shv}(U, \mathcal{C}) \to \mathsf{PShv}(X, \mathcal{C})$ defined by

$$(j_!^{\mathrm{pre}}\mathcal{F})(V) = \begin{cases} \mathcal{F}(V) & V \subseteq U \\ \varnothing & V \nsubseteq U \end{cases} \tag{2.7.1.1}$$

where $\varnothing \in \mathcal{C}$ denotes the initial object. More formally, $j_!^{\mathrm{pre}} : \mathsf{PShv}(U, \mathcal{C}) \to \mathsf{PShv}(X, \mathcal{C})$ is the left Kan extension functor which is left adjoint to $j^* : \mathsf{PShv}(X, \mathcal{C}) \to \mathsf{PShv}(U, \mathcal{C})$; it is a pointwise left Kan extension (for $\mathsf{Op}(U) \to \mathsf{Op}(X)$, the relevant comma category for $V \subseteq X$ not contained in U is empty, so has colimit the initial object $\varnothing \in \mathcal{C}$).

It is evident that $\mathrm{Hom}(j_!^{\mathrm{pre}}\mathcal{F}, \mathcal{G}) = \mathrm{Hom}(\mathcal{F}, j^*\mathcal{G})$ (the first Hom space involves no data over open sets $V \nsubseteq U$, and for $V \subseteq U$, it is precisely a Hom on the right; more formally, this is just the adjoint property of Kan extensions). Now note that $\mathrm{Hom}((j_!^{\mathrm{pre}}\mathcal{F})^{\mathrm{sh}}, -) = \mathrm{Hom}(j_!^{\mathrm{pre}}\mathcal{F}, -)$ when the second argument is a sheaf, so we may set $j_!\mathcal{F} = (j_!^{\mathrm{pre}}\mathcal{F})^{\mathrm{sh}}$. ☐

2.7.2 Exercise. Show that for any open inclusion $j : U \to X$ the unit map $\mathbf{1} \Rightarrow j^*j_!$ is an isomorphism, hence $j_!$ is fully faithful.

2.7.3 Definition (Proper pushforward $f_!$ and pullback $f^!$)**.** Suppose $\mathcal{C}$ is complete, admits filtered colimits, and has a zero object. For a locally proper map $f : X \to Y$ (**??**), the functor

$$f_! : \mathsf{Shv}(X, \mathcal{C}) \to \mathsf{Shv}(Y, \mathcal{C}) \tag{2.7.3.1}$$

is defined as $q_*j_!$ for any factorization of f as a composition $X \overset{j}{\hookrightarrow} \bar{X} \overset{q}{\to} Y$ of an open inclusion j and a proper map q. The right adjoint of $f_!$ is denoted $f^!$ (if it exists; a sufficient condition for the existence of $f^!$ is given in (**??**)).

To see that $f_!$ is independent, up to natural isomorphism, of the choice of factorization, suppose given two factorizations $X \overset{j}{\hookrightarrow} \bar{X} \overset{q}{\to} Y$ and $X \overset{j'}{\hookrightarrow} \bar{X}' \overset{q'}{\to} Y$. Let Z be the closure of the image of X inside $\bar{X} \times_Y \bar{X}'$. Evidently $Z \to \bar{X}$ and $Z \to \bar{X}'$ are proper. There is an open subset $X \times_Y X \subseteq \bar{X} \times_Y \bar{X}'$, and the map $X \to X \times_Y X$ is a closed embedding since $X \to Y$ is separated, so we conclude that $X \to Z$ is an open inclusion. To show that $q_*j_! = q'_*j'_!$, it thus suffices to show that for any composition $X \overset{j}{\hookrightarrow} \bar{X} \overset{q}{\to} Y$ (j open embedding, q proper) for which qj is also an open embedding, that $q_*j_! = (qj)_!$. (This may require additional hypotheses on X, Y, and/or $\mathcal{C}$.)

2.8 Fundamental cycles

2.8.1 Definition (Cycles)**.** Let X be a locally compact Hausdorff space, and let $\pi : X \to *$. For a cosheaf $\mathcal{F}$ of spectra on X, we will regard $\pi_!\mathcal{F}$ as the *space of (Borel–Moore) $\mathcal{F}$-valued cycles on X*. An *$\mathcal{F}$-valued cycle on X* is a map $S^0_{\mathrm{cosh}} \to \pi_!\mathcal{F}$ (equivalently, a map $\pi^! S^0_{\mathrm{cosh}} \to \mathcal{F}$).

2.8.2 Definition (Fundamental cycles). The *fundamental cycle* of a locally compact Hausdorff space X is the $\pi^! S^0_{\text{cosh}}$-valued cycle on X given by the unit map $S^0_{\text{cosh}} \to \pi_! \pi^! S^0_{\text{cosh}}$ (equivalently, the identity map $\pi^! S^0_{\text{cosh}} \to \pi^! S^0_{\text{cosh}}$).

The map $S^0_{\text{cosh}} \to \pi_! \mathcal{F}$ corresponding to a given map $c : \pi^! S^0_{\text{cosh}} \to \mathcal{F}$ is, by definition the composition $S^0_{\text{cosh}} \to \pi_! \pi^! S^0_{\text{cosh}} \xrightarrow{\pi_! c} \pi_! \mathcal{F}$. We can make a stronger statement: the cosheaf $\pi^! S^0_{\text{cosh}}$ on X equipped with the fundamental cycle is the initial object of the category of cosheaves of spectra $\mathcal{F}$ equipped with an $\mathcal{F}$-valued cycle. That is, every cycle on X is the pushforward of the fundamental cycle on X by a unique map. This justifies the term 'fundamental cycle'.

2.8.3 (Computation of $\pi^! S^0_{\text{cosh}}$). To understand the cosheaf $\pi^! S^0_{\text{cosh}}$, we appeal to Spanier–Whitehead duality. The Spanier–Whitehead dual of $\pi^! S^0_{\text{cosh}}$ is $\pi^! S^0_{\text{sh}}$, which sends U to $U/\infty := \varprojlim_{K \subseteq U} U/(U - K)$. Thus $(\pi^! S^0_{\text{cosh}})(U) = D(U/\infty)$. Applying $\pi_!$ cancels the $/\infty$, so $\pi_! \pi^! S^0_{\text{cosh}} = DX$, and the map $S^0 \to DX$ is simply the Spanier–Whitehead dual of the map $X \to *$. Alternatively, $\pi_! \pi^! S^0_{\text{cosh}}$ is the colimit over $p \in X$ rel infinity of $D(U_x/\infty)$. In the case $X = M$ is a topological manifold, this is $(M/\infty)^{-TM}$. We conclude an identification of $DM = (M/\infty)^{-TM}$, and a fundamental cycle $S^0 \to (M/\infty)^{-TM}$, the Spanier–Whitehead dual of the map $M \to *$. In the case $X = M$ is a topological manifold-with-boundary, this cosheaf is $j_!$ of $-TM$ on the interior. So, in this case, $\pi_! \pi^! S^0_{\text{cosh}}$ is the same for M as for $\bar{M}$.

2.8.4 Definition (Relative cycles). Let $f : X \to Y$ be a locally proper map of topological spaces, and let $\mathcal{F}$ and $\mathcal{G}$ be cosheaves of spectra on X and Y, respectively. A *relative cycle* is a map $\mathcal{G} \to f_! \mathcal{F}$ (equivalently, $f^! \mathcal{G} \to \mathcal{F}$).

A relative cycle $\mathcal{G} \to f_! \mathcal{F}$ evidently induces a map $\pi_{Y!} \mathcal{G} \to \pi_{X!} \mathcal{F}$ (i.e. from the space of $\mathcal{G}$-valued cycles on Y to the space of $\mathcal{F}$-valued cycles on X). More generally, relative cycles compose, in the sense that for $X \xrightarrow{f} Y \xrightarrow{g} Z$, the composition of $g_!(\mathcal{G} \to f_! \mathcal{F})$ and $\mathcal{H} \to g_! \mathcal{G}$ is a map $\mathcal{H} \to g_! f_! \mathcal{F}$. Relative cycles can also be pulled back. Namely, fix a pullback square

$$\begin{array}{ccc} X' & \xrightarrow{a_X} & X \\ \downarrow{\scriptstyle f'} & & \downarrow{\scriptstyle f} \\ Y' & \xrightarrow{a_Y} & Y \end{array} \tag{2.8.4.1}$$

where f is locally proper, let $\mathcal{F}$ and $\mathcal{G}$ be cosheaves of spectra on X and Y, and denote by $\mathcal{F}'$ and $\mathcal{G}'$ their pullbacks to X' and Y'. Proper base change (**??**) identifies the pullback of $f_! \mathcal{F}$ to Y' with $f'_! \mathcal{F}'$. Thus the pullback of a map $\mathcal{G} \to f_! \mathcal{F}$ is a map $\mathcal{G}' \to f'_! \mathcal{F}'$.

2.8.5 Example. Let $f : X \to Y$ be locally proper, and fix a cosheaf of spectra $\mathcal{F}$ on X. We have $\pi_{X!} \mathcal{F} = \pi_{Y!} f_! \mathcal{F}$; that is, an $\mathcal{F}$-valued cycle on X is the same as an $f_! \mathcal{F}$-valued cycle on Y. This identification is induced by the relative cycle $\mathbf{1} : f_! \mathcal{F} \to f_! \mathcal{F}$ (equivalently, the unit map $f^! f_! \mathcal{F} \to \mathcal{F}$).

2.8.6 Example. Given a cosheaf of spectra $\mathcal{G}$ on X and a closed embedding $i : Z \to X$, we may understand the cosheaf of spectra $i^! \mathcal{G}$ as follows. It assigns to an open set $U \subseteq Z$ the inverse limit of $\mathcal{G}(V)/\mathcal{G}(V-U)$ over all open subsets $V \subseteq X$ with $V \cap Z = U$. For example, considering the inclusion $0 : * \to E$ where E is a finite-dimensional real vector space, we have that $0^! S^0_{\mathrm{cosh}} = S^E_{\mathrm{cosh}}$.

2.8.7 Definition (Relative fundamental cycles)**.** The *relative fundamental cycle* of a locally proper map $f : X \to Y$ with respect to a cosheaf of spectra $\mathcal{G}$ on Y is the unit map $\mathcal{G} \to f_! f^! \mathcal{G}$.

2.8.8 Example (Thom map). Let $\pi : E \to X$ be a vector bundle, $s : X \to E$ a section, and $i : Z \to X$ the inclusion of the zero locus of s. We have a diagram

$$\begin{array}{ccc} Z & \longrightarrow & X \\ \big\downarrow & & \big\downarrow{\scriptstyle 0} \\ X & \xrightarrow{s} & E \end{array} \tag{2.8.8.1}$$

Fix any cosheaf of spectra $\mathcal{F}$ on X. The map $0 : X \to E$ has a relative fundamental cycle $\pi^*\mathcal{F} \to 0_! 0^! \pi^* \mathcal{F}$. Its pullback under $s : X \to E$ is a relative cycle $\mathcal{F} \to i_! i^* 0^! \pi^* \mathcal{F}$. The cosheaf $0^! \pi^* \mathcal{F}$ may be identified with $\mathcal{F} \wedge S^E$. We thus obtain a relative cycle

$$\mathcal{F} \to i_! i^* (\mathcal{F} \wedge S^E). \tag{2.8.8.2}$$

This relative cycle gives rise to the Thom/Gysin/Umkehr map from $\mathcal{F}$-valued cycles on X to $(\mathcal{F} \wedge S^E)|_Z$-valued cycles on Z.

2.8.9 Example. Let X be a topological space, and let V be a vector bundle on X. We search for a map of spectra $S^0 \to (X/\infty)^{-V}$, where $(X/\infty)^{-V}$ is defined as $\pi_!(-V)_{\mathrm{cosh}}$, where $(-V)_{\mathrm{cosh}}$ denotes the cosheaf on X associated to the local system given by the Spanier–Whitehead dual of the spherical fibration associated to V. This is just a map $S^0_{\mathrm{cosh}} \to \pi_!(-V)_{\mathrm{cosh}}$, which by adjunction is the same thing as a map $\pi^! S^0_{\mathrm{cosh}} \to (-V)_{\mathrm{cosh}}$. Passing to Spanier–Whitehead duals, this is the same as a map $V_{\mathrm{sh}} \to \pi^* S^0_{\mathrm{sh}}$. Recalling that $(\pi^* S^0_{\mathrm{sh}})(U) = U/\infty$, such a map would be encoded by a compatible family of 'exponential maps'

$$V_p/(V_p - 0) \to X/(X-p) \tag{2.8.9.1}$$

for $p \in X$. When X is a moduli space of pseudo-holomorphic curves, it will be more convenient to construct instead the inverse to this map and then argue that it is an isomorphism.

2.9 Verdier duality

Verdier duality is an equivalence $\mathbb{D}$ on any locally compact Hausdorff space between the categories of sheaves and cosheaves valued in a stable ∞-category. Concretely, Verdier duality may be defined by $(\mathbb{D}\mathcal{F})(U) = \varprojlim_{K \subseteq U} \mathcal{F}(U)/\mathcal{F}(U-K)$. Verdier duality exchanges the adjoint pairs (f^*, f_*) and $(f_!, f^!)$.

2.9.1 Verdier Duality ([26] [18, 5.5.5.1]). *For any stable ∞-category $\mathcal{C}$ which is complete and cocomplete and any locally compact Hausdorff space X, there is an equivalence of categories*

$$\mathbb{D} : \mathsf{Coshv}(X, \mathcal{C}) \xleftrightarrow{\sim} \mathsf{Sh}(X, \mathcal{C}) \tag{2.9.1.1}$$

squaring to the identity, defined by

$$(\mathbb{D}\mathcal{F})(U) = \varprojlim_{K \subseteq U} \mathcal{F}(U)/\mathcal{F}(U - K) \tag{2.9.1.2}$$

Verdier duality exchanges the adjoint pairs (f^*, f_*) and $(f_!, f^!)$ in the sense that $\mathbb{D}f^* = f^!\mathbb{D}$ and $\mathbb{D}f_* = f_!\mathbb{D}$. Note that this is consistent with adjunctions since on cosheaves the adjunctions are reversed (as cosheaf categories are the opposite of sheaf categories).

2.9.2 Example. Let M be a topological manifold and $\pi : M \to *$. There are four (co)sheaves of particular interest, related by Spanier–Whitehead duality D and Verdier duality $\mathbb{D}$ as follows:

$$\begin{array}{ccccc} \pi^* S^0_{\mathrm{cosh}} = U & \overset{D}{\text{———}} & (U/\infty)^{-TU} = \pi^* S^0_{\mathrm{sh}} \\ \Big| \mathbb{D} & & \Big| \mathbb{D} \\ \pi^! S^0_{\mathrm{sh}} = U/\infty & \overset{D}{\text{———}} & U^{-TU} = \pi^! S^0_{\mathrm{cosh}} \end{array} \tag{2.9.2.1}$$

where S^0_{sh} and S^0_{cosh} denote the constant sheaf and cosheaf on $*$ with value S^0. We begin at the upper left corner $\pi^* S^0_{\mathrm{cosh}}(U) = \mathrm{colim}_U\, S^0 = U$, which holds provided U is paracompact. The lower left corner $\pi^! S^0_{\mathrm{sh}}(U) = U/\infty$ then follows by definition of Verdier duality and the fact that $\mathbb{D}\pi^* = \pi^!\mathbb{D}$ and $\mathbb{D}S^0_{\mathrm{sh}} = S^0_{\mathrm{cosh}}$. Now we note that the sheaf $U \mapsto U/\infty$ has another name: it is the sheaf associated to the local system TM (the tangent spherical fibration of M). Applying Spanier–Whitehead dual thus gives the cosheaf associated to the local system $-TM = D(TM)$, and the notation U^{-TU} simply means by definition the sections over U of this cosheaf (or, rather, the colimit over U of $-TU$, which is the same thing). This equals $\pi^! S^0_{\mathrm{cosh}}$ since $DS^0_{\mathrm{cosh}} = S^0_{\mathrm{sh}}$. Finally, applying Verider duality again and again using $\mathbb{D}\pi^* = \pi^!\mathbb{D}$ and $\mathbb{D}S^0_{\mathrm{sh}} = S^0_{\mathrm{cosh}}$, we see that $\pi^* S^0_{\mathrm{sh}}$ is the Verdier dual of $U \mapsto U^{-TU}$, which by definition is $U \mapsto (U/\infty)^{-TU}$. Going around the commuting diagram above in this way proves the identity $DU = (U/\infty)^{-TU}$.

The adjunctions between π_* and π^* and between $\pi_!$ and $\pi^!$ give us maps, corresponding to the diagram above,

$$\pi_*\pi^* S^0_{\mathrm{cosh}} \to S^0_{\mathrm{cosh}} \qquad S^0_{\mathrm{sh}} \to \pi_*\pi^* S^0_{\mathrm{sh}} \tag{2.9.2.2}$$

$$\pi_!\pi^! S^0_{\mathrm{sh}} \to S^0_{\mathrm{sh}} \qquad S^0_{\mathrm{cosh}} \to \pi_!\pi^! S^0_{\mathrm{cosh}} \tag{2.9.2.3}$$

The maps on the left are $M \to *$, and the maps on the right are both its Spanier–Whitehead dual $S^0 \to DM = (M/\infty)^{-TM}$. This latter map is the fundamental cycle of M.

2.10 Derived category of vector bundles

2.10.1 Definition. A *differential graded category* is a category C enriched over cochain complexes of abelian groups.

2.10.2 Definition. The *differential graded nerve* $\mathcal{C}$ of a differential graded category C has n-simplices the tuples of objects $X_0, \ldots, X_n \in \mathsf{C}$ together with maps $f_{ij} : \mathbb{Z}(\Delta^1)^{\{i+1,\ldots,j-1\}} \to \mathsf{C}(X_i, X_j)$ such that $f_{ik}|_{\{t_j=1\}} = f_{ij} \times f_{jk}$. This data pulls back under a map $s : \Delta^m \to \Delta^n$ exactly as in the definition of the simplicial nerve (2.3.2).

2.10.3 Exercise. Describe explicitly the 0-simplices (objects), 1-simplices (morphisms), and 2-simplices of the differential graded nerve $\mathcal{C}$ of a differential graded category C. Consider the subcategory of $\mathcal{C}$ consisting of those simplices in which every f_{ij} is constant; how is this related to C? What is the homotopy category of $\mathcal{C}$ in terms of C?

2.10.4 Exercise. Show that the differential graded nerve of a differential graded category is an ∞-category (compare (2.3.4)).

2.10.5 Definition (Derived category of vector bundles $\mathcal{D}^b_{\mathrm{Vect}}(X)$)**.** We define $\mathcal{D}^b_{\mathrm{Vect}}(X)^{\mathrm{pre}}$ to be the differential graded nerve of the differential graded category whose objects are bounded complexes of vector bundles $V^\bullet$ on X and whose morphisms are global sections of $\mathrm{Hom}^\bullet(V^\bullet, W^\bullet)$. Now $\mathcal{D}^b_{\mathrm{Vect}}(-)^{\mathrm{pre}}$ is a presheaf of categories on Top, and we define $\mathcal{D}^b_{\mathrm{Vect}}(-)$ to be its sheafification.

A complex of vector bundles specializes, at every point, to a complex of vector spaces, whose cohomology is known as the fiberwise cohomology of complex. The category $\prod_{p\in X} \mathrm{Vect}$ ($\mathbb{Z}$-graded vector spaces) is evidently a sheaf, so objects of $\mathcal{D}^b_{\mathrm{Vect}}(X)$ also have fiberwise cohomology groups (use the universal property of sheafification).

2.10.6 Lemma. *An object of $\mathcal{D}^b_{\mathrm{Vect}}(X)$ is zero iff it is fiberwise acyclic.*

Proof. Since $\mathcal{D}^b_{\mathrm{Vect}}$ is a sheaf, an object is zero iff it is zero locally. We may thus assume that our given object is a complex of vector bundles $V^\bullet$. The loci $\{\dim \operatorname{im} d \geq r\}$ and $\{\dim \ker d \leq r\}$ are open; since $\ker d = \operatorname{im} d$, this implies that $\dim \ker d = \dim \operatorname{im} d$ is locally constant, which in turn implies that $\ker d = \operatorname{im} d$ is a vector subbundle of $V^\bullet$. The extension $0 \to \ker d \to V^\bullet \to \operatorname{im} d \to 0$ is degreewise split locally on X. A choice of splitting expresses $V^\bullet$ as a direct sum of two-term complexes $\mathbf{1}_A : A \to A$. Such objects are evidently zero. Thus $V^\bullet$ is locally zero, hence zero by the sheaf property. □

2.10.7 Corollary. *A morphism in $\mathcal{D}^b_{\mathrm{Vect}}(X)$ is an isomorphism iff it is a fiberwise quasi-isomorphism.*

Proof. Apply (2.10.6) to mapping cones. □

2.10.8 Definition (Dualization on $\mathcal{D}^b_{\mathrm{Vect}}$)**.** Dualizing complexes of vector bundles is a contravariant involution of the dg-category $(\mathcal{D}^b_{\mathrm{Vect}})^{\mathrm{pre}}$, hence passes to its sheafification $\mathcal{D}^b_{\mathrm{Vect}}$.

2.11 Global spaces and orbispaces

Our objective is to generalize the homotopy theories of spaces and topological spaces $\mathcal{S} \subseteq \mathsf{Top}^\infty$ to the setting of stacks. Stacks on the site of topological spaces $\mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})$ form a homotopy theory $\mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})^\infty$ by a direct generalization of the definition of Top^∞. Just as $\mathcal{S} \subseteq \mathsf{Top}^\infty$ is freely generated by the tiny object $* \in \mathsf{Top}^\infty$, we are more interested in a certain full subcategory $\mathcal{S}_{\mathrm{gl}} \subseteq \mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})^\infty$ which we call *global spaces*. It is freely generated under colimits by the tiny objects $\mathbb{B}G \in \mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})^\infty$ for all compact Lie groups G, namely

$$\mathcal{S}_{\mathrm{gl}} = \mathcal{P}(\mathbb{B}\mathcal{G}) \subseteq \mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})^\infty. \tag{2.11.0.1}$$

This construction is based on one introduced by Gepner–Henriques [7]; another presumably equivalent construction is given by Schwede [24]. The proof of cocompleteness of $\mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})^\infty$ is reasonably explicit, hence so are the objects comprising $\mathcal{S}_{\mathrm{gl}}$: they can be built out of cells $(D^k, \partial D^k) \times \mathbb{B}G$. We can also consider the homotopy theory $\mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})^\infty_{\mathsf{rep}}$ where all maps are required to be representable. The objects $\mathbb{B}G \in \mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})^\infty_{\mathsf{rep}}$ remain tiny, so they freely generate under colimits a full subcategory

$$\mathcal{S}_{\mathrm{orb}} = \mathcal{P}(\mathbb{B}\mathcal{G}_{\mathsf{rep}}) \subseteq \mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})^\infty_{\mathsf{rep}} \tag{2.11.0.2}$$

which we call the ∞-category of *orbispaces*. The natural functor $\mathcal{S}_{\mathrm{orb}} \to \mathcal{S}_{\mathrm{gl}}$ is not fully faithful, but at least on Hom spaces it is an inclusion of components. One can also replace the set $\mathcal{G}$ of isomorphism classes of compact Lie groups with a subset thereof (for instance, finite groups) and obtain homotopy theories which we denote $\mathcal{S}_{\mathcal{G}\mathrm{gl}}$ and $\mathcal{S}_{\mathcal{G}\mathrm{orb}}$, respectively.

2.11.1 Definition (∞-categories $\mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})^\infty$ and $\mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})^\infty_{\mathsf{rep}}$)**.** An n-simplex in $\mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})^\infty$ consists of the data of objects $X_0, \ldots, X_n \in \mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})$ and maps $f_{ij} : X_i \times I^{\{i+1,\ldots,j-1\}} \to X_j$ for $i < j$ satisfying the same conditions as in the definition of Top^∞ (2.3.14). The ∞-category $\mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})^\infty_{\mathsf{rep}}$ has the same definition except with the additional requirement that every map f_{ij} must be representable.

2.11.2 Definition. We denote by $\mathbb{B}\mathcal{G}$ and $\mathbb{B}\mathcal{G}_{\mathsf{rep}}$ the full subcategories of $\mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})^\infty$ and $\mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})^\infty_{\mathsf{rep}}$, respectively, spanned by the objects $\mathbb{B}G$ for compact Lie groups G.

2.11.3 Exercise. Let G and H be compact Lie groups. Show that $\mathrm{Hom}(\mathbb{B}G, \mathbb{B}H)$ in $\mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})$ is the groupoid whose objects are homomorphisms $\varphi : G \to H$ and whose isomorphisms $\varphi \to \varphi'$ are elements $h \in H$ satisfying $h\varphi h^{-1} = \varphi'$ (this groupoid has finitely many objects, and the automorphism of a given object $\varphi : G \to H$ is the centralizer of its image, equipped with the discrete topology). Show that $\mathrm{Hom}(\mathbb{B}G, \mathbb{B}H)$ in $\mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})^\infty$ is the homotopy quotient $\mathrm{Hom}(G, H)/H$ where H acts by conjugation. Conclude that the components of the latter are bijection with isomorphism classes in the former, and that whereas the automorphism group of $\mathbb{B}\varphi$ in the former is $Z_H(\varphi(G))$ (as an abstract group), the corresponding component in the latter is the classifying space of $Z_H(\varphi(G))$ as a topological group.

2.11.4 Lemma. *Disjoint unions are coproducts in* $\mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})^\infty$ *and* $\mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})^\infty_{\mathsf{rep}}$.

Proof. The argument used to prove the corresponding assertion for Top^∞ (2.3.17) applies without change. □

2.11.5 Lemma. *Let $C \in \mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})$, let $B_+, B_- \subseteq C$ be open substacks, and let $A := B_+ \cap B_-$. If there exists a map $f : C \to \mathbb{R}$ such that $f^{-1}((-\varepsilon, \infty)) \subseteq B_+$ and $f^{-1}((-\infty, \varepsilon)) \subseteq B_-$, then the diagram*

$$\begin{array}{ccc} A & \longrightarrow & B_+ \\ \downarrow & & \downarrow \\ B_- & \longrightarrow & C \end{array} \tag{2.11.5.1}$$

is a pushout in $\mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})^\infty$.

Proof. The proof of the corresponding assertion for Top^∞ (2.3.18) applies without change. □

2.11.6 Proposition. $\mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})^\infty$ *is cocomplete and the functor* $\mathsf{Top}^\infty \to \mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})^\infty$ *is cocontinuous.*

Proof. The only nontrivial part is to show that any diagram $X \leftarrow A \to Y$ in $\mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})^\infty$ may be completed to one of the form considered in (2.11.5). Define a stack $\hat{X}$ as follows: a map $f : Z \to \hat{X}$ is a map $Z \to X$, a function $g : Z \to (-\infty, \varepsilon)$, a map $h : g^{-1}((-\varepsilon, \varepsilon)) \to A$, and an isomorphism between $f|_{g^{-1}((-\varepsilon,\varepsilon))}$ and h composed with $A \to X$. Similarly define $\hat{Y}$ using $(-\varepsilon, \infty)$. There is now an evident diagram

$$\begin{array}{ccccc} \hat{X} & \longleftarrow & A \times (-\varepsilon, \varepsilon) & \longrightarrow & \hat{Y} \\ \downarrow & & \downarrow & & \downarrow \\ X & \longleftarrow & A & \longrightarrow & Y \end{array} \tag{2.11.6.1}$$

The vertical maps are all isomorphisms in $\mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})^\infty$ (for $\hat{X} \to X$, note that there is a map $X \to \hat{X}$ defined by taking $g \equiv -\varepsilon$, and the composition $\hat{X} \to X \to \hat{X}$ is homotopic to the identity map by taking a linear homotopy of the map g). The upper row may be completed to a diagram of the form considered in (2.11.5) since the upper two horizontal arrows are open inclusions and the resulting gluing as an evident map to $\mathbb{R}$ with the desired properties. □

2.11.7 Lemma. *Any $X \in \mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})$ with $|X| = *$ is a tiny object of* $\mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})^\infty$.

2.11.8. Let $\mathcal{Q} \subseteq \mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})^\infty$ be the full subcategory spanned by a set of stacks with coarse space $*$. Recall that $\mathcal{P}(\mathcal{Q})$ is the free cocompletion of $\mathcal{Q}$ (**??**), so since $\mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})^\infty$ is cocomplete, there is an induced functor

$$\mathcal{P}(\mathcal{Q}) \to \mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})^\infty \tag{2.11.8.1}$$

which is fully faithful since the objects of $\mathcal{Q}$ are tiny (2.11.7) (**??**). In other words, we have identified the full subcategory of $\mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})^\infty$ spanned by colimits of objects of $\mathcal{Q}$ with the presheaf category $\mathcal{P}(\mathcal{Q})$. The inclusion above moreover has a right adjoint given by pullback: an object of $\mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})^\infty$ can be tested against (morphisms from) objects of $\mathcal{Q}$ to obtain an object of $\mathcal{P}(\mathcal{Q})$.

2.12 Global ∞-categories

Given a discrete group G, there is a natural notion of a G-equivariant object in a category $\mathcal{C}$, namely an object $X \in \mathcal{C}$ together with a homomorphism $G \to \mathrm{Aut}(X)$; these form a category $G\mathcal{C}$. For instance, a G-equivariant object in the category of vector spaces is a G-representation. When $\mathcal{C}$ is an ∞-category, a 'naive G-equivariant object' is a functor $BG \to \mathcal{C}$ where $BG \in \mathcal{S} \subseteq \mathsf{Cat}_\infty$ denotes the classifying space of G, which can be any compact Lie group. As suggested by the choice of adjective 'naive', one is usually interested in other notions of equivariant objects. For instance, G-equivariant homotopy theory takes place not in $\mathsf{Fun}(BG, \mathcal{S})$, but rather in the ∞-category $G\mathcal{S} := (\mathcal{S}_{\mathrm{orb}})_{/\mathbb{B}G}$. It would appear this ∞-category has no simple construction in terms of just the ∞-category $\mathcal{S}$ and the compact Lie group G. Given a family of compact Lie groups $\mathcal{G}$, the categories of G-equivariant objects for $G \in \mathcal{G}$ (in all of the above senses) are related by pullback functors associated to group homomorphisms $\varphi : G \to H$.

A *global ∞-category*, encodes a 'collection of ∞-categories of G-equivariant objects for all compact Lie groups G'.

2.12.1 Definition (∞-category of G-spaces)**.** We define $G\mathcal{S} = (\mathcal{S}_{\mathrm{orb}})_{/\mathbb{B}G}$ for any compact Lie group G.

A representable map $X \to \mathbb{B}G$ is 'the same' as a topological space Y carrying a continuous G-action (namely $Y = X \times_{\mathbb{B}G} *$). This explains how to think of $(\mathcal{S}_{\mathrm{orb}})_{/\mathbb{B}G}$ as the ∞-category of G-spaces.

A homomorphism $\rho : H \to G$ induces a triple of adjoint functors $(\rho_!, \rho^*, \rho_*)$ between the categories $H\mathcal{S}$ and $G\mathcal{S}$. The middle functor $\rho^* : G\mathcal{S} \to H\mathcal{S}$ is pullback (fiber product) under $\mathbb{B}\rho : \mathbb{B}H \to \mathbb{B}G$ (equivalently, this regards a G-space as an H-space via ρ). It has a left adjoint $\rho_! : H\mathcal{S} \to G\mathcal{S}$ given by composition with $\mathbb{B}\rho$ (this is $G \times_H -$ from H-spaces to G-spaces). It also has a right adjoint $\rho_* : H\mathcal{S} \to G\mathcal{S}$ given by $\underline{\mathrm{Hom}}_H(G, -)$ from H-spaces to G-spaces.

2.13 Equivariant site

2.13.1 Definition (Equivariant site)**.** Let X be a stack on topological spaces. The *equivariant site* of X is the arrow category

$$X_{\mathrm{eq}} = \mathsf{Arr}(\mathsf{Op}(X) \to \mathbb{B}\mathcal{G}_{\mathsf{rep}}) \tag{2.13.1.1}$$

(arrows in $\mathsf{Shv}(\mathsf{Top}, \mathsf{Grpd})^{\infty}_{\mathsf{rep}}$). An object of X_{eq} is thus an open substack $U \subseteq X$ and a representable map $U \to \mathbb{B}G$ for some compact Lie group G. A sieve on an object $(U \to \mathbb{B}G) \in X_{\mathrm{eq}}$ is deemed a covering sieve iff it contains an open cover of U (the maps to $\mathbb{B}\mathcal{G}$ are thus completely ignored in determining whether a given sieve is a covering sieve).

2.13.2 Definition (Sheaves on the equivariant site)**.** Let X be a stack on topological spaces. A *precosheaf* on X_{eq} valued in an inner fibration $\mathcal{C} \to \mathbb{B}\mathcal{G}_{\mathsf{rep}}$ is a lift

$$\begin{array}{ccc} & & \underline{\mathcal{C}} \\ & \nearrow & \downarrow \\ X_{\mathrm{eq}} & \longrightarrow & \mathbb{B}\mathcal{G}_{\mathsf{rep}} \end{array} \tag{2.13.2.1}$$

These form an ∞-cateogry $\mathsf{PCoshv}(X_{\mathrm{eq}}, \mathcal{C})$ (2.2.72). Similarly, $\mathsf{PShv}(X_{\mathrm{eq}}, \mathcal{C})$ for an inner fibration $\mathcal{C} \to \mathbb{B}\mathcal{G}_{\mathsf{rep}}^{\mathsf{op}}$ defined as $\mathsf{PCoshv}(X_{\mathrm{eq}}, \mathcal{C}^{\mathsf{op}})^{\mathsf{op}}$.

The inner fibrations $\mathcal{C} \to \mathbb{B}\mathcal{G}_{\mathsf{rep}}$ which occur in practice are moreover both cartesian and cocartesian. If one is satisfied with restricting to cocartesian fibrations, then one could speak of cosheaves valued in functors $\mathcal{C} : \mathbb{B}\mathcal{G}_{\mathsf{rep}}^{\mathsf{op}} \to \mathsf{Cat}_\infty$. A sheaf valued in $\mathcal{C} : \mathbb{B}\mathcal{G}_{\mathsf{rep}}^{\mathsf{op}} \to \mathsf{Cat}_\infty$ is then a cosheaf valued in $\mathsf{op} \circ \mathcal{C}$.

A precosheaf $\mathcal{F}$ on X_{eq} is called a *cosheaf* iff for every covering sieve on an object $(U \to \mathbb{B}G) \in X_{\mathrm{eq}}$, the restriction of $\mathcal{F}$ to the slice category over $(U \to \mathbb{B}G)$ is left Kan extended from its restriction to (the full subcategory spanned by) the sieve.

Chapter 3

Analysis

3.1 Function spaces

We recall various standard function spaces and their basic properties. References include [23, 1, 12]. By 'manifold' we mean 'paracompact smooth manifold', and by 'vector bundle' we mean 'finite-dimensional smooth real vector bundle'.

3.1.1 Definition (C^∞ and C^k spaces)**.** Given a manifold M and a vector bundle $E \to M$, its space of smooth sections is denoted $C^\infty(M,E)$, and the subspace consisting of those sections of compact support is denoted $C^\infty_c(M,E)$. Similarly C^k for integers $k \geq 0$ indicates sections which are k times continuously differentiable, so $C^\infty = \bigcap_{k\geq 0} C^k$.

The space C^k is equipped with seminorms $\|\cdot\|_{C^k(K)}$ for $K \subseteq M$ compact defined as

$$\|f\|_{C^k(K)} = \sum_i \sum_{j=0}^{k} \sup_{x\in U_i\cap K} \left|D^j(\varphi_i f|_{U_i})(x)\right|, \tag{3.1.1.1}$$

where $U_i \subseteq \mathbb{R}^n$ is a finite collection of charts of M covering K and $\varphi_i : M \to \mathbb{R}_{\geq 0}$ is supported inside U_i with $\sum_i \varphi_i$ everywhere positive on K (the seminorm $\|\cdot\|_{C^k(K)}$ is independent of the choice of charts and functions φ_i up to commensurability).

The spaces C^k and C^∞ are complete with respect to the family of seminorms $\|\cdot\|_{C^k(K)}$ parameterized by compact $K \subseteq M$ (and, in the latter case, all $k < \infty$). Their subspaces C^k_K and C^∞_K of functions supported inside a fixed compact $K \subseteq M$ are closed (and hence complete). The subspaces C^k_c and C^∞_c are not closed (in fact, they are dense); rather than giving them the subspace topology, it is most relevant to give them the directed colimit topology from their presentation as the ascending union of the spaces C^k_K and C^∞_K.

Denote by Ω_M the bundle of densities on M. It is a smooth real line bundle defined by the existence of a canonical integration map $\int : C^\infty_c(M,\Omega_M) \to \mathbb{R}$. In fact, it is the line bundle associated to a principal $\mathbb{R}_{>0}$-bundle, so it has powers Ω^t_M for any $t \in \mathbb{R}$.

3.1.2 Definition (Distributions $C^{-\infty}$)**.** A distribution $f \in C^{-\infty}(M,E)$ is a continuous linear functional $C^\infty_c(M, E^* \otimes \Omega_M) \to \mathbb{R}$ denoted $g \mapsto \int fg$. Continuity means, concretely, that for every compact $K \subseteq M$, there exists $k < \infty$ such that $|\int fg| \leq \text{const}\|g\|_{C^k_K}$ for $g \in C^\infty_K$. There is a natural inclusion $C^\infty(M,E) \hookrightarrow C^{-\infty}(M,E)$. The product of a distribution f and a smooth function g is defined to be $h \mapsto \int f(gh)$.

We shall regard $C^{-\infty}$ as equipped with the subspace topology inside $\prod_{f \in C^\infty_c} \mathbb{R}$ (this is a very weak topology). In other words, if V is a topological vector space, a map $T : V \to C^{-\infty}$ is continuous iff $\int T(\cdot)f : V \to \mathbb{R}$ is continuous for every $f \in C^\infty_c$.

3.1.3 Example (Delta function). The delta function $\delta_p \in C^{-\infty}(\mathbb{R}^n)$ is the distribution given by the linear functional 'evaluate at $p \in \mathbb{R}^n$'. On a manifold, the delta distribution is naturally a distribution valued in densities $\delta_p \in C^{-\infty}(M, \Omega_M)$.

3.1.4 Exercise ($C^k \subseteq C^{-\infty}$). Show that for fixed $f \in C^\infty_c$, we have $|\int fg| \leq \text{const}\|g\|_{C^k}$, and conclude from this that the inclusion $C^\infty \to C^{-\infty}$ admits a unique continuous extension $C^k \to C^{-\infty}$.

3.1.5 Exercise. Show that $C^\infty_c \subseteq C^{-\infty}$ is dense.

3.1.6 Exercise (Support of a distribution). Let $f \in C^{-\infty}(M,E)$. Define a closed set $\operatorname{supp} f \subseteq M$ by the property that $p \notin \operatorname{supp} f$ iff there exists an open set U containing p such that $\int fg = 0$ for every g supported inside U. Show that if $g \in C^\infty_c(M,E)$ and $\operatorname{supp} g \cap \operatorname{supp} f = \varnothing$, then $\int fg = 0$. Conclude that if $\varphi : M \to \mathbb{R}$ is smooth and identically 1 in a neighborhood of $\operatorname{supp} f$, then $\varphi f = f$. Denote by $C^{-\infty}_c(M,E)$ the distributions of compact support.

3.1.7 Definition (L^2-norm)**.** The space $L^2(\mathbb{R}^n)$ is the completion of $C^\infty_c(\mathbb{R}^n)$ (equivalently, $C^0_c(\mathbb{R}^n)$) with respect to the L^2-norm $\|f\|_{L^2} = (\int |f|^2)^{1/2}$. This norm evidently comes from the inner product $\langle f, g\rangle = \int fg$. On a manifold M, the L^2-norm is most naturally defined for sections of the bundle of half-densities $\Omega^{1/2}_M$, giving us the space $L^2(M, \Omega^{1/2}_M)$. To define spaces $L^2(M,E)$ where E is a vector bundle over M, the necessary data is that of a positive definite inner product $E \otimes E \to \Omega_M$. Commensurable such metrics define commensurable norms, so everything is unique up to commensurability if M is compact.

Similarly, define $L^2_{\text{loc}}(M)$ as the completion of $C^\infty(M)$ with respect to the family of seminorms $\|\cdot\|_{L^2(K)}$ for $K \subseteq M$ compact.

3.1.8 Exercise ($L^2_{\text{loc}} \subseteq C^{-\infty}$). Show that the inclusion $C^\infty \to C^{-\infty}$ admits a unique continuous extension $L^2_{\text{loc}} \to C^{-\infty}$.

3.1.9 Definition (H^s-norm for $s \in \mathbb{Z}_{\geq 0}$)**.** For integers $s \geq 0$, the H^s-norm is defined as

$$\|f\|^2_{H^s(\mathbb{R}^n)} = \int \sum_{|\alpha| \leq s} |D^\alpha f|^2. \tag{3.1.9.1}$$

The space H^s is the completion of C^∞_c with respect to this norm.

3.1.10 Definition (Schwartz functions)**.** The space of Schwartz functions $\mathcal{S}(\mathbb{R}^n)$ consists of those infinitely differentiable functions all of whose norms $\sup(1+|x|^A)\sum_{|\alpha|\leq B}|D^\alpha_x f(x)|$ are finite; the space $\mathcal{S}$ is complete with respect to this family of norms. We can also consider Schwartz functions valued in any vector space.

3.1.11 Definition (Fourier transform)**.** The Fourier transform is a linear map $\mathcal{S}(\mathbb{R}^n, \mathbb{C}) \to \mathcal{S}(\mathbb{R}^n, \mathbb{C})$ denoted $f \mapsto \hat{f}$ and given by

$$\hat{f}(\xi) = \int e^{-2\pi i\langle \xi, x\rangle} f(x)\, dx. \tag{3.1.11.1}$$

The Fourier transform is continuous: the required decay properties of $\hat{f}$ and its derivatives follow from integration by parts. The Fourier transform also makes sense for Schwartz functions valued in a complex vector space V. When V is a real vector space, the Fourier transform maps $\mathcal{S}(\mathbb{R}^n, V)$ to the subspace of $\mathcal{S}(\mathbb{R}^n, V \otimes_{\mathbb{R}} \mathbb{C})$ consisting of functions g satisfying $g(-\xi) = \overline{g(\xi)}$ (and conversely, the Fourier transform of such a function g lies in the subspace $\mathcal{S}(\mathbb{R}^n, V) \subseteq \mathcal{S}(\mathbb{R}^n, V \otimes_{\mathbb{R}} \mathbb{C})$).

3.1.12 Exercise. Show that $\int_{\mathbb{R}^n} e^{-\pi|x|^2}\, dx = 1$ by reducing to the case $n = 2$ and using polar coordinates.

3.1.13 Proposition (Fourier inversion)**.** *For $f \in \mathcal{S}(\mathbb{R}^n)$, we have $\hat{\hat{f}}(x) = f(-x)$.*

Proof. Note that $g(\xi)e^{-\pi(\xi/N)^2} \to g(\xi)$ in $\mathcal{S}(\mathbb{R}^n)$ as $N \to \infty$. It thus suffices to show that the Fourier transform of $\hat{f}(\xi)e^{-\pi(\xi/N)^2}$ approaches $f(-x)$ as $N \to \infty$. This is the expression

$$\iint e^{2\pi i\langle \xi, x-y\rangle} f(y) e^{-\pi(\xi/N)^2}\, dy\, d\xi = \int \left(\int e^{-\pi(\xi/N)^2 + 2\pi i\langle \xi, z\rangle}\, d\xi \right) f(x+z)\, dz \tag{3.1.13.1}$$

Now we may compute the inner integral of ξ by completing the square and moving the contour and appealing to the identity $\int e^{-\pi x^2}\, dx = 1$. The result is $N^n e^{-\pi(Nz)^2}$, making the desired convergence to $f(x)$ clear. □

3.1.14 Exercise (Fourier transform and convolution)*.* Show that for $f, g \in \mathcal{S}(\mathbb{R}^n)$, we have $\widehat{f*g} = \hat{f}\hat{g}$, where $(f*g)(x) = \int f(y)g(x-y)\, dy$ denotes convolution. Using Fourier inversion, conclude that $\widehat{fg} = \hat{f}*\hat{g}$, and specialize this to conclude that $\int fg = \int \hat{f}\hat{g}$, so in particular $\|f\|_{L^2} = \|\hat{f}\|_{L^2}$ (Plancherel's formula).

3.1.15 Definition (Sobolev norms H^s for $s \in \mathbb{R}$)**.** The H^s-norm of f is defined in terms of its Fourier transform by the formula

$$\|f\|_{H^s(\mathbb{R}^n)} = \left\| (1 + |\xi|^2)^{s/2} \hat{f} \right\|_{L^2(\mathbb{R}^n)}. \tag{3.1.15.1}$$

Note that when $s \in \mathbb{Z}_{\geq 0}$ is a nonnegative integer, this agrees with the previous definition of H^s-norm (up to commensurability) by Plancherel.

3.1.16 Lemma (Sobolev embedding)**.** *For integer $k \geq 0$ and $s > k + \frac{n}{2}$, we have $\|u\|_{C^k} \leq \mathrm{const}_{s,k} \|u\|_{H^s}$.*

Proof. It is enough to treat the case $k = 0$, where we are supposed to show that

$$|u(0)|^2 = \left|\int e^{2\pi i\langle \xi,x\rangle}\hat{u}(\xi)\,d\xi\right|^2 \leq \mathrm{const}_s \int |\hat{u}(\xi)|^2(1+|\xi|)^{2s}\,d\xi \tag{3.1.16.1}$$

which is simply Cauchy–Schwarz provided that $\int (1+|\xi|)^{-2s}\,d\xi < \infty$ which is the case for $s > \frac{n}{2}$ (using polar coordinates, it is equivalent to $\int_1^\infty r^{-2s}r^{n-1}\,dr < \infty$).

Alternatively, boundedness of $u \mapsto u(0)$ in the H^s-norm is equivalent, by duality, to the statement $\delta \in H^{-s}$. The Fourier transform of δ is simply $\hat{\delta} \equiv 1$, so $\delta \in H^{-s}$ iff $\int (1+\xi^2)^{-s} < \infty$. □

3.1.17 Lemma (Sobolev restriction estimate)**.** *The restriction of smooth functions* $\mathbb{R}^n$ *to* $\mathbb{R}^{n-1} \times 0$ *is bounded* $H^s \to H^{s-\frac{1}{2}}$ *provided* $s > \frac{1}{2}$.

Of course, this implies Sobolev embedding by interating.

Proof. We inspect Fourier transforms in coordinates $\xi \in \mathbb{R}^{n-1}$ and $\eta \in \mathbb{R}$ dual to $x \in \mathbb{R}^{n-1}$ and $y \in \mathbb{R}$; the Fourier transform of $f(x,0)$ is the integral of $\hat{f}(\xi,\eta)\,d\eta$. Cauchy–Schwarz gives

$$\left|\int \hat{f}(\xi,\eta)\,d\eta\right|^2 \leq \int \hat{f}(\xi,\eta)^2(1+\xi^2+\eta^2)^s\,d\eta \int (1+\xi^2+\eta^2)^{-s}\,d\eta \tag{3.1.17.1}$$

Now doing a change of variables scaling η by $(1+|\xi|)^{-1}$, we see that the right most integral is of order $(1+\xi^2)^{-(s-\frac{1}{2})}$ (times $\int (1+\eta^2)^s\,d\eta$ which is finite provided $s > \frac{1}{2}$). We thus conclude that

$$\left|\int \hat{f}(\xi,\eta)\,d\eta\right|^2 (1+\xi^2)^{s-\frac{1}{2}} \leq \int \hat{f}(\xi,\eta)^2(1+\xi^2+\eta^2)^s\,d\eta \tag{3.1.17.2}$$

Now integrating $d\xi$ we conclude (note that $\int \hat{f}(\xi,\eta)\,d\eta$ is the Fourier transform of the restriction of f forgetting the single dimensional variable dual to η). □

The next result allows us to define integer Sobolev regularity classes H^s for maps to non-linear targets (i.e. manifolds) whenever $s \geq 0$ is an integer and $H^s \subseteq C^0$.

3.1.18 Proposition. *Suppose* $F : \mathbb{R}^n \to \mathbb{R}$ *is smooth and vanishes at the origin. Let* $s \geq 0$ *be an integer such that* $H^s \to C^0$ *is bounded. If* $\|g\|_{C^0} \leq M$*, then* $\|F(g)\|_{H^s} \leq \mathrm{const}_{F,M}\|g\|_{H^s}$.

Proof. A general term in a derivative of $F(g)$ up to order k takes the form $D^\alpha F(g)\prod_{i=1}^{|\alpha|} D^{\beta_i}g$ where $|\beta_i| \geq 1$ and $\sum_i |\beta_i| \leq k$. The factor $D^\alpha F(g)$ is bounded since F is smooth and $\|g\|_{C^0} \leq M$ by hypothesis. Using Hölder's inequality, it thus suffices to show that $D^\beta g \in L^{2s/|\beta|}$. That is, we should show that $H^{s-r} \to L^{2s/r}$ is bounded for integers $r = 1,\ldots,s$. For $r = 0$, this is the assumption that $H^s \to C^0$ is bounded, and for $r = s$ this is the definition $H^0 = L^2$. It thus follows for general $r \in [0,s]$ by interpolation. □

3.1.19 Corollary. *In* (3.1.18)*, if* F *vanishes to order* $m \geq 1$ *at the origin, then* $\|F(g)\|_{H^s} \leq \mathrm{const}_{F,M}\|g\|_{H^s}^m$ *for* $\|f\|_{C^0} \leq M$.

Proof. If $\|f\|_{C^0} \geq 1$, there is nothing to prove beyond (3.1.18), so we assume that $\|f\|_{C^0} \leq 1$. Consider the rescaling $F_\varepsilon(x) = \varepsilon^{-m}F(\varepsilon x)$; since F vanishes to order m at the origin, the derivatives of F_ε over the unit ball are bounded uniformly as $\varepsilon \to 0$ (since F vanishes to order m at the origin, it follows that $D^\alpha F$ vanishes to order $m - |\alpha|$ at the origin, which gives uniform boundedness of $D^\alpha F_\varepsilon = \varepsilon^{|\alpha|-m}(D^\alpha F)(\varepsilon x)$ over the unit ball as $\varepsilon \to 0$). Now we have $F(g) = \varepsilon^m F_\varepsilon(\varepsilon^{-1}g)$. As long as $\|f\|_{C^0} \leq \varepsilon$, we can apply the previous result (3.1.18) to conclude that $\|F_\varepsilon(\varepsilon^{-1}g)\|_{H^s} \leq \text{const}_F\|\varepsilon^{-1}g\|_{H^s}$. This is the same as $\|F(g)\|_{H^s} \leq \text{const}_F\varepsilon^{m-1}\|f\|_{H^s}$. Taking $\varepsilon = \|f\|_{C^0}$, we conclude that $\|F(g)\|_{H^s} \leq \text{const}_F\|f\|_{C^0}^{m-1}\|f\|_{H^s}$, which implies the desired result since $H^s \to C^0$ is assumed to be bounded. □

3.1.20 Exercise. Suppose $A : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$ vanishes to order two at the origin and satisfies $A(x, 0) = 0$. Using a rescaling of the form $A(\alpha x, \beta y)$, show that if $\|f\|_{C^0} \leq 1$ and $\|g\|_{C^0} \leq 1$, then

$$\|A(f, g)\|_{H^s} \leq \text{const}(\|g\|_{C^0}\|f\|_{H^s} + (\|f\|_{C^0} + \|g\|_{C^0})\|g\|_{H^s}) \tag{3.1.20.1}$$

for a constant depending on the derivatives of A. Make a change of variables to conclude that for $B : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}$ vanishing to order two at the origin and satisfying $B(x, x) = 0$, if $\|f\|_{C^0} \leq 1$ and $\|g\|_{C^0} \leq 1$, then

$$\|B(f, g)\|_{H^s} \leq \|f - g\|_{H^s}(\|f\|_{C^0} + \|g\|_{C^0}) + \|f - g\|_{C^0}(\|f\|_{H^s} + \|g\|_{H^s}) \tag{3.1.20.2}$$

3.2 Differential operators

3.2.1 Definition (Differential operator)**.** On a manifold M carrying vector bundles E and F, a differential operator $L : C^\infty(M, E) \to C^\infty(M, F)$ of order $\leq m$ is a map which is given in local coordinates $M = \mathbb{R}^n$ by an expression of the form

$$Lf = \sum_{|\alpha| \leq m} c_\alpha D^\alpha f \tag{3.2.1.1}$$

where c_α are smooth functions taking values in $\text{Hom}(E, F)$ (an operator being of this form is evidently preserved by diffeomorphisms). The order m terms transform under diffeomorphisms independently of the others, so a differential operator L of order $\leq m$ has a well-defined order m term which takes the form of a homogeneous degree m polynomial map $T^*M \to \text{Hom}(E, F)$.

Evidently, for any differential operator of order $\leq m$, we have a pointwise estimate $\|Lf\|_{C^k} \leq \text{const}_{L,k}\|f\|_{C^{k+m}}$ where the constant depends on the derivatives up to order k of the coefficients of L. To globalize this estimate over a non-compact space requires control of the constant at infinity.

3.2.2 Definition (Formal adjoint)**.** For a differential operator $L : C^\infty(M, E) \to C^\infty(M, F)$, its formal adjoint is the differential operator

$$L^* : C^\infty(M, F^* \otimes \Omega_M) \to C^\infty(M, E^* \otimes \Omega_M) \tag{3.2.2.1}$$

defined by the property $\int_M \langle u, Lv\rangle = \int_M \langle L^* u, v\rangle$ (say for u and v of compact support), where Ω_M denotes the bundle of densities on M. In other words, L^* is obtained from L by formally integrating by parts.

3.2.3 Exercise. Show that a differential operator $L : C^\infty(M, E) \to C^\infty(M, F)$ admits a unique continuous extension $L : C^{-\infty}(M, E) \to C^{-\infty}(M, F)$.

3.2.4 Exercise. Let $w \in C^{-\infty}(M, \Omega_M)$ be a distribution supported at a single point p (3.1.6). Show that w is a linear combination of the delta function δ_p and its derivatives.

3.2.5 Exercise. Denoting by $\Delta = \partial_x^2 + \partial_y^2$ the Laplace operator on $\mathbb{R}^2$, prove that $\Delta(\log r) = 2\pi\delta_0$.

3.2.6 Proposition. *Let $L : C^\infty(\mathbb{R}^n, E) \to C^\infty(\mathbb{R}^n, F)$ be a compactly supported differential operator of order $\leq m$. We have $\|Lf\|_{H^s} \leq \mathrm{const}_{L,s}\|f\|_{H^{s+m}}$.*

Proof. When $s \in \mathbb{Z}_{\geq 0}$, this follows immediately from the pointwise C^s-estimate on Lf. For general $s \in \mathbb{R}$ we will use the Fourier transform.

We begin by looking at the case that L has *constant coefficients*, namely $L : C^\infty(\mathbb{R}^n, E) \to C^\infty(\mathbb{R}^n, F)$ for vector spaces E and F takes the form (3.2.1.1) for constants $c_\alpha \in \mathrm{Hom}(E, F)$ (admittedly this operator is not compactly supported). In other words $L = P(D)$ for the polynomial $P : \mathbb{R}^n \to \mathrm{Hom}(E, F)$ given by $P(\xi) = \sum_{|\alpha|\leq m} c_\alpha \xi^\alpha$ (the domain of P is most naturally the *dual* of $\mathbb{R}^n$, which we are implicitly identifying with $\mathbb{R}^n$ using the standard basis and its dual). The action of L on Fourier transforms now takes the following simple form:

$$\widehat{Lu}(\xi) = P(i\xi)\hat{u}(\xi), \tag{3.2.6.1}$$

from which the bound $\|Lf\|_s \leq \mathrm{const}_{L,s}\|f\|_{s+m}$ is immediate.

In the desired case of interest, namely variable coefficients of compact support, we will also use the Fourier transform. Let $P(x, \xi)$ denote the degree $\leq m$ polynomial in ξ given by the coefficients of L at $x \in \mathbb{R}^n$, so we have

$$Lu(x) = \int e^{2\pi i\langle \xi, x\rangle} P(x, i\xi)\hat{u}(\xi)\, d\xi. \tag{3.2.6.2}$$

by differentiating under the integral sign. We may thus calculate

$$\widehat{Lu}(\zeta) = \int e^{-2\pi i\langle \zeta, x\rangle} \int P(x, \xi) e^{2\pi i\langle \xi, x\rangle}\hat{u}(\xi)\, d\xi\, dx \tag{3.2.6.3}$$

$$= \int \hat{u}(\xi) \int e^{2\pi i\langle \xi-\zeta, x\rangle} P(x, \xi)\, dx\, d\xi \tag{3.2.6.4}$$

where we note that the x-support of P is compact and we take $u \in \mathcal{S}(\mathbb{R}^n, E)$ so that the interchange of integrals is valid. Now the kernel $K(\zeta, \xi) = \int e^{2\pi i\langle \xi-\zeta, x\rangle} P(x, \xi)\, dx$ is bounded by $\mathrm{const}_{L,N} \cdot (1 + |\xi - \zeta|)^{-N} \cdot (1 + |\xi|)^m$ (integrate by parts N times in the direction of $\xi - \zeta$ if $|\xi - \zeta| \geq 1$). The desired bound

$$\int \widehat{Lu}(\zeta)(1 + |\zeta|^2)^{2s}\, d\zeta \leq \mathrm{const}_{L,s} \int \hat{u}(\xi)(1 + |\xi|^2)^{s+m}\, d\xi \tag{3.2.6.5}$$

follows from this estimate on $K(\zeta, \xi)$. □

The basic local bound (3.2.6) globalizes on any compact manifold by a partition of unity argument.

3.3 Elliptic operators

We review the theory of elliptic operators. References include [15].

3.3.1 Definition (Elliptic operator)**.** An elliptic (differential) operator of order m is a differential operator of order $\leq m$ whose order m symbol $T^*M \to \operatorname{Hom}(E, F)$ sends nonzero elements $\xi \in T^*M$ to invertible elements of $\operatorname{Hom}(E, F)$.

3.3.2 Exercise. Show that the operator $C^\infty(\mathbb{R}^n) \to C^\infty(\mathbb{R}^n)$ given by $\sum_i \partial_{x_i}^2$ is elliptic. Show that the operator $C^\infty(\mathbb{C}, \mathbb{C}) \to C^\infty(\mathbb{C}, \mathbb{C})$ given by $f(x+iy) \mapsto \frac{1}{2}(\partial_x f(x+iy) - i\partial_y f(x+iy))$ is elliptic.

The significance of ellipticity is that such operators are 'close' to being invertible. This and the next couple of sections are devoted to making this assertion precise. The main step is to construct, for any order m elliptic operator L, a *parametrix* Q, which is an operator of 'order $-m$' such that both operators $\mathbf{1} - LQ$ and $\mathbf{1} - QL$ have 'order $-\infty$' (having 'order r' means, in particular, being bounded $H^s \to H^{s-r}$). The existence of a parametrix for L has many corollaries; for instance, it implies 'elliptic estimates' of the form $\|u\|_s \leq \operatorname{const}_{L,s} \|Lu\|_{s-m} + \operatorname{const}_{L,s,N} \|u\|_{s-N}$ for any $N < \infty$.

3.3.3 Construction (Parametrices for elliptic operators with constant coefficients)**.** Fix an order m elliptic operator $L = P(D) : C^\infty(\mathbb{R}^n, E) \to C^\infty(\mathbb{R}^n, F)$ with constant coefficients. Ellipticity implies that

$$\|P(i\xi)^{-1}\| \leq \operatorname{const}_L \cdot \|\xi\|^{-m}. \tag{3.3.3.1}$$

Recalling that the action of L on Fourier transforms is to multiply by $P(i\xi)$ (3.2.6.1), this estimate gives a natural way to (almost) invert L. Namely, we define an operator $Q : \mathcal{S}(V, F) \to \mathcal{S}(V, E)$ by the property that

$$\widehat{Qu}(\xi) = A(\xi)\hat{u}(\xi), \tag{3.3.3.2}$$

where $A : \mathbb{R}^n \to \operatorname{Hom}(F, E) \otimes_{\mathbb{R}} \mathbb{C}$ is smooth with $A(-\xi) = \overline{A(\xi)}$ and satisfies $A(\xi) = P(i\xi)^{-1}$ for sufficiently large ξ. This operator Q is an example of a *pseudo-differential operator* of order $-m$. As such, it enjoys bounds of the form $\|Qu\|_s \leq \operatorname{const}_{L,s} \|u\|_{s-m}$ (an immediate consequence of the decay bound (3.3.3.1)). Now we have

$$((\mathbf{1} - QL)u)^\wedge(\xi) = (\mathbf{1} - A(\xi)P(i\xi))\hat{u}(\xi), \tag{3.3.3.3}$$

$$((\mathbf{1} - LQ)u)^\wedge(\xi) = (\mathbf{1} - P(i\xi)A(\xi))\hat{u}(\xi). \tag{3.3.3.4}$$

Noting that both $\mathbf{1} - A(\xi)P(i\xi)$ and $\mathbf{1} - P(i\xi)A(\xi)$ are of compact support, we see that the operators $\mathbf{1} - QL$ and $\mathbf{1} - LQ$ are both smoothing operators, in the sense that we have bounds

$$\|(\mathbf{1}-QL)u\|_s \leq \text{const}_{L,s,N}\cdot\|u\|_{s-N}, \tag{3.3.3.5}$$
$$\|(\mathbf{1}-LQ)u\|_s \leq \text{const}_{L,s,N}\cdot\|u\|_{s-N}, \tag{3.3.3.6}$$

for any $N<\infty$.

3.3.4 Exercise (Elliptic estimate from a parametrix). Let L be a differential operator of order $\leq m$, and suppose there exists an operator Q satisfying $\|Qu\|_s \leq \text{const}_{L,s}\|u\|_{s-m}$ and

$$\|(\mathbf{1}-QL)u\|_s \leq \text{const}_{L,s,N}\cdot\|u\|_{s-N} \tag{3.3.4.1}$$

for any $N<\infty$. Show that

$$\|u\|_s \leq \text{const}_{L,s}\cdot\|Lu\|_{s-m}+\text{const}_{L,s,N}\cdot\|u\|_{s-N} \tag{3.3.4.2}$$

for any $N<\infty$.

3.3.5 Exercise (Patching together parametrices). Let L be an elliptic operator of order m on a compact manifold. Let $M=\bigcup_i U_i$ be a finite cover by coordinate charts $U_i\subseteq\mathbb{R}^n$, and let $\varphi_i: M\to\mathbb{R}_{\geq 0}$ be a partition of unity $\sum_i\varphi_i\equiv 1$ subordinate to the cover $\operatorname{supp}\varphi_i\subseteq U_i$. Further let $\psi_i:U_i\to\mathbb{R}$ have compact support and be identically equal to 1 over a neighborhood of $\operatorname{supp}\varphi_i$. Suppose moreover that the restriction of L to each $U_i\subseteq\mathbb{R}^n$ has constant coefficients (the existence of an open cover with this property is a very strong condition on L), and let Q_i be a parametrix (3.3.3) for this constant coefficient operator (this is the only use of the assumption that each $L|_{U_i}$ has constant coefficients). Using the fact that the commutator $[L,\varphi]$ has order $\leq m-1$, show that the operator $Q:=\sum_i\psi_iQ_i\varphi_i$ satisfies bounds

$$\|(\mathbf{1}-QL)u\|_s \leq \text{const}_s\cdot\|u\|_{s-1}, \tag{3.3.5.1}$$
$$\|(\mathbf{1}-LQ)u\|_s \leq \text{const}_s\cdot\|u\|_{s-1}. \tag{3.3.5.2}$$

These bounds are evidently weaker than (3.3.3.5)–(3.3.3.6), but we will see next how to improve them.

3.3.6 Exercise (Improving an approximate inverse). Let $L:A\to B$ be a linear map. Given a linear map in the reverse direction $Q:B\to A$, show that

$$Q(\mathbf{1}_B+\cdots+(\mathbf{1}_B-LQ)^{N-1})=(\mathbf{1}_A+\cdots+(\mathbf{1}_A-QL)^{N-1})Q, \tag{3.3.6.1}$$

and denote this expression by Q_N. Show that

$$\mathbf{1}_A-Q_NL=(\mathbf{1}_A-QL)^N, \tag{3.3.6.2}$$
$$\mathbf{1}_B-LQ_N=(\mathbf{1}_B-LQ)^N. \tag{3.3.6.3}$$

The significance of this construction is that if Q is an 'approximate inverse' to L in the sense that one or both of $\mathbf{1}_A-QL$ or $\mathbf{1}_B-LQ$ is 'small', then Q_N is a 'better' approximate inverse in the sense that $\mathbf{1}_A - Q_N L$ or $\mathbf{1}_B - LQ_N$ is 'smaller' (this is also comparable to inverting an upper triangular matrix). For example, applying it to the 1-parametrix from (3.3.5) produces an N-parametrix for any $N < \infty$, i.e. an operator Q_N satisfying bounds

$$\|(\mathbf{1} - Q_N L)u\|_s \leq \text{const}_s \cdot \|u\|_{s-N} \tag{3.3.6.4}$$
$$\|(\mathbf{1} - LQ_N)u\|_s \leq \text{const}_s \cdot \|u\|_{s-N} \tag{3.3.6.5}$$

which is sufficient to produce elliptic estimates (3.3.4).

3.3.7 Exercise. Let V be a complete normed vector space. Any endomorphism $A : V \to V$ satisfying $\|A - \mathbf{1}\| < 1$ is invertible with inverse given by the infinite series

$$A^{-1} = \mathbf{1} + (\mathbf{1} - A) + (\mathbf{1} - A)^2 + \cdots \tag{3.3.7.1}$$

which converges in operator norm.

3.3.8 Discussion. To construct a parametrix for an elliptic operator with variable coefficients, one may attempt a continuous version of the patching argument (3.3.5). Let L be an elliptic operator on an open set $\Omega \subseteq \mathbb{R}^n$. For any $x \in \Omega$, let L_x denote the constant coefficient operator which agrees with L at x. Each L_x thus has a parametrix Q_x constructed in (3.3.3); moreover we may take Q_x to be a Fourier multiplication operator whose symbol $A(x,\xi)$ equals $P_x(i\xi)^{-1}$ for $\|\xi\| \geq M$ for some $M = M(L) < \infty$ and is jointly smooth in x and ξ (with bounds on its derivatives depending only on L). It is therefore reasonable to consider

$$(Qu)(x) = \int e^{2\pi i\langle \xi, x\rangle} A(x,\xi)\hat{u}(\xi)\, d\xi \tag{3.3.8.1}$$

as an operator on $C_c^\infty(\Omega) \to C^\infty(\Omega)$. Note that if L has constant coefficients, we may take A to be independent of x, and this reduces to the construction given previously for constant coefficient operators (3.3.3). The integral (3.3.8.1) can be viewed as a continuous version of the patched operator from (3.3.5). We will eventually see that (3.3.8.1) is only a 1-parametrix; to prove it involves a general study of operators of the form (3.3.8.1) which has the added benefit of allowing us to construct a true parametrix rather than just a 1-parametrix.

To construct parametrices for elliptic operators with variable coefficients, we study a general class of operators known as *pseudo-differential operators*. These are (for us) operators $T_A : C_c^\infty(\mathbb{R}^n) \to C_c^\infty(\mathbb{R}^n)$ of the form

$$(T_A u)(x) = \int e^{2\pi i\langle \xi, x\rangle} A(x,\xi)\hat{u}(\xi)\, d\xi. \tag{3.3.8.2}$$

where A has compact x-support and is a *symbol of order* $\leq m$, meaning that $|D_x^\alpha D_\xi^\beta A(x,\xi)| \leq \text{const}_{\alpha,\beta} \cdot (1 + |\xi|)^{m-|\beta|}$.

3.3.9 Example. A compactly supported differential operator of order $\leq m$ is of the form (3.3.8.2) where A is a polynomial in ξ whose coefficients are smooth functions of x (and conversely).

3.3.10 Exercise. Suppose that L is an elliptic operator of order m, and let $\varphi(x)$ be a cutoff function. Define $B(x,\xi)$ to be smooth, of compact x-support, and satisfy $B(x,\xi) = \varphi(x)P_x(i\xi)^{-1}$ for $|\xi|$ sufficiently large (where $P_x(i\xi)$ is the symbol of L). Show that B is a symbol of order $\leq -m$.

3.3.11 Lemma. *If A is a symbol of order $\leq m$, then the operator T_A (3.3.8.2) satisfies $\|T_A u\|_s \leq \mathrm{const}_{A,s}\|u\|_{s+m}$.*

Proof. Basically the reasoning from (3.2.6) used to prove the special case when T_A is a differential operator applies without change. We have

$$\widehat{T_A u}(\zeta) = \int K(\zeta,\xi)\hat{u}(\xi)\,d\xi \qquad \text{where } K(\zeta,\xi) = \int e^{2\pi i\langle \xi-\zeta,x\rangle}A(x,\xi)\,dx. \tag{3.3.11.1}$$

We again have $|K(\zeta,\xi)| \leq \mathrm{const}_{A,N}\cdot(1+|\xi-\zeta|)^{-N}\cdot(1+|\xi|)^m$ by integrating by parts and appealing to the given bounds on A. This bound on K implies the desired estimate. ☐

3.3.12 Proposition (Composition of pseudo-differential operators)**.** *Consider two operators T_A and T_B of the form* (3.3.8.2) *where A and B are symbols of order m_A and m_B, respectively. Then we have $T_A \circ T_B = T_C$ where C is a symbol of order m_A+m_B and has asymptotic expansion*

$$C(x,\xi) \sim \sum_\alpha \frac{D_\xi^\alpha A(x,\xi) D_x^\alpha B(x,\xi)}{\alpha!} \tag{3.3.12.1}$$

meaning that the difference between C and the sum of terms on the right with $|\alpha| < N$ is a symbol of order $m_A + m_B - N$.

Proof. As we have seen, the action of an operator of the form T_A on Fourier transforms is given by integration against an appropriate kernel (3.3.11.1). Moreover, the decay properties of these kernels imply that the composition of operators is given by composing the kernels. That is, we have $(T_A T_B u)^\wedge(\eta) = \int K_C(\eta,\xi)\hat{u}(\xi)\,d\xi$ where

$$K_C(\eta,\xi) = \int K_A(\eta,\zeta)K_B(\zeta,\xi)\,d\zeta \tag{3.3.12.2}$$

$$= \iiint e^{2\pi i\langle \zeta-\eta,y\rangle}A(y,\zeta)e^{2\pi i\langle \xi-\zeta,x\rangle}B(x,\xi)\,dy\,dx\,d\zeta \tag{3.3.12.3}$$

$$= \int e^{2\pi i\langle \xi-\eta,y\rangle}\left[\iint e^{2\pi i\langle \zeta-\xi,y-x\rangle}A(y,\zeta)B(x,\xi)\,dx\,d\zeta\right]dy \tag{3.3.12.4}$$

$$= \int e^{2\pi i\langle \xi-\eta,y\rangle}\left[\iint e^{-2\pi i\langle \beta,t\rangle}A(y,\xi+\beta)B(y+t,\xi)\,dt\,d\beta\right]dy \tag{3.3.12.5}$$

At least formally, the bracketed expression in the middle will be our new symbol C, however we still need to justify the interchange of order of integration. We begin with the first triple integral. It is not absolutely convergent: it is defined (so that it equals the first integral of K_A against K_B) by first integrating with respect to x and y to get something with rapid decay in ζ, and then integrating $d\zeta$. However, merely doing first the integral dx already gives us rapid decay in ζ, so we can interchange the y integral and ζ integral. Thus if we define

$$C(y,\xi) = \iint e^{-2\pi i\langle \beta,t\rangle} A(y,\xi+\beta)B(y+t,\xi)\,dt\,d\beta \tag{3.3.12.6}$$

(note the order of integration: after doing the integral dt, we have rapid decay in β) then we have $T_A \circ T_B = T_C$ provided we show that C is a symbol of some order.

The key to showing that C admits the desired asymptotic expansion (3.3.12.1) (and thus is a symbol of order $m_A + m_B$) is to consider the Taylor approximation

$$A(y,\xi+\beta) = \sum_{|\alpha|<N} \frac{D_\xi^\alpha A(y,\xi)}{\alpha!}\beta^\alpha + \text{error}. \tag{3.3.12.7}$$

If in the definition of C we replace $A(y,\xi+\beta)$ by this Taylor approximation, we obtain precisely the asymptotic expansion (3.3.12.1) (the integral $\iint e^{-2\pi i\langle\beta,t\rangle}\beta^\alpha B(y+t,\xi)\,dt\,d\beta$ is a y-derivative of $B(y,\xi)$ by Fourier inversion and integration by parts). It thus suffices to show that the error in the Taylor expansion above contributes a symbol of order $\le m_A + m_B - N$ to C.

We wish to bound

$$\int \left[\int e^{-2\pi i\langle\beta,t\rangle} B(y+t,\xi)\,dt\right]\left(A(y,\xi+\beta) - \sum_{|\alpha|<N}\frac{D_\xi^\alpha A(y,\xi)}{\alpha!}\beta^\alpha\right)d\beta, \tag{3.3.12.8}$$

or, rather, show it is a symbol of order $\le m_A + m_B - N$ in the variables (y,ξ). The condition of being a symbol of this order involves *derivatives* with respect to y and x, but we note a simplification: such derivatives fall on A and B, and $D_x^\alpha D_\xi^\beta A$ is itself a symbol of order $m_A - |\beta|$, so the derivatives of the above are of the same for, with A and B replaced by other symbols (namely their derivatives). It therefore suffices to bound the expression above by $\text{const}(1+|\xi|)^{m_A+m_B-N}$.

The bracketed integral dt is bounded by $\text{const}_M(1+|\xi|)^{m_B}(1+|\beta|)^{-M}$ for any $M<\infty$ (use the fact that B is a symbol of order m_B), and the difference in parentheses is, by the Taylor theorem, bounded by $\text{const}_N|\beta|^N(1+|\xi+\beta|)^{m_A-N}$ (the last factor bounds the order N ξ-derivatives of A since A is a symbol of order m_A). We are therefore left with

$$\text{const}_M(1+|\xi|)^{m_B}\int \beta^{-M}(1+|\xi+\beta|)^{m_A-N}\,d\beta. \tag{3.3.12.9}$$

The integral over the locus $|\beta| \ge \frac{1}{2}|\xi|$ has rapid decay (faster than any polynomial in ξ) due to the β^{-M} factor. The integral over the locus $|\beta| \le \frac{1}{2}|\xi|$ contributes at most $(1+|\xi|)^{m_A-N}$, so we are done. □

3.3.13 Definition. Denote by S^m the vector space of symbols of order $\le m$. Let $S = \bigcup_m S^m$ be the ascending union of the spaces S^m, and let $S^{-\infty} := \bigcap_m S^m$ be their intersection.

Let $\hat{S} = \varprojlim_m S/S^m$ denote the 'completion'; an element of $\hat{S}$ is an 'asymptotic expansion'. We will write $\sim$ for equality in $\hat{S}$ following the asymptotic composition formula (3.3.12.1) (which is exactly the statement of an equality in $\hat{S}$). Evidently $\ker(S \to \hat{S}) = S^{-\infty}$. Order is additive under composition of operators, so composition descends to $\hat{S}$; evidently the formula (3.3.12.1) continues to apply.

3.3.14 Exercise ($S/S^{-\infty}$ is complete). Show that the map $S \to \hat{S}$ is surjective (show that given a sequence of symbols $A_i \in S$ with $A_{i+1} - A_i \in S^{-i}$, there exists another such sequence $\tilde{A}_i$ with $\tilde{A}_i - A_i$ of compact ξ-support such that $\tilde{A}_i \in S$ converges to $A \in S$).

3.3.15 Corollary (Inverse symbol)**.** *Let L be an elliptic operator of order m on an open set $\Omega \subseteq \mathbb{R}^n$. There exists a symbol Q defined on $\Omega \times \mathbb{R}^n$ which is inverse to L in $\hat{S}$, meaning that $LQ \sim \mathbf{1} \sim QL$.*

Proof. This is like inverting an upper triangular matrix. Since the symbol of L is invertible for large ξ, we may take Q to be smooth and equal to the pointwise inverse of the symbol of L, and Q will have order $-m$ (3.3.10). By the asymptotic composition formula (3.3.12.1), the difference $\mathbf{1} - LQ$ is (in $\hat{S}$) a symbol of order -1. We may now add to Q a symbol of order $-m-1$ (namely the symbol of $\mathbf{1} - LQ$ multipled on the left by the pointwise inverse to L) to ensure that $\mathbf{1} - LQ$ has order -2. Iterating, we produce an element $Q \in \hat{S}$ such that $LQ \sim \mathbf{1}$ (equality in $\hat{S}$). Similarly, we can produce Q' such that $Q'L \sim \mathbf{1}$. Now $Q \sim Q'LQ \sim Q'$. To conclude, note that $Q \in \hat{S}$ lifts to S by completeness (3.3.14). □

3.3.16 Construction (Parametrices for elliptic operators)**.** Let L be an elliptic operator of order m on an open set $\Omega \subseteq \mathbb{R}^n$, and let Q denote the inverse to L in $\hat{S}(\Omega)$ (3.3.15). Let $\varphi : \Omega \to \mathbb{R}$ be compactly supported, and let $\psi : \Omega \to \mathbb{R}$ be compactly supported and equal 1 on $\operatorname{supp} \varphi$. We have $\psi Q \varphi \sim Q\varphi$ (equality in $\hat{S}$) since both sides are supported inside $\operatorname{supp} \varphi$ and over this set $\psi \equiv 1$; similarly $\varphi Q \psi \sim \varphi Q$. Multiplying by L and appealing to the fact that $LQ \sim \mathbf{1} \sim QL$, we conclude that

$$\varphi Q \psi L \sim \varphi \sim L \psi Q \varphi \tag{3.3.16.1}$$

Both φQ and ψQ are compactly supported symbols in S, so both $\varphi Q \psi$ and $\psi Q \varphi$ therefore define operators $C^\infty \to C^\infty_c$ on Ω.

We now consider an elliptic operator L on a manifold M. Let $M = \bigcup_i U_i$ be an open cover by Euclidean charts, let $\varphi_i : M \to \mathbb{R}$ be a subordinate partition unity, and let $\psi_i : M \to \mathbb{R}$ be supported inside U_i and $\equiv 1$ on a neighborhood of $\operatorname{supp} \varphi_i$. On each chart U_i, we have a symbol Q_i inverse to $L|_{U_i}$, which thus satisfies (3.3.16.1). We claim that both

$$Q = \sum_i \psi_i Q_i \varphi_i \qquad Q' = \sum_i \varphi_i Q_i \psi_i \tag{3.3.16.2}$$

(evidently operators of order $-m$) are parametrices for L. Indeed, it is evident from (3.3.16.1) that $\mathbf{1} - QL$ and $\mathbf{1} - LQ'$ are smoothing operators. As Q is a "left inverse modulo smooth operators" and Q' is a "right inverse modulo smoothing operators", we should expect

$Q - Q'$ is a smoothing operator, and indeed $Q - Q' = Q(\mathbf{1} - LQ') - (\mathbf{1} - QL)Q'$. Thus $\mathbf{1} - LQ = (\mathbf{1} - LQ') + L(Q' - Q)$ and $\mathbf{1} - Q'L = (\mathbf{1} - QL) + (Q - Q')L$ are also smoothing operators.

3.3.17 Corollary (Elliptic regularity)**.** *Let $L : C^\infty(M, E) \to C^\infty(M, F)$ be an elliptic operator of order m.*

(3.3.17.1) *For $u \in C^{-\infty}(M, E)$, if $Lu \in H^s_{\mathrm{loc}}$ then $u \in H^{s+m}_{\mathrm{loc}}$.*

(3.3.17.2) *For $f \in C^{-\infty}(M, F)$, there exists $u \in C^{-\infty}(M, E)$ with $f - Lu \in C^\infty$.*

Proof. Suppose $u \in C^{-\infty}(M, E)$ and $Lu \in H^s_{\mathrm{loc}}$. Fix a coordinate chart on M, a smooth cutoff function φ supported in this chart, and an operator Q of order $-m$ such that $\varphi - QL$ gains one derivative. Then we have $\varphi u = (\varphi - QL)u + QLu$, so $u \in H^r_{\mathrm{loc}}$ (in our chosen coordinate chart) implies $u \in H^{\min(r+1,s+m)}$; eventually we get $u \in H^{s+m}_{\mathrm{loc}}$.

Now let $f \in C^{-\infty}(M, F)$. Let Q be an operator of order $-m$ so that $\mathbf{1} - LQ$ is a smoothing operator. Thus $f - L(Qf) = (\mathbf{1} - LQ)f \in C^\infty$, so we may take $u = Qf$. □

3.3.18 Exercise (Kernel and cokernel of an elliptic operator). Let L be an elliptic operator of order m. Use elliptic regularity (3.3.17) to show that each of the natural inclusions between the two-term complexes

$$C^\infty(M, E) \xrightarrow{L} C^\infty(M, F) \tag{3.3.18.1}$$

$$H^s(M, E) \xrightarrow{L} H^{s-m}(M, F) \tag{3.3.18.2}$$

$$C^{-\infty}(M, E) \xrightarrow{L} C^{-\infty}(M, F) \tag{3.3.18.3}$$

is a quasi-isomorphism. Conclude that $\ker L$ and $\operatorname{coker} L$ have an invariant meaning, and $\ker L \subseteq C^\infty(M, E)$ and $C^\infty(M, F) \twoheadrightarrow \operatorname{coker} L$.

3.3.19 Remark (Parametrices as chain homotopies). Not only are the inclusions of two-term complexes above quasi-isomorphisms, a choice of parametrix Q shows this in a very clean way. For an economy of notation, we write out only the case of $H^s \hookrightarrow H^t$ for $s \geq t$, but the rest are the same. The mapping cone of this inclusion is the total complex of the double complex

$$\begin{array}{ccc} H^s(M, E) & \xrightarrow{L} & H^{s-m}(M, F) \\ {\scriptstyle \mathbf{1}}\downarrow & & \downarrow{\scriptstyle \mathbf{1}} \\ H^t(M, E) & \xrightarrow{L} & H^{t-m}(M, F) \end{array} \tag{3.3.19.1}$$

The endomorphism of this double complex given by

$$\begin{array}{ccc} H^s(M, E) & \xleftarrow{Q} & H^{s-m}(M, F) \\ {\scriptstyle \mathbf{1}-QL}\uparrow & & \uparrow{\scriptstyle \mathbf{1}-LQ} \\ H^t(M, E) & \xleftarrow{Q} & H^{t-m}(M, F) \end{array} \tag{3.3.19.2}$$

is a chain homotopy between the identity map and the zero map, where we use crucially the fact that $\mathbf{1} - LQ$ and $\mathbf{1} - QL$ are smoothing operators. This is a somewhat stronger property than merely being acyclic, as this chain homotopy is a continuous linear map.

The next properties of elliptic operators require compactness of the manifold in question, or at the very least some sort of bounded geometry.

3.3.20 Corollary (Kernel finiteness)**.** *For an elliptic operator L on a compact manifold,* $\ker L$ *is finite-dimensional.*

Proof. It suffices to show that there exists a finite collection of points $P \subseteq M$ such that an element of $\ker L$ which vanishes on P must be zero. Take P to be any set of points such that the ε-balls centered at P cover M. Thus $\varphi|_P \equiv 0$ implies that $\|\varphi\|_{C^0} \le \varepsilon \|\varphi\|_{C^1}$. On the other hand, $\varphi \in \ker L$ implies that $\|\varphi\|_{C^1} \le \mathrm{const}_L \|\varphi\|_{C^0}$ by ellipticity of L. Together these imply $\|\varphi\|_{C^0} \le \varepsilon \mathrm{const}_L \|\varphi\|_{C^0}$, so taking $\varepsilon > 0$ sufficient small implies that $\|\varphi\|_{C^0} = 0$. □

3.3.21 Exercise. Adapt the argument from (3.3.20) to show that for an elliptic operator $L : C^\infty(M, E) \to C^\infty(M, F)$ on a manifold M, there exists a function $\varepsilon : M \to \mathbb{R}_{>0}$ depending only on the local geometry of L such that if $P \subseteq M$ is any set of points with $M = \bigcup_{p \in P} B(p, \varepsilon(p))$, then $\ker L \to \prod_{p \in P} E_p$ is injective (the key is a local elliptic estimate of the form $\|\varphi\|_{C^1(K)} \le \mathrm{const}_{L,K,U} \|\varphi\|_{C^0(U)}$ for $U \subseteq M$ open and $K \subseteq U$ compact).

3.4 Ellipticity in cylindrical ends

The results on elliptic operators from the previous section take their strongest form when the manifold under consideration is compact. We now recall how to extend these results to manifolds with cylindrical ends following Lockhart–McOwen [16].

3.4.1 Definition (Cylindrical operator)**.** We consider the cylinder $\mathbb{R} \times N$ for a compact manifold N. There is a tautological action of $\mathbb{R}$ on $\mathbb{R} \times N$, and we consider $\mathbb{R}$-equivariant (aka cylindrical) vector bundles E and F on $\mathbb{R} \times N$ (equivalently, E and F are identified as pullbacks of vector bundles on N). A cylindrical operator on $\mathbb{R} \times N$ is an $\mathbb{R}$-equivariant operator $C^\infty(\mathbb{R} \times N, E) \to C^\infty(\mathbb{R} \times N, F)$.

3.4.2 Definition (Reduction and twisting of cylindrical operators)**.** Given a cylindrical operator $L : C^\infty(\mathbb{R} \times N, E) \to C^\infty(\mathbb{R} \times N, F)$, its reduction is the operator $L_0 : C^\infty(N, E) \to C^\infty(N, F)$ defined by restricting L to $\mathbb{R}$-invariant sections. For any complex number $z \in \mathbb{C}$, we may consider the conjugation $e^{zt} L e^{-zt}$ which is again a cylindrical operator. Its reduction is denoted $L_z : C^\infty(N, E) \to C^\infty(N, F)$, which may also be viewed as the restriction of L to sections which transform by the character e^{zt} under $\mathbb{R}$-translation. The operator L_z depends polynomially on z, and in fact specifying a cylindrical operator $C^\infty(\mathbb{R} \times N, E) \to C^\infty(\mathbb{R} \times N, F)$ is the same as specifying a polynomial in z valued in operators $C^\infty(N, E) \to C^\infty(N, F)$.

3.4.3 Exercise. Consider $\mathbb{R} \times S^1$ with the complex structure $J(\partial_t) = \partial_s$ and the cylindrical operator $L = \frac{1}{2}(\partial_t - i\partial_s)$ acting on functions valued in $\mathbb{C}$. What are the operators L_z?

3.4.4 Remark. The definition of twisting of cylindrical operators implicitly assumes the setting of complex vector bundles. We leave it to the reader to elaborate on the case of real vector bundles (in a word, one passes to their complexifications and imposes hypotheses such as $f(x) = \overline{f(x)}$, $g(-\xi) = \overline{g(\xi)}$, $L_{\bar{z}} = \overline{L_z}$, etc. as appropriate); this applies to the rest of this section as well.

3.4.5 Definition (Cylindrical Schwartz space $\mathcal{S}(\mathbb{R} \times N)$)**.** The Schwartz space $\mathcal{S}(\mathbb{R} \times N)$ consists of those smooth functions all of whose seminorms $\sup(1 + |t|)^n |D^\alpha f|$ are finite, where $n < \infty$ and D^α stands for any $\mathbb{R}$-invariant differential operator; it is complete with respect to this family of seminorms. This makes sense for functions valued in any cylindrical vector bundle.

3.4.6 Exercise (Cylindrical Fourier transform). Define the Fourier transform $\mathcal{S}(\mathbb{R} \times N) \to \mathcal{S}(\mathbb{R} \times N)$ by the usual integral in just the $\mathbb{R}$-coordinate $\hat{f}(\xi, n) = \int e^{-2\pi i \langle \xi, x\rangle} f(x, n)\, dx$. Show that the Fourier transform is continuous and that Fourier inversion holds. Show that for a cylindrical differential operator L, we have $\widehat{Lu}(\xi, n) = L_{i\xi}\hat{u}(\xi, n)$.

3.5 Families of elliptic operators

We study how the results of the previous sections on function spaces (3.1) and elliptic operators (3.2)–(**??**) apply in families. Much of this is tautological, as most of these results consist of explicit estimates, which are evidently uniform in parameters. One less trivial fact about how elliptic operators vary in families is that the pushforward of a proper family of elliptic operators is (locally) quasi-isomorphic to a two-term complex of finite-dimensional vector bundles on the base.

3.5.1 Definition (Relative smooth manifold)**.** Let B be a topological space, and let $M \to B$. A *smooth atlas* on $M \to B$ consists of a set of diagrams

$$\begin{array}{ccc} U \times V & \longrightarrow & M \\ \downarrow & & \downarrow \\ U & \longrightarrow & B \end{array} \tag{3.5.1.1}$$

in which the horizontal arrows are both open embeddings, U is a topological space, and $V \subseteq \mathbb{R}^n$ is open, subject to the condition that the transition maps are *relatively smooth*. A transition map between two charts of the above form takes the form of a diagram

$$\begin{array}{ccc} A & \longrightarrow & A' \\ \downarrow & & \downarrow \\ U \cap U' & = & U \cap U' \end{array} \tag{3.5.1.2}$$

where $A \subseteq U \cap U' \times \mathbb{R}^n$ and $A' \subseteq U \cap U' \times \mathbb{R}^{n'}$ are open subsets. Such a map is called relatively smooth iff all its derivatives in the $\mathbb{R}^n$ direction exist and are continuous.

On any such $M \to B$, we may talk about relatively smooth functions (maps to $\mathbb{R}$ which in any chart are relatively smooth), relatively smooth vector bundles (transition functions are relatively smooth), relatively smooth sections thereof, and relatively smooth differential operators (given in local coordinates by relatively smooth sections), which act on relatively smooth sections.

On a proper relative smooth manifold $M \to B$, locally on B, we can cover with finitely many charts, and there is a relatively smooth partition of unity subordinate to them. This means that locally on B, we have H^s-norms, C^k-norms, etc., which are well-defined up to commensurability.

3.5.2 Proposition. *A proper relative smooth manifold is locally trivial.*

Proof. Let $M \to B$ be a proper relative smooth manifold. Fix a point $0 \in B$, and denote by M_0 the fiber over it. Using a partition of unity argument, we can define a map $M \to M_0 \times B$ over B in a neighborhood of $M_0 \subseteq M$ on which it is the identity map. This map is defined on the inverse image of a neighborhood of $0 \in B$ by properness. In fact, its derivative is invertible over a neighborhood of M_0 (hence over a neighborhood of $0 \in B$) since this is an open condition and it holds on M_0. It is therefore a relative diffeomorphism over the inverse image of a neighborhood of 0. □

3.5.3. Let $M \to B$ be a proper relative smooth manifold carrying vector bundles E and F, and let $L : E \to F$ be a relatively smooth elliptic operator. In a neighborhood of any fixed basepoint $0 \in B$, properness of $M \to B$ implies (3.5.2) the existence of trivializations

$$\begin{array}{ccc} M_0 \times B & \xrightarrow{\varphi} & M \\ & \searrow \quad \swarrow & \\ & B & \end{array} \qquad \begin{array}{ccc} E_0 \times B & \xrightarrow{\varphi_E} & E \\ \downarrow & & \downarrow \\ M_0 \times B & \xrightarrow{\varphi} & M \end{array} \qquad \begin{array}{ccc} F_0 \times B & \xrightarrow{\varphi_F} & F \\ \downarrow & & \downarrow \\ M_0 \times B & \xrightarrow{\varphi} & M \end{array} \tag{3.5.3.1}$$

With respect to such trivializations, we simply have a family of elliptic operators $L_b : E_0 \to F_0$ on M_0 depending continuously on $b \in B$ in a neighborhood of $0 \in B$ (depending continuously means, concretely, the coefficients of L_b and all their derivatives are continuous functions on $M_0 \times B$).

The constants in the estimates from (3.3) are given in terms of bounds on the coefficients of the operator in question, so these constants are uniform locally on B. For example, the dimension of $\ker L_b$ is a locally bounded function on B.

To continue studying families of elliptic operators, we introduce some notation. Let $C^\infty(M, E)$ denote the sheaf of relatively smooth sections of E, so L defines a map of sheaves $C^\infty(M, E) \to C^\infty(M, F)$, whose kernel is denoted $\ker L$. We can push forward any of these sheaves to B.

3.5.4 Lemma. *Let $L : E \to F$ be a relatively smooth elliptic operator on a proper relative smooth manifold $\pi : M \to B$. The locus of $b \in B$ for which D_b is surjective is open, and $\pi_* \ker D = \ker \pi_* D$ is a vector bundle over it.*

Proof. We work in local coordinates (3.5.3.1) near a fixed basepoint $0 \in B$.

Let Q_0 denote a right inverse to L_0. Via the trivializations (3.5.3.1), the operator Q_0 determines an operator $Q^{\mathrm{pre}} : \pi_* C^\infty(M, F) \to \pi_* C^\infty(M, E)$ defined in a neighborhood of 0. Now the operator $\mathbf{1}_F - LQ^{\mathrm{pre}}$ vanishes at $b = 0$, so by continuity, its restriction to nearby fibers has very small (s, s)-norm in a neighborhood of $0 \in B$ (depending on s). It follows that the series expansion

$$Q = Q^{\mathrm{pre}}(\mathbf{1}_F + (\mathbf{1}_F - LQ^{\mathrm{pre}}) + (\mathbf{1}_F - LQ^{\mathrm{pre}})^2 + \cdots) \tag{3.5.4.1}$$

provides a genuine right inverse $LQ = \mathbf{1}_F$ (3.3.6). This shows that surjectivity is open.

To show that $\pi_* \ker L$ is a vector bundle, we construct a local trivialization of it near 0 as follows. The trivialization (3.5.3.1) provides an identification

$$C^\infty(M_0, E_0)_B \rightleftarrows \pi_* C^\infty(M, E). \tag{3.5.4.2}$$

The right inverses Q_0 and Q define projections

$$\mathbf{1}_E - Q_0 L_0 : C^\infty(M_0, E_0) \to \ker L_0 \tag{3.5.4.3}$$

$$\mathbf{1}_E - QL : \pi_* C^\infty(M, E) \to \pi_* \ker L \tag{3.5.4.4}$$

Composing the identification maps (3.5.4.2) in either direction with these projections defines maps $(\ker L_0)_B \to \pi_* \ker L$ and $\pi_* \ker L \to (\ker L_0)_B$. Theses are mutual inverses at 0, hence are approximate inverses in a neighborhood, hence isomorphisms by (3.3.7). $\square$

Chapter 4

Riemann surfaces

4.1 Riemann surfaces

4.1.1 Definition (Riemann surface)**.** A Riemann surface is a one-dimensional complex manifold. This means that it consists of a paracompact Hausdorff topological space C together with a complex analytic atlas, namely a set of open embeddings $\varphi_i : U_i \to C$ (called charts) from open subsets $U_i \subseteq \mathbb{C}$ such that the transition maps $\varphi_i^{-1}(\varphi_j(U_j)) \to \varphi_j^{-1}(\varphi_i(U_i))$ are holomorphic.

4.2 Families of Riemann surfaces

4.2.1 Definition. A family of Riemann surfaces over a topological space B is a separated map $C \to B$ together with a relative complex analytic atlas. Such an atlas consists of charts of the form $U \times D^2 \to U$ whose transition maps are holomorphic on each fiber (which implies they vary continuously in the C^∞ topology). In a family of nodal Riemann surfaces over a topological space B, another sort of chart is allowed, namely an open subset of a pullback of the multiplication map $\mathbb{C}^2 \to \mathbb{C}$. Such a chart can be expressed alternatively as $(U \times [0, \infty) \times S^1)\#_\lambda(U \times (-\infty, 0] \times S^1) \to U$ for a continuous function $\lambda : U \to D^2$ which specifies how to glue together $[0, \infty) \times S^1$ and $(-\infty, 0] \times S^1$.

It will be important to note that the transition maps between such charts necessarily converge exponentially to translation/rotation maps. In fact, any holomorphic map $I \times S^1 \to I' \times S^1$ (where $I, I' \subseteq \mathbb{R}$ are intervals) preserving the homotopy class of S^1 necessarily takes the form $(s, t) \mapsto (s + a, t + b)$ for constants a and b plus a term whose kth order derivatives are bounded by $M_k e^{-s}$ for some absolute constants $M_k < \infty$ and s denotes the distance from the boundary of I (the domain of the map).

4.3 Cauchy–Riemann operators

4.3.1 Definition (Cauchy–Riemann operator)**.** Let C be a Riemann surface. A real Cauchy–Riemann operator on a smooth complex vector bundle E/C is a first order $\mathbb{R}$-linear differential operator $D : C^\infty(C, E) \to C^\infty(C, E \otimes_{\mathbb{C}} \Omega_C^{0,1})$ satisfying

$$D(f \cdot s) = f \cdot Ds + \bar{\partial} f \cdot s. \tag{4.3.1.1}$$

In local coordinates, a real Cauchy–Riemann operator is given (acting on functions on an open subset of $\mathbb{C}$ valued in a complex vector space E_0) by $(Ds)(z) = (\bar{\partial} s)(z) + a(z)(s(z))$ for some smooth map $a : \mathbb{C} \to \operatorname{End}_{\mathbb{R}}(E_0)$.

4.3.2 Definition (Cylindrical vector bundle)**.** Let C be a Riemann surface, and let E/C be a smooth complex vector bundle. A cylindrical chart for E near a puncture of C is an identification of the restriction of E to $[0, \infty) \times S^1$ with the pullback of a complex vector bundle F on S^1. A cylindrical structure on E is an equivalence class of cylindrical charts, where two cylindrical charts are called equivalent iff the transition function between them is, up to error $O(e^{-\delta s})$, the pullback of a smooth isomorphism of vector bundles over S^1. A cylindrical structure on E/C at a given puncture induces an extension of E to the circle at infinity at that puncture.

4.3.3 Exercise. Fix a complex vector bundle $F \to S^1$, and let E denote its pullback to $\mathbb{R} \times S^1$. Show that $\mathbb{R}$-equivariant real Cauchy–Riemann operators on E are in natural bijection with real connections on $F \to S^1$.

4.3.4 Definition (Cylindrical Cauchy–Riemann operator)**.** Let C be a Riemann surface, and let E/C be a smooth complex vector bundle equipped with a cylindrical structure. A Cauchy–Riemann operator D on E is called cylindrical iff there exists a (necessarily unique) real connection on the associated vector bundle over the circle at infinity, such that with respect to some (equivalently, any) cylindrical chart, it coincides with the induced real Cauchy–Riemann operator up to error all of whose derivatives are $O(e^{-\delta s})$.

4.4 Families of Cauchy–Riemann operators

Let $C \to B$ be a family of Riemann surfaces. Given a (relatively) smooth vector bundle E/C, we may consider a family D of Cauchy–Riemann operators acting on E/C. This means that the coefficients of D and their vertical derivatives are all defined and continuous. When the family $C \to B$ has punctures, the cylindrical charts (4.3.2) (4.3.4) must be relatively smooth, with the estimates holding uniformly. When the family $C \to B$ includes nodal resolutions, cylindricity is imposed with respect to the nodal resolution charts (note that the transition maps between these converge exponentially to translation/rotation).

4.4.1 Lemma. *Let $\pi : C \to B$ be a family of nodal Riemann surfaces equipped with a family of Cauchy–Riemann operators D. The locus of $b \in B$ for which D_b is surjective is open, and $\pi_* \ker D = \ker \pi_* D$ is a vector bundle over it.*

Chapter 5

Pseudo-holomorphic curves

5.1 Pseudo-holomorphic maps

5.1.1 Definition (Pseudo-holomorphic maps)**.** A map $u : C \to X$ from a Riemann surface C to an almost complex manifold (X, J) is called *pseudo-holomorphic* iff its differential $du : TC \to u^*TX$ is complex linear.

We may decompose $du = (du)^{1,0} + (du)^{0,1}$ into its holomorphic and anti-holomorphic components, which take values in $\Omega_C^{1,0}$ and $\Omega_C^{0,1}$ tensor u^*TX, respectively. The equation asserting pseudo-holomorphicity of u may thus be written

$$(du)^{0,1} = 0. \tag{5.1.1.1}$$

Given an almost complex submersion $\pi : W \to C$ over a Riemann surface C (meaning W is an almost complex manifold and π is complex linear), we may consider pseudo-holomorphic sections $u : C \to W$. Since $\pi \circ u = \mathbf{1}_C$ is pseudo-holomorphic, the anti-holomorphic derivative $(du)^{0,1}$ of any section u takes values in $\Omega_C^{0,1}$ tensor $u^*T_{W/C}$.

5.2 A priori estimates

We present here the fundamental *a priori* estimates on pseudo-holomorphic maps. These are non-linear analogues of the linear elliptic estimates from (3.3)–(**??**). A straightforward elliptic bootstrapping argument shows that a gradient bound on a pseudo-holomorphic map u (i.e. a bound on the supremum of $|du|$) imply C^∞ bounds (i.e. bounds on all derivatives of u). A more subtle question is what can be derived from C^0 bounds (i.e. bounds on the image of u) or energy bounds (i.e. bounds on $\int |du|^2$). Even taken together, C^0 bounds and energy bounds do not imply gradient bounds, due to the bubbling phenomenon. We will see, however, that energy bounds or C^0 bounds which are sufficiently small do imply a gradient bound. These results are the first step towards *Gromov compactness*, which provides a complete description of the consequences of C^0 bounds and energy bounds.

The results in this section are known (except for possibly a few marginal improvements of no consequence). We give the simplest proofs we know, but there are quite a number of other approaches; see [6, 13, 25, 21, 10, 11, 5, 19, 2].

We work in the smooth category except when stated otherwise.

5.2.1 Proposition (Gradient bounds imply C^∞ bounds). *Let $u : D^2 \to (B(1), J)$ be pseudo-holomorphic with $\|u\|_{C^1} \leq M$. We then have $|D^k u(0)| \leq \operatorname{const}_{k,J,M} \|u\|_{C^1}$.*

Proof. Applying $\frac{d}{dx} - \frac{d}{dy}J(u)$ to the equation $u_x + J(u)u_y = 0$ yields

$$u_{xx} + u_{yy} = \dot{J}(u, u_y)u_x - \dot{J}(u, u_x)u_y. \tag{5.2.1.1}$$

The L^2-norm of the right hand side is bounded by $\operatorname{const}\|u\|_{C^1}$, so by elliptic regularity we have $\|u\|_{W^{2,2}(K)} \leq \operatorname{const} \cdot \|u\|_{C^1}$ for compact $K \subseteq (D^2)^\circ$. Using this, we see that the $W^{1,2}$-norm of the right hand side is bounded (over compact subsets of $(D^2)^\circ$) by $\operatorname{const} \cdot \|u\|_{C^1}$ (inspect its derivative, applying the C^1-bound to first derivatives of u and the $W^{2,2}$ bound to second derivatives), so we have $\|u\|_{W^{3,2}(K)} \leq \operatorname{const} \cdot \|u\|_{C^1}$. For the remaining bootstrapping, observe right hand side is a smooth function vanishing at zero applied to (u, Du). For $k \geq 3$, a $W^{k,2}$ bound on u is a $W^{k-1,2}$ bound on (u, Du), which implies (since $W^{k-1,2} \subseteq C^0$) a $W^{k-1,2}$ bound on its post-compostion with any smooth function vanishing at zero, linear in the original bound (3.1.18). We therefore have $\|u\|_{W^{k,2}(K)} \leq \operatorname{const}_k \cdot \|u\|_{C^1}$. □

5.2.2 Exercise. Deduce from (5.2.1) that $|D^k u(p)| \leq \operatorname{const}_{k,J,M} \|u\|_{C^1}(1 + d(p, \partial D^2)^{-(k-1)})$.

5.2.3 Definition (Energy). The *energy* of a pseudo-holomorphic map is its area $\int |du|^2$. This requires a measurement of 2-planes in the target (but no data on the source). We will only ever require the energy functional $u \mapsto \int |du|^2$ to be well-defined up to a constant factor, so a metric on the target is not needed unless the target is non-compact. The energy can also be written $\int u^*\omega$ for any 2-form ω taming J.

Energy is scale invariant, whereas the gradient is not. We therefore should not expect energy to control the size of the gradient. The following example makes this precise.

5.2.4 Exercise (Bubbling: energy bounds do not imply gradient bounds). Equip the Riemann sphere $S^2 = \mathbb{C}P^1 = \widehat{\mathbb{C}}$ with the round metric. Consider the maps $u_N : D^2 \to \mathbb{C}P^1$ given by $\times N : D^2 \to \mathbb{C} \subseteq \widehat{\mathbb{C}}$. Show that the energy of u_N is bounded uniformly in N, yet as $N \to \infty$, it concentrates near $0 \in D^2$ and the gradient $du_N(0)$ is unbounded. Show that u_N converges uniformly over compact subsets of $D^2 \setminus 0$ to the constant map to $\infty \in \widehat{\mathbb{C}}$. This is an example of bubbling: in a certain precise sense, the limit of the maps $u_N : D^2 \to \widehat{\mathbb{C}}$ as $N \to \infty$ is a map $D^2 \vee \mathbb{C}P^1 \to \widehat{\mathbb{C}}$ (where $\mathbb{C}P^1$ is glued to $0 \in D^2$) which is constant on D^2 and is a biholomorphism on $\mathbb{C}P^1$.

Given a pseudo-holomorphic map $u : \mathbb{C} \to X$ and a point $p \in \mathbb{C}$ with $|du(p)|$ very large, one can identify the disk of radius $M \cdot |du(p)|^{-1}$ centered at p with the disk $D^2_M \subseteq \mathbb{C}$ of radius M to obtain a map $\tilde{u} : D^2_M \to X$ with $|d\tilde{u}(0)| = 1$. If $\tilde{u}$ satisfies a gradient bound over all of D^2_M (and hence enjoys C^∞ bounds (5.2.1)), then the failure of u to satisfy a gradient bound can be intepreted, at least near p, as due to a bubble modelled on the reasonably nice map $\tilde{u}$. For a general point p for which $|du(p)|$ is large, the rescaled map $\tilde{u}$ need not satisfy a gradient bound. The purpose of Hofer's Lemma (5.2.5) below is to show that a gradient bound $|d\tilde{u}| \leq 2$ can always be achieved by moving p by no more than $2M \cdot |du(p)|^{-1}$.

5.2.5 Hofer's Lemma ([9, Lemma 3.3])**.** *Let (X, d) be a complete metric space, let $f : X \to \mathbb{R}_{\geq 0}$ be locally bounded, and let $M < \infty$. For every $p_0 \in X$, there exists $p \in X$ with $f(p) \geq f(p_0)$ and $d(p, p_0) \leq 2M \cdot f(p_0)^{-1}$ such that $d(x, p) \leq M \cdot f(p)^{-1} \implies f(x) \leq 2f(p)$.*

Proof. If p_0 does not satisfy the desired property, then there exists a violation point p_1, i.e. $d(p_0, p_1) \leq M \cdot f(p_0)^{-1}$ and $f(p_1) \geq 2f(p_0)$. If p_1 does not satisfy the desired property, there is a subsequent violation point p_2. We have $f(p_i) \geq 2^i f(p_0)$, hence $d(p_i, p_{i+1}) \leq 2^{-i} M \cdot f(p_0)^{-1}$, so $d(p_0, p_i) \leq 2M \cdot f(p_0)^{-1}$. This process $p_0, p_1, \ldots$ will eventually terminate at a suitable point p, since otherwise it would converge (since X is complete) to a point p_∞ near which f is not locally bounded. ☐

5.2.6 Definition (Bounded geometry)**.** Let (X, g) be a Riemannian manifold. A *bound on the geometry* of (X, g) at a point $p \in X$ is a collection of constants $\varepsilon > 0$ and $M_0, M_1, \ldots < \infty$ such that there exists a map $\Phi : B(1) \to X$ such that $\Phi^* g \geq \varepsilon g_{\mathrm{std}}$ and $\|\Phi^* g\|_{C^k} \leq M_k$ for every $k < \infty$. A bound on the geometry of (X, g) over a subset $A \subseteq X$ is a collection $\varepsilon, M_0, M_1, \ldots$ which bound the geometry at every $p \in A$. A bound on the geometry of (X, g) means $A = X$. A bound on the geometry and injectivity radius means that in addition Φ is required to be injective (beware that in standard terminology, a 'bound on the geometry' is usually taken to mean what we have decided to call a 'bound on the geometry and injectivity radius').

For an additional structure τ on the tangent bundle (e.g. a symplectic form, almost complex structure, or any combination thereof), bounded geometry of (X, g, τ) means that $\|\Phi^* \tau\|_{C^k} \leq M_k$ as well. When the data τ itself determines a Riemannian metric (e.g. tame pair (J, ω) determining the metric $\omega(v, Jw) + \omega(w, Jv)$), we will simply say (X, τ) has bounded geometry, to mean (X, g_τ, τ) has bounded geometry.

To say that a constant depends on the geometry of X over a given subset $A \subseteq X$ means that there exists $k < \infty$ such that for every $\varepsilon > 0$ and $M_0, \ldots, M_k < \infty$, there is a constant which works whenever the geometry of X over $A \subseteq X$ is bounded by $(\varepsilon, M_0, \ldots, M_k)$ (and $M_i = \infty$ for $i > k$). We will usually not concern ourselves with specifying such a k.

5.2.7 Proposition (Small energy bounds imply gradient bounds)**.** *Let $u : D^2 \to (X, J, g)$ be a pseudo-holomorphic map. If $\int_{D^2} |du|^2 < \varepsilon$ then $|du(0)|^2 \leq \mathrm{const} \int_{D^2} |du|^2$, where $\varepsilon > 0$ and $\mathrm{const} < \infty$ depend on bounds on the geometry of (X, J, g) over the image of u.*

Proof. If $\sup |du| \geq 5$ over the disk of radius $\frac{1}{2}$, then we can use Hofer's Lemma (5.2.5) to find p such that $|du(p')| \leq 2|du(p)|$ for p' in a ball of radius $|du(p)|^{-1}$ around p (which is entirely contained in D^2). Rescale so that this is our disk D^2 (it only decreases the energy!) so now $|du(0)| = 1$ and $|du| \leq 2$. We therefore have C^∞ bounds (5.2.1) on u over D^2. It follows that $|du(0)| = 1$ implies that energy is at least some $\varepsilon > 0$. We may therefore define $\varepsilon > 0$ so that this case does not occur.

We thus have $\sup|du| \leq 5$ over the disk of radius $\frac{1}{2}$. By rescaling this to the whole disk, we have $\sup|du| \leq 5/2$ over D^2, and hence we have C^∞ bounds (5.2.1) on u, depending on bounds on the geometry of the target. We now bootstrap using (5.2.1.1) and the energy bound. The right hand side is bounded in L^2 by $\int|du|^2$ (bound $\dot{J}(u, u_x)$ and $\dot{J}(u, u_y)$ using the C^∞ bounds on u and J). Thus $\|\Delta u\|_{L^2}$ is bounded by a constant times energy. Hence $\|\Delta u_x\|_{H^{-1}}$ is bounded by a constant times energy, and same for Δu_y, i.e. we have bounds on $\|\nabla u\|_{L^2}$ and $\|\Delta \nabla u\|_{H^{-1}}$ which by elliptic regularity combine to give a bound on $\|\nabla u\|_{H^1}$ linear in energy. Iterating, we obtain an estimate on $\|\nabla u\|_{H^2}$, which controls $\|\nabla u\|_{C^0}$, thus is the desired gradient estimate linear in energy. □

5.2.8 Exercise. Deduce from (5.2.7) that $|du(p)| \leq \text{const}(1 + d(p, \partial D^2)^{-1}) \int_{D^2}|du|^2$.

5.2.9. The *action* of a map $\gamma : S^1 \to \mathbb{C}^n$, denoted $a(\gamma)$, is the integral of $\omega_{\text{std}} = \sum_i dx_i \wedge dy_i$ over a disk with boundary γ. Using the expression $\omega_{\text{std}} = \frac{i}{2}\sum_i dz_i \wedge d\bar{z}_i = \frac{i}{2}\sum_i d(z_i\, d\bar{z}_i)$ and expanding $\gamma(t) = \sum_n a_n e^{int}$ in Fourier series, we have

$$a(\gamma) = \frac{i}{2}\int_{S^1} \gamma\, d\bar{\gamma} = \pi \sum_n n|a_n|^2 \leq \pi \sum_n n^2|a_n|^2 = \frac{1}{2}\int_{S^1} |\gamma'(t)|^2\, dt. \tag{5.2.9.1}$$

5.2.10 Exercise. Use the fact that $\sum_n n^{-2} < \infty$ to show that for every $\varepsilon > 0$ there exists $N = N(\varepsilon) < \infty$ such that there exists no smooth function $f : [0, N] \to \mathbb{R}$ satisfying $f(0) = \varepsilon$ and $f'(x) \geq f(x)^2$.

5.2.11 Gromov–Schwarz Lemma (Exact C^0-bounds imply gradient bounds)**.** *Let (X, J, λ) be tame. For any pseudo-holomorphic map $u : D^2 \to X$, we have*

$$|du(0)| \leq \text{const}, \tag{5.2.11.1}$$

where the constant depends on the geometry of (X, λ, J).

Proof. Assume $|du(0)| \geq 1$. Consider $\tilde{u} : D^2 \to X$ obtained by rescaling D^2 to the disk of radius $|du(0)|$ around zero and composing with u. Now $|d\tilde{u}(0)| = 1$, so since small energy bounds imply gradient bounds (5.2.7), we obtain a constant lower bound on the energy of u over the disk of radius $|du(0)|^{-1}$ centered at zero.

Consider cylindrical coordinates $[0, \infty) \times S^1 = D^2 \setminus 0 \subseteq D^2$, and let $E(s)$ be the energy of u over $[s, \infty) \times S^1$. We saw just above that $E(\log|du(0)|) \geq \varepsilon$ for some $\varepsilon > 0$ depending on the geometry of (X, λ, J).

We now establish a differential inequality for E which, together with the lower bound $E(\log|du(0)|) \geq \varepsilon$ and finiteness of $E(0)$, gives an upper bound on $\log|du(0)|$. We have

$$-E'(s) = \int_{s\times S^1} |du|^2 \geq \frac{1}{2\pi}\left(\int_{s\times S^1} |du|\right)^2 \geq \frac{1}{2\pi\|\lambda\|}\left(\int_{s\times S^1} u^*\lambda\right)^2 = \frac{E(s)^2}{2\pi\|\lambda\|}. \tag{5.2.11.2}$$

This differential inequality blows up in finite time (5.2.10), which gives the desired upper bound on $\log|du(0)|$. □

5.2.12 Exercise. Conclude from (5.2.11) that for a pseudo-holomorphic map $u : I \times S^1 \to (X, J, \lambda)$, we have $|du(s,t)| \leq \mathrm{const} \cdot d(s, \partial I)^{-1}$.

5.2.13 Exercise. Fix (X, J, g). Show that there exist $\varepsilon : X \to \mathbb{R}_{>0}$ and $M_k : X \to \mathbb{R}_{>0}$ depending only on bounds on the geometry and injectivity radius such that for any pseudo-holomorphic map $u : D^2 \to X$ with $u(D^2) \subseteq B_{u(0)}(\varepsilon(u(0)))$ satisfies $\|D^k u(0)\| \leq M_k(u(0))$.

5.2.14 Exercise. Let $\Omega \subseteq \mathbb{R}^n$ be open and let $f : \Omega \to \mathbb{R}$ satisfy $\|f\|_{C^{k+1}} \leq M$. Show that for every $\varepsilon > 0$ there exists $\delta = \delta(\varepsilon, M) > 0$ such that if $\|f\|_{C^0} \leq \delta$ then $\|f\|_{C^k} \leq \varepsilon$. (Prove the case $k = 1$ directly and then use induction.) Conclude using (5.2.13) that if pseudo-holomorphic maps are C^0-close then they are C^∞-close. We will see later (5.7.5) that the C^k distance between pseudo-holomorphic curves is bounded linearly by their C^0-distance.

5.2.15. In the action inequality (5.2.9.1) for maps $\gamma : S^1 \to \mathbb{C}^n$, the left hand side $a(\gamma)$ is invariant under reparameterization, whereas the right hand side is not. By reparameterizing γ to have constant speed, we conclude that

$$a(\gamma) \leq (4\pi)^{-1} L(\gamma)^2 \tag{5.2.15.1}$$

where $L(\gamma) = \int_{S^1} |\gamma'(t)|\, dt$ is the length. This inequality continues to hold for $\gamma : X \to \mathbb{C}^n$ for compact 1-manifolds X since $x^2 + y^2 \leq (x+y)^2$ for $x, y \geq 0$.

5.2.16 Proposition (Monotonicity)**.** *For $u : C \to (X, J, g)$ pseudo-holomorphic and proper over $B_r(u(p))$, we have*

$$E(u \cap B_r(u(p))) \geq (1 - \mathrm{const}\, r)\pi r^2 \tag{5.2.16.1}$$

for const $< \infty$ *depending on a bound on the geometry and injectivity radius of (X, J, g) at $u(p)$, and assuming that J is an isometry on the tangent space at $u(p)$.*

Proof. Let $B_r = B_r(u(p))$, and consider the area $A(r) = E(u \cap B_r)$ as a function of r. Fix local coordinates $\Phi : (\mathbb{C}^n, 0) \to (X, u(p))$ which identify almost complex structures and metrics at the basepoint. The set of r for which $u \pitchfork \partial B_r$ is open and has full measure by Sard's theorem [20]. For such r, the derivative $A'(r)$ is the length of $u \cap \partial B_r$ which is bounded below by $(1 - \mathrm{const}\, r)\sqrt{4\pi a(u \cap \partial B_r)}$ by (5.2.15.1). The action of the boundary is, by pseudo-holomorphicity the same as the area $A(r)$, again up to a relative error linear in r. We therefore have $A'(r) \geq (1 - \mathrm{const}\, r)\sqrt{4\pi A(r)}$, which is equivalent to $\frac{d}{dr}\sqrt{A(r)} \geq (1 - \mathrm{const}\, r)\sqrt{\pi}$. Since this is valid for an open set of r with full measure, we can integrate to obtain the desired result. □

5.2.17 Exercise. Show that removing the assumption that J is an isometry on the tangent space at $u(p)$ simply replaces π with a constant $\varepsilon > 0$ depending on the geometry at p.

5.2.18 Exercise. For a pseudo-holomorphic map $u : S \to (X, J, g)$ from a compact Riemann surface with boundary S, show that $\sup_{p \in S} d(u(p), u(\partial S))$ is bounded above by const $\cdot \max(E, E^{1/2})$ where $E = E(u)$ is the energy and const $< \infty$ depends on a bound on the geometry of (X, J, g) over the image of u. Show by example that this cannot be improved to $\leq \mathrm{const} E$ or $\leq \mathrm{const} E^{1/2}$.

5.2.19 Proposition (Removable singularity). *A pseudo-holomorphic map $u : D^2 \backslash 0 \to (X, J)$ extends smoothly to D^2 iff its image is relatively compact in X and it has finite area.*

Proof. Consider instead the coordinates $D^2 \setminus 0 = [0, \infty) \times S^1$. Let $E(s)$ denote the energy of u over $[s, \infty) \times S^1$; evidently $E(s) \to 0$ as $s \to \infty$. Since sufficiently small energy bounds imply gradient bounds (5.2.7), we have

$$|du(s,t)| \leq \text{const}E(s-1)^{1/2} \tag{5.2.19.1}$$

for sufficiently large s. In particular, $|du(s,t)| \to 0$ as $s \to \infty$. Monotonicity (5.2.18) implies moreover that the distance between $u(s \times S^1)$ and $u(s' \times S^1)$ for $s < s'$ is also bounded above by $\text{const}E(s-1)^{1/2}$ for s sufficiently large. It follows that u extends continuously to D^2.

Our next step is to show that $|du|$ is bounded on $D^2 \setminus 0$. In view of the energy controlling the gradient (5.2.19.1), it suffices to establish an exponential upper bound on energy $E(s) \leq \text{const}e^{-s}$. Fix local coordinates $(\mathbb{C}^n, 0) \to (X, u(0))$ (defined near zero) in which the almost complex structure J agrees with the standard one on $\mathbb{C}^n$ at the origin. We may measure energy with respect to the standard flat metric on $\mathbb{C}^n$. We now have $-E'(s) = \int_{s \times S^1} |du|^2$ which is bounded below by $2a(u(s, \cdot))$ by (5.2.9.1). The action $a(u(s, \cdot))$ is commensurable with $E(s)$ up to a factor of $(1 - \text{const}d)$ where d is the distance from the origin (coming from deviation of J from standard complex structure on $\mathbb{C}^n$). To fully justify this assertion, we must argue that $\int_{S^1} u(s, \cdot)^* \lambda = \int_{[s,\infty) \times S^1} u^* d\lambda$; certainly the difference between the right and left sides is independent of s by Stokes, and both sides approach zero as $s \to \infty$. The distance d is bounded by $\text{const} \cdot E(s-1)^{1/2}$ since small energy controls gradient (5.2.7). We conclude that

$$-E'(s) \geq (1 - \text{const}E(s-1)^{1/2}) 2E(s) \tag{5.2.19.2}$$

for s sufficiently large. Write this as $-\frac{d}{ds} \log E(s) \geq (1 - \text{const}E(s-1)^{1/2})2$. Integrating this once shows that E is bounded above by $\text{const}_\mu e^{-\mu s}$ for any $\mu < 2$. Integrating again, noting that this exponential upper bound for say $\mu = 1$ implies that the integral of $\text{const}E(s-1)^{1/2}$ is finite, we get the desired inequality $E(s) \leq \text{const}e^{-s}$.

Now we would like to use elliptic bootstrapping (5.2.1) to conclude that u is smooth on D^2. That result assumed smooth u, whereas here u is merely continuous on D^2 and smooth on $D^2 \backslash 0$, with $|du|$ bounded. So, let us check that the proof goes through in this setting. We first check that the key equation $u_{xx} + u_{yy} = \dot{J}(u, u_y)u_x - \dot{J}(u, u_x)u_y$ remains true, where the left hand side is taken in the sense of distributions and the right hand side is meant in the sense of multiplication of L^∞ functions. Begin with $\partial_x(u_x + J(u)u_y) - \partial_y(J(u)u_x - u_y) = 0$, where the outer derivatives are in the sense of distributions. This yields $u_{xx} + u_{yy} = \partial_y(J(u)u_x) - \partial_x(J(u)u_y)$. To show that the right hand side coincides with $\dot{J}(u, u_y)u_x - \dot{J}(u, u_x)u_y$, it is enough to integrate against compactly supported smooth test functions φ. Integrating by parts it is enough to show that $\int (\varphi J(u))_x u_y = \int (\varphi J(u))_y u_x$. This holds since $J(u), u \in W^{1,2}$ and both sides are $W^{1,2}$-bounded, so it is enough to check for smooth u, where it is just integration by parts. Now we still need to do the actual bootstrapping. For the second step, we need to know that $(L^\infty \cap W^{1,2}) \cdot (L^\infty \cap W^{1,2}) \subseteq W^{1,2}$, which can be shown as follows. Let $f, g \in L^\infty \cap W^{1,2}$. We just need to show that the distribution derivative $(fg)_x$ coincides with

$f_x g + f g_x$ (multiplication of L^∞ and L^2 functions). In other words, we need to show that for any compactly supported smooth test function φ, that $\int \varphi_x f g + \varphi f_x g + \varphi f g_x = 0$. This is, by inspection, a $W^{1,2}$-bounded functional of f, so since C^∞ is dense in $W^{1,2}$ we may assume that f is smooth. By symmetry, we may also assume that g is smooth, so then it just follows from integration by parts. □

5.2.20 Definition (Hofer energy [8])**.** Consider $\mathbb{R} \times Y$ with cylindrical ($\mathbb{R}$-invariant) almost complex structure J satisfying $J(\partial_s) = R$ for some vector field R on Y. Such an almost complex structure induces a Levi distribution $\xi = TY \cap JTY$. Let $u : C \to \mathbb{R} \times Y$ be pseudo-holomorphic. The *ξ-energy* of u is the integral of $u^*\omega$, where ω is any 2-form on Y which tames J on ξ and vanishes on R (the ξ-energy is independent, up to constant factor, of the choice of ω). The *R-energy* of u is the supremum over $\varphi : \mathbb{R} \to [0, 1]$ satisfying $\varphi'(s) \geq 0$ of the integral of $d\varphi \wedge \lambda$ for a 1-form λ which is positive on R and zero on ξ (again this is independent of λ up to constant factor). Note that both integrands ω and $d\varphi \wedge \lambda$ are non-negative on complex subspaces of $T(\mathbb{R} \times Y)$.

5.2.21 Lemma (Bounded R-energy and small ξ-energy implies bounded gradient)**.** *For $u : D^2 \to \mathbb{R} \times Y$ pseudo-holomorphic with $E_R(u) \leq N$ and $E_\xi(u) \leq \varepsilon_{N,Y}$, we have $|du(0)| \leq \mathrm{const}_{N,Y}$, where the subscript indicates dependence on N and on bounds on the geometry of Y over the image of u.*

Proof. Given large $|du(0)|$, we apply Hofer's Lemma (5.2.5) with $M = |du(0)|^{1/2}$ to find a rescaled map $\tilde{u}$ defined on a large disk (radius $|du(0)|^{1/2}$) satisfying $|d\tilde{u}| \leq 2$ and $|d\tilde{u}(0)| = 1$. It therefore suffices to bound above the radius R for which there exists a map $\tilde{u} : D^2_R \to \mathbb{R} \times Y$ with $E_R(\tilde{u}) \leq N$, $E_\xi(\tilde{u}) \leq \varepsilon_{N,Y}$, $|d\tilde{u}| \leq 2$, and $|d\tilde{u}(0)| = 1$. The gradient bounds on $\tilde{u}$ implies C^∞ bounds (5.2.1), so small ξ-energy energy implies C^∞-close to the leaf of the foliation of $\mathbb{R} \times Y$ by $\partial_s \oplus R$ passing through $\tilde{u}(0)$. By projection, we thus obtain a map $D^2_R \to \mathbb{C}$ (where the universal cover of the leaf is identified with $\mathbb{C}$ in the obvious way). This map has C^∞-small $\bar{\partial}$, so is C^∞-close to a holomorphic map (say over D^2_{R-1}); note that this is a linear equation, so the linear elliptic theory applies!. This holomorphic map $w : D^2_R \to \mathbb{C}$ has R-energy bounded by $M + 1$, so it must miss at least one point in the interval $w(0) + i[1, M + 2]$ and one point in the interval $w(0) - i[1, M + 2]$. Thus w lifts to the universal cover of $\mathbb{C}$ based at $w(0)$ with these two points removed, which is D^2. Now the resulting lift $\tilde{w} : D^2_R \to D^2$ has gradient at zero at most R^{-1}, which combined with $|dw(0)| \approx 1$ gives the desired upper bound on R. □

5.3 Families of pseudo-holomorphic maps

5.3.1 Definition (Family of pseudo-holomorphic sections)**.** Let B be a topological space, let $C \to B$ be a family of Riemann surfaces (4.2.1). Let $W \to C$ be an almost complex submersion (meaning $W \to B$ is a relative smooth manifold with a smooth almost complex structure and $W \to C$ is a submersion over B). A family of pseudo-holomorphic sections of $W \to C$ over B is a continuous section $u : C \to W$ whose restriction to every fiber is pseudo-holomorphic.

5.3.2 Exercise. Use (5.2.14) to show that every family of pseudo-holomorphic sections u is smooth over B.

5.4 Moduli stacks of pseudo-holomorphic maps

5.4.1 Definition (Moduli stack of pseudo-holomorphic sections)**.** Let $W \to C$ be an almost complex submersion over a relative Riemann surface $C \to B$ over a topological space B. The moduli stack $\underline{\mathrm{Hol}}(W/C/B)$ assigns to a topological space Z the set of maps $Z \to B$ together with a family of pseudo-holomorphic sections of the pullback $W_Z \to C_Z \to Z$. More generally, we can allow B to be a stack on topological spaces, in which case the moduli problem $W \to C \to B$ is, by definition, the assignment to every $Z \to B$ of a moduli problem $W_Z \to C_Z \to Z$, compatible with pullback.

5.4.2 Example. Given a Riemann surface C and an almost complex manifold X, the moduli stack of pseudo-holomorphic maps $C \to X$ is denoted $\underline{\mathrm{Hol}}(C, X)$ and is defined by the property that a map $Z \to \underline{\mathrm{Hol}}(C, X)$ from a topological space Z is a continuous map $Z \times C \to X$ whose restriction to $z \times C$ is pseudo-holomorphic for every $z \in Z$. Evidently $\underline{\mathrm{Hol}}(C, X) = \underline{\mathrm{Hol}}((C \times X)/C/*)$.

5.5 Stability

The moduli space $\underline{\mathrm{Hol}}(W/C/B)$ is usually not what we are ultimately interested in, rather we are interested in a nice open subset $\underline{\mathrm{Hol}}(W/C/B)^s \subseteq \underline{\mathrm{Hol}}(W/C/B)$ given by the stable locus (1.8.3). To understand this stable locus, the first step is to construct a stable diagonal (1.8.1) for $\underline{\mathrm{Hol}}(W/C/B)$ to which we can apply (1.8.5) to conclude that $\underline{\mathrm{Hol}}(W/C/B)^s \subseteq \underline{\mathrm{Hol}}(W/C/B)$ is an open substack with proper diagonal.

5.5.1 Construction (Stable diagonal for $\underline{\mathrm{Hol}}$)**.** Let $W \to C \to B$ be a holomorphic curve problem. Fix the following data:

(5.5.1.1) A stable diagonal $B \to \widehat{B} \to B \times B$.

(5.5.1.2) A family of Riemann surfaces $\widehat{C} \to \widehat{B}$ fitting into a diagram

$$\begin{array}{ccccc} C & \longrightarrow & \widehat{C} & \longrightarrow & C \times C \\ \downarrow & & \downarrow & & \downarrow \\ B & \longrightarrow & \widehat{B} & \longrightarrow & B \times B \end{array}$$

in which the left square is a pullback $C = \widehat{C} \times_{\widehat{B}} B$ and in which the two resulting maps $\widehat{C} \to C \times_B \widehat{B}$ are maps of families of Riemann surfaces.

(5.5.1.3) An isomorphism between the two resulting pullbacks $W \times_C \widehat{C} \to \widehat{C}$, which we denote by $\widehat{W}$. This isomorphism should be the identity over B.

Such data determines a pair of maps $\underline{\mathrm{Hol}}(W/C/B)\times_B\widehat{B}\to\underline{\mathrm{Hol}}(\widehat{W}/\widehat{C}/\widehat{B})$. We define

$$\widehat{\underline{\mathrm{Hol}}}(W/C/B)=\big(\underline{\mathrm{Hol}}(W/C/B)\times_B\widehat{B}\big)\times_{\underline{\mathrm{Hol}}(\widehat{W}/\widehat{C}/\widehat{B})}\big(\widehat{B}\times_B\underline{\mathrm{Hol}}(W/C/B)\big). \tag{5.5.1.4}$$

5.5.2 Lemma. *The factorization*

$$\underline{\mathrm{Hol}}(W/C/B)\to\widehat{\underline{\mathrm{Hol}}}(W/C/B)\to\underline{\mathrm{Hol}}(W/C/B)\times\underline{\mathrm{Hol}}(W/C/B) \tag{5.5.2.1}$$

is a stable diagonal of $\underline{\mathrm{Hol}}(W/C/B)$.

Proof. Pulling back $\underline{\mathrm{Hol}}(\widehat{W}/\widehat{C}/\widehat{B})$ under $B\to\widehat{B}$ yields $\underline{\mathrm{Hol}}(W/C/B)$, which shows that $\underline{\mathrm{Hol}}(W/C/B)\to\widehat{\underline{\mathrm{Hol}}}(W/C/B)$ is the inclusion of the open substack $\widehat{\underline{\mathrm{Hol}}}(W/C/B)\times_{\widehat{B}}B$.

The map

$$\widehat{\underline{\mathrm{Hol}}}(W/C/B)\to\underline{\mathrm{Hol}}(W/C/B)\times_B\widehat{B}\times_B\underline{\mathrm{Hol}}(W/C/B) \tag{5.5.2.2}$$

is a pullback of the diagonal of $\underline{\mathrm{Hol}}(\widehat{W}/\widehat{C}/\widehat{B})\to\widehat{B}$, hence is a closed embedding since $\widehat{C}\to\widehat{B}$ is open (1.7.23). The map from the right hand side above to $\underline{\mathrm{Hol}}(W/C/B)\times\underline{\mathrm{Hol}}(W/C/B)$ is a pullback of $\widehat{B}\to B\times B$, hence proper. □

5.6 Linear deformation theory

The linear deformation theory of pseudo-holomorphic curves is the study of deformation complexes. These are complexes of vector bundles on moduli spaces of pseudo-holomorphic curves which play the role of the tangent bundle. The zeroth cohomology T^0 is the space of first order deformations of a pseudo-holomorphic map, so can be thought of as the tangent space to the moduli space. The first cohomology T^1 is sometimes called the 'obstruction space', and it measures the failure of the moduli space to be smooth/transverse (when T^1 vanishes, the moduli space is called 'unobstructed' or 'regular'). The vector space T^{-1} is the tangent space to the group of automorphisms of points of the moduli space (it is present only when working over a stack, and in that case it vanishes over precisely the stable locus).

5.6.1 Definition (Relative deformation complex $T_{\mathrm{Hol}(W/C/B)/B}$)**.** Consider an almost complex fibration $W\to C$ over a Riemann surface with a pseudo-holomorphic section $u:C\to W$. First order deformations of u are given by sections of $u^*T_{W/C}$. The first order change in $(du)^{0,1}$ is a real Cauchy–Riemann operator

$$C^\infty(C,u^*T_{W/C})\to C^\infty(C,u^*T_{W/C}\otimes_{\mathbb{C}}\Omega^{0,1}_C). \tag{5.6.1.1}$$

Given a pseudo-holomorphic curve problem $W\to C\to B$ over a topological space B and a section $u:C\to W$, the deformation operators vary continuously (a relatively smooth family of elliptic operators). In particular, this is a family of real Cauchy–Riemann operators over $\underline{\mathrm{Hol}}(W/C/B)$. The resulting two-term complex of vector bundles on $\underline{\mathrm{Hol}}(W/C/B)$ is denoted $T_{\mathrm{Hol}(W/C/B)/B}$, with its cohomology denoted $T^i_{\mathrm{Hol}(W/C/B)/B}$.

5.6.2 Exercise. Show that the deformation complex of a constant map to $p \in X$ in $\underline{\mathrm{Hol}}(C, X)$ is the trivial holomorphic vector bundle with fiber T_pX over C.

5.6.3 Exercise. Show that the deformation complex of the identity map $\mathbf{1}_C \in \underline{\mathrm{Hol}}(C, C)$ is the holomorphic line bundle TC over C.

5.6.4 Definition (Deformation complex $T_{\mathrm{Hol}(W/C/B)}$)**.** Consider a smooth pseudo-holomorphic curve problem $W \to C \to B$, namely B is a smooth manifold, $C \to B$ is a smooth family of Riemann surfaces, and $W \to C$ is a smooth almost complex family. The space of first order deformations of a point $(b, u : C_b \to W_b) \in \mathrm{Hol}(W/C/B)$ is the space of sections of $u^*TW/du(T_{C/B})$ over C_b lying over a fixed element of T_bB. The first order change in $(du)^{0,1}$ induced by such a deformation of (b, u) is a real Cauchy–Riemann operator

$$C^\infty(C_b, u^*TW/du(T_{C/B}))_{\equiv T_bB} \to C^\infty(C_b, u^*T_{W/C} \otimes_{\mathbb{C}} \Omega^{0,1}_{C_b}) \tag{5.6.4.1}$$

(the subscript on the domain indicates sections which project to a specified vector in T_bB).

5.7 Non-linear deformation theory

5.7.1 Definition (Deformation setup)**.** Let $u : C \to X$ be a smooth map from a Riemann surface to an almost complex manifold. Fix a smooth map $\exp : TX \to X$ which to first order along the zero section is $\mathbf{1}_{TX}$ (defined in a neighborhood of the zero section). Also fix a $\mathbb{C}$-linear isomorphism $\tau : \exp^* TX \to \pi^*TX$ which is the identity over the zero section (defined in a neighborhood of the zero section); $\pi : TX \to X$ being the tautological projection. We can then define

$$\hat{D}_{u,\exp,\tau} : C^\infty(C, u^*TX) \to C^\infty(C, u^*TX \otimes_{\mathbb{C}} \Omega^{0,1}_C) \tag{5.7.1.1}$$
$$\xi \mapsto \tau((d \exp \xi)^{0,1}) \tag{5.7.1.2}$$

The derivative of $\hat{D}_{u,\exp,\tau}$ at zero is a real Cauchy–Riemann operator denoted

$$D_{u,\exp,\tau} : C^\infty(C, u^*TX) \to C^\infty(C, u^*TX \otimes_{\mathbb{C}} \Omega^{0,1}_C) \tag{5.7.1.3}$$

This operator is independent of exp and τ when u is pseudo-holomorphic, but otherwise not.

This setup generalizes in a straightforward way to an almost complex fibration $W \to C$. Given a smooth section $u : C \to W$, we may define

$$\hat{D}_{u,\exp,\tau} : C^\infty(C, u^*T_{W/C}) \to C^\infty(C, u^*T_{W/C} \otimes_{\mathbb{C}} \Omega^{0,1}_C) \tag{5.7.1.4}$$
$$\xi \mapsto \tau((d \exp \xi)^{0,1}) \tag{5.7.1.5}$$

for $\exp : T_{W/C} \to W$ a smooth map which to first order along the zero section is $\mathbf{1}_{T_{W/C}}$ and $\tau : \exp^* T_{W/C} \to \pi^*T_{W/C}$ an isomorphism equalling the identity over the zero section. Its derivative is then a real Cauchy–Riemann operator $D_{u,\exp,\tau}$ with the same domain and target.

5.7.2 Lemma (Quadratic estimate)**.** *Fix a deformation setup* (5.7.1)*. If C is compact, then for integer $s > 1$ we have an estimate*

$$\|\hat{D}f - \hat{D}g - D(f-g)\|_{H^s} \leq \text{const}\|f-g\|_{H^{s+1}}(\|f\|_{H^{s+1}} + \|g\|_{H^{s+1}}) \tag{5.7.2.1}$$

where const *depends on bounds on the geometry of C, the geometry of X over the image of C, and the derivatives of u.*

Proof. The expression $\hat{D}f - \hat{D}g - D(f-g)$ is a non-linear first order differential operator. That is, it takes the form $A(J^1f, J^1g)$ for some smooth function A (depending on exp and τ and their derivatives). Now apply (3.1.20) to bound $A(J^1f, J^1g)$ in terms of H^s-norms of J^1f and J^1g and their difference, and reexpress in terms of H^s-norms of f and g and their difference. □

5.7.3 Exercise. Generalize (5.7.2) to the setting that C is non-compact but of bounded geometry and injectivity radius. For this, the H^s-norm should be defined via a partition of unity of bounded geometry, namely $\|f\|_{H^s}^2 = \sum_i \|\varphi_i f \circ \Phi_i\|_{H^s}^2$ for Φ_i a collection of coordinate charts of bounded geometry, and φ_i a partition of unity, also of bounded geometry.

5.7.4 Exercise. Generalize further to the case of Sobolev spaces weighted by e^w for some function $w : C \to \mathbb{R}$ with a bound on dw and all its derivatives.

5.7.5 Exercise. Suppose $u, u' : C \to X$ are two pseudo-holomorphic curves and $u' = \exp_u \xi$. Recall from (5.2.14) that if ξ is C^0-small then it is C^∞-small. Prove a linear bound $\|\xi\|_{C^k} \leq \text{const}\|\xi\|_{C^0}$ by establishing the following chain of inequalities (implied constants omitted):

$$\|\xi\|_{C^k} \leq \|\xi\|_{H^s} \leq \|(\mathbf{1} - Q_u D_u)\xi\|_s \leq \|\xi\|_{L^2} \leq \|\xi\|_{C^0}. \tag{5.7.5.1}$$

by using the quadratic estimate (5.7.2) and an elliptic estimate from (3.3), where Q_u denotes a parametrix of D_u.

5.8 Inverse exponential maps

The inverse exponential map takes two nearby pseudo-holomorphic maps (or sections) u and u', and produces an element $\exp_u^{-1}(u') \in \ker D_u$ of the kernel of the deformation operator of u. The idea of how to define $\log_u(u')$ is simple: we consider the section ξ defined by the property $\exp_u \xi = u'$, we note that pseudo-holomorphicity of u' implies an estimate of the form $\|D_u\xi\|_s \leq \|\xi\|_s^2$, and we define $\exp_u^{-1} u'$ to be $(\mathbf{1} - Q_u D_u)\xi$ for a suitable choice of right inverse Q_u for D_u. The goal of this section is to make this definition precise and to formulate a sense in which the inverse exponential map is unique. In fact, this uniqueness is crucial for showing existence, as it means that it suffices to prove existence locally. Uniqueness is also crucial for comparing the inverse exponential maps of different but related moduli spaces. The eventual use of the inverse exponential map is to express the fundamental cycles of transverse moduli spaces as cycles twisted by the analytic tangent bundle T^0.

5.8.1 Definition (Equality up to quadratic error)**.** Let $\{V_\alpha\}_{\alpha\in A}$ be a family of real vector spaces indexed by a set A. Also fix a set of families of norms $\{\|\cdot\|_\alpha\}_\alpha$ on $\{V_\alpha\}_\alpha$ called 'admissible', such that if $\{\|\cdot\|_\alpha\}_\alpha$ is admissible then so is $\{r\|\cdot\|_\alpha\}_\alpha$ for every positive real number r. We consider families of elements $\{x_\alpha \in V_\alpha\}_\alpha$. We say two such families $\{x_\alpha\}$ and $\{y_\alpha\}$ are *equal up to quadratic error*, written $\{x_\alpha\} \sim \{y_\alpha\}$, iff there exists an admissible family of norms $\{\|\cdot\|_\alpha\}$ such that $\|x_\alpha - y_\alpha\|_\alpha \leq \min(\|x_\alpha\|_\alpha^2, \|y_\alpha\|_\alpha^2)$ for every $\alpha \in A$.

5.8.2 Exercise. Using the fact that admissibility is preserved by overall scaling, show that $\sim$ is an equivalence relation.

5.8.3 Exercise. Show that if $\{x_\alpha\}_\alpha \sim \{y_\alpha\}_\alpha$, then the families $\{tx_\alpha+(1-t)y_\alpha\}_\alpha$ are equivalent for all $t \in \mathbb{R}$. Conclude that equivalence classes are convex.

5.8.4 Exercise. Show that if $\{x_\alpha\}_\alpha \sim \{y_\alpha\}_\alpha$, then $x_\alpha = 0 \iff y_\alpha = 0$. Show the converse when all families of norms are admissible.

5.8.5 Exercise. Take $V_\alpha = \mathbb{R}$, and take the admissible families of norms to be those which are determined by a norm on $\mathbb{R}$. Supposing our index set is $\mathbb{Z}_{\geq 1}$, partition into equivalence classes the following functions: 1, n, $(-1)^n$, $\frac{1}{n}$, $\frac{1}{n} + \frac{3}{n^2}$, $\frac{1}{2^n}$.

5.8.6 Definition (Inverse exponential map of a smooth manifold)**.** Let M be a smooth manifold, and let $p \in M$. We consider germs of maps

$$(M, p) \to (T_pM, 0) \tag{5.8.6.1}$$

up to quadratic error. That is, we consider maps $f : U \to T_pM$ where $U \subseteq M$ is a neighborhood of p satisfying $f(p) = 0$, and we regard two such (U, f) and (V, g) as equivalent iff there exists a neighborhood $W \subseteq U \cap V$ of p and a norm $\|\cdot\|$ on T_pM such that $\|f(x)-g(x)\| \leq \min(\|f(x)\|^2, \|g(x)\|^2)$ for all $x \in W$. After the restriction to W, this is a special case of (5.8.1), so it is indeed an equivalence relation, and equivalence classes are convex. There is a canonical map

$$\log_0 : (\mathbb{R}^n, 0) \to (T_0\mathbb{R}^n, 0), \tag{5.8.6.2}$$

namely the tautological identification $\mathbb{R}^n = T_0\mathbb{R}^n$. Choosing a local coordinate chart $\Phi : (\mathbb{R}^n, 0) \to (M, p)$, we can transport $\log_0$ to a map $\Phi_* \log_0 : (M, p) \to (T_pM, 0)$. If Φ' is any other such coordinate chart, the maps $\Phi_* \log_0$ and $\Phi'_* \log_0$ are equivalent up to quadratic error. We denote by

$$\log_p : (M, p) \to (T_pM, 0) \tag{5.8.6.3}$$

this distinguished equivalence class of germs of maps up to quadratic error. Beware that $\log_p$ is not a specific map (5.8.6.3) as the notation might suggest, rather it is an equivalence class of such maps.

We now consider a simultaneous specification of inverse exponential maps at all basepoints of M, varying continuously. That is, we consider germs of continuous maps

$$(M \times M, M) \to (TM, M) \tag{5.8.6.4}$$

over M up to quadratic error. In other words, f is defined in a neighborhood of the diagonal, $f(p,q) \in T_pM$, and $f(p,p) = 0$. Two such maps f are regarded as equivalent iff there exists a neighborhood of the diagonal over which they agree up to quadratic error, measured with respect to any choice of continuously varying norm on TM. As before, this is an equivalence relation and equivalence classes are convex. Even better, if $f \sim g$ and $t : M \to \mathbb{R}$ is any continuous function, then f and g are both equivalent to $tf + (1-t)g$. For $M = \mathbb{R}^n$, we can define $\log_p(q) = q - p$. For any coordinate chart $\Phi : \mathbb{R}^n \to M$, we may push forward to obtain a map $\Phi_* \log$ over the image of Φ, and these maps are equivalent on overlaps. In view of the sheaf property (5.8.7), this determines a unique equivalence class of map $\log : (M \times M, M) \to (TM, M)$ up to quadratic error.

5.8.7 Exercise. Consider germs of continuous maps $(M \times M, M) \to (TM, M)$ over M up to quadratic error. Using paracompactness of M, show that if two maps f and g are locally equivalent (i.e. there exists an open cover $M = \bigcup_i U_i$ such that $f|_{U_i} \sim g|_{U_i}$) then f and g are equivalent. Using a partition of unity argument, show that if $M = \bigcup_i U_i$ is an open cover and f_i are germs on U_i which are equivalent on overlaps ($f_i|_{U_i \cap U_j} \sim f_j|_{U_i \cap U_j}$), then there exists a germ f on M whose restriction to every U_i is equivalent to f_i. Conclude that germs of continuous maps $(M \times M, M) \to (TM, M)$ up to quadratic error form a sheaf on M (note that if a manifold is paracompact, it is metrizable, hence every open subset is paracompact).

5.8.8 Definition (Inverse exponential map of $\underline{\mathrm{Hol}}(C,X)$)**.** We consider germs of continuous maps

$$\underline{\mathrm{Hol}}(C,X) \times \underline{\mathrm{Hol}}(C,X) \to T^0_{\underline{\mathrm{Hol}}(C,X)} \tag{5.8.8.1}$$

over $\underline{\mathrm{Hol}}(C,X)$ (the domain maps via the projection to the first factor) defined in a neighborhood of the diagonal, sending the diagonal to zero. More generally, for any open substack $U \subseteq \underline{\mathrm{Hol}}(C,X)$, we can consider maps $U \times U \supseteq N \to T^0_U$ over U. We declare two such maps f and g over U to be equivalent iff there exists an open cover $U = \bigcup_i U_i$ such that for every i, there exists a neighborhood of the diagonal of U_i and a norm on $T^0|_{U_i}$ such that $\|f(x) - g(x)\| \le \min(\|f(x)\|^2, \|g(x)\|^2)$ for all $x \in U_i$. Note that equivalence is, by definition, a local property.

Now the inverse exponential map of $\underline{\mathrm{Hol}}(C,X)$ is the equivalence class of germ (5.8.8.1) (if it exists) defined by the property

$$\sup_{p \in C} \left|(\log_u u')(p) - \log_{u(p)}(u'(p))\right| \le \|\log_u u'\|^2 \tag{5.8.8.2}$$

for some (equivalently, any) inverse exponential map on X, locally on $\underline{\mathrm{Hol}}(C,X)$ for some choice of local norm on T^0.

5.8.9 Exercise. Show that $(\log_u u')(p) \sim \log_{u(p)}(u'(p))$ (equality up to quadratic error) implies (5.8.8.2). Explain, however, why this condition is too strong to be useful in practice (there would usually exist no map satisfying it).

5.9 Local structure of moduli spaces

5.9.1 Proposition. *The regular locus* $\underline{\mathrm{Hol}}(C,X)^{\mathrm{reg}} \subseteq \underline{\mathrm{Hol}}(C,X)$ *is a topological manifold.*

Proof. Let $p \in \underline{\mathrm{Hol}}(C,X)$ be given by $u : C \to X$. Maps nearby to u take the form $\exp_u \xi$ for $\xi : C \to u^*TX$ (fix some map $\exp : TX \to X$ whose derivative along the zero section is the identity map of TX). To measure pseudo-holomorphicity of such maps, introduce the non-linear map

$$\hat{D} : C^\infty(C, u^*TX) \to C^\infty(C, u^*TX \otimes_{\mathbb{C}} \Omega_C^{0,1}) \tag{5.9.1.1}$$

sending ξ to the image of $(d\exp_u \xi)^{0,1} \in C^\infty(C, (\exp_u \xi)^*TX \otimes_{\mathbb{C}} \Omega_C^{0,1})$ under a local trivialization of TX relating it to the target above ($\hat{D}$ is defined for sections of sufficiently small C^0-norm). Then $\exp_u \xi$ is pseudo-holomorphic iff $\hat{D}(\xi) = 0$. The moduli space $\underline{\mathrm{Hol}}(C,X)$ near p is thus identified with the zero set of $\hat{D}$ near 0.

When $p \in \underline{\mathrm{Hol}}(C,X)^{\mathrm{reg}}$, its deformation operator

$$D : C^\infty(C, u^*TX) \to C^\infty(C, u^*TX \otimes_{\mathbb{C}} \Omega_C^{0,1}) \tag{5.9.1.2}$$

is surjective. Elements of $\ker D$ are first order deformations of u which preserve pseudo-holomorphicity. Let Q be a right inverse to D. The projection $\mathbf{1} - QD$ from $C^\infty(C, u^*TX)$ to $\ker D_u$ defines a map $\underline{\mathrm{Hol}}(C,X) \to \ker D_u$ (identifying $\underline{\mathrm{Hol}}(C,X)$ with $\hat{D}^{-1}(0) \subseteq C^\infty(C, u^*TX)$ locally around p, as above). It suffices to show that this map is an isomorphism of stacks in a neighborhood of p.

We consider the endomorphism

$$R : C^\infty(C, u^*TX) \to C^\infty(C, u^*TX) \tag{5.9.1.3}$$

$$\xi \mapsto \xi - Q\hat{D}\xi \tag{5.9.1.4}$$

A fixed point of R is evidently a pseudo-holomorphic map (note Q is injective). We wish to show that its restriction to any slice $\xi + \operatorname{im} Q$ is a contraction mapping. To prove this, we should bound

$$\|R\xi - R\zeta\| = \|(QD\xi - Q\hat{D}\xi) - (QD\zeta - Q\hat{D}\zeta)\| \tag{5.9.1.5}$$

$$\leq \|Q\| \|D(\xi - \zeta) - (\hat{D}\xi - \hat{D}\zeta)\| \tag{5.9.1.6}$$

(note the use of $\xi - \zeta \in \operatorname{im} Q$ to write $\xi - \zeta = QD(\xi - \zeta)$.) This is now exactly the quantity to which the bound (5.7.2) applies, which tells us that over the locus where $\|\xi\|_s \leq \varepsilon$, the map R is indeed a contraction on these slices $\xi + \operatorname{im} Q$. It follows that on each such slice there is a unique pseudo-holomorphic map. This means that the map $\underline{\mathrm{Hol}}(C,X) \to \ker D_u$ is bijective on points.

The iteration $\xi \mapsto R(\xi)$ moreover converges uniformly in a family of ξ. Note that the convergence is in H^s-norm, not necessarily higher, but pseudo-holomorphicity tells us that the limiting family is indeed relatively smooth (with effective estimates on its smoothness). Taking the family $\xi \in \ker D$, we have thus constructed an inverse mapping to $\underline{\mathrm{Hol}}(C,X) \to \ker D_u$. □

Bibliography

[1] Robert A. Adams, *Sobolev spaces*, Academic Press [A subsidiary of Harcourt Brace Jovanovich, Publishers], New York-London, 1975, Pure and Applied Mathematics, Vol. 65. MR 0450957 60

[2] F. Bourgeois, Y. Eliashberg, H. Hofer, K. Wysocki, and E. Zehnder, *Compactness results in symplectic field theory*, Geom. Topol. **7** (2003), 799–888. MR 2026549 80

[3] Albrecht Dold, *Partitions of unity in the theory of fibrations*, Ann. of Math. (2) **78** (1963), 223–255. MR 155330 19

[4] Samuel Eilenberg and J. A. Zilber, *Semi-simplicial complexes and singular homology*, Ann. of Math. (2) **51** (1950), 499–513. MR 35434 21

[5] Kenji Fukaya and Kaoru Ono, *Arnold conjecture and Gromov-Witten invariant*, Topology **38** (1999), no. 5, 933–1048. MR 1688434 (2000j:53116) 80

[6] M. Gromov, *Pseudo holomorphic curves in symplectic manifolds*, Invent. Math. **82** (1985), no. 2, 307–347. MR 809718 iii, 80

[7] Andre Henriques and David Gepner, *Homotopy Theory of Orbispaces*, Preprint **arxiv:math/0701916** (2007), 1–52. 56

[8] H. Hofer, *Pseudoholomorphic curves in symplectizations with applications to the Weinstein conjecture in dimension three*, Invent. Math. **114** (1993), no. 3, 515–563. MR 1244912 85

[9] H. Hofer and C. Viterbo, *The Weinstein conjecture in the presence of holomorphic spheres*, Comm. Pure Appl. Math. **45** (1992), no. 5, 583–622. MR 1162367 81

[10] H. Hofer, K. Wysocki, and E. Zehnder, *Properties of pseudoholomorphic curves in symplectisations. I. Asymptotics*, Ann. Inst. H. Poincaré Anal. Non Linéaire **13** (1996), no. 3, 337–379. MR 1395676 80

[11] ———, *Finite energy cylinders of small area*, Ergodic Theory Dynam. Systems **22** (2002), no. 5, 1451–1486. MR 1934144 80

[12] Lars Hörmander, *The analysis of linear partial differential operators. I*, Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 256, Springer-Verlag, Berlin, 1983, Distribution theory and Fourier analysis. MR 717035 60

[13] Christoph Hummel, *Gromov's compactness theorem for pseudo-holomorphic curves*, Progress in Mathematics, vol. 151, Birkhäuser Verlag, Basel, 1997. MR 1451624 80

[14] A. Joyal, *Quasi-categories and Kan complexes*, vol. 175, 2002, Special volume celebrating the 70th birthday of Professor Max Kelly, pp. 207–222. MR 1935979 26

[15] H. Blaine Lawson, Jr. and Marie-Louise Michelsohn, *Spin geometry*, Princeton Mathematical Series, vol. 38, Princeton University Press, Princeton, NJ, 1989. MR 1031992 66

[16] Robert B. Lockhart and Robert C. McOwen, *Elliptic differential operators on noncompact manifolds*, Ann. Scuola Norm. Sup. Pisa Cl. Sci. (4) **12** (1985), no. 3, 409–447. MR 837256 73

[17] Jacob Lurie, *Higher topos theory*, Annals of Mathematics Studies, vol. 170, Princeton University Press, Princeton, NJ, 2009. MR 2522659 26, 30, 33, 35, 46, 47, 49

[18] ______, *Higher algebra*, Book manuscript, 2017. 54

[19] Dusa McDuff and Dietmar Salamon, *J-holomorphic curves and symplectic topology*, American Mathematical Society Colloquium Publications, vol. 52, American Mathematical Society, Providence, RI, 2004. MR 2045629 (2004m:53154) 80

[20] John W. Milnor, *Topology from the differentiable viewpoint*, Based on notes by David W. Weaver, The University Press of Virginia, Charlottesville, Va., 1965. MR 0226651 83

[21] Marie-Paule Muller, *Gromov's Schwarz lemma as an estimate of the gradient for holomorphic curves*, Holomorphic curves in symplectic geometry, Progr. Math., vol. 117, Birkhäuser, Basel, 1994, pp. 217–231. MR 1274931 80

[22] Daniel G. Quillen, *Homotopical algebra*, Lecture Notes in Mathematics, No. 43, Springer-Verlag, Berlin-New York, 1967. MR 0223432 24, 26

[23] Michael Reed and Barry Simon, *Methods of modern mathematical physics. I*, second ed., Academic Press, Inc. [Harcourt Brace Jovanovich, Publishers], New York, 1980, Functional analysis. MR 751959 60

[24] Stefan Schwede, *Global homotopy theory*, New Mathematical Monographs, vol. 34, Cambridge University Press, Cambridge, 2018. MR 3838307 56

[25] Jean-Claude Sikorav, *Some properties of holomorphic curves in almost complex manifolds*, Holomorphic curves in symplectic geometry, Progr. Math., vol. 117, Birkhäuser, Basel, 1994, pp. 165–189. MR 1274929 80

[26] Jean-Louis Verdier, *Dualité dans la cohomologie des espaces localement compacts*, Séminaire Bourbaki, Vol. 9, Soc. Math. France, Paris, 1995, pp. Exp. No. 300, 337–349. MR 1610971 54