Elements of Categorical Metaphysics

The Compositional Organization of Space, Logic, and Structure

June 4, 2021

Contents

Part I

Categories

Chapter 1

Category Theory

This chapter, which introduces category theory and covers the study of spaces from many categorically oriented points of view, is a blend of many sources. Our sources for category theory include [Mac Lane, 2013, Riehl, 2017, Aluffi, 2009]. The section on homotopy theory borrows from [May, 1999, Hatcher, 2005, Munkres, 2018], in roughly that order.

1.1 Category Theory

1.1.1 Categories

A ***category*** C is a class $\mathrm{Ob}(\mathsf{C})$ of ***objects*** and, for every two objects $X, Y \in \mathrm{Ob}(\mathsf{C})$, a class of ***morphisms*** denoted variously as $\mathsf{C}(X, Y)$ or $\mathrm{Hom}_{\mathsf{C}}(X, Y)$. (We will have occasion to use both notations – while $\mathsf{C}(X, Y)$ is more concise and easier on the eyes, the Hom notation is sometimes more enlightening). For every triple of objects X, Y, Z, there is a composition function $\mathsf{C}(Y, Z) \times \mathsf{C}(X, Y) \to \mathsf{C}(X, Z)$ sending g, f to the ***composition morphism*** $g \circ f$, often abbreviated to gf, whose existence we require. We also require that composition is associative, in the sense that $(h \circ g) \circ f = h \circ (g \circ f)$, as well as the existence of ***identity morphisms*** id_X for each $X \in \mathrm{Ob}(\mathsf{C})$ such that $g \circ \mathrm{id}_X = g$ and $\mathrm{id}_X \circ f = f$.

If $\mathrm{Ob}(\mathsf{C})$ is a set rather than a proper class, C is said to be *small*. If $\mathsf{C}(X, Y)$ is a set for all $X, Y \in \mathsf{C}$, then C is *locally small*, and we often refer to $\mathsf{C}(X, Y)$ as a *hom-set*[1].

[1]"Hom" is an abbreviation of homomorphism, a relic from category theory's origins in algebraic topology.

Many common "types" of mathematical objects can be assembled into categories:

- There is a category Set whose objects are sets and whose morphisms are functions (a function $f : X \to Y$ being a selection of an element in Y for every element of X). Composition of functions is defined in the usual sense, and there is an obvious identity morphism $\mathrm{id}_X : X \to X$, $x \mapsto x$.
- The category Top consists of topological spaces and continuous functions.
- The category Ab consists of abelian groups and group homomorphisms.
- The category CRing consists of commutative rings and ring homomorphisms.
- The category R-Mod consists of modules over a commutative ring R and their homomorphisms[2].
- The category Man^p consists of C^p manifolds and maps. For instance, $\mathsf{Diff} := \mathsf{Man}^\infty$ consists of smooth manifolds and maps.

Set is a locally small category, as are all categories whose objects and morphisms can be thought of as sets and set functions, including all of the above examples.

Monomorphisms and Epimorphisms In Set, we can classify morphisms into injective, surjective, and bijective maps. This generalizes in the following manner: A morphism $f : X \to Y$ in a category C is an ***epimorphism*** if, for all $g, h : Y \to Z$, we have $gf = hf$ if and only if $g = h$. f is a ***monomorphism*** if, for all $g, h : W \to X$, we have $fg = fh$ if and only $g = h$. f is an ***isomorphism*** if there is an inverse morphism $g : Y \to X$ such that $fg = \mathrm{id}_Y$ and $gf = \mathrm{id}_X$. Two objects in C are ***isomorphic*** if there is an isomorphism between them. The isomorphisms in Grp, Set, Top, and Diff, for instance, are the group isomorphisms, bijections, homeomorphisms, and diffeomorphisms, respectively; for nearly all intents and purposes, isomorphic objects are to be regarded as equivalent. Note: we often shorten epimorphism to *epi*, or in its adjectival form, an *epic* morphism, whereas monomorphism is shortened to *mono*, or a *monic* morphism.

In Set, (i) epimorphisms are equivalent to surjections, (ii) monomorphisms are equivalent to injections, and (iii) isomorphisms are equivalent to bijections. To prove this, take a map of sets $f : X \to Y$.

[2] We often write $R(X, Y)$ and $\mathrm{Hom}_R(X, Y)$ instead of R-$\mathsf{Mod}(X, Y)$ and $\mathrm{Hom}_{R\text{-}\mathsf{Mod}}(X, Y)$.

(i) Suppose that there is some $y \in Y$ not contained in $f(X)$. Let $Z = \{0, 1\}$, and let $g, h : Y \to Z$ send $Y - y$ to 0 and y to 0 or 1, respectively. $gf = hf$, but $g \neq h$. So if f is an epimorphism, it must be a surjection. Conversely, suppose that f is a surjection, and let $g, h : Y \to Z$ satisfy $gf = hf$. For every $y \in Y$ there is an x_y such that $f(x_y) = y$, so $g(y) = gf(x_y) = hf(x_y) = h(y)$, and $g = h$. Obviously, if $g = h$ then $gf = hf$ as well, so surjections are epimorphisms.

(ii) Similar to (i).

(iii) Bijections obviously have inverses. Conversely, let $f : X \to Y$ admit an inverse $g : Y \to X$ such that $g(f(x)) = x$ and $f(g(y)) = y$. If f is not surjective, then there is some $y \in Y$ mapped to by no $f(x)$, so we cannot have $f(g(y)) = y$, and if f is not injective, then there are $x \neq x' \in X$ with $f(x) = f(x')$ and therefore $x = g(f(x)) = g(f(x')) = x'$, a contradiction. So isomorphisms are injective and surjective, and hence bijective. Importantly, this proof hinges on the fact that injective surjections are bijections; in an arbitrary category, it is *not* necessarily true that a morphism which is both monic and epic is an isomorphism. A category where this is true is known as a ***balanced category***.

Most of our example categories are balanced, but CRing is not. To see this, take the inclusion $i : \mathbb{Z} \to \mathbb{Q}$. First, let $f, g : R \to \mathbb{Z}$ be such that $if = ig$. Since i is an inclusion, $f(r) = g(r)$ for all r, making i monic. Now let $h, k : \mathbb{Q} \to S$ be such that $hi = ki$. $h(p/q) = h(p)h(q^{-1}) = h(p)h(q)^{-1}$, so h and likewise k are completely determined by where they send the integers, and hence $hi = ki$ implies $h = k$. Despite being monic and epic, i fails to be an isomorphism.

Naturality In general, the vast majority of types of mathematical objects assemble into categories, the main concern being what the morphisms between objects of a certain type should be; generally, there is a natural notion of morphism between such objects (as in the above examples) which, when equipped to their category, allow that category to "encapsulate" the nature of that type of object. This natural notion is generally one that preserves precisely the structure associated to that type of object; given enough information about what is needed to define an object of that type, the structure we want morphisms to preserve generally becomes obvious.

For instance, we may define a natural notion of a morphism between categories: a morphism $F : \mathsf{C} \to \mathsf{D}$ should map objects $X \in \mathsf{C}$ to objects $FX \in \mathsf{D}$ and morphisms $f : X \to Y$ to morphisms $Ff : FX \to FY$ in a manner that preserves composition, identity, and associativity. Such a map has a special name: Given two categories C, D, a ***functor*** $F : \mathsf{C} \to \mathsf{D}$ consists of a map $\mathrm{Ob}(\mathsf{C}) \to \mathrm{Ob}(\mathsf{D})$, as well as, for every $X, Y \in \mathsf{C}$, a map $\mathsf{C}(X, Y) \to \mathsf{D}(FX, FY)$. We require $F(g \circ f) = (Fg) \circ (Ff)$ and $F\mathrm{id}_X = \mathrm{id}_{FX}$. (Associativity is trivial).

Given two functors $F, G : \mathsf{C} \to \mathsf{D}$, a ***natural transformation*** $\alpha : F \Rightarrow G$ is a family $\{\alpha_X : FX \to GX\}_{X \in \mathsf{C}}$ of maps in D such that, for any $f : X \to Y$, we have $(Gf) \circ \alpha_X = \alpha_Y \circ (Ff)$. If each α_X is an isomorphism, α is known as a ***natural isomorphism***.

We can define two new categories: the category Cat of small categories and functors, and, for any $\mathsf{C}, \mathsf{D} \in \mathsf{Cat}$, a category D^{C} (also written as $[\mathsf{C}, \mathsf{D}]$) whose objects are functors $\mathsf{C} \to \mathsf{D}$ and whose morphisms are natural transformations between functors. Both of these are subject to set-theoretic issues[3]. We will handwave these issues away, though especially curious/bored readers may see Appendix A for a discussion on the problems this can lead to, and the mechanisms for fixing them.

$$
\begin{tikzcd}
X \arrow[r] \arrow[d, "f"'] \arrow[dr, phantom, "\Longrightarrow"] & FX \arrow[d, "Ff"] \\
Y \arrow[r] & FY
\end{tikzcd}
\qquad
\begin{tikzcd}
& Y \arrow[dl] \arrow[dr] & \\
FY \arrow[rr, bend right, "\alpha_Y"] & & GY \\
& X \arrow[uu, "f"] \arrow[dl] \arrow[dr] & \\
FX \arrow[uu, "Ff"] \arrow[rr, bend right, "\alpha_X"] & & GX \arrow[uu, "Gf"']
\end{tikzcd}
$$

The data associated to a functor and natural transformation

All of our example categories are locally small, and their objects are sets equipped with extra structure. Such locally small categories which can be modeled on sets are called ***concrete***, and they admit functors $\mathsf{C} \to \mathsf{Set}$ which "forget" the structure on their objects, conveniently known as ***forgetful functors***. For instance, the forgetful functor $\mathsf{Ab} \to \mathsf{Set}$ just maps abelian groups to their underlying sets, and group homomorphisms to their underlying set functions. In general, for a category to be concrete we require the existence of a forgetful functor which is injective on

[3] It is for this reason that Cat consists of *small* categories; the set-theoretically problematic CAT is defined as the category of all categories.

hom-sets, as otherwise two different maps in C will be sent to the same set map, so we cannot speak of their "underlying" set maps.

A functor F for which each map $\mathsf{C}(X, Y) \to \mathsf{C}(FX, FY)$ is injective is known as a ***faithful functor***; in contrast, functors which are surjective on hom-sets are called ***full***. Faithfully full functors are bijections on hom-sets. On objects, F is ***essentially surjective*** if every object $Y \in \mathsf{D}$ is isomorphic to some FX, $X \in \mathsf{C}$.

1.1.2 Limits and Colimits

To see how categorical thinking can encapsulate the nature of certain types of mathematical objects, consider the product of topological spaces: given a pair of topological spaces X_1, X_2, we define their product to be a space X equipped with canonical projection maps $\pi_i : X \to X_i$, and give X the smallest topology that makes the π_i continuous. Every open set in this initial topology is required for continuity, making this the "most efficient" space with continuous morphisms into X_1 and X_2. This can be made rigorous by the following observation: any space Y equipped with a pair of functions $(f_1 : Y \to X_1, f_2 : Y \to X_2)$ admits a continuous map $f : Y \to X, y \mapsto (f_1(y), f_2(y))$ such that $\pi_1 f = f_1$ and $\pi_2 f = f_2$; in fact, this f is *uniquely* determined by f_1 and f_2. Pictorially, there is a unique arrow $f : Y \to X$ such that the triangles in the following diagram commute:

$$\begin{tikzcd}
& X \arrow[dl, "\pi_1"'] \arrow[dr, "\pi_2"] & \\
X_1 & & X_2 \\
& Y \arrow[ul, "f_1"] \arrow[ur, "f_2"'] \arrow[uu, dashed, "\exists! f"] &
\end{tikzcd}$$

In particular, if we set $Y = X$, we get $f = \mathrm{id}_X$. We see that $X = X_1 \times X_2$ encodes pairs of morphisms $(f_1 : Y \to X_1, f_2 : Y \to X_2)$ in the most efficient possible way; in fact, if any other space X' with morphisms $(\pi'_1 : X' \to X_1, \pi'_2 : X' \to X_2)$ satisfies this property, then the diagram

$$\begin{tikzcd}
 & X \arrow[dl, "\pi_1"'] \arrow[d, dashed, "f"] \arrow[dr, "\pi_2"] & \\
X_1 & X' \arrow[l, "\pi_1'"'] \arrow[r, "\pi_2'"] \arrow[d, dashed, "f'"] & X_2 \\
 & X \arrow[ul, "\pi_1"] \arrow[ur, "\pi_2"'] &
\end{tikzcd}$$

demonstrates that the unique morphism $f'f : X \to X$ satisfies $\pi_1 = \pi_1 f'f$ and $\pi_2 = \pi_2 f'f$; since id_X also satisfies this property, we have $f'f = \mathrm{id}_X$, and by the same reasoning $ff' = \mathrm{id}_{X'}$, making X' and X homeomorphic to one another. It follows that the product of topological spaces can be *defined* (up to homeomorphism) by this category-theoretic requirement, which takes place abstractly in Top. We can generalize this to an arbitrary category C:

The product of two objects X, Y is, if it exists, an object $X \times Y$ equipped with morphisms $\pi_X : X \times Y \to X$ and $\pi_Y : X \times Y \to Y$ such that for every Z equipped with a pair of morphisms $f : Z \to X$ and $g : Z \to Y$, there is a unique morphism $h : Z \to X \times Y$ with $\pi_X h = f$ and $\pi_Y h = g$.

The product in Top is the topological product, as we've seen; in $\mathsf{Ab}, \mathsf{Set}$, and CRing, it's the product of abelian groups, cartesian product of sets, and product of rings, respectively. All of these share the same property of being unique up to isomorphism. In general, suppose two objects X, Y in a category C have two products, Z_0 and Z_1. Then Z_0 and Z_1 are isomorphic.

Proof. Let ϕ_X, ϕ_Y be the canonical projections from Z_0 and ψ_X, ψ_Y the canonical projections from Z_1. By the universal property of the product, Z_1 has an arrow $\Psi : Z_1 \to Z_0$ such that $\phi_X \circ \Psi = \psi_X$ and $\phi_Y \circ \Psi = \psi_Y$, and Z_0 has an arrow $\Phi : Z_0 \to Z_1$ such that $\psi_X \circ \Phi = \phi_X$ and $\psi_Y \circ \Phi = \phi_Y$. It follows that $\phi_X \circ \Psi \circ \Phi = \psi_X \circ \Phi = \phi_X$, and $\phi_Y \circ \Psi \circ \Phi = \phi_Y$. Likewise, $\psi_X \circ \Phi \circ \Psi = \psi_X$ and $\psi_Y \circ \Phi \circ \Psi = \psi_Y$. It follows that both the morphisms $\Psi \circ \Phi$ and id_{Z_0} satisfy the required factorization identities in the product diagram for Z_0, and likewise for Z_1, as indicated in the following diagrams:

$$\begin{tikzcd}
 & Z_0 \arrow[dl, "\phi_X"'] \arrow[dr, "\phi_Y"] & \\
X & & Y \\
 & Z_0 \arrow[ul, "\phi_X"] \arrow[ur, "\phi_Y"'] \arrow[uu, bend left, "\mathrm{id}_{Z_0}"] \arrow[uu, bend right, "\Psi\Phi"'] &
\end{tikzcd}
\qquad
\begin{tikzcd}
 & Z_1 \arrow[dl, "\psi_X"'] \arrow[dr, "\psi_Y"] & \\
X & & Y \\
 & Z_1 \arrow[ul, "\psi_X"] \arrow[ur, "\psi_Y"'] \arrow[uu, bend left, "\mathrm{id}_{Z_1}"] \arrow[uu, bend right, "\Phi\Psi"'] &
\end{tikzcd}$$

So $\mathrm{id}_{Z_0} = \Psi\Phi$ and $\mathrm{id}_{Z_1} = \Phi\Psi$, making Φ and Ψ isomorphisms between Z_0 and Z_1. ■

This manner of thinking about categorical constructions can be vastly generalized: for instance, we may ask for an object that classifies morphisms into *no* objects, i.e. an object T that has a unique morphism $f : X \to T$ for all $X \in \mathsf{C}$. Such an object is known as a ***terminal object***. We may even throw morphisms into the mix: given a diagram $f, g : X_1 \rightrightarrows X_2$, we may ask for an object Y equipped with morphisms $i : Y \to X_1, j : Y \to X_2$ such that $fi = gi = j$ any other object equipped with commuting morphisms to X_1 and X_2 bears a unique morphism to Y; such a Y, when it exists, is known as the ***equalizer*** of f and g, $\mathrm{Eq}(f, g)$.

$$\begin{tikzcd}
& X_1 \times X_2 \arrow[dl] \arrow[dr] & \\
X_1 & & X_2 \\
& X \arrow[ul] \arrow[ur] \arrow[uu, dashed] &
\end{tikzcd}
\qquad
\begin{tikzcd}
T \\
\\
X \arrow[uu, dashed]
\end{tikzcd}
\qquad
\begin{tikzcd}
\mathrm{Eq}(f,g) \arrow[d] \arrow[dr] & \\
X_1 \arrow[r, shift left, "f"] \arrow[r, shift right, "g"'] & X_2 \\
X \arrow[u] \arrow[ur] \arrow[uu, dashed, bend left] &
\end{tikzcd}$$

Diagrams for products, terminal objects, and equalizers; dashed arrows are unique

This process is generalized in the obvious way to arbitrary diagrams; the object corresponding to a certain diagram is known as the ***limit*** of that diagram. For instance, the limit of the empty diagram is the terminal object, the limit of the diagram $X_1\ X_2$ is the product $X_1 \times X_2$, and the limit of the diagram $f, g : X_1 \rightrightarrows X_2$ is the equalizer $\mathrm{Eq}(f, g)$. The proof of the uniqueness of products up to isomorphism generalizes easily to the uniqueness of any kind of limit. In particular, any category can have at most one terminal object up to isomorphism. In Set, all singletons are terminal objects – for X an arbitrary set, there's only a single function $f : X \to \{*\}$ sending all $x \in X$ to the single object $*$ – and all singletons are isomorphic, allowing us to just speak of "a" terminal set; if we need a specific one, we'll use the ordinal $1 := \{\varnothing\}$.

Duality Given any category C, we can flip all the arrows, obtaining the opposite category C^{op}. For instance, a morphism $X \to Y$ in $\mathsf{Set}^{\mathrm{op}}$ is given by a function $f : Y \to X$. In general, every arrow-theoretic statement and construction has a dual, given by flipping all the arrows and attaching the prefix 'co'; this is known as the principle of duality. For instance, the *co*product of two objects $X_1, X_2 \in \mathsf{C}$ is an object $X_1 \amalg X_2$ equipped with two morphisms $i_1 : X_1 \to X_1 \amalg X_2$,

$i_2 : X_2 \to X_1 \amalg X_2$ such that any X also equipped with such morphisms has a unique morphism *from* $X_1 \amalg X_2$ making everything commute.

$$\begin{tikzcd} & X_1 \times X_2 \arrow[dl] \arrow[dr] & \\ X_1 & & X_2 \\ & X \arrow[ul] \arrow[ur] \arrow[uu, dashed] & \end{tikzcd} \qquad \begin{tikzcd} & X_1 \amalg X_2 \arrow[dd, dashed] & \\ X_1 \arrow[ur] \arrow[dr] & & X_2 \arrow[ul] \arrow[dl] \\ & X & \end{tikzcd}$$

Comparing product and coproduct diagrams

We similarly have ***coequalizers***, *co*terminal (***initial***) objects, and in general, ***colimits***.

An especially ubiquitous notion is given by that of a cofunctor, or a contravariant functor: A ***contravariant functor*** $F : \mathsf{C} \to \mathsf{D}$ is a functor $\mathsf{C}^{\mathrm{op}} \to \mathsf{D}$. Specifically, each arrow $f : X \to Y$ in C is sent to an arrow $Ff : FY \to FX$, and composition works backwards, sending $g \circ f : X \to Y \to Z$ to $F(gf) = (Ff)(Fg) : FZ \to FY \to FX$. Normal functors are often called ***covariant*** when specification is required.

Example. For every object X in a category C, there is a covariant functor $\mathsf{C}(X, -) : \mathsf{C} \to \mathsf{Set}$ sending $Y \in \mathsf{C}$ to the set $\mathsf{C}(X, Y)$, and a morphism $f : Y \to Z$ to the set map $f_* : \mathsf{C}(X, Y) \to \mathsf{C}(X, Z)$ sending a $g : X \to Y$ to $f_*(g) = f \circ g$. The dual, contravariant functor is given by $\mathsf{C}(-, X)$, which sends an object Y to $\mathsf{C}(Y, X)$ and a map $f : Y \to Z$ to $f^* : \mathsf{C}(Z, X) \to \mathsf{C}(Y, X)$, $g \mapsto g \circ f$. $C(X, -)$ and $C(-, X)$ are known as the covariant and contravariant ***representable functors*** for X.

Example. A ***lattice*** is a poset which, as a category, has all binary products and coproducts. The coproduct is to be interpreted as the join (or sup, logical OR) $x \vee y$ and the product as the meet (or inf, logical AND) $x \wedge y$. Since the categorical structure on an arbitrary poset is given by writing an arrow $x \to y$ whenever $x \leq y$, the join of two elements x, y is an element $x \vee y$ satisfying $x, y \leq x \vee y$, and such that any object z satisfying $x, y \leq z$ also satisfies $x \vee y \leq z$. In this way, $x \vee y$ is the least upper bound of x, y, while $x \wedge y$ is the greatest lower bound.

If L has elements 0 and 1 such that $0 \leq x \leq 1$ for all $x \in L$, then 0 and 1 are the initial and terminal objects of L as a category. Equalizers and coequalizers are trivial in lattices, so a lattice with 0 and 1 is a poset which, as a category, has all finite limits and colimits.

We may also define lattices with 0 and 1 equationally: a lattice is a set with two distinguished elements 0 and 1, and two associative, commutative binary operations $\vee$ and $\wedge$ such that $x \wedge x = x \vee x = x$, $1 \wedge x = 0 \vee x = x$, and $x \wedge (y \vee x) = (x \wedge y) \vee x = x$. The partial order is recovered by the relation $x \leq y \iff x = x \wedge y \iff y = x \vee y$. If also $x \wedge (y \vee z) = (x \wedge y) \vee (x \wedge z)$, or, *equivalently*, $x \vee (y \wedge z) = (x \vee y) \wedge (x \vee z)$, we say that the lattice is distributive. If L has for each x an element $\neg x$ such that $x \wedge \neg x = 0$ and $x \vee \neg x = 1$, then such a $\neg x$ is unique, and is known as the ***complement*** of x. A ***Boolean algebra*** is a distributive lattice with 0 and 1 in which every element x has a complement. In such a lattice, the DeMorgan laws hold: $\neg(x \vee y) = \neg x \wedge \neg y$, $\neg(x \wedge y) = (\neg x) \vee (\neg y)$, and $\neg\neg x = x$. For instance, every poset of subsets of a given set is a Boolean algebra under the operations of union, intersection, and complement; in fact, every Boolean algebra can be constructed up to isomorphism in this manner.

Equivalence and Universality As indicated earlier, the notion of naturality plays a large role in category theory; categories and their morphisms serve as a method of organizing objects of a certain type, and basic constructions on categories (taking limits, opposites, etc.) yield natural constructions on the corresponding objects. The key ingredient in all of these constructions is universality, which can be thought of as selecting the "most general" or "best" way of doing something: for instance, the product $X \times Y$ of two objects is the most general object that bears morphisms to both X and Y, in the sense that all other objects with morphisms to X and Y see those morphisms factor *uniquely* through those of $X \times Y$ [4]. Even without the use of category theory, universal properties show up throughout mathematics: for instance, the tensor product $M \otimes N$ of R-modules M and N satisfies the universal property that any bilinear morphism $M \oplus N \to P$ factors uniquely through $M \otimes N$; informally, it is the *most general way* to turn bilinear homomorphisms into linear morphisms. The localization of a ring A at a multiplicatively closed subset $S \subset A$ satisfies the universal property that every ring homomorphism $A \to B$ which sends A to an invertible element of B factors uniquely through $S^{-1}A$; it is the most general way to add inverses to A.

[4]The name "universality" derives from the fact that this property is expressed via *universal properties*, as $\forall \ldots \exists! \ldots$.

Category theory also allows us to weaken the notion of equivalence from strict equality $(=)$ to isomorphism $(\cong)$. Many categories have a natural notion of a "morphism between morphism", or a ***2-morphism***: e.g., natural transformations serve as the 2-morphisms in Cat. In a category with 2-morphisms, known as a ***2-category***, we can further weaken the notion of equivalence: let X, Y be objects of a 2-category C with morphisms $F : X \to Y$ and $G : Y \to X$ such that FG admits a 2-isomorphism $\alpha : FG \cong \mathrm{id}_Y$ and GF a 2-isomorphism $\beta : GF \cong \mathrm{id}_X$. In $\mathsf{C} = \mathsf{CAT}$, this concept bears a special name: An ***equivalence of categories*** $\mathsf{C} \cong \mathsf{D}$ is a pair of functors $F : \mathsf{C} \to \mathsf{D}$, $G : \mathsf{C} \to \mathsf{D}$ equipped with natural isomorphisms $\alpha : FG \cong \mathrm{id}_Y$ and $\beta : GF \cong \mathrm{id}_X$.

Yoneda's Lemma For a category C, we will denote the functor category $\mathsf{Set}^{\mathsf{C}^{\mathrm{op}}}$ of contravariant functors $\mathsf{C} \to \mathsf{Set}$ by $\widehat{\mathsf{C}}$; its elements are known as ***presheaves***. Yoneda's lemma states that C admits a full and faithful embedding into its category of presheaves $\widehat{\mathsf{C}}$.

For a covariant functor $F : \mathsf{C} \to \mathsf{Set}$, the set $\widehat{\mathsf{C}}(\mathsf{C}(X, -), F)$ of natural transformations from $\mathsf{C}(X, -)$ to F is isomorphic to FX. For a contravariant $F : \mathsf{C}^{\mathrm{op}} \to \mathsf{Set}$, $\widehat{\mathsf{C}}(\mathsf{C}(-, X), F) \cong FX$.

Proof. For F covariant, take an arbitrary $a \in FX$. Letting $\alpha_X(\mathrm{id}_X) = a$ defines a unique natural transformation in which any $f : X \to Y$ must be mapped to $(Ff)(a)$. Conversely, any $a \in FX$ defines a unique natural transformation $\alpha_Y(f) = (Ff)(a)$. For F contravariant, flip the direction of f. ∎

Note that when $F = \mathsf{C}(-, Y)$, the contravariant version yields

$$\widehat{\mathsf{C}}(\mathsf{C}(-, X), \mathsf{C}(-, Y)) \cong \mathsf{C}(X, Y)$$

We may use this to define an embedding of C in $\widehat{\mathsf{C}}$: the ***Yoneda embedding*** is the functor よ : $\mathsf{C} \to \widehat{\mathsf{C}}$ sending X to $\mathsf{C}(-, X)$ and $f : X \to Y$ to the natural transformation $\mathsf{C}(-, X) \Rightarrow \mathsf{C}(-, Y)$ corresponding to f. Since the sets of natural transformations between two functors *are* the hom-sets in the functor category $\widehat{\mathsf{C}}$, よ is a full and faithful functor, and hence a proper embedding. Furthermore, $\widehat{\mathsf{C}}$ also contains all colimits in a natural way: (Co-Yoneda lemma) Every element of $\widehat{\mathsf{C}}$ is a colimit of a diagram of contravariant representable functors in a canonical manner. For further details and a proof, see [MacLane and Moerdijk, 2012], pgs. 41-43.

1.1.3 Adjunctions

The "best" relation two functors $F : \mathsf{C} \to \mathsf{D}$ and $G : \mathsf{D} \to \mathsf{C}$ can have is their forming an equivalence of categories $\mathsf{C} \cong \mathsf{D}$. Then, morphisms in C can be mapped to morphisms in D in a natural and reversible manner (up to isomorphism). The *next* best relation F and G can have is a failure of equivalence on objects, but an equivalence on morphisms, in the sense that $\mathsf{D}(FX, Y)$ is in bijection with $\mathsf{C}(X, GY)$ for all $X \in \mathsf{C}$, $Y \in \mathsf{D}$. If this happens in a natural manner, we say that F and G are adjoint functors. Adjunctions show up everywhere, as we will demonstrate.

Given locally small categories C and D, along with functors $F : \mathsf{C} \to \mathsf{D}$ and $G : \mathsf{D} \to \mathsf{C}$, we call F and G ***adjoint*** functors if there's a natural isomorphism Φ between the following functors from $\mathsf{C}^{\mathsf{op}} \times \mathsf{D}$ to Set:

$$\Phi : \mathsf{D}(F-, -) \cong \mathsf{C}(-, G-)$$

Then, F is said to be left adjoint to G, and G is said to be right adjoint to F. This relation is written as $F \dashv G$, with the $\dashv$ symbol pointing towards the *left* adjoint (we could also write $G \vdash F$).

The name "adjoint" comes from linear algebra, where the adjoint of an operator A on an inner product space V is another operator $A^\dagger$ satisfying $\langle Av, w\rangle = \langle v, A^\dagger w\rangle$: we "move" the operator to the other side by taking its adjoint.

Example. The free abelian group on a set S, is defined to be an abelian group $F(S)$ along with an inclusion set map $i_S : S \to F(S)$ such that every set map $u : S \to A$, where A is an abelian group, factors as $u = \varphi \circ i$ for a unique homomorphism φ. A set map $f : S \to T$ generates by composition a map $i_T \circ f : S \to F(T)$, and hence a unique homomorphism $F(S) \to F(T)$; it can be verified that when $f = \mathrm{id}_S$, this homomorphism is $\mathrm{id}_{F(S)}$, and furthermore that composition of these induced maps is associative. This evidences F as a functor $\mathsf{Set} \to \mathsf{Ab}$, known as a ***free functor***. If we let J be the forgetful functor $\mathsf{Ab} \to \mathsf{Set}$, then we see that $\mathsf{Set}(S, JA)$ is in bijection with $\mathsf{Ab}(FS, A)$: the map from set maps to group homomorphisms is given by the definition of the free group, and the map from group homomorphisms to set maps is given by taking $\varphi : F(S) \to A$ to the set map $\varphi \circ i : S \to JA$. This bijection is natural in both S and A, rendering F the left adjoint to J. Free-forgetful adjunctions of this nature are extremely common: in fact, we may define free functors as left adjoints to forgetful functors.

Example. In Set, maps $X \times Y \to Z$ can be identified with maps $X \to \mathsf{Set}(Y, Z)$ by ***currying***:

in lambda notation, we send $\lambda x, y.f(x,y)$ to $\lambda x.\,(\lambda y.f(x,y))$. This yields an adjunction with $- \times Y$ on the left and $\mathrm{Set}(Y, -)$ on the right. As we'll see later, this is the defining feature of a ***cartesian closed category***.

Example. A ***Heyting algebra*** is a lattice H with 0 and 1 which has an right adjoint known as exponentiation associated to the functor $- \wedge y$. That is, there is for every x, y an object, generally written as $x \Rightarrow y$, such that $z \leq (x \Rightarrow y)$ iff $x \wedge x \leq y$, i.e. $x \Rightarrow y$ is a least upper bound for all elements z with $z \wedge x \leq y$. In particular, $y \leq (x \Rightarrow y)$.

The unit and counit of the exponential adjunction give us inclusions $x \leq (y \Rightarrow (x \wedge y))$ and $y \wedge (y \Rightarrow x) \leq x$. The properties $1^X \cong 1$ and $X^1 \cong X$, valid in any category with a right adjoint to its product functor, become $(x \Rightarrow 1) = 1$ and $(1 \Rightarrow x) = x$, and the properties $(y \times z)^x \cong y^x \times z^x$ and $x^{y \times z} \cong (x^y)^z$ become $(x \Rightarrow (y \wedge z)) = ((x \Rightarrow y) \wedge (x \Rightarrow z))$ and $((y \wedge z) \Rightarrow x) = (z \Rightarrow (y \Rightarrow x))$. Heyting algebras are distributive due to the fact that $- \wedge y$ is a left adjoint, and hence preserves coproducts: $((x \vee z) \wedge y) = ((x \wedge y) \vee (z \wedge y))$.

In a Heyting algebra, we may define the negation of x as $\neg x := (x \Rightarrow 0)$, the idea being that "not x" means "x implies falsity". This is not a strict negation: while $x \wedge \neg x = 0$, as evidenced by the identity $x \wedge (x \Rightarrow y) \leq y$, $x \vee \neg x$ isn't necessarily equal to 1. If x does have a strict negation, though, it is $\neg x$. So while $x \leq \neg\neg x$, this isn't a strict equality as in a Boolean algebra. However, $\neg x = \neg\neg\neg x$, and $x \leq y$ implies that $\neg y \leq \neg x$, so we're not totally lost. These features tell us that the logic of a Heyting algebra doesn't necessarily satisfy the law of double negation $x = \neg\neg x$, and as such is an *intuitionistic* logic rather than a classical one.

Given a predicate $S(x,y)$, where $x \in X$ and $y \in Y$ are elements of sets, we may regard S as the subset $S \subseteq X \times Y$ of those pairs for which $S(x,y)$ is true. The statement $(\forall x)S(x,y)$ then picks out a subset $T \subseteq Y$ consisting of all those y such that $X \times y \subseteq S$. Letting p denote the projection $X \times Y \to Y$, we may denote this subset as $\forall_p S$. The statement $(\exists x)S(x,y)$ is equivalent to $y \in p(S)$, and we will denote the corresponding subset by $\exists_p S$. Let $\mathcal{P}Y$ be the Boolean algebra of all subsets $T \subseteq Y$ and $\mathcal{P}(X \times Y)$ the Boolean algebra of all predicates S. Viewing these as categories, we have a pair of functors $\forall_p, \exists_p : \mathcal{P}(X \times Y) \rightrightarrows \mathcal{P}(Y)$. There is a third functor, $p^* : \mathcal{P}(Y) \to \mathcal{P}(X \times Y)$ which sends each subset $T \subseteq Y$ to its inverse image $p^*T = X \times T$. Then, there is an adjoint triple $\exists_p \dashv p^* \dashv \forall_p$. This follows from the fact that $p^*T \subseteq S \iff T \subseteq \forall_p S$ and $S \subseteq p^*T \iff \exists_p S \subseteq T$.

Example. Ab is naturally a subcategory of Grp, so we can define an inclusion functor $i : \mathsf{Ab} \to \mathsf{Grp}$ which just drops the 'abelian' prefix. The left adjoint of this functor is given by abelianization, sending a group G to $G/[G,G]$ and a group homomorphism $\varphi : G \to H$ to the map $\varphi^* : G \to H \to H/[H,H]$, which satisfies $\varphi(xy) = \varphi(x)\varphi(y) = \varphi(y)\varphi(x) = \varphi(yx)$ and hence extends to a morphism $G/[G,G] \to H/[H,H]$. In general, a subcategory C_0 of a category C is ***reflective*** when its inclusion functor has a left adjoint, and ***coreflective*** when the inclusion functor has a right adjoint.

A paramount feature of adjoints which we will state but not prove is their ability to preserve limits and colimits. Let $F : \mathsf{C} \to \mathsf{D}$ be left adjoint to $G : \mathsf{C} \to \mathsf{D}$, let Γ be a diagram in C, and let Δ be a diagram in D. Then, $\mathrm{colim} F\Gamma = F(\mathrm{colim}\Gamma)$ and $\lim G\Delta = G(\lim \Delta)$. Succinctly, *left adjoints preserve colimits and right adjoints preserve limits*.

Units and Counits Given an adjunction $\Phi : \mathsf{C}(X, GY) \cong \mathsf{D}(FX, Y)$, suppose we set $Y = FX$, giving us a bijection $\mathsf{C}(X, GFX) \cong \mathsf{D}(FX, FX)$. Plugging the identity 1_{F_X} in on the right side gives us a unique $\eta_X : X \to GFX$. Doing this for all X gives us a natural transformation $\mathrm{id}_{\mathsf{C}} \to GF$, since an $h : X' \to X$ is translated to a $GFh : GFX' \to GFX$ such that $GFh \circ \eta_{X'} = \eta_X \circ h$ (proof: $GFh \circ \eta_{X'} = GFh \circ \Phi(\mathrm{id}_{FX'}) = \Phi(Fh \circ \mathrm{id}_{FX'}) = \Phi(\mathrm{id}_{FX} \circ Fh) = \Phi(\mathrm{id}_{FX}) \circ h = \eta_X \circ h$). Dually, we can let $X = GY$, so that plugging in id_{GY} into the right hand side of the bijection $\mathsf{C}(GY, GY) \cong \mathsf{D}(FGY, Y)$ gives us a natural transformation $\varepsilon : FG \to \mathrm{id}_{\mathsf{D}}$. Both the composites $G \xrightarrow{\eta G} GFG \xrightarrow{G\varepsilon} G$ and $F \xrightarrow{F\eta} FGF \xrightarrow{\varepsilon F} F$ reduce to the identities 1_G and 1_F; from this, we obtain the adjunction's ***zig-zag identities***

$$(\varepsilon F)(F\eta) = 1_F \qquad (\eta G)(G\varepsilon) = 1_G$$

We call η the ***unit*** of the adjunction and ε the ***counit***.

Monads Consider the iterated composites of an endofunctor $T : \mathsf{C} \to \mathsf{C}$, i.e. $T^2 = TT$, $T^3, \ldots$. If $\mu : T^2 \to T$ is a natural transformation, with μ_X a morphism $T^2X \to TX$, then $T\mu = \{T\mu_X\}_{X \in \mathsf{C}}$ is a natural transformation from T^3 to T^2, defined by $(T\mu)_X$ to $T(\mu_X)$. μT is another natural transformation between T^3 and T^2, defined by $(\mu T)_X \coloneqq \mu_{TX}$.

A ***monad*** in a category C consists of an endofunctor T on C and two natural transformations $\eta : \mathrm{id}_{\mathsf{C}} \to T$ and $\mu : T^2 \to T$ known as the ***unit*** and ***multiplication*** such that the following diagrams commute:

$$\begin{tikzcd} T^3 \arrow[r, "T\mu"] \arrow[d, "\mu T"'] & T^2 \arrow[d, "\mu"] \\ T^2 \arrow[r, "\mu"] & T \end{tikzcd} \qquad \begin{tikzcd} 1T \arrow[r, "\eta T"] \arrow[dr, equal] & T^2 \arrow[d, "\mu"] & T1 \arrow[l, "T\eta"'] \arrow[dl, equal] \\ & T & \end{tikzcd}$$

where 1 is the natural transformation $\{\mathrm{id}_X\}_{X \in \mathsf{C}}$.

The structure is meant to resemble that of a monoid (identity, associative composition), with η the ***unit*** of T and μ the multiplication of T. In this sense, the left diagram just expresses the associativity of multiplication, and the right diagram expresses the left and right unit laws.

Example. As an example, the powerset functor $\mathcal{P} : \mathsf{Set} \to \mathsf{Set}, X \mapsto \mathcal{P}X$, $(\mathcal{P}f)(S) = f(S)$ forms a monad. The unit sends $X \in \mathsf{Set}$ to the map $\eta_X : \mathrm{id}_{\mathsf{Set}}(X) \to \mathcal{P}X, x \mapsto \{x\}$, and the multiplication sends X to the map $\mu_X : \mathcal{P}\mathcal{P}X \to \mathcal{P}X$, $\{S_\lambda\} \mapsto \bigcup_\lambda S_\lambda$.

To verify the coherence laws, let $\mathcal{S} = \{\{S_{\lambda_\xi}\}_{\xi \in \Xi}\}_{\lambda \in \Lambda}$, where each S_{λ_ξ} is a subset of X, be an arbitrary element of $\mathcal{P}\mathcal{P}\mathcal{P}X$. We want to verify that $(\mu_X \mu_{\mathcal{P}X})(\mathcal{S}) = (\mu_X \mathcal{P}\mu_X)(\mathcal{S})$. On one side, $(\mu_X \mu_{\mathcal{P}X})(\mathcal{S}) = \bigcup_{\lambda \in \Lambda}\left(\bigcup_{\xi \in \Xi} S_{\lambda_\xi}\right) = \bigcup_{\lambda,\xi} S_{\lambda_\xi}$. On the other side, note that $\mathcal{P}\mu_X$ is a map $\mathcal{P}\mathcal{P}\mathcal{P}X \to \mathcal{P}\mathcal{P}X$ sending $\mathcal{S}$ to $\{\bigcup_{\xi \in \Xi} S_{\lambda_\xi}\}_{\lambda \in \Lambda}$, so $(\mu_X \mathcal{P}\mu_X)(\mathcal{S}) = \bigcup_{\lambda \in \Lambda}\left(\bigcup_{\xi \in \Xi} S_{\lambda_\xi}\right) = \bigcup_{\lambda,\xi} S_{\lambda_\xi}$ as well. To verify the law for η, we must show that $\mu_X \eta_{\mathcal{P}X} = \mu_X \mathcal{P}\eta_X = \mathrm{id}_{\mathcal{P}X}$, which is evident from the trivial action of μ on singletons.

Every adjunction $F : \mathsf{C} \to \mathsf{D} \dashv G : \mathsf{D} \to \mathsf{C}$ gives rise to a monad in the category C. GF is the endofunctor on C, the unit $\eta : \mathrm{id}_{\mathsf{C}} \to GF$ of the adjunction the unit of the monad, and, given the counit ε, the multiplication is given as $G\varepsilon F : GFGF \to GF$. The coherence laws then look like

$$\begin{tikzcd} GFGFGF \arrow[r, "GFG\varepsilon F"] \arrow[d, "G\varepsilon FGF"'] & GFGF \arrow[d, "G\varepsilon F"] \\ GFGF \arrow[r, "G\varepsilon F"'] & GF \end{tikzcd} \qquad \begin{tikzcd} FGFG \arrow[r, "FG\varepsilon"] \arrow[d, "\varepsilon FG"'] & FG \arrow[d, "\varepsilon"] \\ FG \arrow[r, "\varepsilon"'] & \mathrm{id}_{\mathsf{D}} \end{tikzcd} \qquad \begin{tikzcd} \mathrm{id}_{\mathsf{C}}\, GF \arrow[r, "\eta GF"] \arrow[dr, equal] & GFGF \arrow[d, "G\varepsilon F"] & GF\, \mathrm{id}_{\mathsf{C}} \arrow[l, "GF\eta"'] \arrow[dl, equal] \\ & GF & \end{tikzcd}$$

The middle diagram is just a restatement of the right, obtained by removing the G on the left and the F on the right; it must hold, since $\varepsilon\varepsilon = \varepsilon \cdot (FG\varepsilon) = \varepsilon \cdot (\varepsilon FG)$. The right diagram must hold since $1 = G\varepsilon \cdot \eta G = \varepsilon F \cdot F\eta$.

Example. Consider the free abelian group - forgetful functor adjunction $F \dashv U$. This yields a monad with unit $\eta : \mathrm{id}_{\mathsf{Set}} \to UF$ with $\eta_X : X \to UFX$ sending $x \in X$ to x considered as a basis element of FX and multiplication $U\varepsilon F : UFUF \to UF$, where $\varepsilon : FU \to 1_{\mathsf{Ab}}$ sends an abelian group A to a morphism $FUA \to A$ that takes the elements of an *element* of FUA (a collection of arbitrary un-concatenated elements of A) and multiplies them all together to get an element of A. This is conceptually similar to the power set monad, in that the unit "wraps" a set ($x \mapsto \{x\}$ vs. $x \mapsto \{\text{basis element } x\}$), whereas the multiplication gives us a way to reduce several elements at the same level (set of sets $\mapsto$ set of union of sets vs. set of elements of abelian group $\mapsto$ sum of elements in abelian group). This similarity comes from the fact that both monads involve Set as the base category.

Given a monad $T = (T, \mu, \eta)$ on C, an ***algebra*** over T, or a T-algebra, is an object $X \in \mathsf{C}$ along with a morphism $f : TX \to X$ such that $f\eta_X = \mathrm{id}_X$ and $f(Tf) = f\mu_X$. In the power set monad on Set, for instance, an algebra is an assignment to each subset S of a given object X an element $f(S)$ such that $f(\{x\}) = x$ and $f(\{f(S_\lambda)\}) = f\left(\bigcup_\lambda S_\lambda\right)$. A ***morphism*** of T-algebras $(X, f) \to (Y, g)$ is a morphism $\alpha : X \to Y$ where the obvious square commutes: $g(T\alpha) = \alpha f$. Thus, any monad T on C gives us a category C^T of T-algebras, known as the ***Eilenberg-Moore category*** of T. While there is no natural choice of map $TX \to X$ (we have to choose a T-algebra structure), there is a natural map $\mu_X : T^2X \to TX$ giving TX a T-algebra structure. The functor $F^T : \mathsf{C} \to \mathsf{C}^T$ sending X to the algebra (TX, μ_X) is known as the ***free algebra*** functor, and the subcategory of C^T consisting of the free algebras is known as the ***Kleisli category*** C_T.

The free algebra functor $F^T : \mathsf{C}^T \to \mathsf{C}$ is left adjoint to the forgetful functor $\mathsf{C}^T \to \mathsf{C}$, $(X, f) \mapsto X$. The counit of this adjunction is the natural transformation $\mu : T^2 \to T$ and the unit is $\eta : 1 \to T$. In this way, not only does every adjunction generate a monad, but every monad comes from an adjunction.

1.2 Monoidal Categories

Motivation Many families of objects that naturally assemble into categories can be endowed with additional operations. Some motivating examples, some of which we have already seen:

- (Monoidal structure) Given two R-modules M and N, their tensor product is the module

$M \otimes N$, unique up to isomorphism, such that bilinear maps $\varphi : M \times N \to P$ are naturally in bijection with maps $M \otimes N \to P$. The operator $\otimes$ can be extended to a bifunctor $R\text{-}\mathsf{Mod} \times R\text{-}\mathsf{Mod} \to R\text{-}\mathsf{Mod}$, and equips $R\text{-}\mathsf{Mod}$ with the structure of a *monoid*.

- (Cartesian closed structure) Every function of the form $f : X \times Y \to Z$ in Set is equivalent to a function of the form $X \to \mathrm{Hom}_{\mathsf{Set}}(Y, Z)$ via currying. Similarly, in the category CGWH, the adjunction $- \times X \dashv -^X$ allows us to identify maps $Y \times X \to Z$ with maps $Y \to (X \to Z)$ in a manner entirely internal to CGWH.

- (Model structure) Every morphism in Top can be factored as a fibration followed by a cofibration [Riehl, 2014]. Any morphism which is both a fibration and a cofibration is a weak equivalence, inducing isomorphisms on all higher homotopy groups. The fibrations and cofibrations on Top tell us what we need to know in order to do homotopy theory, and by defining fibrations and cofibrations in arbitrary categories, we may do homotopy theory in categories other than Top.

- (Enriched structure) Every hom-set in $R\text{-}\mathsf{Mod}$ is an abelian group in a natural way: the identity is the zero map $0(m) = 0$, and addition is given by $(\varphi + \psi)(m) = \varphi(m) + \psi(m)$. Composition is a bilinear map $\circ_{XYZ} : R(X, Y) \times R(Y, Z) \to R(X, Z)$ as well, so we say that $R\text{-}\mathsf{Mod}$ is *enriched* over Ab.

- (n-categorical structure) In Cat, morphisms are functors. The set D^{C} of functors $\mathsf{C} \to \mathsf{D}$ is itself a category, with natural transformations as morphisms; we can therefore say that Cat has not just hom-sets but hom-*categories*.

- (Abelian structure) In many categories enriched over Ab, such as $R\text{-}\mathsf{Mod}$, morphisms have kernels, images, cokernels, and coimages; we can correspondingly find quotient objects and speak of the homology of chain complexes. [Weibel, 1995].

- (Topological structure) Diff admits a natural notion of a covering, in which a function family $\{M_i \to M\}$ covers the smooth manifold M if the images of all functions form an open cover of M [MacLane and Moerdijk, 2012]. It is possible to extend this notion of a covering to the notion of a topology on a category, known as a *Grothendieck topology*.

We will use these examples to construct a few hierarchies of structures that can be placed on (arbitrary) categories. Enriched categories, in particular, give us a way to replace the hom-sets of a category C with hom-*objects* in a category V with some additional structure necessary to define composition; n-categories are examples of enriched categories, and abelian categories are categories enriched over Ab with some additional niceness properties.

R-Mod is a very useful case study. Not only does the tensor product give it a monoidal structure, but every R-Mod is enriched over $\mathbb{Z}$-$\mathsf{Mod} = \mathsf{Ab}$ in a manner compatible with the monoidal structure on Ab: the composition map $\circ_{XYZ} : R(X,Y) \times R(Y,Z) \to R(X,Z)$ is a bilinear map in Ab, and hence can be reduced to a single arrow $R(X,Y) \otimes_{\mathbb{Z}} R(Y,Z) \to R(X,Z)$. So, hom-sets in R-Mod are objects in Ab, and composition in R-Mod is described by morphisms in Ab in a manner compatible with Ab's monoidal structure. In general, any category C whose objects and morphisms can be described by a "monoidal category" V in a similar manner is said to be enriched over V.

Our discussion of monoidal categories and enrichment is based largely off of [Mac Lane, 2013,Fong and Spivak, 2018,Riehl, 2014], with extra details pertaining to structures in monoidal categories based off of Coecke's articles [Coecke, 2010, Abramsky and Coecke, 2009].

1.2.1 Definitions

Monoidal Categories A ***monoidal category*** is a category C equipped with a functor $\otimes : \mathsf{C} \times \mathsf{C} \to \mathsf{C}$ known as the tensor, a specific object I known as the unit, and a set of natural isomorphisms:

- $\alpha : -_1 \otimes (-_2 \otimes -_3) \Rightarrow (-_1 \otimes -_2) \otimes -_3$, known as the associator.

- $\lambda : I \otimes - \Rightarrow -$ and $\rho : - \otimes I \Rightarrow -$, known as the left and right unitors.

- Optionally, $\sigma : -_1 \otimes -_2 \Rightarrow -_2 \otimes -_1$, known as the commutator, in which case C is known as *symmetric* monoidal.

We require that the following diagrams commute: first, the pentagon identity

$$\begin{tikzcd}
((W \otimes X) \otimes Y) \otimes Z \arrow[r, "\alpha_{W\otimes X,Y,Z}"] \arrow[d, "\alpha_{W,X,Y}\otimes \mathrm{id}_Z"'] & (W \otimes X) \otimes (Y \otimes Z) \arrow[r, "\alpha_{W,X,Y\otimes Z}"] & W \otimes (X \otimes (Y \otimes Z)) \\
(W \otimes (X \otimes Y)) \otimes Z \arrow[rr, "\alpha_{W,X\otimes Y,Z}"'] & & W \otimes ((X \otimes Y) \otimes Z) \arrow[u, "\mathrm{id}_W \otimes \alpha_{X,Y,Z}"']
\end{tikzcd}$$

and then the triangle identity

$$\begin{tikzcd}
& X \otimes Y & \\
(X \otimes I) \otimes Y \arrow[ur, "\rho_X \otimes \mathrm{id}_Y"] \arrow[rr, "\alpha_{X,I,Y}"'] & & X \otimes (I \otimes Y) \arrow[ul, "\mathrm{id}_X \otimes \lambda_Y"']
\end{tikzcd}$$

If σ exists, we demand that it satisfy the hexagon identity

$$\begin{tikzcd}
(X \otimes Y) \otimes Z \arrow[r, "\alpha_{X,Y,Z}"] \arrow[d, "\sigma_{X,Y}\otimes \mathrm{id}_Z"'] & X \otimes (Y \otimes Z) \arrow[r, "\sigma_{X,(Y\otimes Z)}"] & (Y \otimes Z) \otimes X \arrow[d, "\alpha_{Y,Z,X}"] \\
(Y \otimes X) \otimes Z \arrow[r, "\alpha_{Y,X,Z}"'] & Y \otimes (X \otimes Z) \arrow[r, "\mathrm{id}_Y \otimes \sigma_{X,Z}"'] & Y \otimes (Z \otimes X)
\end{tikzcd}$$

as well as that $\sigma_{X,Y} \circ \sigma_{Y,X} = \mathrm{id}_{X\otimes Y}$. If all of these isomorphisms are in fact equalities, e.g. $(X \otimes Y) \otimes Z$ and $X \otimes (Y \otimes Z)$ are always the exact same object, then C is known as *strong* monoidal.

Two common families of monoidal categories are the *cartesian* monoidal categories, those of the form $(\mathsf{C}, \times, 1)$, and *cocartesian* monoidal categories, those of the form $(\mathsf{C}, \amalg, 0)$.

Monoidal Functors A functor F between symmetric monoidal categories $(\mathsf{C}, \otimes, I)$ and $(\mathsf{D}, \otimes, I)$ is monoidal if there is a natural isomorphism $\Phi : (F-_1) \otimes (F-_2) \Rightarrow F(-_1 \otimes -_2)$ as well as an isomorphism $\phi : I_D \to FI_C$, such that the following diagrams commute for all elements:

$$\begin{tikzcd}
(FX \otimes FY) \otimes FZ \arrow[r, "\Phi_{X,Y}\otimes \mathrm{id}_{FZ}"] \arrow[d, "\alpha_{FX,FY,FZ}"'] & F(X \otimes Y) \otimes FZ \arrow[r, "\Phi_{X\otimes Y,Z}"] & F((X \otimes Y) \otimes Z) \arrow[d, "F\alpha_{X,Y,Z}"] \\
FX \otimes (FY \otimes FZ) \arrow[r, "\mathrm{id}_{FX} \otimes \Phi_{Y,Z}"'] & FX \otimes F(Y \otimes Z) \arrow[r, "\Phi_{X,Y\otimes Z}"'] & F(X \otimes (Y \otimes Z))
\end{tikzcd}$$

$$\begin{tikzcd}
I_D \otimes FX \arrow[r, "\lambda_{FX}"] \arrow[d, "\phi \otimes \mathrm{id}_{FX}"'] & FX \\
FI_C \otimes FX \arrow[r, "\Phi_{I_C,X}"'] & F(I_C \otimes X) \arrow[u, "F\lambda_X"']
\end{tikzcd}
\qquad
\begin{tikzcd}
FX & FX \otimes I_D \arrow[l, "\rho_{FX}"'] \\
F(X \otimes I_C) \arrow[u, "F\rho_X"] & FX \otimes FI_C \arrow[l, "\Phi_{X,I_C}"] \arrow[u, "\mathrm{id}_{FX} \otimes \phi"']
\end{tikzcd}$$

If we want F to preserve symmetry, we require the additional diagram

$$\begin{tikzcd}
FX \otimes FY \arrow[r, "\sigma_{FX,FY}"] \arrow[d, "\Phi_{X,Y}"'] & FY \otimes FX \arrow[d, "\Phi_{Y,X}"] \\
F(X \otimes Y) \arrow[r, "F\sigma_{X,Y}"'] & F(Y \otimes X)
\end{tikzcd}$$

This is sometimes called a *strong* monoidal functor, due to Φ and ϕ being isomorphisms; a triplet (F, Φ, ϕ) in which Φ and ϕ aren't necessarily so is known as a *lax* monoidal functor.

Given monoidal functors $(F, \Phi, \phi), (G, \Psi, \psi) : \mathsf{C} \to \mathsf{D}$, a natural transformation $\alpha : F \Rightarrow G$ is a monoidal natural transformation if the following diagrams commute:

$$\begin{tikzcd}
FX \otimes FY \arrow[r, "\alpha_X \otimes \alpha_Y"] \arrow[d, "\Phi_{X,Y}"'] & GX \otimes GY \arrow[d, "\Psi_{X,Y}"] \\
F(X \otimes Y) \arrow[r, "\alpha_{X \otimes Y}"'] & G(X \otimes Y)
\end{tikzcd}
\qquad
\begin{tikzcd}
I_D \arrow[d, "\phi"'] \arrow[dr, "\psi"] & \\
FI_C \arrow[r, "\alpha_{I_C}"'] & GI_C
\end{tikzcd}$$

Categories of Monoidal Categories Let MonCat be the category of monoidal categories and (strong) monoidal functors, with the variant $\mathsf{SymMonCat}$ having symmetric monoidal categories and symmetric monoidal functors.

The structure of $\mathsf{SymMonCat}$ resembles that of Ab more than that of Cat (a review is given in [Fong and Spivak, 2019]): the terminal category $1 = \{*\}$, equipped with trivial symmetric monoidal structure, has exactly one morphism to every other symmetric monoidal category up to monoidal natural isomorphism (it sends $*$ to something isomorphic to the unit of the codomain), and exactly one morphism from every other symmetric monoidal category, and is therefore a *zero object* of $\mathsf{SymMonCat}$.

Furthermore, the categorical product $\mathsf{C} \times \mathsf{D}$ of symmetric monoidal categories has a natural (strict!) symmetric monoidal structure, with $(X, Y) \otimes (X', Y') = (X \otimes X', Y \otimes Y')$; it is not only the product in $\mathsf{SymMonCat}$ but the *coproduct* as well (for finite collections), being equipped with inclusions $X \in \mathsf{C} \mapsto (X, I_D)$, $Y \in \mathsf{D} \mapsto (I_C, Y)$.

1.2.2 Additional Structures

For this subsection, we fix a monoidal category $(\mathsf{C}, \otimes, I)$.

Closed Categories C is left (alt. right) ***closed monoidal*** if every functor of the form $X \otimes -$ (alt. $- \otimes X$) has a right adjoint, which in either case is written $[X, -]$ and known as the ***internal hom***; we think of it as the set of maps from X to Y represented as an object in C itself[5]. If C is symmetric monoidal, then left and right closedness are equivalent, and C is known simply as closed monoidal. If $\otimes = \times$, i.e. C is cartesian monoidal, then an internal hom renders C a *cartesian closed category* (CCC); in this case, the inner hom $[X, Y]$ is often known as the exponential, and written as Y^X, while the actual adjunction is known as *currying* [6].

The internal hom $[-, -]$ in a closed symmetric monoidal category is functorial in its second argument by definition, but is also functorial in its first argument; the morphism $[Y, Z] \to [X, Z]$ induced by a morphism $X \to Y$ is analogous to precomposition. Hence, the internal hom is a functor $[-, -] : \mathsf{C}^{\mathrm{op}} \times \mathsf{C} \to \mathsf{C}$.

For an object X in a symmetric monoidal C, the adjunction $X \otimes - \dashv [X, -]$ has a counit and a unit:

- The counit is a natural transformation $\mathrm{ev} : X \otimes [X, -] \Rightarrow -$, known as the *evaluation map* whose component $\mathrm{ev}_Y : X \otimes [X, Y] \to Y$ we think of as taking an element of X and a map $X \to Y$ and evaluating the map at the element (in fact, this is exactly what ev does in many examples, such as $(\mathsf{Set}, \times, 1)$).

- The unit is a natural transformation $\mathrm{coev} : - \Rightarrow [X, X \otimes -]$ known as the *coevaluation map* whose component $\mathrm{coev}_Y : Y \to [X, X \otimes Y]$ we think of as sending a $y \in Y$ to the map sending $x \in X$ to $x \otimes y \in X \otimes Y$.

[5]This happens in a great deal of cases: the set of maps between two topological spaces can itself be given a topology, the set of maps between two sets is a set (surprise!), the set of maps between two R-modules...

[6]In computer science, currying is the partial evaluation of functions, e.g. taking the binary function $f : X \times Y \to Z$ and plugging in a fixed x to get a unary function $f_{x_0} : Y \to Z, y \mapsto f(x, y)$; this operation is itself a function $\mathrm{Hom}(X \times Y, Z) \to \mathrm{Hom}(X, \mathrm{Hom}(Y, Z)), \lambda x, y.f(x, y) \mapsto \lambda x.\,(\lambda y.f(x, y))$.

When C has terminal objects and binary products, the category of presheaves $\widehat{\mathsf{C}}$ is cartesian closed: finite products are computed pointwise, and the exponential Q^P is given by

$$Q^P(X) = \mathrm{Hom}_{\widehat{\mathsf{C}}}(h^X \times P, Q)$$

The evaluation counit $\mathrm{ev}_X : \mathrm{Hom}_{\widehat{\mathsf{C}}}(h^X \times P, Q) \times P(X) \to Q(X)$ sends a natural transformation $\alpha : h_X \times P \to Q$ and an element $p \in P(X)$ to $\alpha(\mathrm{id}_X, p) \in Q(X)$.

Local Cartesian Closure We call C ***locally cartesian closed*** if for every $W \in \mathsf{C}$ the slice category C_W is cartesian closed. Equivalently, C is locally cartesian closed if it has pullbacks and every $f^* : \mathsf{C}_Y \to \mathsf{C}_X$ induced by a morphism $f : X \to Y$ (which sends a morphism $g : Z \to Y$ to its pullback along f) has a right adjoint $\Pi_f : \mathsf{C}_X \to \mathsf{C}_Y$. This is known as the ***dependent product***.

It is already true that f^* has a left adjoint, known as the dependent sum $\Sigma_f : \mathsf{C}_X \to \mathsf{C}_Y$; this is simply $f \circ -$, as is not hard to show. Hence, in a locally cartesian closed category we have an adjoint triple

$$\begin{tikzcd} \mathsf{C}_X \arrow[rr, bend left=50, "\Sigma_f"] \arrow[rr, bend right=50, "\Pi_f"'] \arrow[rr, phantom, shift left=4, "\perp"] \arrow[rr, phantom, shift right=4, "\perp"] & & \mathsf{C}_Y \arrow[ll, "f^*"'] \end{tikzcd}$$

associated to every $f : X \to Y$.

If C has a terminal object, then since $\mathsf{C}_* \cong \mathsf{C}$, C itself is cartesian closed.

1.2.3 Enriched Categories

Fix a symmetric monoidal category $(\mathsf{V}, \otimes, I)$. A V-***category***, or category ***enriched*** over the base category V, is a collection $\mathsf{C} = \{X_\lambda\}$ of objects, along with the following data:

- For each pair $X, Y \in \mathsf{C}$, an object $\mathrm{Hom}_{\mathsf{C}}(X, Y) \in \mathsf{V}$ known as the ***hom-object***.
- For each $X \in \mathsf{C}$, a morphism $\mathrm{id}_X : I \to \mathrm{Hom}_{\mathsf{C}}(X, X)$ representing the identity morphism.
- For each triplet $X, Y, Z \in \mathsf{C}$, a morphism $\circ_{XYZ} : \mathrm{Hom}_{\mathsf{C}}(X, Y) \otimes \mathrm{Hom}_{\mathsf{C}}(Y, Z) \to \mathrm{Hom}_{\mathsf{C}}(X, Z)$.

We require composition to be associative, in the sense that

$$\circ_{XYW} \circ (\mathrm{id}_{\mathrm{Hom}_{\mathsf{C}}(X,Y)} \otimes \circ_{YZW}) = \circ_{XZW} \circ (\circ_{XYZ} \otimes \mathrm{id}_{\mathrm{Hom}_{\mathsf{C}}(Z,W)})$$

for all X, Y, Z, W, and we require the identity to play nicely with composition in the usual sense, for which we require

$$\circ_{XXY} \circ (\mathrm{id}_{\mathsf{C}(X,Y)} \otimes \mathrm{id}_X) = \rho_{\mathrm{Hom}_{\mathsf{C}}(X,Y)}$$

and

$$\circ_{YYX} \circ (\mathrm{id}_Y \otimes \mathrm{id}_{\mathrm{Hom}_{\mathsf{C}}(X,Y)}) = \lambda_{\mathrm{Hom}_{\mathsf{C}}(X,Y)}$$

where ρ and λ are the right and left unitors.

Enriched Functors Given two V-categories C, D, a V-***functor*** $F : \mathsf{C} \to \mathsf{D}$ is a map on objects $X \mapsto FX$ along with, for every $\mathrm{Hom}_{\mathsf{C}}(X,Y) \in \mathsf{V}$, a morphism in V, $F_{X,Y} : \mathrm{Hom}_{\mathsf{C}}(X,Y) \to \mathrm{Hom}_{\mathsf{D}}(FX, FY)$. We require that these morphisms commute with composition morphisms in V, in the sense that

$$\circ^{\mathsf{D}}_{FX,FY,FZ} \circ (F_{Y,Z} \times F_{X,Y}) = F_{X,Z} \circ \circ^{\mathsf{C}}_{X,Y,Z}$$

and we require that the identity map $I \to \mathrm{Hom}_{\mathsf{C}}(X,X)$ composed with $F_{X,X}$ be equal to the identity map $I \to \mathrm{Hom}_{\mathsf{D}}(FX, FX)$.

A V-***natural transformation*** $\alpha : F \to G$ between V-enriched $F, G : \mathsf{C} \to \mathsf{D}$ is defined in the usual way, as a family of morphisms $\alpha_X : I \to \mathrm{Hom}_{\mathsf{D}}(FX, GX)$, but we require $(\alpha_Y)_* \circ F_{X,Y} = (\alpha_X)^* \circ G_{X,Y}$. A V-adjunction $F : \mathsf{C} \to \mathsf{D} \dashv G : \mathsf{D} \to \mathsf{C}$ is a natural isomorphism $\mathsf{D}(F-,-) \cong \mathrm{Hom}_{\mathsf{C}}(-, G-)$, or equivalently V-natural transformations $\eta : \mathrm{id} \to GF$ and $\varepsilon : FG \to \mathrm{id}$ satisfying the zig-zag identities.

The set of V-categories along with V-functors forms a category of V-enriched categories, which we will call V-Cat. We may construct a functor $(-)_0 : \mathsf{V}\text{-}\mathsf{Cat} \to \mathsf{Cat}$ sending a V-category C to the ordinary category C_0 which has all the same objects as C, but whose morphisms $X \to Y$ are morphisms $I \to \mathrm{Hom}_{\mathsf{C}}(X,Y)$ in V. Composition of morphisms $f : X \to Y$, $g : Y \to Z$ in C_0

is given by the morphism

$$I \cong I \otimes I \xrightarrow{f \otimes g} \mathrm{Hom}_{\mathsf{C}}(X,Y) \otimes \mathrm{Hom}_{\mathsf{C}}(Y,Z) \xrightarrow{\circ_{XYZ}} \mathrm{Hom}_{\mathsf{C}}(X,Z)$$

The Enriched Yoneda Lemma When V is symmetric monoidal and *closed*, we can consider it as a V-category $\underline{\mathsf{V}}$:

- The hom is given by $\mathrm{Hom}_{\underline{\mathsf{V}}}(X,Y) = [X,Y]$
- The composition $\circ_{XYZ} : [X,Y] \otimes [Y,Z] \to [X,Z]$ is given by the adjunct of the morphism

$$(X \otimes [X,Y]) \otimes [Y,Z] \xrightarrow{\mathrm{ev}_X \otimes \mathrm{id}_{[Y,Z]}} Y \otimes [Y,Z] \xrightarrow{\mathrm{ev}_Y} Z$$

- The identity elements $\mathrm{id}_X : I \to [X,X]$ are the adjuncts of the left unitors $\lambda_X : I \otimes X \to X$

With this special V-category, we can express the representable functors as enriched functors: for any V-category C and any $X \in \mathsf{C}$, there is a V-functor $\mathrm{Hom}_{\mathsf{C}}(X,-) : \mathsf{C} \to \underline{\mathsf{V}}$ sending each $Y \in \mathsf{C}$ to $\mathrm{Hom}_{\mathsf{C}}(X,Y) \in \underline{\mathsf{V}}$, and equipped with $\underline{\mathsf{V}}$-morphisms

$$\mathrm{Hom}_{\mathsf{C}}(X,-)_{Y,Z} : \mathrm{Hom}_{\mathsf{C}}(Y,Z) \to [\mathrm{Hom}_{\mathsf{C}}(X,Y), \mathrm{Hom}_{\mathsf{C}}(X,Z)]$$

adjunct to the composition morphisms $\circ_{XYZ}$ of C. Similarly, the V-functors $\mathrm{Hom}_{\mathsf{C}}(-,X)$ send $Y \in \mathsf{C}$ to $\mathrm{Hom}_{\mathsf{C}}(Y,X) \in \underline{\mathsf{V}}$, and their V-morphisms are those adjunct to $\circ_{YZX}$. Hence, the functor $\mathrm{Hom}_{\mathsf{C}} : \mathsf{C}_0^{\mathrm{op}} \times \mathsf{C}_0 \to \mathsf{V}$ is V-functorial in both arguments.

Take a V-natural transformation α from $\mathrm{Hom}_{\mathsf{C}}(X,-)$ to a V-functor $F : \mathsf{C} \to \underline{\mathsf{V}}$; this is a family of morphisms $\alpha_Y : I \to [\mathrm{Hom}_{\mathsf{C}}(X,Y), FY]$ in V, which by adjunction is equivalently a family of morphisms $\tilde{\alpha}_Y : \mathrm{Hom}_{\mathsf{C}}(X,Y) \to FY$. The composition

$$I \xrightarrow{\mathrm{id}_X} \mathrm{Hom}_{\mathsf{C}}(X,X) \xrightarrow{\tilde{\alpha}_X} FX$$

sends α to a morphism $I \to FX$ in $\underline{\mathsf{V}}$, i.e. an element of FX. The ***enriched Yoneda lemma*** says that this operation is a *bijection* between V-natural transformations $\mathrm{Hom}_{\mathsf{C}}(X,-) \Rightarrow F$ and elements of FX. The proof relies on some technicalities we have neglected to mention here, but which can be found in the first chapter of [Kelly and Kelly, 1982].

1.2.4 2-Categories

Definition and Consequences A 2-***category*** is a Cat-category, or a category enriched over the (cartesian closed) category of small categories[7]. Hence, for every X, Y in a 2-category C, $\mathrm{Hom}_{\mathsf{C}}(X, Y)$ is itself a category, which we will denote $\mathsf{C}(X, Y)$. Objects of this category are 1-morphisms, and morphisms of this category, or morphisms between 1-morphisms, are 2-morphisms; they are denoted as double-tailed arrows $\Rightarrow$, just as natural transformations (the 2-morphisms in Cat when considered as enriched over itself). Composition of morphisms is a functor $\circ_{XYZ} : \mathsf{C}(Y, Z) \times \mathsf{C}(X, Y) \to \mathsf{C}(X, Z)$, and each hom-category $\mathsf{C}(X, X)$ is equipped with an object id_X serving the role of the identity morphism on X.

There is a associativity identity for composition

$$\begin{tikzcd}
\mathsf{C}(Y,Z) \times \mathsf{C}(X,Y) \times \mathsf{C}(W,X) \arrow[r, "\mathrm{id}_{\mathsf{C}(Y,Z)} \times \circ_{WXY}"] \arrow[d, "\circ_{XYZ} \times \mathrm{id}_{\mathsf{C}(W,X)}"'] & \mathsf{C}(Y,Z) \times \mathsf{C}(W,Y) \arrow[d, "\circ_{WYZ}"] \\
\mathsf{C}(X,Z) \times \mathsf{C}(W,X) \arrow[r, "\circ_{WXZ}"'] & \mathsf{C}(W,Z)
\end{tikzcd}$$

$$\circ_{WXZ} \circ (\circ_{XYZ} \times \mathrm{id}_{\mathsf{C}(W,X)}) = \circ_{WYZ} \circ (\circ_{WXY} \times \mathrm{id}_{\mathsf{C}(Y,Z)})$$

for triplets of morphisms $W \xrightarrow{f} X \xrightarrow{g} Y \xrightarrow{h} Z$.

Because composition is now a functor, it is capable of acting on maps as well, immediately leading to new constructions. For instance, $\circ_{XYZ}$ will act not only on a pair of 1-morphisms, or an object $(g, f) \in \mathsf{C}(Y, Z) \times \mathsf{C}(X, Y)$, but on a morphism of the product category, say $(\beta : g \Rightarrow g', \alpha : f \Rightarrow f')$, to yield a 2-morphism $\beta * \alpha : gf \Rightarrow g'f'$. This operation is known as ***horizontal composition***. In the case that $\alpha = \mathrm{id}_f$, we simply write the 2-morphism as βf, and call it composition by ***whiskering***:

$$\begin{tikzcd}
X \arrow[r, "f"] & Y \arrow[r, bend left, "g"] \arrow[r, bend right, "g'"'] \arrow[r, phantom, "\Downarrow\beta"] & Z
\end{tikzcd}
\qquad\qquad
\begin{tikzcd}
X \arrow[r, bend left, "gf"] \arrow[r, bend right, "g'f"'] \arrow[r, phantom, "\Downarrow\beta f"] & Z
\end{tikzcd}$$

[7]Technically, this is a *strict* 2-category, to be distinguished from the *weak* 2-categories where properties of 1-morphisms only hold up to 2-isomorphism rather than on the nose. The term 2-category is often used in the latter sense.

This is associative, as is made clear by the composition associativity identity:

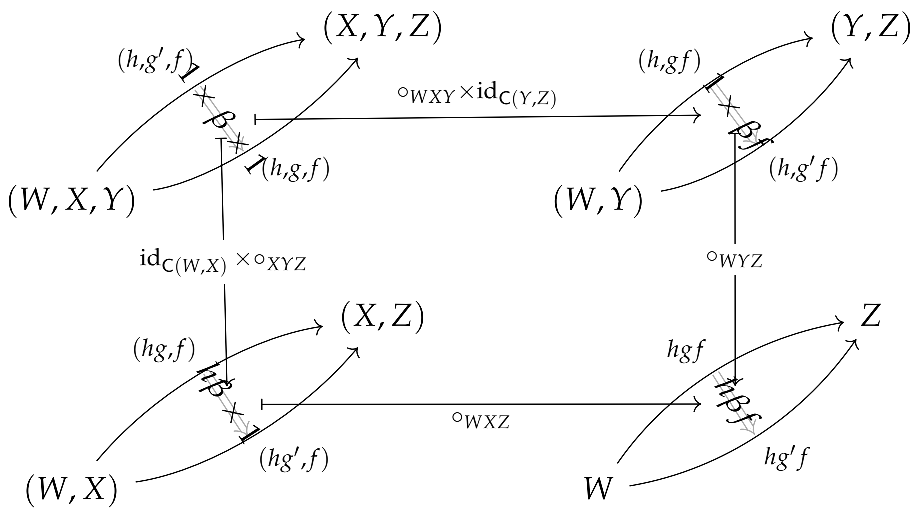

In particular, if given two pairs of 1-morphisms $f, f' : X \to Y$, $g, g' : Y \to Z$ and a pair of 2-morphisms $\alpha : f \Rightarrow f', \beta : g \Rightarrow g'$, whiskering gf along β and then along α will yield the same result as whiskering it along α and then along β. In other words, we have a commutative diagram:

$$\begin{tikzcd} gf \arrow[r, Rightarrow, "\beta f"] \arrow[d, Rightarrow, "g\alpha"'] & g'f \arrow[d, Rightarrow, "g'\alpha"] \\ gf' \arrow[r, Rightarrow, "\beta f'"'] & g'f' \end{tikzcd}$$

There is also a diagram verifying that the identity acts as it should:

$$\begin{tikzcd} * \times \mathsf{C}(X,Y) \arrow[r, equal] \arrow[d, "\mathrm{id}_X \times \mathrm{id}_{\mathsf{C}(X,Y)}"'] & \mathsf{C}(X,Y) \arrow[r, equal] & \mathsf{C}(X,Y) \times * \arrow[d, "\mathrm{id}_{\mathsf{C}(X,Y)} \times \mathrm{id}_X"] \\ \mathsf{C}(X,X) \times \mathsf{C}(W,X) \arrow[ur, "\circ_{WXX}"'] & & \mathsf{C}(X,Y) \times \mathsf{C}(X,X) \arrow[ul, "\circ_{XXY}"] \end{tikzcd}$$

2-Functors When presenting a kind of mathematical object as a collection of types of data subject to certain coherence conditions, the categorical notion of a morphism between objects will be a collection of maps between each type of data which collection preserves all coherence conditions. This is sufficient for instance to come to the proper notions of a functor between categories and a natural transformation between functors, but one can go no further: a morphism between natural transformations $\alpha, \beta : f \Rightarrow g : F \to G$ would have to be either a simple, boring *identification*, or a map between morphisms f_X and $g_X : FX \to GX$, but without a notion of morphism between morphisms there is no natural definition. So we only have two "levels" of morphisms: functors and natural transformations.

2-categories provide us with such a notion, allowing us to define a morphism of natural transformations, and hence find a third level: *modifications* are, roughly, families of 2-morphisms between the families of 1-morphisms defining natural transformations. However, 2-categories have more data than 1-categories, and we have to upgrade the notions of functor and natural transformation to accommodate this.

A ***2-functor*** $F : \mathsf{C} \to \mathsf{D}$ is a mapping of objects of C to objects of D along with, for every pair $X, Y \in \mathsf{C}$, a functor $F_{XY} : \mathsf{C}(X, Y) \to \mathsf{D}(FX, FY)$. We have $F_{XX}\mathrm{id}_X = \mathrm{id}_{FX}$ and $F_{XZ} \circ \circ^{\mathsf{C}}_{XYZ} = \circ^{\mathsf{D}}_{FX,FY,FZ} \circ (F_{YZ} \times F_{XY})$; the second diagram implies functoriality for not only horizontal but vertical 2-morphism compositions.

A ***2-natural transformation*** between 2-functors $F, G : \mathsf{C} \to \mathsf{D}$ is a family of 1-morphisms $\kappa_X \in \mathsf{D}(FX, GX)$ which are not only natural on 1-morphisms $g : X \to Y$ in the sense that $\kappa_Y \circ Fg = Gg \circ \kappa_X$, but natural on 2-morphisms $\alpha : g \Rightarrow h$ between objects X and Y in C in the sense that $(G\alpha)\kappa_X = \kappa_Y(F\alpha)$.

A ***modification*** Ξ between 2-natural transformations $\mu, \nu : F \Rightarrow G : \mathsf{C} \to \mathsf{D}$, written as $\Xi : \mu \Rrightarrow \nu$, is a C-indexed family of 2-morphisms $\Xi_X : \mu_X \to \nu_X$ natural with respect to whiskering, in that $(Gf)\Xi_X = \Xi_Y(Ff)$ for all 1-morphisms $f : X \to Y$, and natural with respect to horizontal composition, in that $(G\alpha) * \Xi_X = \Xi_Y * (F\alpha)$ for all 2-cells $\alpha : f \Rightarrow g : X \to Y$.

Just as we can obtain a new 1-category from the functors and natural transformations between two 1-categories, we can obtain a new 2-category from the 2-functors, 2-natural transformations, and modifications between two 2-categories. In fact, the category of 2-categories, 2-Cat, is cartesian closed; hence, not only is it enriched over itself, but it is possible to enrich other categories over 2-Cat [8]. These are (strict) *3-categories*, the first example of which is the category 3-Cat = 2-Cat-Cat of 2-Cat-enriched categories.

[8] Furthermore, there's a 2-categorical version of the Yoneda lemma and embedding, which gives rise to 2-presheaves and a theory of Grothendieck 2-topoi and so on. We will not study these, preferring instead to skip straight to ∞-topoi.

Adjunctions Recall the counit-unit definition of an adjunction of functors: for $L : \mathsf{C} \to \mathsf{D}$ and $R : \mathsf{D} \to \mathsf{C}$, we write $L \dashv R$ if there are natural transformations $\epsilon : LR \Rightarrow \mathrm{id}_{\mathsf{D}}$ and $\eta : \mathrm{id}_{\mathsf{C}} \Rightarrow RL$ such that $(\epsilon L) \circ (L\eta) = \mathrm{id}_L$ and $(R\epsilon) \circ (\eta R) = \mathrm{id}_R$. This definition, being based entirely on the 1 and 2-morphisms of Cat, easily generalizes to an arbitrary category.

Given 1-morphisms $l : X \to Y$ and $r : Y \to X$ in a 2-category C, we say that l is left adjoint to r and r right adjoint to l, again written $l \dashv r$, if there exist 2-morphisms $\eta : \mathrm{id}_X \Rightarrow rl$ and $\epsilon : lr \Rightarrow \mathrm{id}_Y$ such that $(\epsilon l) \circ (l\eta) = \mathrm{id}_l$ and $(r\epsilon) \circ (\eta r) = \mathrm{id}_r$.

This is often expressed as an equality between the diagrams on the left and right:

$$\begin{tikzcd} & Y \arrow[rr, equal] \arrow[dr, "r"] & & Y & & Y \\ X \arrow[ur, "l"] \arrow[rr, equal] \arrow[r, phantom, "\Uparrow\eta"] & & X \arrow[ur, "l"'] \arrow[u, phantom, "\Uparrow\epsilon"] & & & X \arrow[u, bend left, "l"] \arrow[u, bend right, "l"'] \arrow[u, phantom, "\overset{\mathrm{id}_l}{=}"] \end{tikzcd}$$

$$\begin{tikzcd} & X \arrow[rr, equal] \arrow[dr, "l"] & & X & & X \\ Y \arrow[ur, "r"] \arrow[rr, equal] \arrow[r, phantom, "\Downarrow\epsilon"] & & Y \arrow[ur, "r"'] \arrow[u, phantom, "\Downarrow\eta"] & & & Y \arrow[u, bend left, "r"] \arrow[u, bend right, "r"'] \arrow[u, phantom, "\overset{\mathrm{id}_r}{=}"] \end{tikzcd}$$

The relations between the elements of this diagram are preserved by 2-functors, implying that 2-functors between 2-categories preserve adjunctions. Furthermore, given morphisms $l : X \to Y, l' : Y \to Z, r : Y \to X, r' : Z\ toY$ such that $l \dashv r$ and $l' \dashv r'$, pasting the corresponding diagrams together shows that $r' \circ r \dashv l' \circ l$ as well; that is, adjunctions compose.

There is a notion of morphisms between adjunctions: given adjunctions and morphisms of the form

$$\begin{tikzcd} X \arrow[r, bend left, "l"] \arrow[r, phantom, "\perp"] \arrow[d, bend right, "x"'] & Y \arrow[l, bend left, "r"] \arrow[d, bend left, "y"] \\ X' \arrow[r, bend left, "l'"] \arrow[r, phantom, "\perp"] & Y' \arrow[l, bend left, "r'"] \end{tikzcd}$$

we can define a notion of a 2-morphism going from $l \dashv r$ to $l' \dashv r'$. There seems to be either a left way to do this, as a 2-morphism $\alpha : yl \Rightarrow l'x$, or a right way, as a 2-morphism $xr \Rightarrow r'y$. However, by pre- and post-composing with the unit and counit, we obtain a one-to-one correspondence sending a 2-morphism going in one direction to its *mate* going in the other direction; a 2-morphism $l \dashv r \to l' \dashv r'$ is then a mate pair. This allows us to associate to C a new 2-category of adjunctions in C, denoted $\mathrm{Adj}(\mathsf{C})$, with objects given by objects of C, 1-morphisms $X \to Y$ given by, say, left adjoints $X \to Y$, and 2-morphisms between adjunctions given by mate pairs.

1.2.5 Ends and Coends

CITE: (Co)end Calculus

Dinatural Transformations Given a pair of functors $F, G : \mathsf{C}^{\mathrm{op}} \times \mathsf{C} \to \mathsf{D}$, a natural transformation $\alpha : F \Rightarrow G$ is simply a class of D-arrows $\{\alpha_{X,Y} : F(X,Y) \to G(X,Y)\}_{X,Y \in \mathsf{C}}$ such that, for a pair $f : X' \to X, g : Y \to Y'$ in C, we have $\alpha_{X',Y'} \circ F(f,g) = G(f,g) \circ \alpha_{X,Y}$. In other words, the following diagram commutes:

$$\begin{tikzcd}
F(X,Y) \arrow[r, "\alpha_{X,Y}"] \arrow[d, "F(f,g)"'] & G(X,Y) \arrow[d, "G(f,g)"] \\
F(X',Y') \arrow[r, "\alpha_{X',Y'}"] & G(X',Y')
\end{tikzcd}$$

A ***dinatural transformation*** $F \overset{..}{\Rightarrow} G$, however, is a family $\{\alpha_X : F(X,X) \to G(X,X)\}_{X \in \mathsf{C}}$ such that the following hexagon commutes for any arrow $f : X \to Y$:

$$\begin{tikzcd}
 & F(X,X) \arrow[r, "\alpha_X"] & G(X,X) \arrow[dr, "G(X,f)"] & \\
F(Y,X) \arrow[ur, "F(f,X)"] \arrow[dr, "F(Y,f)"'] & & & G(X,Y) \\
 & F(Y,Y) \arrow[r, "\alpha_Y"'] & G(Y,Y) \arrow[ur, "G(f,Y)"'] &
\end{tikzcd}$$

where $F(f,X) : F(Y,X) \to F(X,X)$ is the arrow obtained by identifying $F : \mathsf{C}^{\mathrm{op}} \times \mathsf{C} \to \mathsf{D}$ with $F' : \mathsf{C} \to [\mathsf{C}^{\mathrm{op}}, D]$ as $F'(X)(-) = F(-,X)$, or equivalently just $F(f, \mathrm{id}_X)$, and likewise for the other arrows.

Denoting by $\Delta_Z : \mathsf{C}^{\mathrm{op}} \times \mathsf{C} \to \mathsf{D}$ the constant functor $(X, Y) \mapsto Z$, we define a ***wedge*** for G to be a dinatural transformation $\Delta_Z \Rightarrow G$, i.e. a family of morphisms $\alpha_X : Z \to G(X, X)$ such that $G(X, f) \circ \alpha_X = G(f, Y) \circ \alpha_Y$ for all $f : X \to Y$. Given two wedges $\alpha : \Delta_Z \Rightarrow G$ and $\beta : \Delta_{Z'} \Rightarrow G$, a morphism $k : Z \to Z'$ suffices to define a morphism of wedges $\alpha \to \beta$ if $\alpha_X = \beta_X \circ k$ for all $X \in \mathsf{C}$. This gives us a category $\mathsf{Wd}(G)$ of wedges for G.

Dually, a ***cowedge*** for a functor $F : \mathsf{C}^{\mathrm{op}} \times \mathsf{C} \to \mathsf{D}$ is a dinatural transformation $F \Rightarrow \Delta_Z$, or a family of morphisms $\alpha_X : F(X, X) \to Z$ such that $\alpha_X \circ F(f, X) = \alpha_Y \circ F(Y, f)$ for all $f : X \to Y$. Defining morphisms of cowedges similarly: $k : Z \to Z'$ defines a morphism $\alpha \to \beta$ if $k \circ \alpha_X = \beta_X$ for all $X \in \mathsf{C}$, we have a category $\mathsf{Cwd}(F)$ of cowedges for F.

$$\begin{tikzcd}
 & G(X,X) \arrow[dr, "G(X,f)"] & \\
Z \arrow[ur, "\alpha_X"] \arrow[dr, "\alpha_Y"'] & & G(X,Y) \\
 & G(Y,Y) \arrow[ur, "G(f,Y)"'] &
\end{tikzcd}
\qquad
\begin{tikzcd}
 & F(X,X) \arrow[dr, "\alpha_X"] & \\
F(Y,X) \arrow[ur, "F(f,X)"] \arrow[dr, "F(Y,f)"'] & & Z \\
 & F(Y,Y) \arrow[ur, "\alpha_Y"'] &
\end{tikzcd}$$

A natural transformation $\varphi : F \Rightarrow G$ of functors $\mathsf{C}^{\mathrm{op}} \times \mathsf{C} \to \mathsf{D}$ yields a functor $\mathsf{Wd}(F) \to \mathsf{Wd}(G)$ as follows: a wedge $\alpha : \Delta_Z \Rightarrow F$ extends to a wedge $(\varphi\alpha)_X = \varphi_{X,X} \circ \alpha_X$, and we can verify that

$$G(X, f) \circ \varphi_{X,X} \circ \alpha_X = \varphi_{X,Y} \circ F(X, f) \circ \alpha_X = \varphi_{X,Y} \circ F(f, Y) \circ \alpha_Y = G(f, Y) \circ \varphi_{Y,Y} \circ \alpha_Y$$

In other words, the large diamond below commutes because not only does the small diamond commute by definition, but the outer trapezoids commute by naturality.

$$
\begin{tikzcd}
 & & G(X,X) \arrow[ddr, "G(X,f)"] & \\
 & F(X,X) \arrow[ur, "\varphi_{X,X}"] \arrow[dr, "F(X,f)"] & & \\
Z \arrow[ur, "\alpha_X"] \arrow[dr, "\alpha_Y"'] & & F(X,Y) \arrow[r, "\varphi_{X,Y}"] & G(X,Y) \\
 & F(Y,Y) \arrow[ur, "F(f,Y)"] \arrow[dr, "\varphi_{Y,Y}"'] & & \\
 & & G(Y,Y) \arrow[uur, "G(f,Y)"'] &
\end{tikzcd}
$$

Definitely if a morphism $\alpha \to \beta$ of wedges is induced by a $k : Z \to Z'$, then $\varphi_{X,X} \circ \alpha_X = \varphi_{X,X} \circ \beta_X \circ k$ for all $X \in \mathsf{C}$, so the map is functorial. Dually, the natural transformation $\varphi : F \Rightarrow G$ yields a functor $\mathsf{Cwd}(G) \to \mathsf{Cwd}(F)$, $(\alpha\varphi)_X = \alpha_X \circ \varphi_{X,X}$.

Coend Calculus The ***end*** of a functor $F : \mathsf{C}^{\mathrm{op}} \times \mathsf{C} \to \mathsf{D}$ is, if it exists, the (unique up to isomorphism) terminal object of $\mathsf{Wd}(F)$. This object, $\mathrm{end}(F)$, includes in its definition an object $Z \in \mathsf{D}$, which we often call the end itself. The ***coend*** of F is the initial object $\mathrm{coend}(F)$ of $\mathsf{Cwd}(F)$. Taking ends yields a functor $[\mathsf{C}^{\mathrm{op}} \times \mathsf{C}, \mathsf{D}] \to \mathsf{D}$, and likewise taking coends yields a functor $[\mathsf{C}^{\mathrm{op}} \times \mathsf{C}, \mathsf{D}]^{\mathrm{op}} \to \mathsf{D}$.

The integral notation for (co)ends is defined as follows:

$$\mathrm{end}(F) = \int_{\mathsf{X} \in \mathsf{C}} F(X,X) \qquad \mathrm{coend}(X) = \int^{X \in \mathsf{C}} F(X,X)$$

though we often just write $\int_X F$ and $\int^X F$. The operation of end on a natural transformation $\varphi : F \Rightarrow G$ is written as $\int_X \varphi : \int_X F \to \int_X G$.

We state without proof the basic rules of the *(co)end calculus*, a set of rules for manipulating and calculating coends.

(Commutativity with Hom) Let $F : \mathsf{C}^{\mathrm{op}} \times \mathsf{C} \to \mathsf{D}$, and $Z \in \mathsf{D}$. We have

$$\mathrm{Hom}_{\mathsf{D}}\left(\int^{X} F(X,X), Z\right) \cong \int_{X} \mathrm{Hom}_{\mathsf{D}}(F(X,X), Z)$$

and

$$\mathrm{Hom}_{\mathsf{D}}\left(Z, \int^{X} F(X,X)\right) \cong \int_{X} \mathrm{Hom}_{\mathsf{D}}(Z, F(X,X))$$

(Commutativity with limit preserving functors in general) If $F : \mathsf{C}^{\mathrm{op}} \times \mathsf{C} \to \mathsf{D}$, and $G : \mathsf{D} \to \mathsf{E}$ commutes with limits, then

$$G \int_{X} F(X,X) = \int_{X} GF(X,X)$$

(Fubini rule) Let $F : (\mathsf{C} \times \mathsf{E})^{\mathrm{op}} \times (\mathsf{C} \times \mathsf{E}) \cong \mathsf{C}^{\mathrm{op}} \times \mathsf{C} \times \mathsf{E}^{\mathrm{op}} \times \mathsf{E} \to \mathsf{D}$, with dummy variables $X \in \mathsf{C}$ and $Y \in \mathsf{E}$. Then

$$\int_{(X,Y)} F(X,X,Y,Y) \cong \int_{X} \int_{Y} F(X,X,Y,Y) \cong \int_{Y} \int_{X} F(X,X,Y,Y)$$

and likewise for coends.

(Natural transformations) Let F, G be functors $\mathsf{C} \to \mathsf{D}$. Then

$$\int_{X} \mathrm{Hom}_{\mathsf{D}}(FX, GX) \cong \mathrm{Nat}(F,G)$$

(Ninja Yoneda lemma) Let $F : \mathsf{C}^{\mathrm{op}} \to \mathsf{Set}$, and $G : \mathsf{C} \to \mathsf{Set}$. Then

$$\int^{Y} FY \times \mathrm{Hom}_{\mathsf{C}}(X,Y) \cong \int_{Y} [FY, \mathrm{Hom}_{\mathsf{C}}(Y,X)] \cong FX$$

$$\int^{Y} GY \times \mathrm{Hom}_{\mathsf{C}}(Y,X) \cong \int_{Y} [GY, \mathrm{Hom}_{\mathsf{C}}(X,Y)] \cong GX$$

Let C be a V-category, for instance with $\mathsf{V} = \mathsf{Set}$ (so that C is locally small). ***Tensoring***, or ***copowering***, is a functor $\otimes : \mathsf{V} \times \mathsf{C} \to \mathsf{C}$, traditionally written infix, such that for $A \in \mathsf{V}$, $X \in \mathsf{C}$, $\mathrm{Hom}_{\mathsf{C}}(A \otimes X, Y) \cong \mathrm{Hom}_{\mathsf{V}}(A, \mathrm{Hom}_{\mathsf{C}}(X,Y))$, this isomorphism being natural in A, X, and Y. ***Cotensoring***, or ***powering***, is a functor $\pitchfork : \mathsf{V}^{\mathrm{op}} \times \mathsf{C} \to \mathsf{C}$ such that $\mathrm{Hom}_{\mathsf{C}}(X, A \pitchfork Y) \cong \mathrm{Hom}_{\mathsf{V}}(A, \mathrm{Hom}_{\mathsf{C}}(X,Y))$, these isomorphisms again being natural.

In the case that $\mathsf{V} = \mathsf{Set}$ and the requisite limits exist in C, it is always possible to choose the following tensor and cotensor:

$$A \pitchfork X = \prod_A X \qquad A \otimes X = \coprod_A X$$

1.2.6 Kan Extensions

Let $F : \mathsf{C} \to \mathsf{D}$ and $G : \mathsf{C} \to \mathsf{E}$. In many natural cases we would like to extend G along F to obtain a functor $H : \mathsf{D} \to \mathsf{E}$ such that $HF = G$, though this isn't always possible. The next best thing is to come up with a functor K bearing some universal property with respect to F and G. There are two natural universal properties, which result in the left and right *Kan extensions*.

A ***left Kan extension*** of G along F is a functor $K : \mathsf{D} \to \mathsf{E}$ along with a natural transformation $\alpha : G \Rightarrow KF$ such that for any other $K' : \mathsf{D} \to \mathsf{E}$ and $\alpha' : G \Rightarrow K'F$, there is a unique natural transformation $\beta : K \Rightarrow K'$ such that α' factors through α and β. This pair (K, α) is unique by the universal property, and is denoted as $\mathrm{Lan}_F G$.

$$\begin{tikzcd}
\mathsf{C} \arrow[rr, "G"] \arrow[dr, "F"'] & \arrow[d, phantom, "\alpha' \Downarrow"] & \mathsf{E} \\
& \mathsf{D} \arrow[ur, "K'"'] &
\end{tikzcd}
\quad = \quad
\begin{tikzcd}
\mathsf{C} \arrow[rr, "G"] \arrow[dr, "F"'] & \arrow[d, phantom, "\alpha \Downarrow"] & \mathsf{E} \\
& \mathsf{D} \arrow[ur, "K'"] \arrow[ur, bend right=60, "K = \mathrm{Lan}_F G"'] \arrow[ur, phantom, shift right=3, "\Uparrow \beta"] &
\end{tikzcd}$$

Dually, a ***right Kan extension*** of G along F is a functor K along with a natural transformation $\alpha : KF \Rightarrow G$ universal in the dual sense.

$$\begin{tikzcd}
\mathsf{C} \arrow[rr, "G"] \arrow[dr, "F"'] & \arrow[d, phantom, "\alpha' \Uparrow"] & \mathsf{E} \\
& \mathsf{D} \arrow[ur, "K'"'] &
\end{tikzcd}
\quad = \quad
\begin{tikzcd}
\mathsf{C} \arrow[rr, "G"] \arrow[dr, "F"'] & \arrow[d, phantom, "\alpha \Uparrow"] & \mathsf{E} \\
& \mathsf{D} \arrow[ur, "K'"] \arrow[ur, bend right=60, "K = \mathrm{Ran}_F G"'] \arrow[ur, phantom, shift right=3, "\Downarrow \beta'"] &
\end{tikzcd}$$

When everything that must exist does, we can compute the left and right Kan extensions as

$$(\mathrm{Lan}_F G)Y \cong \int^X \mathrm{Hom}_{\mathsf{D}}(FX, Y) \otimes GX \qquad (\mathrm{Ran}_F G)Y \cong \int_X \mathrm{Hom}_{\mathsf{D}}(Y, FX) \pitchfork GX$$

Consider for instance the case $\mathsf{E} = \mathsf{C}$, $G = \mathrm{id}_{\mathsf{C}}$. Left Kan extension of id_{C} along F must yield a functor $K : \mathsf{D} \to \mathsf{C}$ along with a universal natural transformation $\alpha_X : X \to KFX$; K will be a right adjoint of F. Dually, $\mathrm{Ran}_F \mathrm{id}_{\mathsf{C}}$ will be a right adjoint of F if it exists. So, if $F : \mathsf{C} \to \mathsf{D}$ has a left adjoint L or a right adjoint R, we can compute L or R as

$$LY = \int^X \mathrm{Hom}_{\mathsf{D}}(FX, Y) \otimes X = \int^X \coprod_{FX \to Y} X \qquad RY = \int_X \mathrm{Hom}_{\mathsf{D}}(Y, FX) \pitchfork X = \int_X \prod_{FX \to Y} X$$

Kan Extensions as Adjoints There is another sense in which we may view Kan extensions: fix $F : \mathsf{C} \to \mathsf{D}$. For any category E, precomposition with F gives us a functor $F^* : \mathsf{E}^{\mathsf{D}} \to \mathsf{E}^{\mathsf{C}}$ sending $H : \mathsf{D} \to \mathsf{E}$ to $HF : \mathsf{C} \to \mathsf{E}$. The left Kan extension along F is, if it exists, identical to the left adjoint $L : \mathsf{E}^{\mathsf{C}} \to \mathsf{E}^{\mathsf{D}}$ of F^*, such that natural transformations $\mathrm{Lan}_F G \Rightarrow H$ are in natural bijection with natural transformations $G \Rightarrow HF$. The right Kan extension along F is the right adjoint $R : \mathsf{E}^{\mathsf{C}} \to \mathsf{E}^{\mathsf{D}}$ of F^*, with natural transformations $HF \Rightarrow G$ being in natural bijection with natural transformations $H \Rightarrow \mathrm{Ran}_F G$. The natural transformations α used in the above definition are the unit and counit of these adjunctions.

$$\mathrm{Lan}_F \dashv F^* \dashv \mathrm{Ran}_F$$

The basic idea may be expressed as follows: We have $F : \mathsf{C} \to \mathsf{D}$, and $G : \mathsf{C} \to \mathsf{E}$, and we want to find the $H : \mathsf{D} \to \mathsf{E}$ that solves the "equation" $G = HF$. So, we have to invert the act of precomposing with F! If F were an equivalence, we could do this as in basic algebra, writing $H = GF^{-1}$ (up to natural isomorphism). It usually isn't invertible, though, so we have to settle for the next best thing to the inverse of precomposing with F, which would be an adjoint to precomposition. Hence, given the equation $G = HF$ for known G and F, Kan extensions try to find the H that is closest to a solution. If we want the "most general" solution, the one that every other possible solution must necessarily factor through (such that this solution solves the "totality" of the problem), we go with the right Kan extension, and if we want the "most efficient" solution, the one that must necessarily factor through every other solution (such that this solution solves the "core" of the problem), we go with the left Kan extension.

Kan Lifts The dual to Kan extensions is given by posing the problem $G = FH$, the next best thing to a solution for which would be finding an adjoint to the *postcomposition* functor $F_* : \mathsf{C}^{\mathsf{B}} \to \mathsf{D}^{\mathsf{B}}$, which sends $G : \mathsf{B} \to \mathsf{C}$ to $FG : \mathsf{B} \to \mathsf{D}$. Its left and right adjoints, should they exist, are known as ***Kan lifts*** along F, denoted $\mathrm{Lift}_F\, G$ and $\mathrm{Rift}_F\, G$.

So we have a pair of dualities: between Kan extensions and Kan lifts, and between left and right Kan extensions/lifts. This is best understood by noting that Cat is a 2-category, and that being a 2-category is all that is necessary to define Kan extensions and lifts (a left Kan extension of a 1-cell $G : X \to Z$ along a 1-cell $F : X \to Y$ is a 1-cell $K : Y \to Z$ equipped with a *2-cell* $\alpha : G \Rightarrow KF$ such that ...). In a 2-category C, there are two kinds of opposites: the first flips the 1-cells, and is denoted as C^{op}, whereas the second flips the 2-cells, and is denoted as C^{co}. Kan extensions in C are Kan lifts in C^{op}, and vice versa, whereas left Kan extensions/lifts in C are right Kan extensions/lifts in C^{co}.

$$\begin{tikzcd} \mathsf{C} \arrow[rr, "G"] \arrow[dr, "F"'] & {} \arrow[d, phantom, "\Downarrow"] & \mathsf{E} \\ & \mathsf{D} \arrow[ur, dashed, "-\mathrm{an}_F G"'] & \end{tikzcd} \qquad \begin{tikzcd} \mathsf{C} \arrow[dr, "F"'] & {} \arrow[d, phantom, "\Uparrow"] & \mathsf{B} \arrow[ll, dashed, "-\mathrm{ift}_F G"'] \arrow[dl, "G"] \\ & \mathsf{D} & \end{tikzcd}$$

Day Convolution We now give an important example of a Kan extension. Given a small V-enriched monoidal category $(\mathsf{C}, \otimes, 1)$, we may define an external tensor product $-\overline{\otimes}- : \mathsf{V}^{\mathsf{C}} \times \mathsf{V}^{\mathsf{C}} \to \mathsf{V}^{\mathsf{C}\times\mathsf{C}}$ by $(F\overline{\otimes}G)(X, Y) := F(X) \otimes_V G(Y)$. $\otimes$ is itself a functor $\mathsf{C} \times \mathsf{C} \to \mathsf{C}$, and precomposition by it defines a functor $\mathsf{V}^{\mathsf{C}} \to \mathsf{V}^{\mathsf{C}\times\mathsf{C}}$. This functor's left adjoint is known as the ***Day convolution*** $* : \mathsf{V}^{\mathsf{C}} \times \mathsf{V}^{\mathsf{C}} \to \mathsf{V}^{\mathsf{C}}$.

We can fix $F, G \in \mathsf{V}^{\mathsf{C}}$ to get a functor $F\overline{\otimes}G : \mathsf{C} \times \mathsf{C} \to \mathsf{V}$; taking the left Kan extension of this functor along $\otimes$ gives us the Day convolution $F * G$, and we can therefore compute it as

$$F * G = \int^{X,Y} \mathrm{Hom}_{\mathsf{C}}(X \otimes Y, -) \otimes_V F(X) \otimes_V G(Y)$$

Day convolution turns V^{C} into a monoidal category, with the Yoneda (co)embedding $\text{よ} : \mathsf{C} \to \mathsf{V}^{\mathsf{C}}, X \mapsto (Y \mapsto \mathrm{Hom}_{\mathsf{C}}(X, Y))$ a monoidal functor.

Yoneda Extension One particularly important class of Kan extensions is given by extensions of functors $F : \mathsf{C} \to \mathsf{D}$ along $\text{よ}_{\mathsf{C}} : \mathsf{C} \to \widehat{\mathsf{C}}$, yielding functors of the form $\mathrm{Lan}_{\text{よ}_{\mathsf{C}}} F : \widehat{\mathsf{C}} \to \mathsf{D}$. $\mathrm{Lan}_{\text{よ}_{\mathsf{C}}}$ extends F from a functor on *objects* of C to a functor on *presheaves* on C. When C is small and D cocomplete, this extension always exists, and is known as the ***Yoneda extension*** of F.

1.2.7 Accessibility and Presentability

We are often prevented from studying certain properties of categories due to set-theoretic constraints, generally largeness. In many cases, though, a large category can be generated or otherwise determined by a proper set of its objects, and we may study these sets to get around size obstruction. The two most common cases are that of *accessibility*, where all objects can be generated via filtered colimits over a certain set of sufficiently nice objects, and *local presentability*, where all objects can be generated by all colimits *and* all colimits exist.

These notions have a significant interplay with cardinal properties in ordinary set theory, and we must explicate these properties first.

Cardinality First, we make a *definition*: A ***set*** is an object whose existence can be deduced from an axiomatic set theory.

Clearly, this definition is useless without an axiomatic set theory to plug in. The most commonly used theory is ZFC, or Zermelo-Fraenkel set theory with the axiom of choice. The alphabet of the first-order language $\mathcal{L}_\in$ of ZFC consists of

- The logical symbols for universal and existential quantification, $\forall$ and $\exists$, as well as those for conjunction ($\wedge$), disjunction ($\vee$), negation ($\neg$), and one/two-sided implication ($\implies$ and $\iff$).
- The *non*-logical symbols $=$ and $\in$ denoting equality and set membership. These binary relations are the primitives of ZFC.

The axioms of ZFC are as follows:

1. (Extensionality) If two sets X and Y have the same elements, then $X = Y$.
2. (Pairing) For any two sets a and b, there is a a pair set $\{a, b\}$.
3. (Separation Schema) For any formula $\phi(x)$ in $\mathcal{L}_\in$ with one free variable x, and any set X, there is a set $\{x \in X \mid \phi(x)\}$.
4. (Power Set) For any set X, there is a power set $\mathcal{P}(X)$ whose elements are subsets of X.
5. (Union) For any set X, there is a set $\cup_{x \in X} x$ given by taking the union of all elements of X.

6. (Infinity) There exists an infinite set.

7. (Replacement Schema) The image of a set under a set function is also a set.

8. (Regularity) Every non-empty set X contains an element disjoint from X.

9. (Choice) We can pick a single representative for each set in a family of arbitrarily large sets through a choice function.

(The schemata each represent infinitely many axioms, one for each formula ϕ; this works around the fact that we cannot directly iterate over the formulae of $\mathcal{L}_\in$). For instance, the existence of the empty set $\varnothing$ can be deduced from the infinite set X postulated by the axiom of infinity and the axiom of separation for the fallacious formula $\phi(x) := (x \in x) \wedge \neg(x \in x)$ applied to X. Any class (collection of sets) whose existence cannot be proved by ZFC is known as a proper class. The prototypical example is the "set of all sets" S, whose existence is contradicted by ZFC: the pair "set" $\{S, S\}$ obviously has no elements disjoint from itself, violating the axiom of regularity.

The Von Neumann Universe An especially important family of sets is given by the *ordinals*: an ordinal is a set α such that every $x \in \alpha$ is a subset of α, and α is well-ordered by $\in$. The successor of an ordinal is given by $\alpha + 1 := \alpha \cup \{\alpha\}$; an ordinal which is the successor of another ordinal is known as a successor ordinal, and an ordinal which is neither empty nor a successor ordinal is known as a limit ordinal.

The class Ord of all ordinals is well ordered by the relation $\alpha < \beta := \alpha \in \beta$, so limit ordinals can be thought of as "jumps" in this ordinal hierarchy. In fact, an arbitrary ordinal α is *equivalent* to the set of all ordinals β that are less than α. The first ordinal is trivially $\varnothing$, and we can proceed to define the von Neumann ordinals as $0 = \varnothing$, $1 = \{0\} = \{\varnothing\}$, $2 = \{0, 1\} = \{\varnothing, \{\varnothing\}\}$, and so on. The first limit ordinal is the limit of the von Neumann ordinals, $\omega = \{0, 1, 2, \ldots\}$.

Using ordinals, we can construct a *cumulative hierarchy* $\{V_\alpha\}$ of sets, which is built up in stages, one stage for each ordinal number. We start by defining V_0 as $\varnothing$ and, for each successor ordinal $\alpha + 1$, define $V_{\alpha+1} := \mathcal{P}(V_\alpha)$. For each limit ordinal β, we define $V_\beta := \bigcup_{\alpha<\beta} V_\alpha$. Finally, we define the (proper) class V to be the union of all stages: $V := \bigcup_\alpha V_\alpha$. The rank of a set is defined to be the ordinal at which it is introduced in this hierarchy. This is the standard set-theoretic approach to building a universe of sets, and is useful in discussing the category Set of sets – which, by definition, is dependent on one's idea of what a "set" is supposed to be. In other set theories, e.g. ZFC with additional axioms, we will have a different Set.

Large Cardinals Bijection is an equivalence relation on the proper class of all sets; naively, we may quotient the proper class of sets by this relation to obtain a notion of the cardinality, or size, of a set. Unfortunately, the equivalence classes are not in general sets. A slightly subtler definition which relies on the axiom of choice fixes this: a cardinal is an ordinal that is not in bijection with any of its proper subsets. The cardinality $|S|$ of a set S is the least ordinal α admitting a bijection with S.

The natural numbers are all cardinals, and ω is the first infinite cardinal; since $|\omega| = |\omega + 1| = \ldots$, we write this cardinal as $\aleph_0$ rather than ω, though cardinals still admit well-orderings as ordinals.

An important property of a cardinal κ is its cofinality $\mathrm{cf}(\kappa)$, defined to be the smallest cardinality among the subsets of κ all of whose sets have maximal cardinality in κ; the definition generalizes to any well-ordered set, ordinals in particular. Example: the cofinality of any nonzero finite ordinal is 1. An ordinal α such that $\mathrm{cf}(\alpha) = \alpha$ is known as a ***regular ordinal***; for instance, all successor ordinals are regular.

Cantor's theorem states that $|S| < |\mathcal{P}(S)|$ for every set S [9], giving us an infinite hierarchy of cardinals $\beth_0 := \aleph_0, \beth_n := 2^{\beth_{n-1}} := |\mathcal{P}(\beth_{n-1})|$. Another infinite hierarchy is given by the *successor cardinal* operation, which associates to a cardinal κ the next largest cardinal κ^+; we have $\aleph_{n+1} := \aleph_n^+$. $\aleph_0$ and the natural numbers are the only countable cardinals; all other cardinals are called uncountable. A successor cardinal is a cardinal which is some cardinal's successor. As with ordinals, we can define limit cardinals, but we must define two flavors: a weak limit cardinal κ is a cardinal which is neither a successor cardinal nor zero. A strong limit cardinal λ is a cardinal such that $\rho < \lambda \implies 2^\rho < \lambda$.

Strong limit cardinals are weak limit cardinals, since obviously $\rho^+ \leq 2^\rho$, and $\aleph_0$ is the first

[9] Proof: suppose there were a bijection f, use replacement to construct the set $T = \{s \in S \mid s \notin f(s)\} \in \mathcal{P}(S)$, and attempt to find an $s \in S$ with $f(s) = T$; we have $s \in T \iff s \notin T$, a contradiction.

strong limit cardinal. For limit ordinals λ, we define $\aleph_\lambda := \bigcup_{\rho<\lambda} \aleph_\rho$, which is in general a weak limit cardinal.

So far, we have stayed within what is provable from ZFC alone. However, weak limit cardinals are as far as ZFC can go; in this sense, such cardinals measure the "strength" of ZFC. We may postulate stronger conditions on the size of a cardinal κ, but there is no guarantee that ZFC can prove the existence of κ. Such cardinals are known as large cardinals. The first condition, or large cardinal property, is given by inaccessibility: a cardinal κ is weakly inaccessible if it is an uncountable regular weak limit cardinal, and strongly inaccessible if it is an uncountable regular strong limit cardinal.

ZFC can neither prove nor disprove the existence of weakly or strongly inaccessible cardinals; in fact, the existence of a weakly inaccessible cardinal would prove the consistency of ZFC.

Accessibility and Presentability For a regular cardinal κ, we define a κ-directed set to be a poset P in which every subset of cardinality at most κ has a join. A κ-directed colimit in a category C is defined to be the colimit of a functor from a κ-directed set to C, and an object $X \in \mathsf{C}$ is defined to be κ-compact when $\mathrm{Hom}_{\mathsf{C}}(X, -)$ preserves κ-directed colimits.

A locally small category C is defined to be ***κ-accessible*** if it has all κ-directed colimits, and there is a proper set S of κ-compact objects such that every object $X \in \mathsf{C}$ is the κ-directed colimit of a set of objects in S. C is just ***accessible*** if there exists some regular cardinal κ for which C is κ-accessible. An accessible functor between accessible categories C, D is a functor F such that there exists a regular κ for which C, D are both κ-accessible and F preserves κ-directed colimits.

An accessible category C is ***locally presentable*** if it has all colimits. Hence, there is a regular κ such that all objects in C are generated by a proper set S of κ-compact objects via taking directed κ-colimits. If the cardinal κ is $\aleph_0$, countable, then C is known as ***locally finitely presentable***. Set is a locally finitely presentable category, as every set is the colimit over the directed set of its finite subsets: hence, we can take the $\aleph_0$-compact set to be $\mathbb{N}$. For any locally finitely presentable category C and small category D, the functor category C^{D} is locally finitely presentable as well, as colimits are computed pointwise; in particular, all presheaf categories over small categories are locally finitely presentable.

Chapter 2

Topos Theory

2.1 Topos Theory

CITE: Handbook of Categorical Algebra Vol. 3, Sketches of an Elephant

Notation Throughout, we will let $\mathcal{E}$ be an elementary topos with subobject classifier Ω and true morphism $t : 1 \to \Omega$. Exponentiation will be denoted by the functor $-_1^{-_2} : \mathcal{E}^{\mathrm{op}} \times \mathcal{E} \to \mathcal{E}$, or sometimes by the functor $[-_1, -_2]$. By adjunction we have a system of isomorphisms $\omega_{A,X,B} : \mathcal{E}(A \times X, B) \cong \mathcal{E}(A, [X, B])$ natural in all variables; the counit is an isomorphism $[X, B] \times X \to B$ known as the evaluation morphism $\mathrm{ev}_{X,B}$, while the unit is an isomorphism $A \to [X, A \times X]$ known as the coevaluation morphism $\mathrm{coev}_{A,X}$. The classifying arrow of a monic $f : A \rightarrowtail B$ will be denoted either by χ_f or $\mathrm{char}\, f : B \to \Omega$.

We will also make use of the Iverson bracket $[-]$, which sends a statement to its truth value.

2.1.1 Grothendieck Topoi

Direct Image Functors Consider a topological space X, and its corresponding category $\mathrm{Sh}(X)$ of sheaves of sets. A continuous morphism $f : X \to Y$ generates a pair of adjoint functors:

- On the right, the direct image functor $f_* : \mathrm{Sh}(X) \to \mathrm{Sh}(Y)$, which sends a sheaf F on X to the sheaf $(f^*F)(V) = F(f^{-1}(V))$.

- On the left, the inverse image functor $f^* : \mathrm{Sh}(Y) \to \mathrm{Sh}(X)$, which sends a sheaf G on Y to the sheaf $(f_*G)(U) = \varinjlim_{V \supseteq f(U)} G(V)$.

By their adjunction, f^* preserves all colimits while f_* preserves all limits. f^* preserves finite limits, in fact, as it is a general fact that filtered colimits such as $\varinjlim$ preserve finite limits. If X and Y are sober[1], such that every point $x \in X$ can be deduced from the lattice of open subsets containing x (and likewise for Y), then in fact any such adjunction $f^* \dashv f_* : \mathrm{Sh}(X) \to \mathrm{Sh}(Y)$ whose left adjoint preserves finite limits comes from a continuous map $f : X \to Y$.

An instructive case is given by setting $X = \{*\}$, the vacuously Hausdorff and hence sober one-point space, since the category $\mathrm{Sh}(X)$ is equivalent to Set. Points of Y are equivalent to morphisms $X \to Y$, and hence equivalent to limit preserving left adjoints $f^* : \mathrm{Sh}(Y) \to \mathsf{Set}$. On the other hand, the fact that X is terminal in Top gives us a unique functor $f_* : \mathrm{Sh}(Y) \to \mathsf{Set}$ for any morphism $f : Y \to X$; this is the global sections functor, and its inverse image is the constant sheaf functor.

Geometric Morphisms Let $\mathcal{E} = \mathrm{Sh}(\mathsf{C}, J)$ and $\mathcal{F} = \mathrm{Sh}(\mathsf{D}, K)$ be Grothendieck topoi. An adjunction $f^* \dashv f_* : \mathcal{E} \to \mathcal{F}$ with f^* preserving finite limits is known as a ***geometric morphism*** $\mathcal{E} \to \mathcal{F}$, with f^* and f_* being called the direct and inverse images, respectively. This will be the topos-theoretic generalization of the above observation that morphisms $f : X \to Y$ generate adjoints $f^* \dashv f_* : \mathrm{Sh}(X) \to \mathrm{Sh}(Y)$. Similarly, we define a ***point*** of $\mathcal{E}$ to be a geometric morphism $p : \mathsf{Set} \to \mathcal{E}$. We form Grothendieck topoi and their geometric morphisms into a category Topos, whose terminal object is Set; the unique morphism $\Gamma : \mathcal{E} \to \mathsf{Set}$ has as its direct image the global sections functor.

If f^*, which preserves finite limits, preserves all small limits, then by the special adjoint functor theorem it has a further left adjoint $f_! : \mathcal{E} \to \mathcal{F}$, which we can compute as $f_! Y = \int^{X \in \mathcal{E}} \coprod_{f^* X \to Y} X$; an adjunction $f_! \dashv f^* \dashv f_* : \mathcal{E} \to \mathcal{F}$ characterizes an ***essential geometric morphism***.

[1]Sobriety is a relatively weak condition, as it is implied by Hausdorffness (and hence present for manifolds, CW complexes, and so on); all affine schemes (and hence all schemes) are sober as well. So it holds in *most* practical cases.

Many useful properties of Grothendieck topos are defined by analogy to topological spaces[2]. For instance, take X sober, and let $p : \mathsf{Sh}(X) \to \mathsf{Set}$. Connectedness of X is equivalent to fullness and faithfulness of $p^* : \mathsf{Set} \to \mathsf{Sh}(X)$. Hence, we call an arbitrary geometric morphism $f : \mathcal{E} \to \mathcal{F}$ ***connected*** if f^* is full and faithful, and call $\mathcal{E}$ itself connected if $\Gamma : \mathcal{E} \to \mathsf{Set}$ is connected (so that $\Gamma^* : \mathsf{Set} \to \mathcal{E}$ is full and faithful). Connected morphisms are necessarily essential, their identifying property being that $f_!$ preserves the terminal object.

2.1.2 Elementary Topoi

An ***elementary topos*** is a category $\mathcal{E}$ which is cartesian closed, has finite limits, including a terminal object 1, and a subobject classifier Ω. We define the contravariant ***power object*** functor as $\mathcal{P} := \Omega^-$, which due to the hom-exponential adjunction satisfies

$$\mathrm{Sub}_{\mathcal{E}}(X \times Y) = \mathcal{E}(X \times Y, \Omega) \cong \mathcal{E}(X, \mathcal{P}Y)$$

As with Grothendieck topoi, the canonical elementary topos is Set; as we will see, constructions in Set directly inspire many definitions of structures in elementary topoi.

2.1.3 Set-like Properties of Topoi

Set as a Topos Set is a topos with the following data:

- The subobject classifier is given by $\Omega = 2 = \{0, 1\}$.
- The true morphism is given by the inclusion $1 \hookrightarrow 2$.
- The exponential $[X, Y]$ is simply the set of all maps from X to Y. Hence, $[-,-] = \mathsf{Set}(-,-)$.
- The evaluation morphism $\mathrm{ev}_{X,Y} : [X, Y] \times X \to Y$ takes a map $\varphi : X \to Y$ and element $x \in X$ and sends it to $\varphi(x) \in Y$ (hence the name evaluation).

[2] Or, more technically, *locales*, though we will note that sober topological spaces embed fully and faithfully into locales.

- The coevaluation morphism $\mathrm{coev}_{X,Y} : X \to [Y, X \times Y]$ sends x to the map sending y to $x \times y$.
- The classifying arrow of an inclusion $f : X \hookrightarrow Y$ is given by $\chi_f(y) = [y \in \mathrm{im} f]$.

These examples will serve as our intuition for how these gadgets work in arbitrary elementary topoi; they will also serve as a foundation for us to characterize more "Set-like" gadgets.

Membership In Set, subsets of a set X are in bijection with morphisms $X \to 2$: an $S \subseteq X$ is mapped to the morphism $\underline{S}(x) = [x \in S]$. Hence, in any topos $\mathcal{E}$ we define the ***power object functor*** $\mathcal{P} = [-, \Omega] : \mathcal{E}^{\mathrm{op}} \to \mathcal{E}$. In Set, the contravariant action sends a morphism $f : X \to Y$ to the morphism $\mathcal{P}f$ sending a $\underline{V} : Y \to 2$ to the composition $\underline{V} \circ f : X \to Y \to 2$, which is equivalent to the inverse image $f^{-1}(V)$; it therefore gives us an inverse image in $\mathcal{E}$.

Now, $\mathrm{ev}_{X,\Omega}$ gives a map $\mathcal{P}X \times X \to \Omega$ which in Set sends $U \subseteq X$ and $x \in X$ to $[x \in U]$; in $\mathcal{E}$ we denote $\mathrm{ev}_{X,\Omega}$ by $\in_X$, calling it the ***membership map*** (or predicate). Note that this map is obtained by adjunction from $\mathrm{id}_{\mathcal{P}X}$, and we therefore call it the ***$\mathcal{P}$-transpose*** of $\mathrm{id}_{\mathcal{P}X}$; the $\mathcal{P}$-transpose of a general map $f : X \times Y \to \Omega$ is the adjunct map $\omega_{X,Y,\Omega}(f) : X \to \mathcal{P}Y$, and the $\mathcal{P}$-transpose of a map $g : X \to \mathcal{P}Y$ is similarly $\omega^{-1}_{X,Y,\Omega}(g) : X \times Y \to \Omega$. For convenience we simply denote transposition by $\widehat{\cdot}$.

Equality Given an $X \in \mathcal{E}$, the universal property of the product $X \times X$ ensures for any pair of arrows $f, g : Y \to X$ an arrow $h : Y \to X \times X$ yielding f and g upon projection. If $f = g = \mathrm{id}_X$, we get an arrow $\Delta_X : X \to X \times X$ with $\pi_X \Delta_X = \mathrm{id}_X$. This is known as the ***diagonal morphism***; if for $f, g : Y \to X$ we have $\Delta_X f = \Delta_X g$, then $\pi_X \Delta_X f = \pi_X \Delta_X g$ and therefore $f = g$, forcing Δ_X monic. A similar construction gives us the epic ***codiagonal*** $\nabla_X : X \amalg X \to X$.

The classifying map of Δ_X is written as $\delta_X : X \times X \to \Omega$. In Set, $\delta_X(x, x') = [x = x']$, so δ_X is in general referred to as the ***equality map*** (or predicate). Its $\mathcal{P}$-transpose $\widehat{\delta}_X : X \to \mathcal{P}X$ will in Set send $x \in X$ to $\{x\}$, and is in general referred to as the ***singleton map***.

Images Given a monic $f : X \to Y$, we will construct a direct image morphism $\exists_f : \mathcal{P}X \to \mathcal{P}Y$. Pull $t : 1 \to \Omega$ back along $\in_X$ to obtain a monic $g : Z \to \mathcal{P}X \times X$. Compose g with the monic $\mathrm{id}_{\mathcal{P}X} \times f$ to get a monic $Z \to \mathcal{P}X \times Y$, take the characteristic map $\mathcal{P}X \times Y \to \Omega$, and transpose to get a map $\exists_f : \mathcal{P}X \to \mathcal{P}Y$. In Set, $Z = \{(U,x) \in \mathcal{P}X \times X \mid x \in U\}$, so the monic $Z \to \mathcal{P}X \times Y$ sends (U,x) to $(U,f(x))$, and its characteristic map sends (U,y) to $[y \in f(U)]$; the transpose of this map sends U to $\{y \in Y \mid y \in f(U)\}$, justifying our interpretation of $\exists_f$ as a ***direct image*** map.

Now we will construct the image of an arbitrary morphism $f : X \to Y$ as a subobject of Y. First, push f out along itself to get a pair of morphisms $g, g' : Y \to Y +_X Y$ with $gf = g'f$. Take the equalizer of g with g' to get a monic $h : Z \to Y$ with $gh = g'h$; its universal property yields for any $h' : Z' \to Y$ with $gh' = g'h'$ a morphism $k : Z' \to Z$ with $h' = hk$. For f, this universal property gives an epic $k : X \to Z$ with $f = hk$. By the fact that this construction involves only universal properties, this gives a factorization of any morphism $f : X \to Y$ into an epic $X \to Z$ followed by a monic $Z \to Y$, the latter of which is known as the ***image*** of f.

Logic We can construct many logical operators using the categorical properties of Ω. While $\mathtt{true} : 1 \to \Omega$ is given by definition, we may define $\mathtt{false} : 1 \to \Omega$ to be the classifying arrow of the monic initial arrow $0 \to 1$. Negation $\neg : \Omega \to \Omega$ is given by $\chi_{\mathtt{false}}$, $\wedge : \Omega \times \Omega \to \Omega$ by δ_Ω, $\implies : \Omega \times \Omega \to \Omega$ by $\chi_{\leq}$, and $\vee : \Omega \times \Omega \to \Omega$ by $(\mathtt{true} \times \mathrm{id}_\Omega) \amalg (\mathrm{id}_\Omega \times \mathtt{true})$.

Furthermore, we may define the existential and universal quantifiers $\exists$ and $\forall$ as "internal" adjoints to the power object functor $\mathcal{P}$. Given $f : X \to Y$, we can construct for each $Z \in \mathcal{E}$ a map $\mathrm{Hom}_{\mathcal{E}}(Z, \mathcal{P}Y) \to \mathrm{Hom}_{\mathcal{E}}(Z, \mathcal{P}X)$ in the functorial manner; an internal left (right) adjoint is a left (right) natural inverse. By Yoneda, existence of such inverses implies existence of natural maps $\exists_f, \forall_f : \mathcal{P}X \to \mathcal{P}Y$ (internally) adjoint to $\mathcal{P}f : \mathcal{P}Y \to \mathcal{P}X$.

In Set, this works as follows: $\exists_f(S)$ is the set $\{y \in Y \mid \exists x \in X \text{ with } f(x) = y \text{ and } x \in S\}$, i.e. the direct image of S. $\forall_f(S)$ is the set $\{y \in Y \mid \forall x \in X, \text{ if } f(x) = y \text{ then } x \in S\}$; there can be no element of $\forall_f(S)$ that is mapped to by an element outside of S. Consider for instance the mapping $f : \mathbb{Z} \to \mathbb{Z}, n \mapsto n^2$. $\exists_f(\mathbb{N})$ will return the non-negatives, while $\forall_f(\mathbb{N})$ will return $\{0\}$, as 0 is the only integer for which $x^2 = 0 \implies x \in \mathbb{N}$.

To summarize, we have defined:

- The *power object functor* $\mathcal{P} = [-, \Omega] : \mathcal{E}^{\mathrm{op}} \to \mathcal{E}$
- The *membership map* $\in_X = \mathrm{ev}_{X,\Omega}$

- *Transposition* $\widehat{\cdot} : \mathcal{E}(X \times Y, \Omega) \cong \mathcal{E}(X, \mathcal{P}Y)$.
- The *diagonal morphism* $\Delta_X : X \to X \times X$
- The *equality map* $\delta_X = \chi_{\Delta_X} : X \times X \to \Omega$
- The *singleton map* $\{\cdot\}_X = \widehat{\delta}_X : X \to \mathcal{P}X$
- The *direct image map* $\exists_f : \mathcal{P}X \to \mathcal{P}Y$
- The *image factorization* $X \twoheadrightarrow \operatorname{im} f \rightarrowtail Y$
- The *logical operators* $\wedge, \vee, \implies : \Omega \times \Omega \to \Omega$ and $\neg : \Omega \to \Omega$.
- The *existential quantifiers* $\forall_f, \exists_f : \mathcal{P}X \to \mathcal{P}Y$ induced by an $f : X \to Y$.

2.1.4 Mitchell-Bénabou Language

The language of an elementary topos $\mathcal{E}$ consists of the following data:

- For every $1 \to X$, a constant c of *type* X. This is often written $c : X$.
- For every X, variables $\{x_n : X\}_{n \in \mathbb{N}}$.

In the interpretation of this language, a term of type X with free variables of type $X_1, \ldots, X_n$ will be given by a morphism $X_1 \times \ldots \times X_n \to X$. The terms of the language are defined inductively: first, we proclaim every constant and variable of type X to be a term of type X, variables being terms with one free variable. We shall write terms as $\alpha, \beta, \ldots$.

- `true` and `false` are terms of type Ω, also known as ***formulas***; they have no free variables, and are interpreted as their corresponding constants.
- (Membership predicate) If $\alpha : X$ and $\beta : \mathcal{P}X$ have the same free variables $x_1, \ldots, x_n$, $\alpha \in \beta$ is a formula with the same free variables $x_1, \ldots, x_n$, interpreted as the arrow $\in_X \circ (\beta \times \alpha)$.
- (Equality predicate) If $\alpha, \beta : X$ have the same free variables $x_1, \ldots, x_n$, then $\alpha = \beta$ is a formula with the same free variables $x_1, \ldots, x_n$, interpreted as the arrow $\delta_X \circ (\alpha \times \beta)$.

- (Application) If α is a term of type X and $f : X \to Y$ a morphism, then $f(\alpha)$ is a term of type Y, interpreted as $f \circ \alpha$.
- (Composition) If α is a term of type X with free variables $x_1, \ldots, x_n$ of types $X_1, \ldots, X_n$, and $y_1, \ldots, y_n$ are terms of types $X_1, \ldots, X_n$ sharing no bound variables with α, and each with free variables $y_1^1, \ldots, y_1^{m_1}, \ldots, y_n^1, \ldots, y_n^{m_n}$, then $\alpha(y_1, \ldots, y_n)$ is a term of type X with free variables $y_1^1, \ldots, y_n^{m_n}$, interpreted as $\alpha \circ (\Pi_i y_i)$.
- (Evaluation) Given $\alpha : X$ and $\beta : Y^X$, $\beta(\alpha)$ is a term of type Y, interpreted as $\mathrm{ev}_{X,Y} \circ (\beta \times \alpha)$. ($\in_X$ is a special case of this).
- (Currying) Given a term α of type X with a free variable y of type Y, $\lambda y.\alpha$ is a term of type X^Y, interpreted as the transpose of α.
- (Logic) If ϕ, ψ are formulas, then so are $\phi \implies \psi, \phi \wedge \psi, \phi \vee \psi, \neg\phi$, and so on. These are interpreted in the obvious way.
- (Quantification) If ϕ is a formula with free variables $y, x_1, \ldots, x_n$ of types $Y, X_1, \ldots, X_n$, then $(\exists y \in Y)\,\phi$ and $(\forall y \in Y)\,\phi$ are formulas with free variables $x_1, \ldots, x_n$. These are interpreted by binding y via $\lambda y.\phi : X_1 \times \ldots \times X_n \to \mathcal{P}Y$, and composing with the $\forall_p$ and $\exists_p : \mathcal{P}Y \to \Omega = \mathcal{P}1$ generated by the terminal morphism $p : Y \to 1$.

We can define further shortcuts using these symbols, such as the uniqueness quantifier $\exists!$:

$$(\exists! x \in X)(\phi(x)) \Leftrightarrow (\exists x \in X)\left(\phi(x) \wedge (\forall x' \in X)(\phi(x') \implies x = x')\right)$$

the $\notin$ and $\neq$ predicates ($x \notin X \Leftrightarrow \neg(x \in X), x \neq x' \Leftrightarrow \neg(x = x')$) (though $\neg(x \notin x)$ isn't necessarily equivalent to $x \in X$ and likewise for $\neq$), and so on. We may also rewrite quantifiers when they are obvious from convention or usage, e.g. rewriting $(\forall x \in X)(\exists y \in Y)$ as $\forall x \exists y$ and $(\forall x_1 \in X)(\forall x_2 \in X)$ as $\forall x_1, x_2$.

A formula ϕ with free variable $x : X$, which we may also write as $\phi(x)$, is equivalent via interpretation to a morphism $X \to \Omega$, and therefore (by $\mathrm{Sub}_{\mathcal{E}}(X) \cong \mathrm{Hom}_{\mathcal{E}}(X, \Omega)$) a *subobject* of X. We write this subobject as $\{x \in X \mid \phi(x)\}$, or just $\{x \mid \phi\}$. Consider for instance the subobject of X^Y given by

$$\mathrm{Inj}(Y, X) = \{f \in X^Y \mid (\forall y, y')(f(y) = f(y') \implies y = y')\}$$

which nominally classifies "injective" maps $Y \to X$. We will translate this: the term $f(y) = f(y') \implies y = y'$ is the arrow

$$(\Rightarrow) \times ((\delta_X \circ (\mathrm{ev}_{X,Y} \times \mathrm{ev}_{X,Y})) \times \delta_Y) \circ \Gamma : X^Y \times Y \times Y \to \Omega$$

where Γ is the purely logistical morphism morally sending (f, y, y') to (f, y, f, y', y, y'). Call this arrow ϕ. We transpose ϕ to get a morphism $X^Y \times Y \to \mathcal{P}Y$, apply $\forall_p$ to get a morphism $X^Y \times Y \to \Omega$, transpose to get $X^Y \to \mathcal{P}Y$, apply $\forall_p$ to get $X^Y \to \Omega$, and then take the fibered product with $\mathtt{true} : 1 \to \Omega$ to get the desired subobject $\mathrm{Inj}(Y, X) \rightarrowtail X^Y$.

We will consider two other examples: for $A, B : \mathcal{P}X$, let $A \cup B$ be the subobject $\{S \in \mathcal{P}X \mid (\forall s \in S)(s \in A \vee s \in B)\}$.

In Set, for instance, ϕ takes a map $f : Y \to X$ and two elements y, y' of Y. It turns this triplet into the sextuplet (f, y, f, y', y, y') via Γ, applies $\mathrm{ev}_{X,Y}$ to the first two pairs to obtain the quadruplet $(f(y), f(y'), y, y')$, then applies δ_X and δ_Y to each pair to obtain the pair $([f(y) = f(y')], [y = y'])$ of truth values, which it applies $\Rightarrow$ to. Transposition and application of $\forall_p$ returns the morphism sending an $f : X \to Y$ to the truth of whether it satisfies $\phi(f, y, y')$ for all $y, y' \in Y$, and pullback returns the subset of all $f : X \to Y$ that do satisfy this. The internal language allows us to reason about things such as injective functions as though they "really" existed.

A ***first-order formula*** in $\mathcal{E}$ is any formula that can be formed via these rules. We may include rules allowing for infinitary conjunction and disjunction, leading to the ***infinitary first-order formulas***. A ***geometric formula*** is an infinitary first-order formula that does not involve negation, implication, or infinitary conjunction; these are called geometric because their truth is preserved by pullback along geometric morphisms $f^* \dashv f_* : \mathcal{E} \to \mathcal{F}$. Logical morphisms preserve the truth of all first-order formulas.

2.1.5 Kripke-Joyal Semantics

Semantics Every formula $\phi(x)$ with free variable $x : X$ has a corresponding subobject $\{x \mid \phi\}$. Every morphism $f : U \to X$ also has a corresponding subobject $\mathrm{im} f$; if $\mathrm{im} f \leq \{x \mid \phi\}$, such that f factors through the subobject $\{x \mid \phi\}$, we say that U ***forces*** ϕ on the "generalized element" f, written as $U \Vdash \phi(f)$, where $\phi(f) := \phi \circ f$. Given this, the following relations on $\Vdash$, which state the ***Kripke-Joyal semantics*** of $\mathcal{E}$, hold:

1. $U \Vdash \phi(f) \wedge \psi(f)$ iff $U \Vdash \phi(f)$ and $U \Vdash \psi(f)$.
2. $U \Vdash \phi(f) \vee \psi(f)$ iff there are arrows $g : V \to U$, $h : W \to U$ such that $g \amalg h : V \amalg W \to U$ is epi, with $V \Vdash \phi(fg)$ and $W \Vdash \phi(fh)$.
3. $U \Vdash \phi(f) \implies \psi(f)$ iff for any $g : V \to U$ such that $V \Vdash \phi(fg)$, V also forces $\psi(fg)$.
4. $U \Vdash \neg\phi(f)$ if for any $g : V \to U$ such that $V \Vdash \phi(fg)$, V is the initial object.
5. $U \Vdash \exists y\, \phi(f, y)$ (for some formula $\phi : X \times Y \to \Omega$ and generalized element $f : U \to X$) iff there's an epic $e : V \to U$ and generalized element $g : V \to Y$ such that $V \Vdash \phi(fe, g)$.
6. $U \Vdash \forall y\, \phi(f, y)$ iff for *every* arrow $h : V \to U$ and generalized element $g : V \to Y$ we have $V \Vdash \phi(fh, g)$.

We say that a formula $\phi(x_1, \ldots, x_n)$ is *true* in $\mathcal{E}$, writing $\mathcal{E} \models \phi$, if the morphism $1 \to \Omega$ given by $\forall x_1, \ldots, \forall x_n\ \phi(x_1, \ldots, x_n)$ is equal to the arrow $\texttt{true} : 1 \to \Omega$, or equivalently if we have $1 \Vdash \forall x_1, \ldots, \forall x_n\ \phi(x_1, \ldots, x_n)$.

The language and semantics of a topos admit several rules for inference that we can use in order to think about this language independent from its arrow-theoretic nature: for instance, we have a *modus ponens rule*: if $U \Vdash \phi(f)$ and $U \Vdash \phi(f) \implies \psi(f)$, then, since $\mathrm{id}_U : U \to U$ has $U \Vdash \phi(f \circ \mathrm{id}_U) = \phi(f)$, it follows that $U \Vdash \psi(f)$. In general, we can carry out *intuitionistic logic*, which is more or less the same as classical logic save for a lack of the PEM. So it is not generally true in a non-Boolean topos $\mathcal{E}$ that $\mathcal{E} \models \phi \vee \neg\phi$, nor is it true that $\mathcal{E} \models \neg\neg\phi \implies \phi$.

Axioms in Topoi There are many useful axioms we can assume our topos $\mathcal{E}$ to have, which using $\mathcal{E}$'s internal logic we can state precisely. We may have, for instance, the *(internal) principle of excluded middle* (PEM):

$$\mathcal{E} \models (\forall p \in \Omega)(p \vee \neg p)$$

If this holds, we call $\mathcal{E}$ a *Boolean topos*; in such a topos we can obtain for every subobject $S \rightarrowtail X$ a complement $S^c \rightarrowtail X$.

The *internal axiom of choice* (IAC) is the internal statement that "every surjection has a section", which in Set really is equivalent to the axiom of choice:

$$\mathcal{E} \models (\forall f \in Y^X) \left[(\forall y \in Y)(\exists x \in X)(f(x) = y) \implies (\exists s \in X^Y) \left(\forall y' \in Y\right) \left(f(s(y')) = y'\right) \right]$$

This is strictly stronger than the PEM, but *weaker* than the external AC: the IAC can be true in $\mathcal{E}$ without the actual statement "every surjection has a section" being true in $\mathcal{E}$.

The *axiom of infinity* is not phrased in the internal language, but is far-reaching nevertheless: it postulates the existence of a ***natural numbers object*** (n.n.o.), or an object $\mathbb{N} \in \mathcal{E}$ equipped with two morphisms $s : \mathbb{N} \to \mathbb{N}$, $z : 1 \to \mathbb{N}$ which is universal in the sense that for any $1 \xrightarrow{x} X \xrightarrow{f} X$, there's a unique $h : \mathbb{N} \to X$ with $hz = x$ and $hs = fh$.

Given an n.n.o. $\mathbb{N}$, we can define an addition map $+ : \mathbb{N} \times \mathbb{N} \to \mathbb{N}$: this is the unique map such that the following diagram is commutative:

$$\begin{array}{ccccc}
1 \times \mathbb{N} & \xrightarrow{z \times \mathrm{id}_{\mathbb{N}}} & \mathbb{N} \times \mathbb{N} & \xrightarrow{s \times \mathrm{id}_{\mathbb{N}}} & \mathbb{N} \times \mathbb{N} \\
\| & & \downarrow{+} & & \downarrow{+} \\
\mathbb{N} & \xrightarrow{\mathrm{id}_{\mathbb{N}}} & \mathbb{N} & \xrightarrow{s} & \mathbb{N}
\end{array}$$

To get this map, apply the universal property of $\mathbb{N}$ to the diagram $1 \to \mathbb{N}^{\mathbb{N}} \to \mathbb{N}^{\mathbb{N}}$, where the first map is the transpose of the identity and the second is $s^{\mathbb{N}}$; this gives us a map $\widehat{+} : \mathbb{N} \to \mathbb{N}^{\mathbb{N}}$ with $\widehat{+} \circ z = \mathrm{id}_{\mathbb{N}}$ and $s^{\mathbb{N}} \circ \widehat{+} = \widehat{+} \circ s$, which by transpose corresponds to a map $+ : \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ making the above diagram commutative.

Given an n.n.o. $\mathbb{N}$, it is straightforward to mimic the construction of $\mathbb{Z}$ and $\mathbb{Q}$. Recall that in Set, $\mathbb{Z}$ is defined to be $\mathbb{N} \times \mathbb{N}$ modulo the relation that $(a, b) \sim (c, d)$ if $a + d = b + c$. In $\mathcal{E}$, we can take the pullback of $+$ along itself to get an object X morally representing all pairs of pairs of integers with equal sums, along with projections $\pi_1, \pi_2 : X \to \mathbb{N} \times \mathbb{N}$. Taking the two projections $\pi_1', \pi_2' : \mathbb{N} \times \mathbb{N} \to \mathbb{N}$, we quotient by the equivalence relation by taking the coequalizer of $\pi_1'\pi_1 \times \pi_2'\pi_2$ with $\pi_2'\pi_1 \times \pi_1'\pi_2$, giving us an integers object $\mathbb{Z}$. We can similarly define a multiplication $* : \mathbb{Z} \to \mathbb{Z}$ and use it to create a rational numbers object $\mathbb{Q} \in \mathcal{E}$.

It is not as easy to get a real numbers object $\mathbb{R}$, though; there are many different possible constructions, and while these are equivalent in Set, they are not generally equivalent in elementary topoi. We shall use the Dedekind real numbers, which is the "largest" among many popular constructions. A Dedekind cut in a topos $\mathcal{E}$ with rational numbers object $\mathbb{Q}$ is a pair of subobjects $L, U \rightarrowtail \mathbb{Q}$ such that the following hold in $\mathcal{E}$:

- (Non-emptiness) $(\exists x \in \mathbb{Q})(x \in L)$ and $(\exists x \in \mathbb{Q})(x \in R)$
- (Disjointness) $(\forall x)(\neg(x \in L \wedge x \in U))$
- (Order) $(\forall x, y)(x < y \wedge y \in L \implies y \in L)$ and $(\forall x, y)(x < y \wedge x \in U \implies y \in U)$
- (Dichotomy) $(\forall x, y)(x < y \implies (x \in L \vee y \in U))$
- (Openness) $(\forall x)(x \in L \implies (\exists y)(y \in L \wedge x < y))$ and $(\forall x)(x \in U \implies (\exists y)(y \in U \wedge y < x))$.

Taking the conjunction of all of these gives a formula φ on $\mathcal{P}\mathbb{Q} \times \mathcal{P}\mathbb{Q}$, the corresponding subobject $\{(L, U) \mid \varphi\}$ of which is known as the (Dedekind) real numbers object $\mathbb{R}$.

Objects in Topoi Given an object $G \in \mathcal{E}$, we may stipulate internal axioms amounting to the existence of an algebraic structure on G: for instance, suppose we equip G with a morphism $0 : 1 \to G$ and a morphism $+ : G \times G \to G$ written infix, and assume that $\mathcal{E}$ models the following sentences:

- $(\forall g \in G)(0 + g = g + 0 = g)$
- $(\forall g, h, k)((g + h) + k = g + (h + k))$.
- $(\forall g \exists h)(g + h = 0)$.
- $(\forall g, h)(g + h = h + g)$.

This will be an abelian group from $\mathcal{E}$'s point of view, and since the theory of abelian groups can be expressed intuitionistically, objects which are abelian groups according to the internal logic are also internal abelian groups; this holds for most similar theories, including rings and modules.

We shall make particular use of a certain kind of object known as a *Weil algebra*. Given a ring object R in a topos $\mathcal{E}$ (or a ***ringed topos*** $(\mathcal{E}, R)$), a ***Weil algebra*** is a local ring $(W, \mathfrak{m})$ with an R-algebra structure, such that W is finite-dimensional as an R-module and can be written as the direct sum $R \oplus \mathfrak{m}$. In the ringed topos $(\mathsf{Set}, \mathbb{R})$, Weil algebras are equivalent to $\mathbb{R}$-algebras, finite-dimensional as vector spaces, of the form $C_0^\infty(\mathbb{R}^n)/I$, where C_0^∞ denotes smooth functions vanishing at 0. For instance, $C_0^\infty(\mathbb{R})/(x^2)$ is the ring of dual numbers $\mathbb{R}[\varepsilon] := \mathbb{R}[x]/(x^2)$. With R-algebra homomorphisms mapping maximal ideals into maximal ideals, Weil algebras form a category $\mathsf{W}(\mathcal{E})$.

2.2 Classifying Topoi

Given a mathematical structure T, a classifying topos for T is a topos $\mathcal{T}$ representing the structure, in the sense that instantiations of the structure T in an arbitrary topos $\mathcal{E}$ are equivalent to geometric morphisms from $\mathcal{E}$ to $\mathcal{T}$. This structure simultaneously generalizes the notion of an algebraic theory, which is a category $\mathbb{T}$ for some sort of algebraic structure such as groups such that product preserving functors $\mathbb{T} \to \mathsf{C}$ are equivalent to instantiations of the algebraic structure in C, and the notion of a classifying space, which is a space representing a topological structure in a similar way.

2.2.1 Algebraic Theories

Motivation Many objects studied in algebra, such as groups, can be presented as sets S along with n-ary functions $f_n^i : S^n \to S$, satisfying certain conditions on the f_n^i. For instance, a group is given by a set G along with an identity, or a function $f_0 : G^0 = \{*\} \to G$ (the existence of which implies that $G \neq \varnothing$), a unary inversion operation $f_1 : G^1 \to G$, and a binary composition operation $f_2 : G^2 \to G$ representing addition, all satisfying the group axioms. These are:

1. Associativity: $f_2 \circ (\mathrm{id}_G \times f_2) = f_2 \circ (f_2 \times \mathrm{id}_G)$
2. Identity: $f_2 \circ (f_0 \times \mathrm{id}_G) = f_2 \circ (\mathrm{id}_G \times f_0) = \mathrm{id}_G$.
3. Inversion: $f_2 \circ (f_1 \times \mathrm{id}_G) \circ \Delta_{G,G} = \mathrm{id}_G$, where $\Delta_{G,G}$ sends x to (x, x).

The natural isomorphism $\gamma : -_1 \times -_2 \Rightarrow -_2 \times -_1$ and diagonal functor Δ exist in any category with finite products. Hence, these elements and axioms can be turned into a category G generated by the natural numbers $[n], n \in \mathbb{N}$ and morphisms $f_0^0, f_0^1 : [0] \to [1], f_1 : [1] \to [1], f_2^0, f_2^1 : [2] \to [1]$. A product preserving functor $F : \mathsf{G} \to \mathsf{Set}$ will not only send $[n]$ to a set $G^n = F([1])^n$, but send the morphisms in G to functions between sets that preserve the given relations, and therefore define a group structure on G. All groups arise in this way, and a natural transformation between functors is precisely a group homomorphism. Therefore, the category of all finite product preserving functors $\mathsf{G} \to \mathsf{Set}$ is isomorphic to Grp, and, for an arbitrary category C with finite products, this category is isomorphic to the category of group objects in C, denoted $\mathsf{Grp}(\mathsf{C})$.

Algebraic Theories An ***algebraic theory*** is a category which is essentially small (has at most a set's worth of isomorphism classes)[3], has finite products, and is such that every object X is isomorphic to a finite product of a specific object X_1. By tradition, we shall write such theories as $\mathbb{T}, \mathbb{S}$, and so on.

An ***algebra*** of an algebraic theory $\mathbb{T}$, or a $\mathbb{T}$-algebra, is a product-preserving functor $\mathbb{T} \to \mathsf{Set}$, and a natural transformation between $\mathbb{T}$-algebras is known as a ***homomorphism***. Hence, we have a category $\mathbb{T}\text{-}\mathsf{Alg} \subset \mathsf{Set}^{\mathbb{T}}$. The contravariant Yoneda embedding[4] $よ^{\mathrm{op}}(X) = h^X = \mathrm{Hom}_{\mathbb{T}}(X, -)$ sends every object $X \in \mathbb{T}$ to a limit-preserving functor, and therefore defines a functor $よ^{\mathrm{op}} : \mathbb{T}^{\mathrm{op}} \to \mathbb{T}\text{-}\mathsf{Alg}$; since $\mathrm{Hom}_{\mathbb{T}}$ preserves limits in both variables, $よ^{\mathrm{op}}$ itself is a product preserving functor.

The simplest algebraic theory is that encoded by $\mathsf{FinSet}^{\mathrm{op}}$, which encodes no operations in particular; it is known as the theory of equality. A product preserving morphism $\mathsf{FinSet}^{\mathrm{op}} \to \mathsf{Set}$ is a coproduct preserving morphism $\mathsf{FinSet} \to \mathsf{Set}$, which since every finite set S is isomorphic to $\amalg_{s \in S}\{*\}$ is determined by the choice of a single set. Hence, the category of algebras over the theory of equality is Set itself.

Free $\dashv$ Forgetful Adjunctions This theory is contained in every other theory – $\mathsf{FinSet}^{\mathrm{op}}$ is generated via products (coproducts in FinSet) of projections – and there is a unique inclusion $\mathsf{FinSet}^{\mathrm{op}} \to \mathbb{T}$ for every algebraic theory $\mathbb{T}$. This yields an inclusion functor $\mathbb{T}\text{-}\mathsf{Alg} \to \mathsf{FinSet}^{\mathrm{op}}\text{-}\mathsf{Alg} = \mathsf{Set}$ sending each algebra to its underlying set. This is known as the ***forgetful***

[3]This is a technical enlargement of the idea of smallness, and is equivalent to saying that the category in question has a small *skeleton*. It is not a particularly egregious enlargement: Set, for instance, is not essentially small – if its skeleton, the set of all cardinals, were a set, the union of its elements would be an even larger cardinal.

[4]Equivalently, the Yoneda embedding for C^{op}.

functor $U_{\mathbb{T}} : \mathbb{T}\text{-}\mathsf{Alg} \to \mathsf{Set}$. $U_{\mathbb{T}}$ has a left adjoint, $F_{\mathbb{T}} : \mathsf{Set} \to \mathbb{T}\text{-}\mathsf{Alg}$, known as the ***free functor***: for a finite set $[n]$, $F_{\mathbb{T}}([n]) = \text{よ}^{\mathrm{op}}([n]) = \mathrm{Hom}_{\mathbb{T}}([n], -)$, and for an infinite set S, $F_{\mathbb{T}}(S)$ sends an object $T \in \mathbb{T}$ to the colimit over all finite subsets $[n]$ of S of $F_{\mathbb{T}}([n])(T)$.

For instance, let $\mathbb{T}_{\mathsf{Grp}}$ be the algebraic theory of groups, described previously. $F_{\mathbb{T}_{\mathsf{Grp}}}([n])([1])$ consists of all morphisms $[n] \to [1]$ in $\mathbb{T}_{\mathsf{Grp}}$. The inclusion $\mathsf{FinSet}^{\mathrm{op}} \to \mathbb{T}_{\mathsf{Grp}}$ gives us one of these for each of the n elements $x_i : [1] \to [n], 0 \mapsto i$, but the identity, addition, and inversion operations give us additional morphisms. In particular, they allow us to freely add and invert the x_i subject to the constraint $x_i x_i^{-1} = 1$. Hence, $F_{\mathbb{T}_{\mathsf{Grp}}}([n])([1])$ is the free group on n generators. So this functor $F_{\mathbb{T}_{\mathsf{Grp}}} : \mathsf{Set} \to \mathbb{T}_{\mathsf{Grp}}\text{-}\mathsf{Alg} \cong \mathsf{Grp}$ reproduces the usual free functor $\mathsf{Set} \to \mathsf{Grp}$.

2.2.2 Classifying Topoi

Algebraic theories are categories which classify algebraic structures in categories; deloopings are topological groups that classify principal bundles over a given group. Classifying topoi generalize both of these examples.

Flat Functors Flat functors are the first and, while not most straightforward, most natural example of classifying topoi. A ***flat functor*** from a nonempty category C to Set is a functor F such that:

1. There is an $X \in \mathsf{C}$ with $FX \neq \varnothing$.
2. There is for every $x \in FX, y \in FY$ a span $X \xleftarrow{\alpha} W \xrightarrow{\beta} Y$ sending some $w \in FW$ to x and y.
3. For every $Ff, Fg : FX \to FY$ agreeing on some $x \in FX$ we can choose α above such that $f\alpha = g\alpha = \beta$.

A functor from a category C to a Grothendieck topos $\mathcal{E}$ is ***internally*** flat if the above conditions hold in $\mathcal{E}$'s internal logic. The conditions above ensure that any finite limits that happen to exist in C will be preserved by a flat functor F: (2) ensures that products will be preserved, and (3) that equalizers will be preserved. Hence, if C has all finite limits, the internally flat functors are the left exact functors.

Diaconesceu's Theorem states that, for $\mathcal{E}$ a Grothendieck topos, the category of internally flat functors from C to $\mathcal{E}$ is equivalent to the category of geometric morphisms from $\mathcal{E}$ to $\widehat{\mathsf{C}}$. That is,

$$\mathrm{Hom}_{\mathsf{Topos}}(\mathcal{E}, \widehat{\mathsf{C}}) \cong \mathrm{Hom}_{\mathsf{FlatFunc}}(\mathsf{C}, \mathcal{E})$$

This equivalence sends a geometric morphism $f^* \dashv f_* : \mathcal{E} \to \widehat{\mathsf{C}}$ is sent to the composition of the inverse image $f^* : \widehat{\mathsf{C}} \to \mathcal{E}$ with the Yoneda embedding よ $: \mathsf{C} \to \widehat{\mathsf{C}}$. Hence, we can say that $\widehat{\mathsf{C}}$ is the classifying topos for flat functors from C.

Theories A ***signature*** Σ is a triplet consisting of:

- A set S of *types* of Σ.
- A set R of *relation symbols*, each with an arity describing what types it relates.
- A set F of *function symbols*, each with a set of types consisting of its domain and a single type consisting of its codomain.

We may define a collection of ***terms*** over a given signature Σ, each of which has a type. We write $t : A$ to say that a term t has type A. The terms over Σ are defined recursively:

- A variable $x : A$ is a term.
- For $f : A_1, \ldots, A_n \to B$ a function symbol and $x_1 : A_1, \ldots, x_n : A_n$, $f(x_1, \ldots, x_n)$ is a term.

Next, we list many procedures for recursively defining a set of formulas over Σ, each with a finite set of free variables:

1. The formulas of truth $\top$ and falsity $\bot$ have no free variables.
2. For formulas ϕ, ψ, the conjunction $\phi \wedge \psi$, disjunction $\phi \vee \psi$, and implication $\phi \Rightarrow \psi$ are formulas with the combined free variables of ϕ and ψ.
3. For a formula ϕ, the negation $\neg\phi$ is a formula with the same free variables.
4. For a formula ϕ with at least one free variable x, the universal and existential quantifications $(\forall x)\phi$ and $(\exists x)\phi$ are formulas with all free variables of ϕ except for x.
5. Given a relation symbol $R : A_1, \ldots, A_n$ and terms $x_1 : A_1, \ldots, x_n : A_n$, $R(x_1, \ldots, x_n)$ is a formula with free variables the union of those in all the x_i.
6. For x, y terms of the same type, $x = y$ is a formula with free variables those of x and y combined.

7. For a set of formulas $\{\phi_\lambda\}_{\lambda\in\Lambda}$ with a combined finite set S of free variables, the infinitary disjunction $\bigvee_\lambda \phi_\lambda$ and infinitary conjunction $\bigwedge_\lambda \phi_\lambda$ are formulas with free variables S.

To form a set of formulas over Σ, we may select a subset of these procedures and build the smallest set of formulas closed under their application. To simplify our exposition, we shall give these procedures labels: relation and equality are *atomic* procedures, truth and conjunction are *conjunctive* procedures, falsity and disjunction are *disjunctive* procedures[5], implication and negation are *conditional* procedures, and infinitary conjunction/disjunction are *infinitary* procedures.

This results in many sets of kinds of formulas [Johnstone, 2002]:

1. ***Atomic formulas*** are built from atomic procedures.
2. ***Horn formulas*** are built from atomic and conjunctive procedures.
3. ***Regular formulas*** are built from atomic and conjunctive procedures as well as existential quantification.
4. ***Coherent formulas*** are built from atomic, conjunctive, and disjunctive procedures, as well as existential quantification.
5. ***Geometric formulas*** are built from atomic, conjunctive, and disjunctive procedures, as well as existential quantification and infinitary disjunction.
6. ***First-order formulas*** are built from all but the infinitary procedures.
7. ***Infinitary first-order formulas*** are built from all of these procedures.

A ***context*** is a finite list of distinct variables $x_1, \ldots, x_n$ ($n = 0$ is possible), traditionally denoted as $\vec{x}$. This is just a list; $\vec{x}, x_{n+1}$ will denote the list $x_1, \ldots, x_n, x_{n+1}$, whereas $\vec{x}, \vec{x}'$ will denote the list $x_1, \ldots, x_n, x'_1, \ldots, x'_m$. If a context $\vec{x}$'s variables contain all free variables in a formula or term ϕ, we may say that $\vec{x}$ is suitable for ϕ, and put ϕ in this context by writing $\vec{x}.\phi$. If $\vec{x}'$ is a context with the same length and sequence of types as a context $\vec{x}$ sharing free variables with a term or formula ϕ, replacing each x_i occurring in ϕ with the corresponding x'_i results in a new formula $\phi[\vec{x}'/\vec{x}]$.

We will use the turnstile $\vdash$ to talk about entailment, writing $\phi \vdash_{\vec{x}} \psi$ to say that, in a context $\vec{x}$, ϕ entails ψ. This is a statement in the metalanguage, **not** the formal language. So while we may

[5] I offer as justification for these pairings the fact that $\top \wedge -$ and $\bot \vee -$ both resolve to the identity in classical logic.

write the English word "implies" to mean the symbol $\Rightarrow$ of the formal language, we write the symbol $\vdash$ to mean the English word "entails". We could technically say that $\phi \dashv_{\vec{x}} \psi$ expresses the validity of $(\forall x_1, \ldots, x_n)(\phi \Rightarrow \psi)$, but the latter does not make sense in cases where we don't have access to the symbols $\forall, \Rightarrow$, such as when studying coherent formulas.

With this in mind, we define a ***sequent*** over a signature Σ to be an expression of the form $\phi \dashv_{\vec{x}} \psi$, where ϕ and ψ are formulas over Σ, and $\vec{x}$ a context suitable for both of them. A ***theory*** $\mathbb{T}$ over Σ is a set of sequents over Σ known as $\mathbb{T}$'s ***axioms***. If all formulas appearing in all sequents of a theory are all of one of the same seven kinds listed above, we say that the theory itself is of that kind.

We have previously discussed the interpretation of the above symbols in an arbitrary elementary topos $\mathcal{E}$. To reiterate:

- There is one type for each object $X \in \mathcal{E}$, with combination of types corresponding to a product of objects.
- There is a relation symbol $R : A_1, \ldots, A_n$ for each subobject of $A_1 \times \ldots \times A_n$ (equivalently, every morphism $A_1 \times \ldots \times A_n \to \Omega$).
- There is a function symbol $f : A_1, \ldots, A_n \to B$ for each morphism $f : A_1 \times \ldots \times A_n \to B$ in $\mathcal{E}$.
- There is a constant term $x : X$ for each object $X \in \mathcal{E}$, and countably many variables $x_n : X$.

The interpretations of the various logical operations are also listed in the previous sections. To say that $\phi \dashv_{\vec{x}} \psi$ for a context $\vec{x} = x_1 : X_1, \ldots, x_n : X_n$ and formulas $\phi(x_1 : X_1, \ldots, x_n : X_n), \psi(x_1 : X_1, \ldots, x_n : X_n)$ is to say that if the morphism $\phi \circ (x_1, \ldots, x_n) : 1 \to X_1 \times \ldots \times X_n \to \Omega$ is equal to $\texttt{true} : 1 \to \Omega$, then so is $\psi \circ (x_1, \ldots, x_n)$. A theory in $\mathcal{E}$ is a collection of such sequents.

Classifying Topoi Given a theory $\mathbb{T}$ and a (Grothendieck) topos $\mathcal{E}$, there are many different ways to interpretations $\mathbb{T}$ in $\mathcal{E}$: we have to find an assignment of the types, relations, and functions of $\mathbb{T}$ to certain objects, subobjects, and morphisms, in a way that satisfies all sequents of $\mathbb{T}$. Once we have some interpretations, though, we can define a homomorphism between interpretations to be a set of morphisms between the objects associated to each to each type which, similar to a natural transformation, preserves all interpretations of relations, function symbols, and constants. Hence, we have for each theory $\mathbb{T}$ and topos $\mathcal{E}$ a category $\mathbb{T}\text{-}\mathsf{Mod}(\mathcal{E})$

of interpretations of $\mathbb{T}$ in $\mathcal{E}$.

Any given functor between topoi will not necessarily preserve theories, but placing specific constraints on the functor, such as preserving certain types of (co)limits, will let it preserve certain types of theories. Our taxonomy of formulas is useful in this regard. If the theory $\mathbb{T}$ is *geometric*, for instance, then it is preserved by inverse images of geometric morphisms: a geometric morphism $f^* \dashv f_* : \mathcal{E} \to \mathcal{F}$ defines a functor $f^* : \mathbb{T}\text{-}\mathsf{Mod}(\mathcal{F}) \to \mathbb{T}\text{-}\mathsf{Mod}(\mathcal{E})$. Hence,

We are now in a position to define classifying topoi: given a geometric theory $\mathbb{T}$, a ***classifying topos*** for $\mathbb{T}$ is a topos, traditionally denoted $S[\mathbb{T}]$, such that the functor $\mathbb{T}\text{-}\mathsf{Mod}(-)$ is naturally isomorphic to $\mathrm{Hom}_{\mathsf{Topos}}(-, S[\mathbb{T}])$.

Every geometric theory has a classifying topos, which we can build in two steps [MacLane and Moerdijk, 2012]. First, construction of the *syntactic category* $B(\mathbb{T})$. We define two geometric formulas ϕ, ψ of $\mathbb{T}$ to be equivalent when they have the same sequences of types $X_1, \ldots, X_n$, and for any interpretation in any topos, the subobjects $\{x \mid \phi(x)\}$ and $\{x \mid \psi(x)\}$ are not just isomorphic but equal. The objects of $B(\mathbb{T})$ are the isomorphism classes of geometric formulas, written as $[\phi], [\psi]$, and so on. A morphism between objects $[\phi]$ and $[\psi]$ with (WLOG disjoint) types $X_1, \ldots, X_n$ and $Y_1, \ldots, Y_n$, respectively, is a geometric formula σ with types $X_1, \ldots, X_n, Y_1, \ldots, Y_n$ such that in every interpretation in every topos, the subobject corresponding to σ is the graph of a morphism $\widehat{\sigma} : \{x \mid \phi(x)\} \to \{x \mid \psi(x)\}$. Two morphisms $[\phi] \to [\psi]$ are identified when they are always the graphs of the same morphism.

This category has all finite limits, and can be equipped with a basis for a Grothendieck topology $J(\mathbb{T})$: a family $\{\sigma_n : [\phi_n] \to [\psi]\}_{n=1}^N$ is a cover when the corresponding map $\amalg_n \widehat{\sigma}_n : \amalg_n \{x \mid \phi_n(x)\} \to \{x \mid \psi(x)\}$ is epic.

Second, take the category of sheaves on this site, obtaining the Grothendieck topos $\mathcal{B}(\mathbb{T}) = \mathrm{Sh}_{J(\mathbb{T})}(B(\mathbb{T}))$. This yields the classifying topos for $\mathbb{T}$. The proof that this is indeed a classifying topos is quite involved, and we will not reproduce it here, instead directing fastidious readers to Chapter X of the above reference.

Examples The ***theory of objects*** $\mathbb{O}$ has as its signature a single type, no relation or function symbols except for equality, and no axioms. The classifying topos for this theory is the presheaf topos on $\mathsf{FinSet}^{\mathrm{op}}$, or $\mathsf{Set}^{\mathsf{FinSet}}$: by Diaconescu's theorem, geometric morphisms from a topos $\mathcal{E}$

into $\mathsf{Set}^{\mathsf{FinSet}}$ are equivalent to finite limit preserving functors from $\mathsf{FinSet}^{\mathrm{op}}$ to $\mathcal{E}$. FinSet is generated by finite colimits on 1, and FinSet is therefore generated by finite *limits* on 1; since these are preserved, a left exact functor $\mathsf{FinSet}^{\mathrm{op}} \to \mathcal{E}$ is precisely a choice of object of $\mathcal{E}$. Hence, $\mathrm{Hom}_{\mathsf{Topos}}(\mathcal{E}, \mathsf{Set}^{\mathsf{FinSet}}) \cong \mathcal{E}$.

Part II

Logic

Chapter 3

The Language of Higher Categories

3.1 Simplices

3.1.1 The Simplex Category

The simplex category $\mathbf{\Delta}$ has as its objects the finite non-empty ordinals $[n] = \{0, \dots, n\}$ and as its morphisms the order-preserving maps $[m] \to [n], f(i) \leq f(j) \iff i \leq j$. All such morphisms can be decomposed into the following two kinds of morphisms:

- ***Face maps*** $\delta_i^n : [n-1] \to [n]$ sending $\{0, \dots, n-1\}$ to $\{0, \dots, i-1, i+1, \dots, n\}$
- ***Degeneracy maps*** $\sigma_i^n : [n+1] \to [n]$ sending $\{0, \dots, n+1\}$ to $\{0, \dots, i, i, \dots, n\}$.

Given a small category B (we're thinking $\mathbf{\Delta}$) and a locally small cocomplete category C, we can turn any functor $F : \mathsf{B} \to \mathsf{C}$ into a functor $\tilde{F} : \mathsf{C} \to \widehat{\mathsf{B}}$ as $(\tilde{F}X)(B) = \mathrm{Hom}_{\mathsf{C}}(FB, X)$. That is, $\tilde{F} = F^* よ_{\mathsf{C}}$, where F^* is precomposition. Under the above conditions, $\tilde{F}$ has a left adjoint $F_! : \widehat{\mathsf{B}} \to \mathsf{C}$, along with a unique natural isomorphism $F \cong F_! \circ よ_{\mathsf{B}}$. In general, this presents $\tilde{F}$ as a right adjoint of the left Kan extension $F_! = \mathrm{Lan}_{よ_{\mathsf{B}}} F$.

For instance, let $\mathsf{B} = \mathbf{\Delta}, \mathsf{C} = \mathsf{Top}$, and F the functor sending $[n]$ to the topological n-simplex

$$|\Delta^n| := \{(x_0, \dots, x_n) \in \mathbb{R}_{\geq 0}^{n+1} \mid \sum_i x_i = 1\}$$

Hence, $F[0] = \{1\}$, $F[1] = \{(i, 1-i) \mid i \in [0,1]\} \cong [0,1]$, and so on: this is the usual notion of an n-simplex. The functor $\tilde{F}$ as above is the ***singular complex*** functor $\mathrm{Sing}(X)([n]) = \mathrm{Hom}_{\mathsf{Top}}(|\Delta^n|, X)$, and there exists a left adjoint to Sing sending presheaves on simplicial sets to topological spaces. This left adjoint, known as ***geometric realization***, is often written as $|\cdot|$; it sends the simplicial set $よ_\Delta([n]) = \mathrm{Hom}_\Delta(-, [n])$ to the topological simplex $|\Delta^n|$ (up to homeomorphism).

Other Shapes There are many other small categories that we can use in place of the simplex category:

- The *globe category* $\mathbb{G}$, with objects the naturals and morphisms $\sigma_n, \tau_n : [n] \to [n+1]$ s.t. $\sigma_{n+1} \circ \sigma_n = \tau_{n+1} \circ \sigma_n, \sigma_{n+1} \circ \tau_n = \tau_{n+1} \circ \tau_n$.
- The *cube category* $(\square, \otimes, [0])$, the strict monoidal category on the naturals freely generated by arrows $i_0, i_1 : [0] \to [1]$ and a right inverse $p : [1] \to [0]$ to both.
- The *tree category* $\mathbf{\Omega}$, with objects the non-planar rooted trees[1], and order-preserving morphisms.
- In contrast, the *cell category* $\mathbf{\Theta}_n$, with objects planar trees with level at most n, and order-preserving morphisms. Θ_1, for instance, is $\mathbf{\Delta}$.

Our next step is to study the category of presheaves on the simplex category, known as simplicial sets; they can be thought of as sets "modeled" on the simplex category.

Presheaves on the previous categories are, respectively, globular sets, cubical sets, dendroidal sets, and n-cellular sets; they form nice categories as well. There is, however, a natural way to obtain simplicial sets from categories via the *nerve* construction, a full and faithful functor embedding the theory of categories into the theory of simplicial sets. For this reason and more, simplicial sets form an especially nice foundation for ∞-category theory.

[1]"Planar" here doesn't mean "embeddable in $\mathbb{R}^2$" (all trees are planar in this sense), but the following: a non-planar tree is a finite poset T with a minimal element/root, a linear order on every sublevel set $T_y = \{x \in T \mid x \leq y\}$, and a distinguished subset of maximal elements.

3.1.2 Simplicial Sets

The category of ***simplicial sets*** is defined to be the category of presheaves on the simplex category; that is, $\mathsf{sSet} = \widehat{\boldsymbol{\Delta}}$. For a simplicial set X, we generally write the set $X[n]$ of n-simplices of X as X_n. Geometric realization allows us to intuit simplicial sets as 'kinds of spaces', and standard n-simplices as ways to probe these spaces: by the Yoneda lemma, we have $X_n = \mathrm{Hom}_{\mathsf{sSet}}(\Delta^n, X)$, so that an element of X_0, or map $\Delta^0 \to X$, is after geometric realization a *point* $|\Delta^0| = * \to |X|$, an element of X_1 is a *path* oriented from the point 0 to the point 1, and so on. In general, we will call elements of X_0 ***objects*** or ***vertices***, and elements of X_1 ***arrows*** or ***morphisms***.

Standard Simplex Anatomy We can completely understand Δ^n: its m-simplices are order-preserving maps $[m] \to [n]$, and the number of these are counted by the famous *stars and bars* combinatorial argument. This goes as follows: gather n stars, labeling them from 1 to n, and gather $m+1$ bars, labeling them from 0 to m. Every order preserving function $f : [m] \to [n]$ corresponds uniquely to the figure you'd get by placing one bar after the star labeled $f(0)$ (so, *before* the star labeled 1 if $f(0) = 0$, one after $f(1)$, and so on, up to $f(n)$. For instance, the map $[4] \to [5]$ given by

$$\begin{tikzcd}
0 \arrow[d] & 1 \arrow[dr] & 2 \arrow[d] & 3 \arrow[d] & 4 \arrow[dr] & \\
0 & 1 & 2 & 3 & 4 & 5
\end{tikzcd}$$

looks like $| * * | | * | * * |$. Every possible figure corresponds uniquely to a function via this algorithm as well, and the number of figures is

$$\frac{(m+n+1)!}{(m+1)!(n)!} = \binom{m+n+1}{n}$$

making this the number of functions as well.

In fact we can pictorially represent the m-simplices of Δ^n via these figures. We will enumerate the n stars and $m+1$ bars when convenient, though. The map $\delta_i^m : [m-1] \to [m]$ yields a map $d_i^m : \Delta_m^n \to \Delta_{m-1}^n$ by precomposition, represented by dropping the ith bar. $\sigma_i^m : [m+1] \to [m]$

similarly yields a map $s_i^m : \Delta^n_m \to \Delta^n_{m+1}$, represented by duplicating the ith bar.

Δ^n has only one "non-degenerate" n-simplex, namely $\mathrm{id}_{[n]} = | * | * | \dots | * |$, and removing this from the set of n-simplices yields a simplicial set known as the ***boundary*** $\partial\Delta^n$. In passing to geometric realization, this amounts to removing the n-dimensional filling of the simplex, leaving only the surrounding crust: while $|\Delta^n|$ is homeomorphic to the closed n-ball, $|\partial\Delta^n|$ is homeomorphic to the $(n-1)$-sphere.

We can continue to strip $\partial\Delta^n$, removing individual faces of $\mathrm{id}_{[n]}$: the ith such face is the $(n-1)$-simplex $| * \dots * |_{i-1} *_i \;\; *_{i+1}|_{i+1} \dots * |$. Removing this face gives us the ith ***horn*** Λ^n_i. If $i = 0$ or n, this is an ***outer horn***, otherwise it is a ***inner horn***.

Maps between simplices As the morphisms in $\mathbf{\Delta}$ are generated by the face and degeneracy maps $\delta^n_i : [n-1] \to [n]$ and $\sigma^n_i : [n] \to [n-1]$, $0 \leq i \leq n \in \mathbb{N}$, we can break the morphisms in a simplicial set down into morphisms of the form $d^n_i : X_n \to X_{n-1}$ (face maps) and $s^n_i : X_n \to X_{n+1}$ (degeneracy maps). Hence, we have a pair of face maps $d^1_0, d^1_1 : X_1 \to X_0$ sending a morphism f to its source and target X and Y, respectively. A 2-simplex has 3 1-simplices g, h, f as its faces, and since in general $\delta^{n+1}_j \circ \delta^n_i = \delta^{n+1}_i \circ \delta^n_{j-1}$ for $i < j$, we have that $d^1_0 \circ d^2_1 = d^1_0 \circ d^2_0$, $d^1_1 \circ d^2_2 = d^1_1 \circ d^2_1$, and $d^1_0 \circ d^2_2 = d^1_1 \circ d^2_0$. That is to say, f and h share a source, g and h share a target, and the source of g is the target of f. It is as though we had a diagram

$$\begin{tikzcd}
\bullet \arrow[rr, "h"] \arrow[dr, "f"'] & & \bullet \\
& \bullet \arrow[ur, "g"'] &
\end{tikzcd}$$

This is not to shout "simplicial sets are categorical!!" at you, as this is not generally true. We need certain relatively simple niceness conditions, and then it will be true. Therefore, we would merely like to whisper "simplicial sets are categorical" to you.

3.1.3 Simplicial Functors

Homotopy categories and nerves $\mathbf{\Delta}$ embeds into Cat in the obvious way: we have a functor F sending $[n]$ to the category $\boldsymbol{n} = 0 \to \dots \to n$. Left Kan extension of this embedding along $\text{よ}_{\mathbf{\Delta}}$ gives us a pair of adjoints between sSet and Cat. The left adjoint $h : \mathrm{sSet} \to \mathrm{Cat}$ sends a simplicial set X to its ***homotopy category***, the category hX whose objects are 0-simplices and whose morphisms are 1-simplices, modulo the relation that any 2-simplex with zeroth, first, and second faces g, h, and f yields $h = gf$, as above; the existence of such a 2-simplex is interpreted as a *composition relation*. The right adjoint $N : \mathsf{Cat} \to \mathsf{sSet}$ sends a category to its ***nerve***, the n-simplices of which are strings of n morphisms $X_1 \to X_2 \to \ldots \to X_n$. Face maps compose morphisms (truncating the outside morphisms), whereas degeneracy maps add identity morphisms.

(Co)skeleta For any $n \in \mathbb{N}$, we may define the full subcategory $\mathbf{\Delta}_{\leq n}$ whose objects are $[0], \ldots, [n]$, and let $\mathsf{sSet}_{\leq n}$ be the corresponding presheaf category. The inclusion $i_n : \mathbf{\Delta}_{\leq n} \to \mathbf{\Delta}$ induces by precomposition a functor $i_n^* : \mathsf{sSet} \to \mathsf{sSet}_{\leq n}$, which decapitates simplicial sets, leaving nothing past n-simplices. This functor has left and right adjoints $L, R : \mathsf{sSet}_{\leq n} \to \mathsf{sSet}$, and precomposition with i_n makes each into an idempotent endofunctor on sSet.

The left endofunctor Li_n is called the ***n-skeleton*** functor sk_n, whereas Ri_n is the ***n-coskeleton*** functor cosk_n. As left and right adjoints, these represent two different approaches to destroying all information a simplex may have above degree n: the n-skeleton $\mathrm{sk}_n X$ has no simplices above degree n which are not created by degeneracy operators, i.e. has only *degenerate* simplices freely created by applying degeneracy operators to $\leq$ n-simplices; it therefore removes all simplices from X that it possibly can. The n-coskeleton $\mathrm{cosk}_n X$ has, for $k \geq 1$, an $(n+k)$-simplex whenever it has all of that simplex's faces; it adds all the simplices it possibly can.

There is furthermore a canonical natural transformation $\mathrm{sk}_n \Rightarrow \mathrm{cosk}_n$. The unit of the $L \dashv i_n$ adjunction is a natural isomorphism $\eta_L : \mathrm{id}_{\mathsf{sSet}_{\leq n}} \Rightarrow i_n L$, as is the counit $\epsilon_R : i_n R \Rightarrow \mathrm{id}_{\mathsf{sSet}_{\leq n}}$ of the $i_n \dashv R$ adjunction, since both L and R are full and faithful. Hence, we have a canonical natural transformation $(R \circ \eta_L^{-1}) \circ (\eta_L \circ L) = (\epsilon_R \circ R) \circ (L \circ \epsilon_R^{-1}) := \tau : L \Rightarrow R$. Precomposing with i_n gives us our natural transformation $\tau_n : \mathrm{sk}_n \Rightarrow \mathrm{cosk}_n$.

We call simplicial sets that are invariant under sk_n ***n-skeletal***, and simplicial sets that are invariant under cosk_n ***n-coskeletal***. n-coskeletality of a simplicial set X implies that, for all $k \geq 1$, every boundary "sphere" $\partial\Delta^{n+k} \to X$ can be filled by a unique $\Delta^{n+k} \to X$, i.e. a unique element of X_{n+k}.

Barycentric subdivision This is our final Kan extension for now. For any $n \in \mathbb{N}$, we can take the poset of nonempty subsets of $[n] = \{0, 1, \ldots, n\}$, ordered by inclusion, to be a category.

The nerve of this category is a simplicial set. Order-preserving maps $f : [m] \to [n]$ define order-preserving maps between posets, as inclusions $A \subseteq B$ result in inclusions $f(A) \subseteq f(B)$. Hence, this construction defines a functor $\mathrm{Sd} : \mathbf{\Delta} \to \mathrm{sSet}$. Left Kan extension of Sd along the Yoneda embedding $よ_{\mathbf{\Delta}}$ allows us to extend Sd to a functor sSet $\to$ sSet known as ***subdivision***. Explicitly, this extension is given by

$$\mathrm{Sd}\, X \cong \varinjlim_{\Delta^n \to X} \mathrm{Sd}\, \Delta^n$$

which encapsulates the general idea of the Kan extensions we've seen so far: if we know how to do something to standard n-simplices, we can extend this knowledge to all simplicial sets via left Kan extension along $よ_{\mathbf{\Delta}}$. Furthermore, the functor we get this way will have a right adjoint.

The right adjoint to subdivision is known as the ***extension*** functor $\mathrm{Ex} : \mathrm{sSet} \to \mathrm{sSet}$. By definition, we have $(\mathrm{Ex}\, X)_n = \mathrm{Hom}_{\mathrm{sSet}}(\Delta^n, \mathrm{Ex}\, X) = \mathrm{Hom}_{\mathrm{sSet}}(\mathrm{Sd}\, \Delta^n, X)$. Hence, the 1-simplices of $\mathrm{Ex}\, X$ are shapes of the form $* \to * \leftarrow *$ among 0 and 1-simplices in X (i.e., cospans), the 2-simplices are towers of cospans filled by compatible 2-cells, and so on. There is a natural map $X \to \mathrm{Ex}\, X$ sending an n-simplex of X to the tower formed by iteratively taking the faces of that simplex. We therefore have a sequence $X \to \mathrm{Ex}\, X \to \mathrm{Ex}^2\, X \to \ldots$, the colimit over which defines the Ex^∞ functor. While this functor gets very complex very quickly, we can characterize the 1-simplices of $\mathrm{Ex}^\infty X$ as "zig-zags" in X, or patterns of the form $* \to * \leftarrow * \to * \leftarrow * \to \ldots \leftarrow *$ among 0 and 1-simplices of X. This functor is known as ***Kan fibrant replacement***.

The fact that 1-simplices in $\mathrm{Ex}^\infty X$ closely resemble morphisms in the "groupoidification" of a category given by inverting all its arrows is not accidental: Kan fibrant replacement is the analogue of groupoidification for simplicial sets. Hence, we shall call $\mathrm{Ex}^\infty X$ the ***∞-groupoidification*** of X.

3.2 Model Category Theory

3.2.1 Homotopical Categories

Weak Equivalences Take a locally small bicomplete category M. A ***model structure*** on M consists of three classes of maps, each of which has its origins in the study of topology:

- Weak equivalences $\mathcal{W}$, which generalize the notion of a homotopy equivalence between spaces.
- Fibrations $\mathcal{F}$, which generalize fiber bundles over a space.
- Cofibrations $\mathcal{C}$, which generalize closed inclusions into spaces.

We can define weak equivalences right off the bat: a wide subcategory (one which contains all objects) $\mathcal{W}$ of a category M is a class of ***weak equivalences*** if it satisfies the 2-of-6 property: for any triplet $f : X \to Y, g : Y \to Z, h : Z \to W \in \mathsf{M}$, if both hg and gf are in $\mathcal{W}$, then so are f, g, h, and hgf. Equipping M with a class of weak equivalences makes it into a ***homotopical category***, and we may construct its ***homotopy category*** $\mathrm{Ho}\mathsf{M}$ by formally inverting M at $\mathcal{W}$, adjoining formal inverses to each weak equivalence.

The most obvious class of weak equivalences is the class of isomorphisms of M, but in this case $\mathrm{Ho}\mathsf{M}$ is equivalent to M. More interestingly, we can choose the weak equivalences in Top to be the weak homotopy equivalences[2]. By the method of CW approximation, given any space X we may construct a CW complex A with a weak homotopy equivalence $A \to X$, giving us a homotopy category composed entirely of CW complexes .

3.2.2 Lifting Problems

Retracts, Saturation, and Lifting Let $\mathbf{2}$ denote the 'walking arrow' category $* \to *$. The functor category $\mathsf{M}^{\mathbf{2}}$ has as its objects arrows in M, and as its morphisms commutative squares between arrows, and is also known as the ***arrow category*** $\mathrm{Arr}(\mathsf{M})$. Given a class $\mathcal{K}$ of arrows

[2] A homotopy equivalence $X \simeq Y$ is a pair $f : X \to Y, g : Y \to X$ with fg and gf homotopic to their respective identities; a weak homotopy equivalence $f : X \to Y$ is a map inducing bijections $\pi_n(X, x) \cong \pi_n(Y, f(x))$ for all $n \geq 0$. Weak homotopy equivalence is *not* symmetric, hence the need for formal inversion.

in M, we say that M is ***closed under retracts*** if, whenever id_f factors through a morphism g, in the sense that there is a diagram

$$\begin{tikzcd}
\bullet \arrow[r] \arrow[rr, bend left, "\mathrm{id}_{\mathrm{dom}(f)}"] \arrow[d, "f"'] & \bullet \arrow[r] \arrow[d, "g"] & \bullet \arrow[d, "f"] \\
\bullet \arrow[r] \arrow[rr, bend right, "\mathrm{id}_{\mathrm{cod}(f)}"'] & \bullet \arrow[r] & \bullet
\end{tikzcd}$$

then $g \in \mathcal{K}$ implies $f \in \mathcal{K}$. So, for instance, if all identity maps are in $\mathcal{K}$, then we can simplify the diagram as

$$\begin{tikzcd}
\bullet \arrow[rr, "\mathrm{id}_{\mathrm{dom}(f)}"] \arrow[dr] \arrow[dd, "f"'] & & \bullet \arrow[dd, "f"] \\
& \bullet \arrow[ur] \arrow[dr] & \\
\bullet \arrow[ur] \arrow[rr, "\mathrm{id}_{\mathrm{cod}(f)}"] & & \bullet
\end{tikzcd}$$

The red-green and blue-green triangles both commute, so that f has both a left and right inverse and is therefore an isomorphism between its domain and codomain; any isomorphism arises in this way, implying that all isomorphisms are contained in $\mathcal{K}$ as well. It follows that wide subcategories closed under retracts contain all isomorphisms.

If $\mathcal{K}$ is a wide subcategory of M which is closed under retract, we say that it is ***left (right) saturated*** if it is closed under coproducts (products), the pushout (pullback) of any arrow in $\mathcal{K}$ along any arrow of M is again in $\mathcal{K}$, and $\mathcal{K}$ is closed under *transfinite composition* (transfinite sequential limits). The final condition means the following: consider a nonempty ordinal α, understood to be a (small) category by virtue of the fact that the ordinals form a total order via inclusion. The coproduct in α is simply taking maxima, and equalizers are trivial, so α is cocomplete. Given a colimit-preserving functor $X : \alpha \to$ M, or an α-sequence, the transfinite composition of X is defined to be $\mathrm{colim}_{\beta<\alpha} X(\beta)$.

Given a pair of maps $f, g \in$ M, we say that (f, g) has the ***lifting property*** if for every morphism $f \to g$ in the arrow category there is an arrow from $\mathrm{cod}(f)$ to $\mathrm{dom}(g)$ making everything commute:

$$
\begin{tikzcd}
\bullet \arrow[r] \arrow[d, "f"'] & \bullet \arrow[d, "g"] \\
\bullet \arrow[r] & \bullet
\end{tikzcd}
\qquad
\begin{tikzcd}
\bullet \arrow[r] \arrow[d, "f"'] & \bullet \arrow[d, "g"] \\
\bullet \arrow[r] \arrow[ur] & \bullet
\end{tikzcd}
$$

This morphism factors both f and g. Given an arbitrary class $\mathcal{K}$ of maps, g satisfies the ***right lifting property*** (RLP) with respect to $\mathcal{K}$ if (e, g) has the lifting property for *all* $e \in \mathcal{K}$; dually, f satisfies the ***left lifting property*** (LLP) with respect to $\mathcal{K}$ if (f, h) has the lifting property for all $h \in \mathcal{K}$. We introduce the amazingly literal symbol $\boxslash$ to summarize this data: if (f, g) has the lifting property, we write $f \boxslash g$. $\mathcal{K}^{\boxslash}$ will be used to denote all arrows with the RLP w.r.t. $\mathcal{K}$, and ${}^{\boxslash}\mathcal{K}$ will be used to denote all arrows with the LLP w.r.t. $\mathcal{K}$. If we have a second class $\mathcal{L}$ such that (f, g) has the lifting property for all $f \in \mathcal{K}$, $g \in \mathcal{L}$, we write $\mathcal{K} \boxslash \mathcal{L}$. We note but do not prove the following important property: for any class $\mathcal{K}$, the class ${}^{\boxslash}\mathcal{K}$ is left saturated, and the class $\mathcal{K}^{\boxslash}$ is right saturated.

A pair $(\mathcal{L}, \mathcal{R})$ of morphisms of M is a ***weak factorization system*** if:

1. Every $f : X \to Y$ in M can be written as a composite $f = R_f \circ L_f$, where $R_f \in \mathcal{R}$ and $L_f \in \mathcal{L}$.

2. $\mathcal{L} = {}^{\boxslash}\mathcal{R}$ and $\mathcal{R} = \mathcal{L}^{\boxslash}$.

If the mappings $f \mapsto L_f, f \mapsto R_f$ come from functors $L, R : \mathrm{Arr}(\mathcal{M}) \to \mathrm{Arr}(\mathcal{M})$ satisfying the conditions $\mathrm{dom} \circ L = \mathrm{dom}$, $\mathrm{cod} \circ R = \mathrm{cod}$, and $\mathrm{cod} \circ L = \mathrm{dom} \circ R$ (so that L preserves the domain of f, R preserves the codomain of f, and the codomain of L_f is always the domain of R_f, and these are all natural), then $(\mathcal{L}, \mathcal{R})$ is furthermore a ***functorial factorization system***.

3.2.3 Model Categories

A ***model category*** is a homotopical category $(\mathsf{M}, \mathcal{W})$ equipped with two additional classes of maps, the ***cofibrations*** $\mathcal{C}$ and the ***fibrations*** $\mathcal{F}$, such that $(\mathcal{C} \cap \mathcal{W}, \mathcal{F})$ and $(\mathcal{C}, \mathcal{W} \cap \mathcal{F})$ are functorial factorization systems for $\mathcal{M}$[3]. Maps simultaneously in $\mathcal{W}$ and $\mathcal{C}$ ($\mathcal{F}$) are known as ***trivial***, or ***acyclic cofibrations (fibrations)***. Notationally, we often just refer to the model category as M, speaking of cofibrations rather than "elements of $\mathcal{C}$" and so on.

In a model category M, an object X is ***fibrant*** if the terminal arrow $X \to *$ is a fibration, and ***cofibrant*** if the initial arrow $\varnothing \to X$ is a cofibration. By definition, every fibrant object's morphism $X \to *$ can be factored as $X \to FX \to *$, with the first arrow acyclic cofibrant and F fibrant; dually, every cofibrant object's morphism $\varnothing \to X$ can be factored as $\varnothing \to CX \to X$, with the second arrow acyclic fibrant and CX cofibrant. Both operations are functorial in arrows (not functorial in the usual sense! In particular, such replacements are not unique!).

Quillen's *small object argument*, applying to any cocomplete category C with technical but readily met size limitations, allows us to construct a functorial factorization system on C from a class of morphisms $\mathcal{I}$ each of whose domains is a *small object*, or an object X such that h^X preserves transfinite directed colimits. We will not make the argument here, but the system generated is $({}^{\boxslash}(\mathcal{I}^{\boxslash}), \mathcal{I}^{\boxslash})$; any system generated in such a manner is said to be ***cofibrantly generated***. We say also that a model category is cofibrantly generated if both of its weak factorization systems are.

Examples of model structures The ***classical model structure***, also known as the Kan-Quillen model structure, is a common model structure placed on sSet. A fibration in this model structure, also known as a ***Kan fibration***, is a morphism $f : X \to Y$ of simplicial sets with the right lifting property with respect to all horn inclusions $\Lambda^n_i \to \Delta^n$ for all $n > 0$ and $0 \leq i \leq n$. That is, for every commuting square

$$\begin{array}{ccc} \Lambda^n_i & \xrightarrow{i_0} & X \\ \big\downarrow & & \big\downarrow f \\ \Delta^n & \xrightarrow[i]{} & Y \end{array}$$

there is a morphism $\Delta^n \to X$, or n-simplex of X, agreeing with the image of the horn in X and

[3]Often, they are simply required to be weak; weak factorization systems which are not functorial are very rare, though, so we shall assume functoriality for convenience.

whose image via f is the n-simplex of Y picked out by i. If f has the RLP with respect only to *inner* horn inclusions, for which $0 < i < n$, then it is known as an ***inner Kan fibration***; if we require $0 < i \leq n$, it is a right fibration, and if $0 \leq i < n$, a left fibration.

The cofibrations of this model structure are the monomorphisms, which, since sSet is a presheaf category, are checked elementwise: a map $f : X \to Y$ is a cofibration/monomorphism iff each $f([n]) : X_n \to Y_n$ is an inclusion.

The weak equivalences are those maps $f : X \to Y$ whose ∞-groupoidifications $\mathrm{Ex}^\infty(f) : \mathrm{Ex}^\infty X \to \mathrm{Ex}^\infty Y$ have the RLP with respect to boundary inclusions $\partial\Delta^n \to \Delta^n$. That is, any n-simplex y in $\mathrm{Ex}^\infty Y$ whose faces are in the image of $\mathrm{Ex}^\infty f$ is itself the image of an n-simplex x in $\mathrm{Ex}^\infty X$, the faces of which are mapped to the faces of y. By adjunction, we can say that the weak equivalences are those maps $f : X \to Y$ with the RLP with respect to all $\mathrm{Sd}\,\partial\Delta^n \to \mathrm{Sd}\,\Delta^n$.

The ***canonical model structure*** on Cat has equivalences as its weak equivalences, functors which are injective on objects as its cofibrations, and whose fibrations are those functors with the right lifting property against the inclusion of the terminal category $\{0\}$ into the "walking isomorphism" $\{0 \cong 1\}$, known as ***isofibrations***. Another way of stating this is that a functor $F : \mathsf{C} \to \mathsf{D}$ is an isofibration iff for any isomorphism $b : FX \cong Y$ there is an isomorphism $a : X \cong X'$ with $FX' = Y$ and $Fa = b$. This is, in fact, the only model structure on Cat whose weak equivalences are the equivalences of categories, hence the name "canonical". Every nonempty category is not only cofibrant (the inclusion is trivial) but fibrant in this model structure: b is necessarily id_*, which all maps get sent to.

The ***classical model structure*** on Top, also by Quillen, has as its weak equivalences the weak homotopy equivalences, and has Serre fibrations, or maps with the right lifting property against all inclusions of disks into cylinders of the form $D^n \to D^n \times [0,1]$, $x \mapsto (x, 0)$, as its fibrations. The cofibrations are retracts of *relative cell complexes*, which are maps $f : X \to Y$ where Y can be formed from X by attaching cells as one does to form a cell complex. The cofibrant spaces are retracts of cell complexes, while all spaces are fibrant.

3.2.4 Quillen Adjunctions

Quillen functors Having described what model categories are, we should like to describe the proper notion of a morphism between them. There are two kinds: left and right ***Quillen functors***. The former preserve colimits, cofibrations, and trivial cofibrations, while the latter preserve limits, fibrations, and trivial fibrations. An adjunction $L : \mathsf{C} \to \mathsf{D} \dashv R : \mathsf{D} \to \mathsf{C}$ generally preserves the $\boxslash$ operator, in the sense that for any classes of arrows $\mathcal{K}, \mathcal{L}$ of C and D respectively, $L\mathcal{K} \boxslash \mathcal{L} \iff \mathcal{K} \boxslash R\mathcal{L}$. Hence, if C, D are model categories and L is left Quillen, then its preserving cofibrations and acyclic cofibrations implies that R preserves acyclic fibrations and fibrations, and vice-versa. We call such a pair of adjunct Quillen functors a (you guessed it) ***Quillen adjunction***.

A Quillen adjunction $L \dashv R$ as above is a ***Quillen equivalence*** if, for any cofibrant $C \in \mathsf{C}$ and fibrant $F \in \mathsf{D}$, arrows of the form $LC \to F$ are weak equivalences iff their adjuncts $C \to RF$ are weak equivalences. In such a case, the functors between homotopy categories obtained by applying L and R to cofibrant and fibrant objects are *equivalences*. So, Quillen equivalences are categorical equivalences "only up to" weak equivalence.

We have already seen some examples of Quillen equivalences, the most important one of which is that between Quillen's very own model structures on sSet and Top. This equivalence is given by the $|\cdot| \dashv$ Sing adjunction, a fact which we shall come to know as the *homotopy hypothesis*, for it allows us to understand the simplicial analogue of groupoids as equivalent to topological spaces, up to weak homotopy equivalence.

3.3 Simplicial Objects

By virtue of being a presheaf category over Set, sSet is complete, cocomplete, and cartesian closed. Its exponential is given by

$$(Y^X)_n = \mathrm{Hom}_{\mathsf{sSet}}(\Delta^n, Y^X) = \mathrm{Hom}_{\mathsf{sSet}}(\Delta^n \times X, Y)$$

where the first equality is by definition and the second is by adjunction. sSet is also symmetric monoidal, with product given by $\times$ and unit given by Δ^0. A category with all of these properties is known as a ***cosmos***, and is perfect for enriching new categories over.

3.3.1 Simplicially Enriched Categories

A ***simplicially enriched category*** C is a category enriched over the cosmos sSet. To reiterate, this means that:

- For every $X, Y \in \mathsf{C}$ there is a simplicial set denoted $\mathrm{Hom}_{\mathsf{C}}(X, Y)$.
- For every $X \in \mathsf{C}$ there is a morphism of simplicial sets $\mathrm{id}_X : \Delta^0 \to \mathrm{Hom}_{\mathsf{C}}(X, X)$.
- For every $X, Y, Z \in \mathsf{C}$ there is a morphism $\circ : \mathrm{Hom}_{\mathsf{C}}(Y, Z) \times \mathrm{Hom}_{\mathsf{C}}(X, Y) \to \mathrm{Hom}_{\mathsf{C}}(X, Z)$.

The obvious commutativity conditions are placed on these sets and morphisms. We may refer to the n-simplices of the hom-simplicial sets of a simplicially enriched category as its n-arrows. Generally, the 0-arrows will be the "actual" arrows between objects, the 1-arrows the arrows between arrows, and so on. The first example of a simplicially enriched category is sSet itself, by virtue of its being cartesian closed; the hom-object can be represented by Y^X, the identity on X by the actual identity ($(X^X)_0 = \mathrm{Hom}_{\mathsf{sSet}}(\Delta^0 \times X, X)$, and $\Delta^0 \times X \cong X$), and the composition map is obtained by taking $Z^{\mathrm{ev}_{X,Y}} : Z^Y \to Z^{X \times Y^X}$, applying adjunction to get a double-evaluation arrow $X \times Y^X \times Z^Y \to Z$, then applying adjunction again to get $\circ : Z^Y \times Y^X \to Z^X$.

A functor F between simplicially enriched categories C, D is given by a map of their objects as well as, for each $X, Y \in \mathsf{C}$, a morphism $\mathrm{Hom}_{\mathsf{C}}(X, Y) \to \mathrm{Hom}_{\mathsf{D}}(FX, FY)$ in V satisfying the necessary identity and composition laws. This gives us a category sSet-Cat of simplicially enriched categories. Natural transformations are defined in the only possible way as well.

A simplicial object in an arbitrary category C is simply an element of the functor category $\mathsf{C}^{\mathbf{\Delta}^{\mathrm{op}}}$, or a set $\{X([n])\}_{n \in \mathbb{N}}$ of elements of C along with face and degeneracy maps satisfying relations originating in $\mathbf{\Delta}$. We will again write $X_n = X([n])$; for a simplicially enriched category C, C_n will refer to the normal category whose morphisms $X \to Y$ are the n-simplices of $\mathrm{Hom}_{\mathsf{C}}(X, Y)$.

As regards simplicial objects in Cat, pulling back the constant functor Set $\to \mathbf{\Delta}$ along the forgetful functor $\mathsf{Cat}^{\mathbf{\Delta}^{\text{op}}} \to \mathbf{\Delta}$ yields precisely sSet-Cat, equipped with the functor F : sSet-Cat $\to \mathsf{Cat}^{\mathbf{\Delta}^{\text{op}}}$ defined as $F(\mathsf{C})([n]) = \mathsf{C}_n$; as the constant functor is monic and pullbacks of monics are monic, F is an *embedding*. Simplicially enriched categories are those simplicial objects in Cat whose objects are constant. As the functor F commutes with limits and colimits, which are calculated pointwise in Cat, sSet-Cat is complete and cocomplete.

3.3.2 Homotopy Nerves

Homotopy Coherent Nerves Let's define a basic functor $S : \mathbf{\Delta} \to$ sSet-Cat that works by defining simplicial sets between the elements of $[n]$. Specifically, the simplicial set $\mathrm{Hom}_{S[n]}(i, j)$ is given by constructing the poset P_{ij} of all subsets of $\{i, i+1, \dots, j-1, j\}$ that contain i and j, ordered by inclusion, and then taking the nerve of this poset. We think of a given element of this poset as a path from i to j: for instance, if $i = 1$ and $j = 17$, the subset $\{1, 6, 11, 14, 16, 17\}$ is thought of as the path $1 \to 6 \to 11 \to 14 \to 16 \to 17$. The composition operation $\circ$: $S[n]_{j,k} \times S[n]_{i,j} \to S[n]_{i,k}$ sends chains of inclusions of posets of the same length to their union. It is worth noting that $\mathrm{Hom}_{S[n]}(i, j)$ is isomorphic to $(\Lambda^1)^{j-i-1}$ if $i < j$, the terminal simplex Λ^0 if $i = j$, and the initial simplex $\varnothing$ otherwise.

We now have a functor $S : \mathbf{\Delta} \to$ sSet-Cat; immediately, the voices in our head start chanting three words in unison, and we cannot help but to give them what they desire: a left Kan extension. We extend S along $よ_{\mathbf{\Delta}} : \mathbf{\Delta} \to$ sSet, obtaining a functor $\mathfrak{C}$: sSet $\to$ sSet-Cat, which is left adjoint to a functor $\mathfrak{N}$: sSet-Cat $\to$ sSet. $\mathfrak{C}$ is known as the ***thickening functor***, and while it is equivalent to S on standard n-simplices (by definition of Kan extension), it is (as all left Kan extensions) constructed for an arbitrary simplicial set by taking colimits. $\mathfrak{N}$ is known as the ***homotopy coherent nerve***, and upgrades the ordinary notion of a nerve from categories to simplicially enriched categories. For a simplicial category C, the homotopy coherent nerve is given by $(\mathfrak{N}\mathsf{C})_n = \mathrm{Hom}_{\mathsf{sSet}}(\Delta^n, \mathfrak{N}\mathsf{C}) = \mathrm{Hom}_{\mathsf{sSet\text{-}Cat}}(S[n], \mathsf{C})$.

For X a simplicial set, we can describe the simplicially enriched category $\mathfrak{C}X$ as follows: its objects are simply the vertices of X, $\mathrm{Ob}(\mathfrak{C}X) = X_0$. The 0-arrows $A \to B$ are in correspondence with the elements f of X_1 with faces $d_0^1 f = A, d_1^1 f = B$, i.e. morphisms $f : A \to B$ in X. We will not attempt to describe higher simplices in the simplicial sets $\mathrm{Hom}_{\mathfrak{C}X}(A, B)$ in detail, as they are notoriously complicated, but we will note that these hom-simplicial sets are always at most 3-coskeletal.

The model structure on simplicial categories There is a useful model structure on sSet-Cat known as the ***Bergner model structure*** [Bergner, 2007]. Its weak equivalences are the ***Dwyer-Kan*** (DK) weak equivalences, or those sSet-functors $F : \mathsf{C} \to \mathsf{D}$ that are essentially surjective on the homotopy categories, and for which each map $\mathrm{Hom}_{\mathsf{C}}(X, Y) \to \mathrm{Hom}_{\mathsf{D}}(FX, FY)$ is a weak equivalence in the Quillen model structure on sSet. The fibrations are those sSet-functors $F : \mathsf{C} \to \mathsf{D}$ for which the maps $\mathrm{Hom}_{\mathsf{C}}(X, Y) \to \mathrm{Hom}_{\mathsf{D}}(FX, FY)$ are fibrations, and which are fibrations on homotopy categories in the canonical model structure on Cat. The cofibrations are the maps which have the left lifting property against all acyclic fibrations.

This structure allows us to define a second model structure on simplicial sets, known as the ***Joyal model structure***. The cofibrations are the monomorphisms, as with the Quillen model structure, while the weak equivalences are those maps $f : X \to Y$ such that $\mathfrak{C}f$ is a weak equivalence in Bergner's model structure. The fibrations in this structure can be presented as the maps $f : X \to Y$ which have the right lifting property against all inner horn inclusions, or inclusions $\Lambda^n_i \to \Delta^n$ with $n \geq 2$ and $0 < i < n$; these are known as ***inner fibrations***. All objects are cofibrant, but the fibrant objects satisfy the special condition of having all inner horn inclusions.

3.3.3 Enriched Model Categories

Pushout-Pullback Consider maps $f : X \to Y$ and $f' : X' \to Y'$ of simplicial sets. These assemble into a pushout diagram

$$\begin{tikzcd}
X \times X' \arrow[r, "\mathrm{id}_X \times f'"] \arrow[d, "f \times \mathrm{id}_{X'}"'] & X \times Y' \arrow[d] \\
Y \times X' \arrow[r] & (X \times Y') +_{X \times X'} Y \times X' \arrow[ul, phantom, "\ulcorner", very near start]
\end{tikzcd}$$

However, $Y' \times Y'$ has maps $\mathrm{id}_Y \times f'$ and $f \times \mathrm{id}_{Y'}$ coming in from $Y \times X'$ and $X \times Y'$, respectively, which form into a commutative square, and therefore has a unique arrow coming from the cartesian product $Z = (X \times Y') +_{X \times X'} Y \times X'$ making the diagram below commute:

$$\begin{tikzcd}
X \times X' \arrow[r, "\mathrm{id}_X \times f'"] \arrow[d, "f \times \mathrm{id}_{X'}"'] & X \times Y' \arrow[d, "g'"] \arrow[ddr, bend left, "f \times \mathrm{id}_{Y'}"] & \\
Y \times X' \arrow[r, "g"'] \arrow[drr, bend right, "\mathrm{id}_Y \times f'"'] & (X \times Y') +_{X \times X'} Y \times X' \arrow[ul, phantom, "\ulcorner", very near start] \arrow[dr, dashed] & \\
 & & Y \times Y'
\end{tikzcd}$$

We call this map the ***pushout-product*** $f \widehat{\times} g : Z \to Y \times Y'$. The dual construction is defined by the pullback diagram

$$\begin{tikzcd}
X'^Y \arrow[dr, dashed] \arrow[drr, bend left, "X'^f"] \arrow[ddr, bend right, "f'^Y"'] & & \\
 & X'^X \times_{Y'^X} Y'^Y \arrow[r] \arrow[d] \arrow[dr, phantom, "\lrcorner", very near start] & X'^X \arrow[d, "f'^X"] \\
 & Y'^Y \arrow[r, "Y'^f"'] & Y'^X
\end{tikzcd}$$

and is known as the ***pullback-hom*** $\widehat{f'^f} : X'^Y \to X'^X \times_{Y'^X} Y'^Y$. By virtue of their universal constructions, these are functorial in arrow categories, and by virtue of their not interchanging the positions of their objects, they extend to give generalized constructions. Specifically, for a bifunctor $\otimes : \mathsf{C} \times \mathsf{D} \to \mathsf{E}$ (where E has the appropriate pullbacks), the pushout-product construction gives us an arrow bifunctor $\widehat{\otimes} : \mathrm{Arr}(\mathsf{C}) \times \mathrm{Arr}(\mathsf{D}) \to \mathrm{Arr}(\mathsf{E})$, and for a bifunctor $[-,-] : \mathsf{C}^{\mathrm{op}} \times \mathsf{D} \to \mathsf{E}$, the pullback-hom construction gives us an arrow bifunctor $\widehat{[-,-]} : \mathrm{Arr}(\mathsf{C})^{\mathrm{op}} \times \mathrm{Arr}(\mathsf{D}) \to \mathrm{Arr}(\mathsf{E})$.

Now, suppose we have a category enriched over V. We should want two "niceness properties" from C: that it is tensored and cotensored.

- C is ***tensored*** over V when there is a functor $- \otimes - : \mathsf{V} \times \mathsf{C} \to \mathsf{C}$ equipped with natural isomorphisms $\mathrm{Hom}_{\mathsf{C}}(V \otimes X, Y) \cong \mathrm{Hom}_{\mathsf{V}}(V, \mathrm{Hom}_{\mathsf{C}}(X, Y))$ for each V.
- C is ***cotensored*** over V when there's a $[-,-] : \mathsf{V}^{\mathrm{op}} \times \mathsf{C} \to \mathsf{C}$ equipped with natural isomorphisms $\mathrm{Hom}_{\mathsf{C}}(X, [V, Y]) \cong \mathrm{Hom}_{\mathsf{V}}(V, \mathrm{Hom}_{\mathsf{C}}(X, Y))$ for each V.

This suffices to define a ***two-variable adjunction***

$$\mathrm{Hom}_{\mathsf{C}}(X,[V,Y]) \cong \mathrm{Hom}_{\mathsf{V}}(V,\mathrm{Hom}_{\mathsf{C}}(X,Y)) \cong \mathrm{Hom}_{\mathsf{C}}(V\otimes X,Y)$$

to which the pushout-product and pullback-hom constructions associate a two-variable adjunction between arrow categories: for $f,g \in \mathsf{C}$ and $v \in \mathsf{V}$,

$$\mathrm{Hom}_{\mathrm{Arr}(\mathsf{C})}(f,\widehat{[v,y]}) \cong \mathrm{Hom}_{\mathrm{Arr}(\mathsf{V})}(v,\widehat{\mathrm{Hom}_{\mathsf{C}}}(f,g)) \cong \mathrm{Hom}_{\mathrm{Arr}(\mathsf{C})}(v\widehat{\otimes} f,g)$$

Quillen Two-Variable Adjunctions Given a tensored and cotensored V-category C where both V and C have model structures, we say that the two-variable adjunction $(\otimes,[-,-],\mathrm{Hom}_{\mathsf{C}})$ is a ***Quillen two-variable adjunction*** if:

- If f and g are cofibrations, then $f\widehat{\otimes}g$ is, and if either of them are acyclic, $f\widehat{\otimes}g$ is as well.
- If f and g are fibrations, then $\widehat{[f,g]}$ and $\widehat{\mathrm{Hom}_{\mathsf{C}}}(f,g)$ are, and if either of them are acyclic, $\widehat{[f,g]}$ and $\widehat{\mathrm{Hom}_{\mathsf{C}}}(f,g)$ are as well.

A V-model category is a V-category that is both tensored and cotensored, and whose two-variable adjunction is Quillen. Quillen two-variable adjunctions preserve lifting problems: for $\mathcal{A},\mathcal{B}$ classes of maps in C, and V a class of maps in C, we have:

$$(\mathcal{V}\otimes\mathcal{A})\boxslash\mathcal{B} \Leftrightarrow \mathcal{A}\boxslash\widehat{[\mathcal{V},\mathcal{C}]} \Leftrightarrow \mathcal{V}\boxslash\widehat{\mathrm{Hom}_{\mathsf{C}}}(\mathcal{A},\mathcal{B})$$

The pushout-product and pullback-hom constructions associated to the tensor and cotensor are referred to as the ***Leibniz tensor*** and ***Leibniz cotensor***, respectively.

3.3.4 Quasi-categories

Recall that the fibrations in the Quillen model structure on sSet are the Kan fibrations, those maps which satisfy the right lifting property against all horn inclusions. The fibrant objects in sSet, then, are those objects X which satisfy the following conditions: for any map $\Lambda^n_i \to X$, there is a morphism $\Delta^n \to X$ commuting with the horn inclusion. Quillen-fibrant objects in sSet are known as ***Kan complexes***.

For instance, if a Kan complex $\mathcal{X}$ has a 2-cell x with zeroth and first faces f and g, there is a horn inclusion $\Lambda^2_2 \to \mathcal{X}$ sending the central 2-cell id_{Δ^2} to x, and therefore a map $k : \Delta^2 \to \mathcal{X}$

commuting with the horn inclusion. By naturality (these are presheaves, after all), we can identify the second face h of x to be $(k \circ d_2^2)(\mathrm{id}_{\Delta^2})$. What is going on here is essentially an inversion:

$$\begin{tikzcd} & \bullet \arrow[dr, dashed, "h"] & \\ \bullet \arrow[ur, "f"] \arrow[rr, "g"'] & & \bullet \end{tikzcd}$$

In normal category theory, this can only happen in groupoids, and for this reason we often interpret Kan complexes as the higher-categorical equivalent of groupoids. A way to get more normal behavior is to withdraw the request that the map $X \to *$ have the right lifting property against *all* horns: so f cannot be forced to exist given that g and h do, nor can h be forced to exist given that f and g do. The only horn we are allowed to fill is the first one, forcing g to exist contingent on the existence of f and h. We identify this as the composite $h \circ f$.

Making this request restricts us to considering only inner horns, moving from Quillen-fibrant objects to *Joyal-fibrant* objects. These are known as weak Kan complexes, or ***quasi-categories***. Being fundamental to our upcoming discussion, we see fit to denote these objects with the style $\mathcal{C}, \mathcal{D}, \ldots$. Given a quasi-category $\mathcal{C}$, members of the set $\mathcal{C}_0$ are identified with the *objects* of the "category", members of $\mathcal{C}_1$ with the *morphisms*. Note that because simplicial sets are functors into Set, the objects $\mathcal{C}_0 = \mathcal{C}([0])$ of the quasi-category will always be a set; in particular, we cannot have "large" quasi-categories without performing some intricate set-theoretic maneuvers.

Given two objects (0-simplices) $X, Y \in \mathcal{C}$, consider the pullback of the morphism $\mathcal{C}^i : \mathcal{C}^{\Delta^1} \to \mathcal{C}^{\partial\Delta^1} \cong \mathcal{C} \times \mathcal{C}$ along the morphism $* \to \mathcal{C} \times \mathcal{C}$ with image (X, Y). This will be a simplicial set whose 0-simplices are 0-simplices of $\mathcal{C}^{\Delta^1}$, or maps $f : \Delta^1 \times \Delta^0 \cong \Delta^1 \to \mathcal{C}$, such that $f(\{0\}) = X$ and $f(\{1\}) = Y$, i.e. morphisms $X \to Y$ in $\mathcal{C}$. Higher simplices $f : \Delta^1 \times \Delta^n \to \mathcal{C}$ send $f(\{0\}, -)$ and $f(\{1\}, -)$ to degenerate simplices generated by X and Y. This simplicial set is called the ***mapping space*** between X and Y, and will be written as $\mathrm{Map}(X, Y)$, or $\mathrm{Map}_{\mathcal{C}}(X, Y)$ if $\mathcal{C}$ is not already clear; it is the ∞-categorical analogue of the hom-sets possessed by ordinary categories. By degeneracy of the higher simplices, each mapping space is an ∞-groupoid.

The nerve construction on ordinary small categories, i.e. the functor $N : \mathsf{Cat} \to \mathsf{sSet}$, is full and faithful, and sends categories to quasi-categories: the only thing to verify here is that the inner horns $\Lambda^2_1 \to N(\mathsf{C})$ are filled, which is true since composition arrows are guaranteed to exist in C. Indeed, not only do they exist, but they exist uniquely; the nerves of categories are identified among all quasi-categories by the existence of *unique* fillers for inner horns.

Summary

Having gone through many constructions of more or less a few major types, we can taxonomize the basic theory.

Functors on Simplicial Sets

The general principle is as follows: we have a functor F from the simplex category $\mathbf{\Delta}$ to a codomain category C, and left Kan extend it along $よ_{\mathbf{\Delta}} : \mathbf{\Delta} \to \mathsf{sSet}$ to obtain a functor $L : \mathsf{sSet} \to \mathsf{C}$ left adjoint to a second functor $R : \mathsf{C} \to \mathsf{sSet}$. There are many constructions of this type:

Left Kan extensions of functors $F : \mathbf{\Delta} \to \mathsf{C}$ along $よ_{\mathbf{\Delta}} : \mathbf{\Delta} \to \mathsf{sSet}$

Codomain C	$F[n]$	Left adjoint $\mathsf{sSet} \to \mathsf{C}$	Right adjoint $\mathsf{C} \to \mathsf{sSet}$
Top	$\lvert\Delta^n\rvert$	Geometric realization $\lvert \cdot \rvert$	Singular complex Sing
Cat	$\{0 \to \ldots \to n\}$	Homotopy category h	Nerve N
sSet	$N(\mathcal{P}[n])$	Subdivision Sd	Extension Ex
sSet-Cat	$(\Lambda^1)^{j-i-1}$ etc.	Thickening $\mathfrak{C}$	H.c. nerve $\mathfrak{N}$

In addition, left and right Kan extensions along $よ_{\mathbf{\Delta}_{\leq n}}$ yield left and right adjoints to the truncation functor $\mathrm{tr}_n : \mathsf{sSet} \to \mathsf{sSet}_{\leq n}$ induced by precomposition by the inclusion $i : \mathbf{\Delta}_{\leq n} \to \mathbf{\Delta}$; when precomposed with tr_n, these adjoints yield a pair of adjoint idempotent endofunctors $\mathrm{sk}_n \dashv \mathrm{cosk}_n$ on sSet.

The last vertex natural transformation $\mathrm{Lv} : \mathrm{Sd} \Rightarrow \mathrm{id}_{\mathsf{sSet}}$ defined on standard n-simplices and extended to all simplicial sets by colimit yields via adjunction a natural map $\mathrm{id}_{\mathsf{sSet}} \Rightarrow \mathrm{Ex}$, which we can iterate to get chains $X \to \mathrm{Ex}\, X \to \mathrm{Ex}^2\, X \to \ldots$ natural in X; taking colimits gives us a functor $\mathrm{Ex}^\infty : \mathsf{sSet} \to \mathsf{sSet}$ known alternatively as Kan fibrant replacement or ∞-groupoidification.

Model Categories

Classical (Quillen) Model Structure on sSet

- Fibrations are the Kan fibrations (RLP against all horn inclusions). Fibrant objects are Kan complexes.
- Cofibrations are the monomorphisms. All objects are cofibrant.
- Weak equivalences are the weak homotopy equivalences.

Joyal Model Structure on sSet

- Fibrations are the inner fibrations (RLP against all inner horn inclusions). Fibrant objects are quasi-categories.
- Cofibrations are the monomorphisms. All objects are cofibrant.
- Weak equivalences are preimages of Bergner weak equivalences under $\mathfrak{C}$.

Classical (Quillen) Model Structure on Top

- Fibrations are the Serre fibrations (RLP against all disk $\to$ base of cylinder inclusions). All spaces are fibrant.
- Cofibrations are the retracts of relative cell complexes. Retracts of cell complexes are cofibrant.
- Weak equivalences are weak homotopy equivalences.

Canonical Model Structure on Cat

- Fibrations are the isofibrations (RLP against inclusions $\{0\} \to \{0 \cong 1\}$). All categories are fibrant.
- Cofibrations are the functors injective on objects. All categories are cofibrant.
- Weak equivalences are the equivalences of categories.

Bergner Model Structure on sSet**-**Cat

- Fibrations are functors which are Quillen fibrations on all hom-objects. Fibrant objects are the Kan complex-enriched categories.
- Cofibrations are defined by necessity (LLP against acyclic fibrations).

- Weak equivalences are functors which are essentially surjective on homotopy categories and Quillen weak equivalences on all hom-objects.

Quillen Adjunctions and Equivalences

- $|\cdot| \dashv \mathrm{Sing}$ is an equivalence between Quillen's sSet and Quillen's Top.
- $\mathfrak{C} \dashv \mathfrak{N}$ is an equivalence (adjunction) between Joyal's (Quillen's) sSet and Bergner's $\mathsf{sSet}\text{-}\mathsf{Cat}$.

3.4 Models of Higher Category Theory

3.4.1 ∞-Cosmoi

Kan complexes and quasi-categories are often called ∞-groupoids and ∞-categories, respectively. But there are many different models of higher category theory, and we should not like to commit ourselves just yet. Riehl and Verity have pioneered a synthetic approach to higher category theory, that of *∞-cosmoi* [Riehl and Verity, 2016b]; they are defined by a set of abstract properties which any model of higher category theory must satisfy, and their elements are known as ∞-categories.

An ***∞-cosmos*** is an sSet-category equipped with a class of maps $\mathcal{I}$ known as *isofibrations*; isofibrations will be denoted by . We will write ∞-cosmoi in the vertiginous style $\mathfrak{C}, \mathfrak{D}, \ldots$, and demand the following properties from them:

- All of $\mathfrak{C}$'s hom-objects, denoted by the notation Fun, are quasi-categories.
- $\mathfrak{C}$ has a terminal object, small products, pullbacks of isofibrations, and limits of countable towers of isofibrations.
- $\mathfrak{C}$ is cotensored by sSet.
- $\mathcal{I}$ has all isomorphisms and maps into the terminal object.
- $\mathcal{I}$ is closed under composition, product, pullback, inverse limits of towers, and Leibniz cotensors with sSet-cofibrations.
- For all $C \in \mathfrak{C}$, $\mathrm{Fun}(C, -)$ sends isofibrations to isofibrations of quasi-categories, or inner fibrations with the LLP against the inclusion $\{0\} \to \{0 \cong 1\}$.

We demand that all of the above limits in $\mathfrak{C}$ satisfy sSet-enriched notions of their universal properties[4]. If furthermore there is an adjunction $\mathrm{Fun}(\mathcal{C} \times \mathcal{D}, \mathcal{E}) \cong \mathrm{Fun}(\mathcal{C}, [\mathcal{D}, \mathcal{E}])$ with all functors $[\mathcal{C}, -]$ preserving isofibrations, we say that $\mathfrak{C}$ is ***cartesian closed***.

We define an ∞-***category*** to be an object in an ∞-cosmos and an ∞-***functor*** to be a morphism in an ∞-cosmos. An ∞-functor $F : \mathcal{C} \to \mathcal{D}$ is defined to be an equivalence of ∞-categories when $\mathrm{Fun}(\mathcal{E}, F) : \mathrm{Fun}(\mathcal{E}, \mathcal{C}) \to \mathrm{Fun}(\mathcal{E}, \mathcal{D})$ is a weak equivalence of quasi-categories for all $\mathcal{E}$, and an acyclic (trivial) fibration if it is an isofibration and equivalence.

It is not particularly hard to be an ∞-cosmos: Cat and its previously defined isofibrations form an ∞-cosmos, with the understanding that its functor categories are embedded as quasi-categories in sSet via the nerve constructions. Two more interesting examples are given by $\mathfrak{Kan}$ and $\mathfrak{QCat}$, the full subcategories of sSet consisting of the Kan complexes and quasi-categories, respectively. Functors between ∞-cosmoi that preserve isofibrations and cosmological limits are known as ***cosmological functors***. If a cosmological functor $F : \mathfrak{C} \to \mathfrak{D}$ is furthermore surjective up to equivalence and defines for all $\mathcal{C}, \mathcal{D}$ equivalences $\mathrm{Fun}_{\mathfrak{C}}(\mathcal{C}, \mathcal{D}) \to \mathrm{Fun}_{\mathfrak{D}}(F\mathcal{C}, F\mathcal{D})$, it is a ***cosmological biequivalence***. Cosmological biequivalence is the basic standard by which we can compare models of higher category theory.

Unless otherwise specified, we will work in $\mathfrak{QCat}$, so that our ∞-categories are quasi-categories. $\mathrm{Fun}(\mathcal{C}, \mathcal{D})$ will denote the quasi-category of ∞-functors (or, morphisms of simplicial sets) $\mathcal{C} \to \mathcal{D}$, and $\mathrm{Map}_{\mathcal{C}}(X, Y)$ will denote the ∞-groupoid, or Kan complex, whose 0-simplices are morphisms $X \to Y$. However, the remainder of the discussion in this section will be applicable to general ∞-cosmoi.

As all the hom-objects $\mathrm{Fun}(\mathcal{C}, \mathcal{D})$ in an ∞-cosmos $\mathfrak{C}$ are quasi-categories, for which all n-morphisms for $n \geq 2$ are invertible, there are essentially 2 non-trivial levels of morphism in $\mathfrak{C}$: the objects of $\mathrm{Fun}(\mathcal{C}, \mathcal{D})$, or "actual" functors $\mathcal{C} \to \mathcal{D}$, and the (1-)morphisms, or natural transformations between functors. So, the basic constructions on ∞-cosmoi are 2-categorical: they can be expressed in terms of the ***homotopy 2-category*** $\mathrm{h}\,\mathfrak{C}$ has the same objects as $\mathfrak{C}$, but

[4]For instance, morphisms $\mathcal{C} \to \mathcal{D} \times \mathcal{E}$ are in natural bijection with pairs of morphisms $\mathcal{C} \to \mathcal{D}$, $\mathcal{C} \to \mathcal{E}$. We demand that this bijection be sSet-natural instead of merely Set-natural, so that in particular it holds on higher simplices. We can call these "cosmological" limits.

the hom-*category* between $\mathcal{C}$ and $\mathcal{D}$ is given by the homotopy category $\mathrm{h}\,\mathrm{Fun}(\mathcal{C}, \mathcal{D})$. For this reason we give a basic account of 2-category theory in 1.2.4.

3.4.2 Constructions on ∞-Categories

Adjunctions Given ∞-categories $\mathcal{C}, \mathcal{D}$, an ∞-functor $L : \mathcal{C} \to \mathcal{D}$ is left adjoint to an ∞-functor $R : \mathcal{D} \to \mathcal{C}$ if there are a pair of ∞-natural transformations, the unit $\eta : \mathrm{id}_{\mathcal{C}} \Rightarrow RL$ and counit $\epsilon : LR \Rightarrow \mathrm{id}_{\mathcal{D}}$, satisfying the 2-categorical triangle identities. Hence, an adjunction between ∞-categories is an adjunction between their homotopy 2-categories. We have noted that adjunctions between 2-categories are preserved by 2-functors; owing to its quasi-categorical nature, an ∞-functor is more than nice enough to induce 2-functors on the homotopy categories of its source and target.

For instance, an adjunction $L \dashv R$ between ∞-categories as above is sent by every representable functor $\mathrm{Fun}(X, -)$ to an adjunction between quasi-categories, and by every $\mathrm{h}\,\mathrm{Fun}(X, -)$ to an adjunction between 1-categories.

Limits Let's take a closer look at the definition of a limit in a 1-category C, which we have only previously defined in fuzzy, conceptual terms.

1. We take a certain *shape* J, for instance $\{*, *\}$, which defines products, or $\{* \rightrightarrows *\}$, which defines equalizers. This is clearly some sort of category.

2. We move to consider a *diagram* of shape J in C, or a functor $F : \mathsf{J} \to \mathsf{C}$. For instance, we might send the shape $\{*, *\}$ to $\{X, Y\}$.

3. We then consider the category of all *cones* over F, or elements $Z \in \mathsf{C}$ equipped with a morphism $f_i : Z \to FJ_i$ for each $J_i \in J$ such that $(F\lambda_{ij}^k) \circ f_i = f_j$ for any $\lambda_{ij}^k \in \mathrm{Hom}_{\mathsf{J}}(J_i, J_j)$. Cones with summit Z are precisely natural transformations from the constant functor $\Delta_Z : \mathsf{J} \to \mathsf{C}$ to the functor F.

4. Find an element $\lim F \in \mathsf{C}$ such that cones with summit Z over F are in natural bijection with morphisms $Z \to \lim F$. In other words, $\mathrm{Hom}_{\mathsf{C}^{\mathsf{J}}}(\Delta_Z, F) \cong \mathrm{Hom}_{\mathsf{C}}(Z, \lim F)$. This is the *limit* of F.

If limits exist for all functors $F : \mathsf{J} \to \mathsf{C}$, then we have a functor $\lim : \mathsf{C}^{\mathsf{J}} \to \mathsf{C}$, and the equivalence of hom-sets above upgrades to an *adjunction* between the constant diagram functor $\Delta : \mathsf{C} \to \mathsf{C}^{\mathsf{J}}, \Delta(Z) = \Delta_Z$, and lim. So we may say the following: limits of shape J exist in C if and only if there is a right adjoint lim to the constant functor Δ. Dually, colimits of shape J exist if and only if there is a left adjoint colim to Δ. With this construction in hand, it becomes easy to prove many of the previous assertions concerning limits. For instance, *right adjoints preserve limits*: let $L : \mathsf{C} \to \mathsf{D}$ be left adjoint to $R : \mathsf{D} \to \mathsf{C}$, where C and D admit all limits of shape J. Let $F : \mathsf{J} \to \mathsf{D}$. Then, for $X \in \mathsf{C}$,

$$\mathrm{Hom}_{\mathsf{C}}(X, R\lim F) \cong \mathrm{Hom}_{\mathsf{D}}(LX, \lim F) \cong \mathrm{Hom}_{\mathsf{D}^{\mathsf{J}}}(\Delta_{LX}, F) \cong \mathrm{Hom}_{\mathsf{D}^{\mathsf{J}}}(L\Delta_X, F)$$

$$\mathrm{Hom}_{\mathsf{D}^{\mathsf{J}}}(L\Delta_X, F) \cong \mathrm{Hom}_{\mathsf{C}^{\mathsf{J}}}(\Delta_X, RF) \cong \mathrm{Hom}_{\mathsf{C}}(X, \lim RF)$$

which by the Yoneda lemma implies that $R\lim F \cong \lim RF$. That left adjoints preserve colimits is proved in a dual manner.

Therefore, limits can be defined by their being right adjunct to a constant functor, and colimits by being left adjunct. We extend this to ∞-categories trivially: An ∞-category $\mathcal{C}$ has all limits of shape $\mathcal{J}$ if the constant diagram functor $\Delta : \mathcal{C} \to \mathcal{C}^{\mathcal{J}}$ has a right adjoint, and has all colimits of shape $\mathcal{J}$ if Δ has a left adjoint.

For instance, $\mathcal{C}$ has a terminal object, or a limit over $\varnothing$, if the constant functor $\mathcal{C} \to \mathcal{C}^{\varnothing} \cong 1$ has a right adjoint $1 \to \mathcal{C}$, which we identify as an object of $\mathcal{C}$. Dually, an initial object is identified with a left adjoint to the functor $\mathcal{C} \to 1$.

3.4.3 Stable ∞-Categories

We may interpret the theory of spectra, introduced in A.3.1, in the general context of ∞-category theory, obtaining a notion of *stable* ∞-categories, the primary example of which is the ∞-categorical version of Sp. Following Jacob Lurie's approach [Lurie, 2006], we will end up inventing a higher-categorical analogue of homological algebra along the way (as one does).

Recall that a zero object in a category C is an object which is both initial and final. Given a zero object $0 \in \mathsf{C}$, we can define maps $0^X_Y : X \to Y$ given by the composition $X \to 0 \to Y$; these maps contain "no data", akin to the zero homomorphisms between abelian groups. An

∞-category C is ***pointed*** if it contains a zero object 0. We will always refer to the zero object of a category as 0, the map $X \to 0$ as 0^X, and the map $0 \to Y$ as 0_Y (so that $0^X_Y = 0_Y 0^X$).

In a pointed ∞-category C, the maps 0_{XY} may not exist "uniquely": as with all composition, there are multiple possible fillers of the horn $X \to 0 \to Y$. However, all such fillers are homotopic: by the general argument for composition, we have a skeleton

$$\begin{tikzcd}
& 0 \arrow[dr, "0_Y"] \arrow[dd, "0_Y"] & \\
X \arrow[ur, "0^X"] \arrow[rr, "a"] \arrow[dr, "b"'] & & Y \arrow[dl, equal, "s^0 Y"] \\
& Y &
\end{tikzcd}$$

where a and b are competitors, both claiming to be the composition $0_Y \circ 0^X$, and indeed both have 2-simplices witnessing this. We can get a canonical 2-simplex with edges $0_Y, 0_Y, s^0 Y$ by degeneracy, giving us a 3-horn which we fill with a 3-simplex witnessing the equivalence of a and b as composites $0_Y \circ 0^X$.

We define a ***triangle*** in C with edges $X \xrightarrow{f} Z \xrightarrow{g} Y$ to be a commutative diagram of the form

$$\begin{tikzcd}
X \arrow[r, "f"] \arrow[d] & Z \arrow[d, "g"] \\
0 \arrow[r] & Y
\end{tikzcd}$$

If this diagram is a pullback square, the triangle is said to be ***exact***, and f is said to be a ***kernel*** for g. If it is a pushout, the triangle is ***coexact***, and g a ***cokernel***[5] for f. In particular, we may construct $\ker g$ by pulling 0_Y back along g, and $\operatorname{coker} f$ by pushing 0_X out along f; up to equivalence, all kernels and cokernels arise as pushouts and pullbacks.

A ***stable ∞-category*** is a pointed ∞-category C satisfying the additional properties:

- For all $f : X \to Y$ in C, the pushout $\ker f : Y \to Y +_X 0$ and the pullback $\operatorname{coker} f : X \times_Y 0 \to X$ exist. (All kernels and cokernels exist).

[5] In the context of ∞-category theory, these kernels and cokernels are often called fibers and cofibers, as in [Lurie, 2012].

- Triangles in $\mathcal{C}$ are exact if and only if they are coexact. (In particular, the kernel of the cokernel and the cokernel of the kernel are equivalent).

While much of the motivation for this definition comes from homological algebra, kernels and cokernels in a stable ∞-category $\mathcal{C}$ can behave in remarkably different ways due to the presence of homotopical data. To that end, a pair of suggestive definitions:

- The ***loop space object*** ΩX associated to an object $X \in \mathcal{C}$ is defined to be $\ker 0_X$.
- The ***suspension object*** ΣX associated to X is defined to be $\operatorname{coker} 0^X$.

An $f : X \to Y$ yields morphisms $\Omega f : \Omega X \to \Omega Y$ essentially by the universal property of pullbacks, and likewise yields a morphism $\Sigma f : \Sigma X \to \Sigma Y$ for the dual reason. Hence, Ω and Σ are functors $\mathcal{C} \to \mathcal{C}$. These definitions work in any pointed ∞-category with the necessary (co)limits, and yield adjunctions, but in a stable category, that triangles are exact iff they are coexact implies that Σ and Ω are not just adjoint but inverses to one another.

Stabilization Given a pointed ∞-category $\mathcal{C}$, we define a ***prespectrum object*** of $\mathcal{C}$ to be a functor $X : N(\mathbb{Z} \times \mathbb{Z}) \to \mathcal{C}$ such that $X(i,j)$ is a zero object for all $i \neq j$. We will write $X[n]$ for $X(n,n)$. Prespectrum objects assemble into a full ∞-subcategory of $\operatorname{Fun}(N(\mathbb{Z} \times \mathbb{Z}), \mathcal{C})$ written $\operatorname{PSp}(\mathcal{C})$. The diagram

$$\begin{array}{ccc} X[n] & \longrightarrow & X(n,n+1) = 0 \\ \downarrow & & \downarrow \\ X(n+1,n) = 0 & \longrightarrow & X[n+1] \end{array}$$

determines by universal properties a pair of morphisms $\Sigma X[n] \to X[n+1], X[n] \to \Omega X[n+1]$; these are adjunct to one another under the $\Sigma \dashv \Omega$ adjunction. If all maps $X[n] \to \Omega X[n+1]$ are equivalences, we say that X is a ***spectrum object***. The ∞-subcategory of spectrum objects in $\operatorname{PSp}(\mathcal{C})$ is denoted $\operatorname{Sp}(\mathcal{C})$. The ***stabilization*** of an ∞-category $\mathcal{C}$ with a terminal object 1 is the category of spectrum objects on the subcategory of pointed objects of $\mathcal{C}$, the slice category $\mathcal{C}^1$ whose objects are morphisms $1 \to X$ and whose morphisms are morphisms $X \to Y$ fitting in commutative diagrams [6]. This category is denoted $\operatorname{Stab}(\mathcal{C})$.

[6] This is a pointed category, with zero object given by $\mathrm{id}_1 : 1 \to 1$.

The suspension spectrum functor $\Sigma^\infty : \mathcal{C} \to \mathrm{Stab}(\mathcal{C})$ has a right adjoint, denoted Ω^∞: its action on a spectrum X is to yield $X[0]$.

Consider $\mathrm{Stab}(\mathcal{G}rpd)$, the stabilization of the terminal ∞-topos. By the homotopy hypothesis, this is precisely the ∞-categorical equivalent of the previous category Sp of spectra, and we are content to denote it $\mathcal{S}p$. A commutative monoid object in $\mathcal{S}p$ is known as an $\mathbb{E}_\infty$-***ring***. Breaking this down, we have:

- A spectrum E of pointed ∞-groupoids (equivalent to CW complexes), along with structure maps $\Sigma E_n \to E_{n+1}$.
- A multiplication map $m : E \wedge E \to E$, i.e. a set of maps $m_n : E_n \wedge E_n \to E_n$.
- A unit map $e : \mathbb{S} \to E$.

We require the associativity and commutativity conditions to hold. Here is the motivation for considering only a commutative monoidal structure: in Ab, the commutative monoids are the commutative rings. The stabilization of an ∞-category is a model for homological algebra, bearing a natural triangulated structure on its homotopy category, and hence has many of the interesting categorical properties of Ab, albeit in their ∞-categorical form. (This is described in detail in [Lurie, 2006]). Hence, the higher categorical analogue of commutative rings should be commutative monoids on stable ∞-categories.

3.5 ∞-Topoi

3.5.1 Definition

While our definition of Grothendieck topoi was based on sites, there is a much easier way to define them: a Grothendieck topos $\mathcal{E}$ is a category equipped with a geometric morphism $f^* \dashv f_*$ from $\mathcal{E}$ into some presheaf category $\widehat{\mathsf{C}}$, or in other words a reflective subcategory of $\widehat{\mathsf{C}}$ whose reflector is left exact. We will translate this definition to quasi-categories.

First, we note that in $\mathfrak{QCat}$, Set isn't the appropriate base over which to define ∞-presheaves over an ∞-category $\mathcal{C}$; it doesn't take the data of higher simplices into account! Instead, we define the presheaf category $\widehat{\mathcal{C}}$ to be $\mathrm{Fun}(\mathcal{C}^{\mathrm{op}}, \mathcal{Grpd})$[7]. We define an ∞-geometric morphism between ∞-categories $\mathcal{C}, \mathcal{D}$ to be an ∞-functor $f_* : \mathcal{C} \to \mathcal{D}$ right adjoint to an ∞-functor $f^* : \mathcal{D} \to \mathcal{C}$ which preserves limits.

We define a (Grothendieck-Rezk-Lurie) ∞-***topos*** to be an accessible ∞-category $\mathcal{H}$ equipped with an inclusion i_* into some ∞-category of ∞-presheaves $\widehat{\mathcal{C}}$ which forms the direct image part of an ∞-geometric morphism $i^* \dashv i_*$. ∞-topoi and ∞-geometric morphisms form a category $\mathcal{Topos}$.

We may equivalently describe ∞-topoi constructively as follows: take a small ∞-category $\mathcal{C}$, and a set $F = \{F_i : X_i \to Y_i\}$ of morphisms in $\widehat{\mathcal{C}}$. We define the ∞-category $\mathrm{Shv}(\mathcal{C}, F)$ to be the full subcategory of $\widehat{\mathcal{C}}$ on objects Z such that $\mathrm{Map}_{\widehat{\mathcal{C}}}(-, Z)$ sends each $F_i \in F$ to an equivalence of quasi-categories. An ∞-topos is any ∞-category isomorphic to some inclusion of the form $i_* : \mathrm{Shv}(\mathcal{C}, F) \to \widehat{\mathcal{C}}$ which admits a left exact left adjoint i^*. If we take $F = \varnothing$, we immediately see that every ∞-category of the form $\widehat{\mathcal{C}}$ is an ∞-topos, and in particular that $\mathcal{Grpd} = \mathrm{Shv}(\varnothing, 1)$ is an ∞-topos.

The ∞-topos $\mathcal{Grpd}$ is, by the homotopy hypothesis, equivalent to the category CW of CW complexes (up to weak equivalence of simplicial sets and spaces, respectively). In ∞-topos theory, it plays the role that Set does in 1-topos theory: every other ∞-topos $\mathcal{H} \subseteq \mathrm{Fun}(\mathcal{C}^{\mathrm{op}}, \mathcal{Grpd})$ has one geometric morphism (up to equivalence) into $\mathcal{Grpd}$, making it the terminal ∞-topos. The direct image $\Gamma : \mathcal{H} \to \mathcal{Grpd}$ sends an ∞-sheaf X to its ∞-groupoid of global sections, $\mathrm{Fun}(*, X)$, while the inverse image $\Delta : \mathcal{Grpd} \to \mathcal{H}$ sends an ∞-groupoid $\mathcal{G}$ to the (sheafification of the) constant functor $\mathcal{C}^{\mathrm{op}} \ni X \mapsto \mathcal{G}$.

Truncation As ∞-groupoids are weakly equivalent to CW complexes, we can reason about their homotopy groups in an essentially topological way. In particular, we can reason about *n-connected* ∞-groupoids, whose homotopy groups vanish for n or below, and *homotopy n-type* ∞-groupoids, whose homotopy groups vanish above n.

The latter has a special name: an ∞-groupoid $\mathcal{G}$ is ***n-truncated*** if $\pi_m(\mathcal{G}, x) = 0$ for all $m > n$

[7] $\mathcal{Grpd}$ is identical to the ∞-cosmos $\mathfrak{Kan}$, but the difference in notation allows us to distinguish between its roles as an ∞-cosmos and as a mere ∞-category.

and $x \in \mathcal{G}$, i.e. if its corresponding space is a homotopy n-type. Hence, a 0-truncated ∞-groupoid is simply a set, a 1-truncated ∞-groupoid is a set with nontrivial automorphisms – a *groupoid* – and so on. A -1-truncated ∞-groupoid $\mathcal{G}$ has $\pi_0(\mathcal{G}) = 0$ and therefore is either $\varnothing$ or $\{*\}$. We define a -2-truncated ∞-groupoid to be one which is equivalent to $\{*\}$.

In an arbitrary ∞-category $\mathcal{C}$, an object X is n-truncated if all ∞-groupoids of the form $\mathrm{Map}(-, X)$ are n-truncated, and a morphism $f : X \to Y$ in $\mathcal{C}$ is n-truncated if it is so as an element of $\mathcal{C}_Y$. Write $\tau_{\leq n}\mathcal{C}$ for the (full) subcategory on $\mathcal{C}$'s n-truncated objects[8].

In many cases of interest, τ_n is functorial: given two left exact ∞-categories $\mathcal{C}, \mathcal{D}$, a left exact functor $F : \mathcal{C} \to \mathcal{D}$ sends n-truncated objects and morphisms in $\mathcal{C}$ to n-truncated objects and morphisms in $\mathcal{D}$ (HTT, 5.5.6.16), evidencing the functoriality of $\tau_{\leq n}$ on ∞-categories. If $\mathcal{C}$ is presentable, the inclusion ∞-functor $\tau_{\leq n}\mathcal{C} \to \mathcal{C}$ has a right adjoint $\mathcal{C} \to \tau_{\leq n}\mathcal{C}$, which we will also denote by $\tau_{\leq n}$. (HTT, 5.5.6.18). Both presentability and left exactness hold for ∞-topoi, for example, with both the direct and inverse images of a geometric morphism being left exact as well; hence, objects and morphisms of ∞-topoi can be truncated to arbitrary degree $n \geq -2$, as can geometric morphisms between ∞-topoi. We define an ***n*-topos** to be an ∞-topos of the form $\tau_{\leq(n-1)}\mathcal{H}$, for some ∞-topos $\mathcal{H}$; all Grothendieck topoi arise as 1-topoi in this way, though different ∞-topoi may give rise to the same 1-topos. The terminal ∞-topos $\mathcal{G}rpd$ yields a terminal n-topos $\tau_{\leq(n-1)}\mathcal{G}rpd = (n-1)\text{-}\mathsf{Grpd}$; in particular, the terminal object Set in the 1-category of 1-topoi arises as $\tau_{\leq 0}\mathcal{G}rpd$.

3.5.2 Types of ∞-Topoi

Many properties of ∞-topoi can be expressed in terms of their unique geometric morphism into $\mathcal{G}rpd$. There are many similar properties defined in this form, which for convenience we put into a table. Read f as the geometric morphism $f^* \dashv f_*$; the Adjunction column lists all adjunctions that must exist to satisfy the corresponding condition on f, and the Additional column lists conditions on these adjunctions. If the terminal geometric morphism from an ∞-topos $\mathcal{H}$ into $\mathcal{G}rpd$ satisfies the given conditions, the rightmost column tells us what to call

[8] Fixing an n-truncated Y, n-truncation of X implies and is implied by n-truncation of f (HTT, 5.5.6.14), so all morphisms in $\tau_{\leq n}\mathcal{C}$ are automatically n-truncated.

$\mathcal{H}$. For instance, an *essential* geometric morphism is one with whose inverse image has a left adjoint, and an ∞-topos whose terminal geometric morphism is essential is known as *locally connected*.

Condition on f	Adjunction	Additional	For $\Gamma : \mathcal{H} \to \mathcal{Grpd}$
Connected	$f^* \dashv f_*$	f^* f.f.	$\times$
Essential	$f_! \dashv f^* \dashv f_*$	n.a.	Locally connected
$\times$	$f_! \dashv f^* \dashv f_*$	f_* f.f.	Connected
Strongly connected	$f_! \dashv f^* \dashv f_*$	$f_!$ cartesian	Strongly connected
Totally connected	$f_! \dashv f^* \dashv f_*$	$f_!$ lex	Totally connected
$\times$	$f^* \dashv f_* \dashv f^!$	n.a.	Locally local
Local	$f^* \dashv f_* \dashv f^!$	f^* f.f.	Local
Discrete	$f^* \dashv f_* \dashv f^!$	f^*, f_* f.f.	$\times$
Cohesive	$f_! \dashv f^* \dashv f_* \dashv f^!$	f^* f.f., $f_!$ cartesian	Cohesive

$\times$: No recognized name. F.f.: fully faithful[9]. Lex: left exact.

In general, we will denote a right adjoint to the global sections functor $\Gamma : \mathcal{H} \to \mathcal{Grpd}$ by ∇, and a left adjoint to the constant sheaf functor $\Delta : \mathcal{Grpd} \to \mathcal{H}$ by Π. An adjoint quadruple, as seen in for instance a cohesive ∞-topos, will therefore be represented as follows:

$$\begin{tikzcd} \mathcal{H} \arrow[r, shift left=3, "\Gamma"] \arrow[r, shift right=9, "\Pi"'] & \mathcal{Grpd} \arrow[l, shift right=9, "\nabla"'] \arrow[l, shift left=3, "\Delta"] \end{tikzcd} \qquad \Pi \dashv \Delta \dashv \Gamma \dashv \nabla$$

3.5.3 Cohesion, Elasticity, and Solidity

Cohesion Read off the previous table that a geometric morphism $f : \mathcal{E} \to \mathcal{F}$ is ***cohesive*** if its inverse image f^* is fully faithful and has an additional left adjoint $f_!$ which preserves

[9] A left (right) adjoint is fully faithful iff its unit (counit) is a natural isomorphism. Given an adjoint triple $F \dashv G \dashv H$, we have $\mathrm{Hom}(GFX, Y) = \mathrm{Hom}(FX, HY) = \mathrm{Hom}(X, GHY)$, so that GF is left adjoint to GH; adjoints are unique up to natural isomorphism, so one is naturally isomorphic to the identity iff the other is. Therefore F is fully faithful precisely if H is. In particular, in the adjoint triple $f_! \dashv f^* \dashv f_*$, f_* is fully faithful precisely if $f_!$ is.

finite products, *and* its direct image f_* has an additional right adjoint $f^!$, giving us an adjoint quadruple $f_! \dashv f^* \dashv f_* \dashv f^!$.

Take a cohesive ∞-topos $\mathcal{H}$, with adjoint quadruple $\Pi \dashv \Delta \dashv \Gamma \dashv \nabla$ into $\mathcal{Grpd}$. We think of the functors $\Delta, \nabla : \mathcal{Grpd} \to \mathcal{H}$ as sending an ∞-groupoid X to a *discrete space*, in which every object of X is distinguished in the space ΔX, and a *codiscrete space*, in which X is treated as one big "clump" in the space ∇X. The functors $\Pi, \Gamma : \mathcal{H} \to \mathcal{Grpd}$ we think of as sending a space $X \in \mathcal{H}$ to its space of connected components and its space of points, respectively.

The adjoint quadruple induces an adjoint triple of endofunctors on $\mathcal{H}$,

$$\Delta \circ \Pi \dashv \Delta \circ \Gamma \dashv \nabla \circ \Gamma$$

the first and last of which are idempotent monads, and the second of which is an idempotent comonad.

- The leftmost endofunctor acts on a space $X \in \mathcal{H}$ by dropping all information internal to each connected component, rendering the collection of all connected components discrete, keeping the shape of X; it is known as the ***shape modality*** $\int$.

- The middle endofunctor dissolves all spatial structure on X, endowing its points with the discrete topology, while the rightmost endofunctor dissolves spatial structure and then endows points with the codiscrete topology; they are known as the ***flat*** and ***sharp modalities***, $\flat$ and $\sharp$.

  Objects X for which $\flat X \cong X$ or $\sharp X \cong X$ are known as ***discrete*** and ***codiscrete objects***, respectively.

Elasticity Take two cohesive ∞-topoi $\mathcal{E}, \mathcal{F}$, and let $i_{inf} : \mathcal{E} \to \mathcal{F}$ be a functor fitting into a series of adjunctions

$$i_{inf} \dashv \Pi_{inf} \dashv \Delta_{inf} \dashv \Gamma_{inf}$$

such that $\Pi_{\mathcal{F}} = \Pi_{\mathcal{E}} \circ \Pi_{inf}$ and likewise for Δ and Γ. We say that $\mathcal{F}$ is an ***elastic topos*** over $\mathcal{E}$, or ***differentially cohesive***. So, the situation is as follows:

$$\begin{tikzcd}
\mathcal{G}rpd \arrow[r, shift left=4.5, leftarrow, "\Pi_{\mathcal{E}}"] \arrow[r, shift left=1.5, "\Delta_{\mathcal{E}}"] \arrow[r, shift right=1.5, leftarrow, "\Gamma_{\mathcal{E}}"] \arrow[r, shift right=4.5, "\nabla_{\mathcal{E}}"] \arrow[rr, bend right, "\nabla_{\mathcal{F}}"] & \mathcal{E} \arrow[r, shift left=6, "i_{inf}"] \arrow[r, shift left=3, leftarrow, "\Pi_{inf}"] \arrow[r, "\Delta_{inf}"] \arrow[r, shift right=3, leftarrow, "\Gamma_{inf}"] & \mathcal{F}
\end{tikzcd}$$

Again, each of $\mathcal{E}$ and $\mathcal{F}$ have their own co/monads $(\int_{\mathcal{E}}, \flat_{\mathcal{E}}, \sharp_{\mathcal{E}})$, $(\int_{\mathcal{F}}, \flat_{\mathcal{F}}, \sharp_{\mathcal{F}})$, but we now have an additional adjoint triple

$$i_{inf} \circ \Pi_{inf} \dashv \Delta_{inf} \circ \Pi_{inf} \dashv \Delta_{inf} \circ \Gamma_{inf}$$

of endofunctors on $\mathcal{F}$.

- $i_{inf} \circ \Pi_{inf}$ is an idempotent comonad known as the ***reduction modality*** $\Re$. Objects invariant under $\Re$ are known as ***reduced***.
- $\Delta_{inf} \circ \Pi_{inf}$ is an idempotent monad, the ***infinitesimal shape modality*** $\Im$. Objects invariant under $\Im$ are known as ***coreduced***.
- $\Delta_{inf} \circ \Gamma_{inf}$ is an idempotent comonad, the ***infinitesimal flat modality*** &.

We have

$$\& \circ \flat_{\mathcal{F}} = \Delta_{inf} \circ \Gamma_{inf} \circ \Delta_{\mathcal{F}} \circ \Gamma_{\mathcal{F}} = \Delta_{inf} \circ \Gamma_{inf} \circ \Delta_{inf} \circ \Delta_{\mathcal{E}} \circ \Gamma_{\mathcal{E}} \circ \Gamma_{inf}$$

$$= \Delta_{inf} \circ \Gamma_{inf} \circ \Delta_{inf} \circ \flat_{\mathcal{E}} \circ \Gamma_{inf} = \Delta_{inf} \circ \flat_{\mathcal{E}} \circ \Gamma_{inf} = \flat_{\mathcal{F}}$$

So all objects invariant under $\flat_{\mathcal{F}}$ are invariant under & as well, implying that the modality of $\flat_{\mathcal{F}}$ is *subsumed* by that of &; we write this relation as $\& > \flat$. Likewise, $\Im > \int_{\mathcal{F}}$. (Careful factorization suffices to prove all the relations of this form that we will encounter, and will be omitted).

Solidity Let $\mathcal{E}, \mathcal{F}, \mathcal{G}$ be cohesive ∞-topoi with both $\mathcal{F}$ and $\mathcal{G}$ elastic over $\mathcal{E}$. $\mathcal{G}$ is a ***solid*** ∞-topos if it is equipped with a functor even : $\mathcal{G} \to \mathcal{F}$ with a series of right adjoints

$$\text{even} \dashv i_{sup} \dashv \Pi_{sup} \dashv \Delta_{sup} \dashv \Gamma_{sup}$$

such that $\Pi_{\mathcal{G}} = \Pi_{\mathcal{F}} \circ \Pi_{sup} = \Pi_{\mathcal{E}} \circ \Pi_{inf} \circ \Pi_{sup}$, and likewise for $\Delta_{\mathcal{G}}$ and $\Gamma_{\mathcal{G}}$.

The situation is as follows:

$$\begin{tikzcd}[column sep=huge]
\mathcal{G}rpd
\arrow[r, shift left=4.5, leftarrow, "\Pi_{\mathcal{E}}"]
\arrow[r, shift left=1.5, "\Delta_{\mathcal{E}}"]
\arrow[r, shift right=1.5, leftarrow, "\Gamma_{\mathcal{E}}"]
\arrow[r, shift right=4.5, "\nabla_{\mathcal{E}}"]
\arrow[rr, bend right, "\nabla_{\mathcal{F}}"']
\arrow[rrr, bend right=40, "\nabla_{\mathcal{G}}"']
& \mathcal{E}
\arrow[r, shift left=4.5, "i_{inf}"]
\arrow[r, shift left=1.5, leftarrow, "\Pi_{inf}"]
\arrow[r, shift right=1.5, "\Delta_{inf}"]
\arrow[r, shift right=4.5, leftarrow, "\Gamma_{inf}"]
& \mathcal{F}
\arrow[r, shift left=7.5, leftarrow, "\mathrm{even}"]
\arrow[r, shift left=4.5, "i_{sup}"]
\arrow[r, shift left=1.5, leftarrow, "\Pi_{sup}"]
\arrow[r, shift right=1.5, "\Delta_{sup}"]
\arrow[r, shift right=4.5, leftarrow, "\Gamma_{sup}"]
& \mathcal{G}
\end{tikzcd}$$

We again have a triplet of endofunctors on $\mathcal{G}$:

- The idempotent monad $i_{sup} \circ \mathrm{even}$ known as the ***fermionic modality*** $\rightrightarrows$
- The idempotent comonad $i_{sup} \circ \Pi_{sup}$ known as the ***bosonic modality*** $\rightsquigarrow$
- The idempotent monad $\Delta_{sup} \circ \Pi_{sup}$ known as the ***rheonomy modality*** Rh

By being elastic over $\mathcal{E}$ and cohesive over $\mathcal{G}rpd$, it also has the two previous triplets of modalities, and admits the relations $\rightsquigarrow > \Re$ and $\mathrm{Rh} > \Im$. We therefore have three generations of modalities, which [nLab authors, 2020] arranges into the progression

$$\begin{array}{ccccccccl}
\mathrm{id} & \dashv & \mathrm{id} & & & & & & \\
\vee & & \vee & & & & & & \\
\rightrightarrows & \dashv & \rightsquigarrow & \dashv & \mathrm{Rh} & & & & \text{solidity} \\
 & & \vee & & \vee & & & & \\
 & & \Re & \dashv & \Im & \dashv & \& & & \text{elasticity} \\
 & & & & \vee & & \vee & & \\
 & & & & \int & \dashv & \flat & \dashv \ \sharp & \text{cohesion} \\
 & & & & & & \vee & \quad\ \vee & \\
 & & & & & & \varnothing & \dashv \ * &
\end{array}$$

including the trivial adjunctions $\mathrm{id} \dashv \mathrm{id}$ and $\varnothing \dashv *$, where $\varnothing$ and $*$ are the constant endofunctors on the initial and terminal objects, respectively.

Chapter 4

Homotopy Type Theory

4.1 Type Theory

4.2 Homotopy Type Theory

Chapter 5

Modality

5.1 Modal Logic

5.1.1 Idea

5.1.2 Classical Logic

5.1.3 Modal Logic

5.1.4 Modal Axioms

5.2 ∞-Monads

5.2.1 2-Categorical Adjunctions

The 2-Category of Quasi-Categories Recall that the proper notion of a hom-set in an ∞-groupoid is that of the mapping space $\mathrm{Map}_{\mathcal{C}}(X, Y)$, which is the pullback

$$\begin{tikzcd}
\mathrm{Map}_{\mathcal{C}}(X,Y) \arrow[r] \arrow[d] \arrow[dr, phantom, "\lrcorner"] & \mathrm{Fun}(\Delta^1, \mathcal{C}) \arrow[d, "\mathrm{Fun}(i,\mathcal{C})"] \\
* \arrow[r, "\text{to } (X,Y)"'] & \mathcal{C} \times \mathcal{C} \cong \mathrm{Fun}(\partial\Delta^1, \mathcal{C})
\end{tikzcd}$$

This is an ∞-groupoid, the objects[1] of which are equivalent to morphisms $X \to Y$ and the morphisms of which are maps $f : \Delta^1 \times \Delta^1 \to \mathcal{C}$ such that $f(\{0\}, s)$ is degenerate at X and $f(\{1\}, s)$ is degenerate at Y. Homotopically speaking, these are identifications of equivalent morphisms in $\mathcal{C}$, and we may take $\pi_0 \mathrm{Map}_{\mathcal{C}}(X, Y)$ for all X, Y to get a proper category associated to $\mathcal{C}$.

The proper notion of the *functor category* from $\mathcal{C}$ to $\mathcal{D}$ is given by the internal hom in simplicial sets,

$$(\mathcal{D}^{\mathcal{C}})_n = \mathrm{Fun}_{\mathfrak{QCat}}(\Delta^n \times \mathcal{C}, \mathcal{D}) = \mathrm{Hom}_{\mathsf{sSet}}(\Delta^n \times \mathcal{C}, \mathcal{D})$$

a quasi-category itself. Hence, we may consider the category of quasi-categories to be a 2-category, with objects the quasi-categories and hom-categories the homotopy categories of the internal homs. We will denote this 2-category QCat_2.

Whiskering in QCat_2 is given by actual composition of maps of simplicial sets, e.g. $L\eta = L \circ \underline{\eta} : \Delta^1 \times \mathcal{C} \to \mathcal{C} \to \mathcal{D}$. Composition of 2-cells $\alpha : A \Rightarrow B$, $\beta : B \Rightarrow C$ within $\mathrm{Fun}(\mathcal{C}, \mathcal{D})$ is given by taking the 1-simplices represented by $\underline{\alpha}(-, X)$ and $\underline{\beta}(-, X)$, which form an image of the horn Λ^2_1, and filling this horn in to obtain a further 1-simplex, which we associate to $(\underline{\beta} \circ \underline{\alpha})(-, X)$.

The Walking Adjunction 2-categorically, adjunctions merely represent a selection of 1-cells and 2-cells satisfying certain coherence data. Hence, we might expect there to be a 2-category freely built on this data which represents adjunctions. We may construct such a 2-category as follows:

- Consider two objects, denoted $+$ and $-$.
- Add two 1-cells $L : + \to -$, $R : - \to +$.
- Add two 2-cells $\eta : \mathrm{id}_+ \Rightarrow RL$, $\epsilon : LR \Rightarrow \mathrm{id}_-$
- Freely generate the category on these cells, subject to the relations $(\epsilon L) \circ (L\eta) = \mathrm{id}_L$ and $(R\epsilon) \circ (\eta R) = \mathrm{id}_R$.

The resulting 2-category, which consists of an adjunction $L \dashv R$ and no other data, is known as the ***walking adjunction*** Adj. Recalling that the 2-category of adjunctions in a 2-category C,

[1] The terms object, 0-cell, 0-simplex, and vertex are interchangeable here, as are the terms morphism, 1-cell, 1-simplex, and edge; we will use n-cell for higher cells/simplices.

denoted $\mathsf{Adj}(\mathsf{C})$, has as its 1-cells the adjunctions in C and as its 2-cells pairs of mates between adjunctions, we have $\mathrm{Hom}_{\mathsf{2\text{-}Cat}}(\mathsf{Adj}, \mathsf{C}) \cong \mathsf{Adj}(\mathsf{C})$.

The hom-categories of Adj can be characterized as follows [Riehl and Verity, 2016a]:

- $\mathrm{Hom}_{\mathsf{Adj}}(+,+) \cong \mathbf{\Delta}_+$, the category obtained by appending $[-1] = \varnothing$ to $\mathbf{\Delta}$; id_+ is given by $[-1]$.
- $\mathrm{Hom}_{\mathsf{Adj}}(-,+) \cong \mathbf{\Delta}_\infty$, the wide subcategory of $\mathbf{\Delta}$ on those morphisms preserving greatest elements.
- $\mathrm{Hom}_{\mathsf{Adj}}(+,-) \cong \mathbf{\Delta}_{-\infty}$, the wide subcategory of $\mathbf{\Delta}$ on those morphisms preserving 0. (Equivalently, $\mathbf{\Delta}_\infty^{\mathrm{op}}$).
- $\mathrm{Hom}_{\mathsf{Adj}}(-,-) \cong \mathbf{\Delta}_+^{\mathrm{op}}$.

Homotopy Coherent Adjunctions Given ∞-categories (quasi-categories) $\mathcal{C}, \mathcal{D}$ and a pair of ∞-functors $L : \mathcal{C} \to \mathcal{D}$, $R : \mathcal{D} \to \mathcal{C}$, L and R are left and right adjoints when they are so in QCat_2; hence, whenever there is a 2-cell $\eta : \mathrm{id}_{\mathcal{C}} \Rightarrow RL$, or a map of simplicial sets $\underline{\eta} : \Delta^1 \times \mathcal{C} \to \mathcal{C}$ with $\underline{\eta}(\{0\}, X) = X$ and $\underline{\eta}(\{1\}, X) = RLX$, as well as a 2-cell $\epsilon : LR \Rightarrow \mathrm{id}_{\mathcal{D}}$ associated to a map $\underline{\epsilon} : \Delta^1 \times \mathcal{D} \to \mathcal{D}$. We require that $(\epsilon L) \circ (L\eta) = \mathrm{id}_L$ and $(R\epsilon) \circ (\eta R) = \mathrm{id}_R$ as 2-cells in QCat_2, and therefore that they are *homotopic* as 1-simplices in $\mathrm{Fun}(\mathcal{C}, \mathcal{D})$ and $\mathrm{Fun}(\mathcal{D}, \mathcal{C})$.

2-categories embed into simplicially enriched categories by replacing hom-categories with their nerves. In this manner, we replace the walking adjunction $\mathsf{Adj} \in \mathsf{2\text{-}Cat}$ with the ***free homotopy coherent adjunction*** $\mathcal{A}dj \in \mathsf{sSet\text{-}Cat}$, with simplicially enriched functors from $\mathcal{A}dj$ to $\mathsf{sSet\text{-}Cat}$ (for instance, to $\mathfrak{QCat}$, or any other ∞-cosmos) being known as ***homotopy coherent adjunctions***.

Any homotopy coherent adjunction yields an adjunction in QCat_2 via forgetting higher simplicial data. Conversely, given an adjunction $\mathsf{Adj} \to \mathsf{QCat}_2$, the set of homotopy coherent adjunctions $\mathcal{A}dj \to \mathsf{sSet\text{-}Cat}$ extending this adjunction forms a contractible Kan complex [Riehl and Verity, 2016a]; hence, there is, up to homotopy, a unique homotopy coherent adjunction associated to any 2-categorical adjunction.

5.2.2 Homotopy Coherent Monads

Monads Consider the sub-2-category of Adj consisting solely of the object $+$ and its hom-category $\mathbf{\Delta}_+$. A functor F from this 2-category to a 2-category C will have the following data:

- An object $X = F+$ of C.
- For each integer $n \geq -1$, a 1-cell $F^{n+1} = F[n] : X \to X$. In particular, $F^0 = F\mathrm{id}_+ = \mathrm{id}_X$.
- For each order-preserving map $f : [m] \to [n]$, $m, n \geq -1$, a 2-cell $f : F^m \Rightarrow F^n$.

The 1-cells satisfy $F^m \circ F^n = F^{m+n}$, and the 2-cells satisfy the simplicial identities. In particular, if we suggestively label the unique 2-cell $F^0 \Rightarrow F^1$ by η, and the unique 2-cell $F^2 \Rightarrow F^1$ by μ, then $\mu \circ F^1\eta = \mu \circ \eta F^1$:

$$
\begin{tikzcd}
 & 1 \arrow[dr, mapsto] & \\
0 \arrow[r, mapsto] & 0 \arrow[r, mapsto] & 0
\end{tikzcd}
\quad = \quad
\begin{tikzcd}
 & 1 \arrow[dr, mapsto] & \\
0 \arrow[ur, mapsto] & 0 \arrow[r, mapsto] & 0
\end{tikzcd}
$$

and $\mu \circ (T\mu) = \mu \circ (\mu T)$:

$$
\begin{tikzcd}
2 \arrow[dr, mapsto] & & \\
1 \arrow[dr, mapsto] & 1 \arrow[dr, mapsto] & \\
0 \arrow[r, mapsto] & 0 \arrow[r, mapsto] & 0
\end{tikzcd}
\quad = \quad
\begin{tikzcd}
2 \arrow[dr, mapsto] & & \\
1 \arrow[r, mapsto] & 1 \arrow[dr, mapsto] & \\
0 \arrow[r, mapsto] & 0 \arrow[r, mapsto] & 0
\end{tikzcd}
$$

Hence, a functor from $\{+\}$ to C defines a ***monad*** on the object $X \in \mathsf{C}$. We denote the 2-category $\{+\}$ by Mnd. Again, we obtain the free homotopy coherent monad $\mathcal{M}nd \in \mathsf{sSet\text{-}Cat}$ by taking nerves of hom-categories, defining a ***homotopy coherent monad***, or ∞-monad, to be a simplicially enriched functor from $\mathcal{M}nd$. This continues to work in any ∞-cosmos; our primary focus shall be $\mathfrak{QCat}$.

Dually, a functor from $\{-\}$ to C defines a ***comonad*** on the target object, so we denote the 2-category $\{-\}$ by Cmnd; taking nerves, we get ***homotopy coherent comonads***, or ∞-comonads, as simplicially enriched functors from the nerve $\mathcal{C}mnd$ of Cmnd.

5.3 Modal Homotopy Types

Part III

Four Notions of Space

Introduction

One of the core aims of metaphysics is an explication of the nature of the physical dimensions of our experience: of the phenomena of *space* and *time*. Einstein's general theory of relativity demonstrated that the two concepts could be unified in a manner consistent with experience: that time was a special kind of space, differing only due to a metric in which one dimension, that of time, bore a different sign from the others. While we now know that general relativity is, strictly speaking, incorrect (it breaks down at the smallest scales), we believe that any correct theory would still treat time as a kind of space, restricting the study of space and time to that of space and the various structures that may be placed on it.

In this part, we will describe three conceptions of space, each with a different categorical structure:

1. ***Synthetic differential geometry***, a synthetic conception of space taking place within the internal language of a topos. The motto: *a space is an object which satisfies a certain set of axioms*.

2. ***Diffeological spaces***, a sheaf-theoretic conception of space that uses the geometry of $\mathbb{R}^n$ to generate a notion of "smooth structure" going far beyond that of smooth manifolds. The motto: *a space is a set which can be smoothly probed*.

3. ***Noncommutative geometry***, an algebraic conception of space relying on the description of topological spaces by their algebras of functions to develop a spatial theory of algebras of functions in general. The motto: *a space is its set of functions*.

4. ***Structured spaces***, a higher categorical conception of space that generalizes the algebro-geometric notion of a locally ringed space to provide an account of geometric objects in arbitrary ∞-topoi. The motto: *???*.

We will then describe some of the phenomenological aspects of space, from the philosophical, physical, and mathematical perspectives.

Chapter 6

Synthetic Differential Geometry

6.1 Infinitesimals

6.1.1 The Kock-Lawvere Axiom

Given a commutative ring object R in a topos $\mathcal{E}$, we define the subobject of ***infinitesimals*** of R by $D := \{x \in R \mid x^2 = 0\}$. The ***Kock-Lawvere axiom*** for R reads

$$(\forall f \in R^D)(\exists! c \in R)\left((\forall \epsilon \in D)(f(\epsilon) = f(0) + c\epsilon)\right)$$

Clearly $0 \in D$, so $0 : 1 \to R$ factors through D. As a consequence, we have that if $c_1\epsilon = c_2\epsilon$ for all $\epsilon \in D$, then $c_1 = c_2$ (let $f(\epsilon) = c_1\epsilon$). The KL axiom allows us to work with infinitesimals as though they actually exist, using them to define derivatives around points. However, this comes at a cost: we cannot in general exhibit non-zero infinitesimals.

In order to work with the KL axiom, we must explicitly reject the principle of excluded middle: to see this, define a map $f : D \to R$ which sends ϵ to 0 if $\epsilon = 0$ and to 1 otherwise; the KL axiom implies that there's a unique $c \in R$ such that $f(\epsilon) = c \cdot \epsilon$ for all $\epsilon \in D$. Assuming the LEM, either D contains only 0 or D contains other elements. If D contains only 0, then c cannot be unique; hence, it contains an $\epsilon \neq 0$, and a unique c such that $c\epsilon = 1$. It follows that $0 = (c\epsilon)^2 = 1^2 = 1$, a contradiction. Hence, we must throw out the LEM, and work constructively. Another consequence of this is the *undecidability* of R: the sentence $(\forall x, y)(x = y \wedge x \neq y)$ is not true. In particular, $\mathcal{E}$ cannot show that infinitesimals are non-zero.

This is in part because the KL axiom is very strong: fixing an $x \in R$, $f : R \to R$, and $k : D \to R$ sending 0 to $f(x)$ and ϵ to $k(\epsilon) = f(x+\epsilon)$, the KL axiom gives a unique c_x in R such that $f(x+\epsilon) = f(x) + c\epsilon$. We write $f'(x) := c_x$ to get a function $f' : R \to R$ known as the derivative of f, and state ***Taylor's formula***:

$$\forall \epsilon \in D(f(x+\epsilon) = f(x) + \epsilon f'(x))$$

So KL implies that every function $f : R \to R$ is differentiable.

An Alternative Statement Here's another statement equivalent to the KL axiom: take the R-algebra $R[\varepsilon] = R \times R$ with multiplication $(a,b) \cdot (c,d) = (ac, ad+bc)$. Then (KL2), the map $\alpha : R[\varepsilon] \to R^D, \alpha(a,b)(\epsilon) = a + \epsilon b$ is an R-algebra isomorphism.

It's clear that $(\alpha(a,b)\alpha(c,d))(\epsilon) = \alpha(ac, ad+bc)(\epsilon)$, as well as that this statement, KL2, implies the original statement (KL1). To see the converse, assume KL1. Then, not only is every function f of the form $\alpha(f(0), c)$, but for every $\alpha(a,b)$ there is a unique $c \in R$ such that $a + b\epsilon = \alpha(a,b)(\epsilon) = \alpha(a,b)(0) + c\epsilon = a + c\epsilon$ for all ϵ; b obviously satisfies this, and hence is the only element of R that satisfies this, making it, and hence the pair (a,b) recoverable from the function $\alpha(a,b)$. So KL1 is equivalent to KL2.

Spectra Given an arbitrary R-algebra $A \in \mathcal{E}$ and a finitely generated R-algebra $B = R[x_1, \ldots, x_n]/I$, for instance a Weil algebra, the ***spectrum*** $\mathrm{Spec}_A(B)$ is a subobject of A^n consisting of those $a = (a_1, \ldots, a_n)$ such that $P(a) = 0$ for all $P \in I$. For instance, $\mathrm{Spec}_R(R[x]/(x^2)) = \{x \in R \mid x^2 = 0\} = D$. For W a Weil algebra, the object $\mathrm{Spec}_R(W)$ is known as the ***formal infinitesimals*** object of R (with respect to W). The process of taking spectra with respect to R is functorial: a morphism $\psi : W \to W'$ of Weil algebras generates a morphism $\Psi : \mathrm{Spec}_R(W') \to \mathrm{Spec}_R(W)$

A third formulation of the KL axiom states that (KL3) the R-algebra homomorphism $\alpha : W \to R^{\mathrm{Spec}_R(W)}$, $\alpha(P)(x_1, \ldots, x_n) = P(x_1, \ldots, x_n)$, is an isomorphism. In the topos $\mathcal{E}$, every Weil algebra W yields a functor $(-)^{\mathrm{Spec}_R W}$ which is right adjoint to the functor $- \times \mathrm{Spec}_R W$. If each W satisfies the KL axiom and $(-)^{\mathrm{Spec}_R W}$ is always a left adjoint as well, $\mathcal{E}$ is known as a ***smooth topos***. The right adjoint, known as the ***amazing right adjoint***, is denoted $(-)^{1/\mathrm{Spec}_R W}$.

Differentiation The differentiation given by the KL axiom satisfies the usual properties: for instance, consider two functions $g, f : R \to R$. $(gf)(x+\epsilon)$ is equal to $(gf)(x) + \epsilon(gf)'(x)$, but also equal to $g(f(x) + \epsilon f'(x))$, which since $\epsilon f'(x)$ is an infinitesimal is itself equal to $(gf)(x) + \epsilon f'(x)g'(f(x))$, implying that $(gf)'(x) = f'(x)(g'f)(x)$, i.e. the chain rule. Similarly, differentiation satisfies the product rule, is R-linear, sends constants to 0, and sends id_R to 1.

We define D_n to be the set of all nth order infinitesimals, or elements $x \in R$ such that $x^{n+1} = 0$. (In particular, $D = D_1$). D_∞ is defined to be the set of all nilpotent elements, or $x \in R$ such that $x^n = 0$ for some $n \geq 1$. Supposing that $2, 3, \ldots$ are invertible in R, the higher order extensions of the KL axiom are as follows:

$$\forall f \in R^{D_n} \exists! c_1, \ldots, c_n \in R \left(\forall \epsilon \in D_n (f(\epsilon) = f(0) + c_1\epsilon^1 + c_2\epsilon^2 + \ldots + c_n\epsilon^n) \right)$$

and the corresponding Taylor formulas are

$$\forall \epsilon \in D_n \left(f(x+\epsilon) = f(x) + \epsilon f'(x) + \frac{\epsilon^2}{2} f''(x) + \ldots + \frac{\epsilon^n}{n!} f^{(n)}(x) \right)$$

An R-module V satisfying the following vector version of the KL axiom is known as a ***Euclidean R-module***:

$$\forall f \in V^D \exists! v \in V \left(\forall \epsilon \in D (f(\epsilon) = f(0) + \epsilon \cdot v) \right)$$

When $V \cong R^n$, we can write $\vec{x} = (x_1, \ldots, x_n)$, and we have for a function $g : R^n \to R^n$ such that $g(\vec{x} + \epsilon \cdot \vec{y}) = f(\epsilon)$ a $\vec{z} \in R^n$ such that $g(\vec{x} + \epsilon \cdot \vec{y}) = g(\vec{x}) + \epsilon \cdot \vec{z}$. We define the ***directional derivative*** $\partial_{\vec{y}} g$ of g in the direction $\vec{y}$ to be this $\vec{z}$, and the ith ***partial derivative*** $\partial_i f$ to be the directional derivative in the direction of the ith unit vector. The map $\vec{y} \to \partial_{\vec{y}} g$ is known as the ***differential*** g' of g.

6.1.2 Differential Geometry

Microlinear Spaces Given a topos $\mathcal{E}$ and a commutative ring object R satisfying the KL axiom, take the nested categories $\mathsf{Weil} \subseteq R\text{-}\mathsf{Alg}_{\mathsf{FP}} \subseteq R\text{-}\mathsf{Alg}$ of Weil algebras, finitely presented R-algebra objects, and R-algebra objects, respectively. We have a pair of functors $R^- : \mathcal{E}^{\mathrm{op}} \to \mathcal{E}$ and $\mathrm{Spec}_R : R\text{-}\mathsf{Alg}_{\mathsf{FP}}^{\mathrm{op}} \supseteq \mathsf{Weil}^{\mathrm{op}} \to \mathcal{E}$. Given a finite limit diagram $\mathcal{J}$ of Weil algebras, $\mathcal{D} =$

$\mathrm{Spec}_R(\mathcal{J})$ is, while not necessarily a colimit, at least a cocone. An object $M \in \mathcal{E}$ is a ***microlinear space*** if $M^{\mathcal{D}}$ is a limit diagram for every $\mathcal{J}$. Microlinear spaces will serve as our generalized manifolds. These spaces contain R, are closed under limits (e.g., arbitrary products), and contain exponentials: if M is microlinear and X an arbitrary object, M^X is again microlinear. Thus, we already have a rich abundance of microlinear spaces. A ***Lie group*** is a group internal to $\mathcal{E}$ which is also a microlinear space; again, the trivial example is R.

Tangent Vectors Given a microlinear space M, a ***vector bundle*** over M is an epic $E = \pi : E \to M$ such that $\pi^{-1}(x)$ is a Euclidean R-module, and a ***section***, also known as an E-***vector field***, of the vector bundle E is a morphism $s : M \to E$ such that $\pi s = \mathrm{id}_M$. The ***tangent bundle*** of a microlinear space M is the object M^D equipped with a map $\pi : M^D \to M, t \mapsto t(0)$; its elements are ***tangent vectors***, and the ***tangent space*** of M at a point x is the collection M^D_x of $t \in M^D$ with $\pi(t) = t(0) = x$. We write $TM = M^D, T_xM = M^D_x$, and think of elements of M^D as probings of M in infinitesimal directions, hence tangent vectors. A TM-vector field, just known as a vector field, is a map $M \to M^D$ satisfying the above properties; by cartesian closure, we can look at a vector field X not just as a map $M \to M^D$, but as a map $M \times D \to M$, or even as a map $D \to M^M$ taking an infinitesimal d and giving us an infinitesimal deformation X_d of M. Using this definition, the object $\mathfrak{X}(M)$ of all vector fields on M becomes an R-module under the action $(rX)_d = X_{rd}$. This definition also allows isomorphisms φ to act on vector fields X: we define $(\varphi_* X)_d = \varphi X_d \varphi^{-1}$. If φ is an *endo*morphism, we may define $(\varphi^*\omega)(v) = \omega(\varphi \circ v)$.

Given a $v \in M^{D^n}$, which we think of as a function taking in n infinitesimals and outputting an element of the microlinear space M, as well as an $r \in R$, we define $r_k v(d_1, \ldots, d_n) = v(d_1, \ldots, rd_k, \ldots, d_n)$. Given a $\sigma \in S_n$, we define $v^\sigma(d_1, \ldots, d_n) = v(d_{\sigma 1}, \ldots, d_{\sigma n})$. An n***-form*** on M is a map $\omega : M^{D^n} \to R$ such that $\omega(r_k v) = r\omega(v)$ and $\omega(v^\sigma) = (-1)^\sigma \omega(v)$. The object $\Lambda^n(M)$ of all n-forms on M is a microlinear space as well as a Euclidean R-module. We denote by $X * v$ the element of $M^{D^{n+1}}$ given by $(X * v)(d_1, \ldots, d_{n+1}) = X_{d_1}(v(d_2, \ldots, d_{n+1}))$, and by $i_X\omega$ the $(n-1)$-form acting on a $w \in M^{D^{n-1}}$ by $(i_X\omega)(w) = \omega(X * w)$.

For $X, Y \in \mathfrak{X}(M)$, we define $[X, Y]_{d_1 d_2} = Y_{-d_2} X_{-d_1} Y_{d_2} X_{d_1}$; the vector field $[X, Y]$ is also written $L_X Y$, and is equivalently the unique vector field such that $(X_{-d})_* Y - Y = dL_X Y$. The

exterior derivative of an n-form ω is given by

$$(d\omega)(v) = \sum_{i=1}^{n+1} (-1)^{i+1} (F_v^i)'(0)$$

where $F_v^i(e) = \omega(v(d_1, \ldots, d_{i-1}, e, d_{i+1}, \ldots, d_n))$; as expected, it satisfies $d^2 = 0$. With this in mind, we state Cartan's three "magical formulae" without proof: $L_{[X,Y]} = L_{[X,}L_{Y]}$, $i_{[X,Y]} = L_{[X,}i_{Y]}$, and $L_X = di_X + i_X d$.

Formal Manifolds More specific than the microlinear spaces are the *formal manifolds*, which take some effort to set up. A morphism $f : X \to Y$ is ***étale*** if for every element $x : 1 \to X$ and morphism g from an infinitesimal object $\operatorname{Spec}_R W$ to Y, there is a unique arrow $h : \operatorname{Spec}_R W \to X$ which maps $0 \in \operatorname{Spec}_R W$ to x and satisfies $fh = g$, i.e. makes the diagram below commutative.

$$\begin{tikzcd} 1 \arrow[r, "0"] \arrow[d, "x"'] & \operatorname{Spec}_R W \arrow[d, "g"] \arrow[dl, dashed, "h"'] \\ X \arrow[r, "f"'] & Y \end{tikzcd}$$

If $Y = R^n$ and f is monic, X is said to be an n-dimensional ***model object***. An object M is an n-dimensional ***formal manifold*** if there are étale monics $X_i \to M$, where each X_i is an n-dimensional monic object, whose coproduct is a regular epic morphism $\amalg_i X_i \to M$.

6.1.3 Smooth Algebras

Let CartSp be the subcategory of Diff consisting of the cartesian spaces $\{\mathbb{R}^n\}_{n\in\mathbb{N}}$. A ***$C^\infty$-ring***, or a ***smooth algebra***, is a product-preserving functor $\mathsf{CartSp} \to \mathsf{Set}$, and a C^∞-ring ***homomorphism*** is a natural transformation of functors. These form a category which we will denote C^∞-Alg. Intuitively, C^∞-rings are modeled on (but not restricted to) rings of the form $C^\infty(M)$, for some smooth manifold M; for such a ring, we may define $\Phi_f(\varphi_1, \ldots, \varphi_n)(p) = f(\varphi_1(p), \ldots, \varphi_n(p))$ to get a C^∞-ring.

Given a C^∞-ring $A : \mathsf{CartSp} \to \mathsf{Set}$, we may endow $A(\mathbb{R})$, and hence all $A(\mathbb{R}^n)$, with the structure of an $\mathbb{R}$-algebra by using the images of the morphisms $+ : \mathbb{R}^2 \to \mathbb{R}$ and $c \cdot - : \mathbb{R} \to \mathbb{R}$: for $x, y \in A(\mathbb{R})$ and $c \in \mathbb{R}$, we denote by $x + y$ the image of $(x, y) \in \mathbb{R}^2$ under the morphism

$A(+) : A(\mathbb{R}^2) = A(\mathbb{R})^2 \to A(\mathbb{R})$, and we denote by cx the image of x under the morphism $A(c \cdot -) : A(\mathbb{R}) \to A(\mathbb{R})$. That the necessary $\mathbb{R}$-algebra identities hold in CartSp imply that they hold in Set as well. Hence, we may associate to every C^∞-ring an underlying $\mathbb{R}$-algebra $A(\mathbb{R})$. We will often *identify* A with $A(\mathbb{R})$, though we can't identify any given $\mathbb{R}$-algebra X with a C^∞-ring: it's necessary that X lifts morphisms $\mathbb{R}^n \to \mathbb{R}^m$ to morphisms $X^n \to X^m$ in a nice way. Specifically, we require an operation $\Phi_f : X^n \to X$ for every smooth map $f : \mathbb{R}^n \to \mathbb{R}$ such that, for $h(x_1, \ldots, x_n) = g(f_1(x_1, \ldots, x_n), \ldots, f_m(x_1, \ldots, x_n))$, we have $\Phi_h(x_1, \ldots, x_n) = \Phi_g(\Phi_{f_1}(x_1, \ldots, x_n), \ldots, \Phi_{f_m}(x_1, \ldots, x_n))$ as well as $\Phi_{\pi_i}(x_1, \ldots, x_n) = x_i$.

Finitely Generated Ideals Of particular consequence is when A is equivalent to $C^\infty(\mathbb{R}^n)/I$ for some ideal I of $C^\infty(\mathbb{R}^n)$: when this happens, A is said to be ***finitely generated***, and when $I = (i_1, \ldots, i_m)$ is finitely generated as an ideal, A is said to be ***finitely presented***. Every C^∞-ring of the form $C^\infty(M)$ for a smooth manifold M is finitely presented, for instance. If A is local as a normal ring, it's known as a ***local*** C^∞-ring. The primary example is, as encountered in algebraic geometry, the stalk of the sheaf of smooth functions on $\mathbb{R}^n$, written $C^\infty_p(\mathbb{R}^n)$.

We define the category L^{op} to be the subcategory of C^∞-Alg consisting of the finitely generated algebras; the objects of L are known as ***loci***, and written as $\ell A, \ell B, \ldots$ (where A, B are finitely generated smooth algebras). A morphism $\ell B \to \ell A$ of L is a morphism $A \to B$, or, if $B = C^\infty(\mathbb{R}^m)/J$ and $A = C^\infty(\mathbb{R}^n)/I$, an equivalence class $[\varphi]$ of functions $\mathbb{R}^m \to \mathbb{R}^n$ acting as $\varphi(f) = f \circ \varphi$; we require each φ to satisfy $f \in I \implies \varphi(f) \in J$, so that φ extends to a function $C^\infty(\mathbb{R}^n)/I \to C^\infty(\mathbb{R}^m)/J$, $f + (I) \mapsto \varphi(f) + (J)$, and write $\varphi \sim \psi$ if each $\pi_i \circ (\varphi - \psi) : \mathbb{R}^n \to \mathbb{R}$ is in I.

$\mathsf{Set}^{\mathsf{L}^{\mathrm{op}}}$ is a Grothendieck topos (by equipping L with the indiscrete topology in which all presheaves are sheaves). The functor $s : \mathsf{Diff} \to \mathsf{L}$ sending a smooth manifold M to $\ell C^\infty(M)$ is full and faithful, and when combined with the full and faithful Yoneda embedding $よ : \mathsf{L} \to \mathsf{Set}^{\mathsf{L}^{\mathrm{op}}}$ evidences Diff as a subcategory of $\mathsf{Set}^{\mathsf{L}^{\mathrm{op}}}$. So, $\mathsf{Set}^{\mathsf{L}^{\mathrm{op}}}$ can be thought of as a category of "generalized" smooth spaces, and at the same time as a category of "variable" sets. For a functor $P \in \mathsf{Set}^{\mathsf{L}^{\mathrm{op}}}$, we say that a ***element of*** P ***at stage*** ℓA is an element x of the set $P(\ell A)$. By Yoneda, these can be identified with natural transformations from ℓA to P (where we have silenced the Yoneda embedding). A map $\varphi : A \to B$ in L yields a map $\varphi : \ell B \to \ell A$ in $\mathsf{Set}^{\mathsf{L}^{\mathrm{op}}}$, and hence maps elements of P at stage ℓA to elements of P at stage ℓB by composition; this is known as

restriction, and written as $x|_\varphi$.

Smooth Reals In the topos $\mathsf{Set}^{\mathsf{L}^{\mathrm{op}}}$, the smooth real line R can be identified as the functor $R = \ell C^\infty(\mathbb{R})$; elements of R at stage ℓA, or natural transformations $\ell A \to R$, are just called reals at stage ℓA. For $A = C^\infty(\mathbb{R}^n)/I$, this is an equivalence class $f(x)$ mod I, where $f : \mathbb{R}^n \to \mathbb{R}$. The internal ring structure on R derives from a ring structure on each set of reals at a given stage ℓA given by simply taking pointwise addition and multiplication of functions mod I. The terminal object ("point") is given by $1 = \ell(C^\infty(\mathbb{R})/(x))$, and the object of nth order infinitesimals is $\ell(C^\infty(\mathbb{R})/(x^{n+1}))$. The ***smooth interval object*** $[a,b]$ is given by $\ell(C^\infty(\mathbb{R})/m^\infty_{[a,b]})$, where $m^\infty_{[a,b]}$ is the ideal consisting of functions that vanish on $[a,b]$. Again, we may analyze these objects by their elements at stage ℓA for $A = C^\infty(\mathbb{R}^n)/I$: for instance, the nth order infinitesimals are those smooth functions f such that $f^{n+1} \in I$. To prove all of this, we state the Kripke-Joyal semantics for $\mathsf{Set}^{\mathsf{L}^{\mathrm{op}}}$: letting x be an element of X at stage ℓA, we have

- $\ell A \Vdash \psi(x) \wedge \phi(x)$ (resp. $\psi(x) \vee \phi(x)$) iff $\ell A \Vdash \psi(x)$ and (resp. or) $\ell A \Vdash \phi(x)$.
- $\ell A \Vdash \phi(x) \implies \psi(x)$ iff for every $f : \ell B \to \ell A$ in L, $\ell B \Vdash \phi(x|_f)$ implies $\ell B \Vdash \psi(x|_f)$.
- $\ell A \Vdash \exists y \in F\, \phi(x,y)$ iff there's an element y_0 of F at stage ℓA such that $\ell A \Vdash \phi(x, y_0)$.
- $\ell A \Vdash \forall y \in F\, \phi(x,y)$ iff for every $f : \ell A \to \ell B$ in L and element y_0 of F at stage ℓB we have $\ell B \Vdash \phi(x|_f, y_0)$.

This allows us to prove that the KL axiom $\forall f \in R^D \exists! c \in R\, (\forall \epsilon \in D(f(\epsilon) = f(0) + c\epsilon))$ is valid for R, as well as the following ***integration axiom***:

$$\forall f \in R^{[0,1]} \exists! F \in R^{[0,1]} \left(F' = f \wedge F(0) = 0\right)$$

The function F whose derivative is f is known as the integral of f.

While L^{op} consists of the finitely generated smooth algebras, we define G^{op} to consist of finitely generated smooth algebras whose ideals are determined by germs. The category G, then, consists of loci of the form $\ell(C^\infty(\mathbb{R}^n)/I)$, where I is such that $f \in I$ iff the germ of f at an arbitrary point $x \in Z(I)$ (i.e., $g(x) = 0$ for all $g \in I$) is in the germ of I. (The $\Rightarrow$ part is trivial, whereas the $\Leftarrow$ part is the real restriction, and where the name "ideal determined by germs" comes from). A second subcategory $\mathsf{F}^{\mathrm{op}} \subset \mathsf{L}^{\mathrm{op}}$ is given by smooth algebras of the form $C^\infty(\mathbb{R}^n)/I$, where I is ***closed***, or such that if for every $x \in Z(I)$, the Taylor series of a function

f at x resembles the Taylor series of some element of I at x, then $f \in I$. Finally, an ideal I of $C^\infty(\mathbb{R}^n)$ is ***point determined*** if $Z(f) \supseteq Z(I) \implies f \in I$. These generate the subcategory E^{op}.

Since the germ of a function contains its Taylor series, closed ideals are germ determined, so that $\mathsf{F}^{\mathrm{op}} \subset \mathsf{G}^{\mathrm{op}}$ and hence $\mathsf{F} \subset \mathsf{G} \subset \mathsf{L}$; furthermore, since the Taylor series of f in particular tells us about its vanishing points, point determined ideals are closed, and hence $\mathsf{E} \subset \mathsf{F} \subset \mathsf{G} \subset \mathsf{L}$. Every ideal I of $C^\infty(\mathbb{R}^n)$ admits a smallest germ determined ideal $\tilde{I}$ given by the set of all f whose germ is an element of the germ of I at all points $x \in Z(I)$; this assignment is functorial, and is in fact left adjoint to the inclusion $\mathsf{G}^{\mathrm{op}} \to \mathsf{L}^{\mathrm{op}}$. The same formula gives us left adjoints to the inclusions $\mathsf{E}^{\mathrm{op}} \to \mathsf{F}^{\mathrm{op}}$, $\mathsf{F}^{\mathrm{op}} \to \mathsf{G}^{\mathrm{op}}$, and hence a sequence of *co*reflective subcategory inclusions $\mathsf{E} \to \mathsf{F} \to \mathsf{G} \to \mathsf{L}$. The right adjoints $\mathsf{L} \to \mathsf{E}$, $\mathsf{L} \to \mathsf{F}$, $\mathsf{L} \to \mathsf{G}$ are customarily denoted by γ, κ, and λ, respectively; we'll also denote the right adjoints $\mathsf{G} \to \mathsf{E}$, $\mathsf{G} \to \mathsf{F}$, and $\mathsf{F} \to \mathsf{E}$ by γ, κ, and γ, so that γ makes a finitely generated ideal in any of these categories point determined, κ makes an ideal closed, and λ makes an ideal germ determined.

Given a function $f \in C^\infty(\mathbb{R}^n)$, the most general solution to providing $C^\infty(\mathbb{R}^n)$ with an inverse of f is given by the smooth algebra $C^\infty(f^{-1}(\mathbb{R} - \{0\}))$. We write this algebra as $C^\infty(\mathbb{R}^n)\{f^{-1}\}$, and associate to it a canonical morphism $\eta_f : C^\infty(\mathbb{R}^n) \to C^\infty(\mathbb{R}^n)\{f^{-1}\}$ restricting a smooth g on $\mathbb{R}^n$ to the subset of $\mathbb{R}^n$ on which f doesn't vanish. We define $(C^\infty(\mathbb{R}^n)/I)\{f^{-1}\} = C^\infty(\mathbb{R}^n)\{f^{-1}\}/\eta_f(I)$; while this construction doesn't necessarily map elements of G^{op} to elements of G^{op}, $C^\infty(\mathbb{R}^n)/\{f^{-1}\}/\eta_f(I)$ will be finitely generated so long as $C^\infty(\mathbb{R}^n)/I$ is, and hence we can obtain a germ determined locus $\lambda\ell((C^\infty(\mathbb{R}^n)/I)\{f^{-1}\})$ equipped with a canonical morphism into $\ell(C^\infty(\mathbb{R}^n)/I)$.

The Topos We define a Grothendieck topology J on G as follows: a family $\{f_\alpha : \ell A_\alpha \to \ell A\}_{\alpha \in \Omega}$ is a covering family if for every $\alpha \in \Omega$ there's a function $b_\alpha \in A$ such that f_α factors through the canonical map $\lambda\ell(A\{b_\alpha^{-1}\}) \to \ell A$, and the family $\{\gamma f_\alpha\}_{\alpha \in \Omega}$ covers $\gamma \ell A$. J sends ℓA to its collection of covering families. The Grothendieck topos $\mathrm{Sh}(\mathsf{G}, J)$ is denoted $\mathcal{G}$. As usual, we have a sheafification functor $-^{sh} : \mathsf{Set}^{\mathsf{G}^{\mathrm{op}}} \to \mathcal{G}$ left adjoint to the inclusion functor $\mathcal{G} \to \mathsf{Set}^{\mathsf{G}^{\mathrm{op}}}$, as well as a global sections functor $\Gamma : \mathcal{G} \to \mathsf{Set}$, $\Gamma(F) = F(1)$, right adjoint to the sheafification of the constant presheaf functor $\Delta(S)(\ell A) = S$. Writing $A = C^\infty(\mathbb{R}^n)/I$, this sheafification sends ℓA to the set of locally constant functions $Z(I) \to S$. Γ is also left adjoint to the functor B sending a set S to the sheaf sending ℓA to the set of all functions $Z(I) \to S$.

The Kripke-Joyal semantics for $\mathcal{G}$ are equivalent to those of $\mathrm{Set}^{\mathbb{L}^{op}}$ for the operators $\wedge$, $\implies$, and $\forall$, but differ for the other connectives.

- $\ell A \Vdash \varphi(x) \vee \psi(x)$ iff there's a covering family $\{f_\alpha : \ell A_\alpha \to \ell A\}$ such that, for each α, $\ell A_\alpha \Vdash \varphi(x|_{f_\alpha})$ or $\ell A_\alpha \Vdash \psi(x|_{f_\alpha})$.
- $\ell A \Vdash \exists y \in F\, \phi(x, y)$ iff there's a covering family $\{f_\alpha : \ell A_\alpha \to \ell A\}$ such that, for each α, there's an element y_α of F at stage ℓA_α (i.e., $y_\alpha \in F(\ell A_\alpha)$) with $\ell A_\alpha \Vdash \phi(x, y_\alpha)$.
- $\ell A \Vdash \neg\phi(x)$ iff for every $f : \ell B \to \ell A$ such that $\ell B \Vdash \phi(x|_f)$, $B = 0$.

Just as in $\mathrm{Set}^{\mathbb{L}^{op}}$, $R = \mathcal{G}(-, \ell C^\infty(R))$ is a commutative ring object with orders $<, \leq$. The difference is that, in $\mathcal{G}$, R satisfies the following additional properties: $\mathcal{G} \models \neg(0 = 1)$, $\mathcal{G} = \forall x, y \in R(x + y \in U(R) \implies x \in U(R) \vee y \in U(R))$, and $\mathcal{G} \models \forall x \in R\, \exists n \in \mathbb{N}(x < n)$. Here, $\mathbb{N}$ is the natural numbers object/sheaf sending ℓA to the set of locally constant functions $\ell A \to \mathbb{N}$. The first two statements state that R is a local ring, and the third states that R is Archimedean. Furthermore, R satisfies the ***field axiom***

$$\forall x_1, \ldots, x_n \in R\, (\neg(x_1 = 0 \wedge \ldots \wedge x_n = 0) \implies (x_1 \in U(R) \vee \ldots \vee x_n \in U(R)))$$

as well as the Kock-Lawvere and integration axioms from $\mathrm{Set}^{\mathbb{L}^{op}}$. Locality is often studied in the form of an ***apartness relation*** # whereby $x \# y$ if $x - y \in U(R)$, or equivalently if $x < y \vee x > y$.

If we replace G with F and λ in the definition of a covering family with κ, we obtain a Grothendieck topology J on F whose corresponding Grothendieck topos $\mathrm{Sh}(F, J)$ is denoted $\mathcal{F}$; the entirety of the above discussion of $\mathcal{G}$ holds for $\mathcal{F}$.

6.2 Physical Models

6.2.1 General Relativity

Synthetic differential geometry allows us to construct an intuitionistic theory of spacetime in which general relativity can be constructed; we will use the model of SDG provided by the topos $\mathcal{G}$ of sheaves over the site of finitely generated smooth algebras with germ determined ideals. Our plan will be to set up the elements of classical Riemannian geometry (connections, curvature, and so on) in a synthetic manner, and study the interpretation of Einstein's equations in $\mathcal{G}$.

Connections and Curvature An ***infinitesimal*** ***n-cube*** on an object M is an element of $M^{D^n} \times D^n$, and an ***infinitesimal*** ***n-chain*** is an element of the free R-module $C_n(M)$ generated by all infinitesimal n-cubes on M. Writing $I = [0,1]$, a ***finite*** (or "big") ***n-cube*** on M is a morphism $I^n \to M$, and a ***finite*** ***n-chain*** an element of the free R-module $\Gamma_n(M)$ generated by finite n-cubes.

A ***affine connection*** on a microlinear space M is a bilinear morphism $\nabla : TM \times_M TM \to M^{D\times D}$ (where the pullback is taken over the morphisms $v \mapsto v(0)$, so these are two tangent vectors at the same point) such that $\nabla(v,w)(d_1,0) = v(d_1)$ and $\nabla(v,w)(0,d_2) = w(d_2)$. If $\nabla(v,w)(d_1,d_2) = \nabla(w,v)(d_2,d_1)$, ∇ is said to be ***torsion-free***. From a connection ∇ on M, we may define another function τ which associates to each $(v,d) \in TM \times D$ a ***parallel transport*** $\tau_d(v,-) : \pi^{-1}(v(0)) \cong \pi^{-1}(v(d))$; this map is linear in both v and its argument, is the identity for $d = 0$, and $\tau_d(\lambda v,-) = \tau_{\lambda d}(v,-)$. We identify $\tau_d(v,w)$ with the parallel transport of w along v for an infinitesimal period of time d. Specifically, $\tau_{d_1}(v,w)(d_2)$ is defined to be $\nabla(v,w)(d_1,d_2)$.

Given a connection ∇ on a microlinear space M, we would like to define the Riemann curvature tensor in terms of the parallel transport of a vector along the boundary of an infinitesimal 2-chain. Given such a 2-chain $(\gamma,d_1,d_2) \in M^{D^2} \times D^2$ based at a point $x = \gamma(0,0)$, we do this as follows: take a vector v and transport it along $\gamma(-,0)$ for a period of d_1 "seconds". Transport the new vector along $\gamma(d_1,-)$ for a period of d_1 seconds, transport backwards along $\gamma(0,-)$ for d_2 seconds and finally transport backwards along $\gamma(-,d_2)$ for d_2 seconds, before subtracting v from the result. This gives a preliminary map

$$R'(\gamma,d_1,d_2,v) = \tau_{d_2}^{-1}(\gamma(-,d_2),\tau_{d_2}^{-1}(\gamma(0,-),\tau_{d_1}(\gamma(d_1,-),\tau_{d_1}(\gamma(-,0),v)))) - v$$

Being bilinear in both d_1 and d_2, we may define a map $\varphi(d_1,d_2) = R'(\gamma,d_1,d_2,v)$ which induces by microlinearity of T_xM a function $\psi : D \to T_xM$ such that $\psi(d_1d_2) = \varphi(d_1,d_2)$. By KL, this can be written as $\psi(d) = d\hat{v}$ for a unique $v \in T_xM$. We define $R'' : M^{D\times D} \times_M TM \to TM$ to send a pair (γ,v) to this $\hat{v}$, and define the ***Riemann curvature tensor*** $R : TM \times_M TM \times_M TM \to TM$ by $R(v_1,v_2,v_3) = R''(\nabla(v_1,v_2),v_3)$. If M is a formal manifold, we may work in local coordinates: the connection ∇ becomes a function that takes in a point $x \in M$ along with two vectors $v, w \in R^n$, and returns an element of $M \times R^n \times R^n \times R^n$. The fourth component of this tuple is denoted ∇_4, and used to define the ***Christoffel symbols***: in a basis $\{e_1, \ldots, e_n\}$ of R^n, these are given by ${\Gamma^i}_{jk}(x) = \pi_i(\nabla_4(x, e_k, e_j))$. The Riemann curvature tensor decomposes into components in the usual manner: ${R^\ell}_{ijk} = \partial_j {\Gamma^\ell}_{ki} - \partial_k {\Gamma^\ell}_{ji} + {\Gamma^\ell}_{jm} {\Gamma^m}_{ki} - {\Gamma^\ell}_{km} {\Gamma^m}_{ji}$ (again, at every point).

Hence, to a formal manifold $M \in \mathcal{G}$ we may associate a Riemann curvature tensor ${R^\ell}_{ijk}$ to a connection ∇. This gives us a Ricci curvature tensor $R_{ik} = {R^\ell}_{i\ell k}$ and, with a Riemannian metric g_{ij}, a scalar curvature $R = g^{ij} R_{ij}$ and Einstein tensor $G_{ij} = R_{ij} - \frac{1}{2} R g_{ij}$.

Einstein's Equations Consider R^4 filled with dust with 4-velocity u^i and density ρ. The classical Einstein equations read $G_{ij} = T_{ij} = \kappa c^2 \rho u_i u_k$, where κ is Einstein's constant. In $\mathcal{G}$, real numbers become elements of R at stage ℓA for $A = C^\infty(\mathbb{R}^n)/I$; these are natural transformations from $よ(\ell A)$ to $R = よ(\ell C^\infty(\mathbb{R}))$, which by Yoneda are in bijection with smooth functions $\varphi : \mathbb{R}^n \to \mathbb{R}$ modulo I. So, using $\mathcal{G}$ as a model for SDG, an arbitrary real number $r \in R$ at stage ℓA is really a "parametrized" element of $\mathbb{R}$, changing smoothly as we vary the point $v \in \ell A$. Similarly, an event, or element of R^4, at stage ℓA is really a smooth function $\mathbb{R}^n \to \mathbb{R}^4, v \mapsto (x^0(v), x^1(v), x^2(v), x^3(v)) \bmod I$. Taking the reals at stage $1 = よ(C^\infty(\{*\}))$ recovers the usual set $\mathbb{R}$. So, in SDG, the Einstein equations $G_{ij}(x) = T_{ij}(x), x \in R^4$ carry over without modification at stage 1, stating that two pairs of 16 reals coincide at every point in $\mathbb{R}^4$ ($G_{00}(x)(*) = T_{00}(x)(*)$ and so on). At stage $\ell C^\infty(\mathbb{R})$, the equations state that two pairs of 16 *smooth curves* through $\mathbb{R}^4$, assigned to each point in $\mathbb{R}^4$, coincide; at stage $\ell C^\infty(\mathbb{R}^2)/I$, they become surfaces $\varphi : \mathbb{R}^2 \to \mathbb{R}^4$ modulo the ideal I, and so on. [Guts and Zvyagintsev, 2000] interprets the Einstein equations for a dusty universe at various stages.

This interpretation of general relativity can be carried out in any other smooth topos, thereby inheriting its internal logic instead of $\mathcal{G}$'s logic; to quote [Guts and Grinkevich, 1996],

"The resulting space-time theory will be non-classical, different from that of the Minkowski space-time. This is a *new* theory of space-time, created in a purely logical manner. It will reflect the real space-time properties to the same extent as the development of mathematical abstractions accompanies the development of the real world."

6.2.2 Classical Mechanics

Here's where we bring in the language of cohesive topoi. Let $\mathcal{S} = \mathcal{S}mooth\mathcal{S}et$ be the cohesive topos of smooth sets, constructed above as the sheaf topos on CartSp with the differentiably good open cover topology. Letting $\Omega^p_{cl}(M)$ be the set of closed p-forms on a manifold M, we define a smooth set $\mathbf{\Omega}^p$ by $\mathbf{\Omega}^p(\mathbb{R}^n) = \Omega^p(\mathbb{R}^n)$, as well as a morphism $\boldsymbol{d} : \mathbf{\Omega}^p \to \mathbf{\Omega}^{p+1}, d_{\mathbb{R}^n} = d : \Omega^p(\mathbb{R}^n) \to \Omega^{p+1}(\mathbb{R}^n)$. This smooth set is a "universal moduli space" for p-forms, in the sense that for any smooth manifold M, considered as a smooth set, there's a natural bijection between morphisms $M \to \mathbf{\Omega}^p$ and p-forms on M. Note that the machinery of smooth sets is necessary to solve this moduli problem: $\mathbf{\Omega}^p$ is *not* the image of a smooth manifold, nor is it even a diffeology. However, this anomaly allows us to lift the definition of p-forms from manifolds to smooth sets: given an arbitrary smooth set X, a p-form ω on X is a morphism $X \to \mathbf{\Omega}^p$, and if $d\omega := \boldsymbol{d} \circ \omega = 0$, ω is ***closed***. There is an object $\mathbf{\Omega}^p_{cl}$ of closed p-forms given by $\mathbf{\Omega}^p_{cl}(\mathbb{R}^n) = \{\text{closed } p\text{-forms on } \mathbb{R}^n\}$.

Presymplectic Sets A ***presymplectic smooth set*** is a pair (X, ω), where X is a smooth set and ω a closed 2-form on X. (While ω is closed, we haven't said anything about nondegeneracy, hence *pre*symplectic), or equivalently a morphism $X \to \mathbf{\Omega}^2_{cl}$. A p-form on X is really just an assignment to every plot $\phi \in X(\mathbb{R}^n)$ of a p-form $\omega_{\mathbb{R}^n}(\phi)$ on $\mathbb{R}^n$, so we can add and multiply them, and in particular we can take the ***tensor product*** of presymplectic sets $(X, \omega) \otimes (Y, \eta)$, which assigns to every product plot $\phi \times \psi X(\mathbb{R}^n) \times Y(\mathbb{R}^n)$ the sum $\omega_{\mathbb{R}^n}(\phi) + \eta_{\mathbb{R}^n}(\psi)$. A ***symplectomorphism*** between presymplectic sets (X, ω) and (X', ω') is just a morphism $\phi : X \to X'$ such that $\omega'\phi = \omega$. Hence, presymplectic sets assemble into the *slice topos* $\mathcal{S}/\mathbf{\Omega}^2_{cl}$. A presymplectic subset of a presymplectic set (X, ω) is simply a subobject $\phi : X' \rightarrowtail X$, which induces by composition a presymplectic set $(X', \omega|_{X'} := \omega\phi)$. If $(\omega\phi)_{\mathbb{R}^n} : X'(\mathbb{R}^n) \to \mathbf{\Omega}^2_{cl}(\mathbb{R}^n) = \Omega^2_{cl}(\mathbb{R}^n)$ is the constant morphism $x \mapsto 0$, and the dimension of X' is half that of X, we call X' a ***Lagrangian subset*** of X.

Given two objects X, Y, we define a ***correspondence*** to be a diagram of the form $X \leftarrow C \to Y$, and a ***equivalence*** of correspondences to be an isomorphism $C \cong C'$ forming a commutative diagram

$$\begin{tikzcd} & C \arrow[dl] \arrow[dr] \arrow[dd] & \\ X & & Y \\ & C' \arrow[ul] \arrow[ur] & \end{tikzcd}$$

Given two correspondences $X \leftarrow C \rightarrow Y \leftarrow C' \rightarrow Z$, their composition along Y is defined to be the correspondence $X \leftarrow C \times_Y C' \rightarrow Z$. Hence, we can for an arbitrary topos $\mathcal{E}$ define a 2-category $\mathsf{Corr}(\mathcal{E})$ of correspondences whose 1-morphisms $X \rightarrow Y$ are correspondences $X \leftarrow C \rightarrow Y$ and whose 2-morphisms are morphisms between correspondences. The category $\mathsf{Corr}(\mathcal{S}/\mathbf{\Omega}^2_{cl})$, for instance, has as its objects commutative squares

$$\begin{tikzcd} & Z \arrow[dl, "\varphi"'] \arrow[dr, "\psi"] \arrow[dd] & \\ X \arrow[dr, "\omega"'] & & Y \arrow[dl, "\eta"] \\ & \mathbf{\Omega}^2_{cl} & \end{tikzcd}$$

This is a symmetric monoidal category under the tensor product $(X, \omega) \otimes (Y, \eta) = (X \times Y, \omega + \eta)$ and unit $(*, 0)$.

Smooth Groupoids Suppose that instead we would like $X(\mathbb{R}^n)$ to capture not just plots of $\mathbb{R}^n$ in X, but gauge transformations – nontrivial isomorphisms – between plots. To do this, we need a groupoid structure on each $X(\mathbb{R}^n)$. A ***smooth groupoid*** is a functor $X : \mathsf{CartSp}^{\mathrm{op}} \rightarrow \mathsf{Grpd}$ such that both the set of objects of $X(\mathbb{R}^n)$, denoted $X_0(\mathbb{R}^n)$, and the set of morphisms, denoted $X_1(\mathbb{R}^n)$, assemble into smooth sets. The category of smooth groupoids is denoted $\mathcal{SmoothGrpd}$; this is just a "refinement" of $\mathcal{SmoothSet}$, and we'll also denote it $\mathcal{S}$. We may obtain smooth groupoids by taking a smooth set X with an action of a smooth group G, and taking the ***smooth homotopy quotient*** $X//G$, whose objects $(X//G)_0(\mathbb{R}^n)$ are the objects of $X(\mathbb{R}^n)$, and whose morphisms are of the form $x \rightarrow gx$. For X an arbitrary one-point space, $X//G$ is a groupoid with a single object and an automorphism for each $g \in G$, with composition of morphisms given by composition of group elements. This groupoid is known as $\boldsymbol{B}G$. We define $\boldsymbol{B}U(1)_{conn}$

to be the smooth groupoid to send $\mathbb{R}^n$ to the groupoid $\Omega^1(\mathbb{R})//\mathsf{Diff}(\mathbb{R}^n, U(1))$ (where the composition of two smooth functions $f, g : \mathbb{R}^n \to U(1)$ is $(f \cdot g)(v) = f(v) \cdot g(v)$).

6.2.3 Quantum Mechanics

Take a smooth topos $\mathcal{E}$ with smooth real line R, and denote by $U(R)$ the subobject of invertible (non-infinitesmal) elements of R. Assume that R satisfies the ***field axiom***,

$$\forall x_1, \ldots, x_n \in R\,(\neg(x_1 = 0 \wedge \ldots \wedge x_n = 0) \implies (x_1 \in U(R) \vee \ldots \vee x_n \in U(R)))$$

(For instance, we can again let $\mathcal{E} = \mathcal{G}$). In particular, for $n = 1$ we have $\forall x \in R\,(x \neq 0 \implies x \in U(R))$. Denoting by C the complex numbers object (a 2-dimensional R-algebra, which also satisfies the field axiom), we define a ***inner product*** on an R-module V to be a symmetric, bilinear map $\langle -, - \rangle : V \times V \to C$ satisfying $v \neq 0 \implies \langle v, v \rangle > 0$. Note that, for $V = R$, we have for $x \neq 0$ that $\langle x, x \rangle = x^2 \langle 1, 1 \rangle > 0$, implying that $x^2 = 0$ and hence $x \in U(R)$; it follows that the existence of an inner product on R relies on the field axiom for $n = 1$.

We'll analyze the case of a spin 1/2 interaction, first in the classical case studied in [Sakurai et al., 2014], and then in the case of SDG, exposited in [Fearns, 2002].

The Stern-Gerlach Experiment In the Stern-Gerlach experiment, silver atoms are shot at a target, passing through an inhomogeneous magnetic field $\vec{B}$ which splits the silver atoms along the z axis. The electron shell structure of silver is 2, 8, 18, 18, and 1: four full shells, followed by a fifth shell with a single electron. The first four shells cancel each other out magnetically, so the magnetic moment $\vec{\mu}$ of the atom is proportional to the spin $\vec{S}$ of the one electron. If the electron behaved classically, the magnetic moment of the atom along the z axis, μ_z, would be distributed anywhere between $-|\vec{\mu}|$ and $|\vec{\mu}|$, resulting in the silver atoms forming a continuous interval on the target. What we observe in practice is two distinct spots on the target, indicating that the electron spin along the z axis is either fully up, $S_z = \hbar/2$, or fully down, $S_z = -\hbar/2$. The same holds when we reorient the machine to split the atoms along the x or y axes, suggesting that the electron's spin, when measured along a given axis, will take either an up or down spin along that axis. We model this as follows: we have three axes x, y, z and three operators S_x, S_y, S_z, each of which has two eigenvectors with eigenvalues $\pm\hbar/2$. We can model these operators as elements of $\mathbb{C}^{2\times 2}$: recalling the definition of the Pauli matrices

$$\sigma^1 = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \quad \sigma^2 = \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix} \quad \sigma^3 = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$$

we write $S_{x_i} = \frac{\hbar}{2}\sigma^i$. So the spin of an electron with spin up along the z axis is modeled by the ket $|S_z;+\rangle = [1,0]^T$, and likewise $|S_y;+\rangle = [1,i]^T/\sqrt{2}$, $|S_x;+\rangle = [1,1]^T/\sqrt{2}$.

Microlinear Lie Groups Moving to a smooth topos $\mathcal{E}$, define the microlinear group $G = \mathrm{SO}(3)$ to be the subobject of $R^{3\times 3}$ consisting of the orthogonal matrices with determinant 1. With matrix multiplication, this is a Lie group internal to $\mathcal{E}$ with identity $e = I_3$. The fiber T_eG, consisting of all $f : D \to G$ such that $f(0) = e$, then has a bilinear operation $[-,-] : T_eG \times T_eG \to T_eG$ given as $[v,w](d_1d_2) = w(-d_2)v(-d_1)w(d_2)v(d_1)$. This is antisymmetric and satisfies the Jacobi identity, so we call it the Lie algebra $\mathfrak{g}$ associated to the Lie group G. $\mathfrak{so}(3)$ is, in fact, isomorphic to the Lie algebra $\mathfrak{su}(2)$ generated by the Pauli matrices, implying that we can consider these matrices, and hence the spin operators themselves, as elements of T_eG.

Now, suppose we have a system consisting of two interacting electrons, the total energy being encapsulated in a unitary Hamiltonian operator H. The classical time-dependent Schrödinger equation expressing the evolution of a time-dependent state $|\psi;t\rangle$ is $ih\frac{d}{dt}|\psi;t\rangle = H|\psi;t\rangle$. In SDG, we take $t \in R, d \in D$, and instead write $|\psi;t+d\rangle = |\psi;t\rangle - \frac{id}{\hbar}H|\psi;t\rangle$. As proven in the paper [Kock, 1986], if $\mathcal{E}$ is ***well-adapted***, possessing a full and faithful functor $\mathsf{Diff} \to \mathcal{E}$, then we have the following integration axiom for a Lie group G with Lie algebra $\mathfrak{g}$:

$$\forall f \in \mathfrak{g}^R \exists! F \in G^R \left(F(0) = e \wedge \forall t \in R \forall d \in D \left(F(t+d)F(t)^{-1} = f(t)(d)\right)\right)$$

The Hamiltonian is a member of the Lie group $U(4)$, and an infinitesimal perturbation to it, as expressed by the SDG Schrödinger equation, is a member of $\mathfrak{u}(4)$; by the integration axiom, this can be integrated to obtain a unique time evolution of $|\psi\rangle$.

While computing actual results in a well-adapted topos such as $\mathcal{G}$ would be tedious, this result is a proof of concept that well-adapted topoi have the necessary structure required to formulate quantum mechanics.

Chapter 7

Cohesive Topoi

7.0.1 Diffeologies

Smooth Sets Let CartSp be, as before, the category whose objects are the $\mathbb{R}^n$'s for $n \geq 0$, and whose morphisms $\mathbb{R}^m \to \mathbb{R}^n$ are the smooth sets. We equip this category with the *good open cover* coverage: an open cover $\{f_\lambda : U_\lambda \to U\}_{\lambda \in \Lambda}$ is ***good*** if each nonempty finite intersection of the $f_\lambda(U_\lambda)$ is contractible. For instance, the open cover of the circle $\circ$ by two halves $\cup$ and $\cap$ is not good: the intersection is homotopy equivalent to two points, rather than one. A ***smooth set*** is a sheaf on this site, and hence there is a Grothendieck topos $\mathcal{SmoothSet}$ of smooth sets.

Hence, a smooth set is a functor $X : \mathsf{CartSp}^{\mathrm{op}} \to \mathsf{Set}$, with $X(\mathbb{R}^n)$ written as X_n, such that for every good open cover $\{f_\lambda : \mathbb{R}^{m_\lambda} \to \mathbb{R}^m\}_{\lambda \in \Lambda}$, if we have a selection of elements $x_\lambda \in X_{m_\lambda}$ such that for all $\mu, \nu \in \Lambda$ and all diagrams $\mathbb{R}^{m_\mu} \xleftarrow{g} \mathbb{R}^\ell \xrightarrow{h} \mathbb{R}^{m_\nu}$ with $f_\mu g = f_\nu h$ we have $X(g)(x_\mu) = X(h)(x_\nu) \in X_\ell$, there is a unique $x \in \mathbb{R}^m$ such that $X(f_\lambda)(x) = x_\lambda$ for all $\lambda \in \Lambda$.

Hence, a smooth set X is first of all a *set*, X_0, along with *plots* of curves in X, or elements of X_1, plots of surfaces, or elements of X_2, and so on, and for every function $f : \mathbb{R}^m \to \mathbb{R}^n$ a map $Xf : X_n \to X_m$ describing how f is used to construct m-plots from n-plots. Every $\mathbb{R}^n$ forms a smooth set $\underline{\mathbb{R}}^n = \mathrm{Hom}_{\mathsf{CartSp}}(-, \mathbb{R}^n)$, as does every smooth manifold $M \in \mathsf{Diff}$ via $\underline{M} = \mathrm{Hom}_{\mathsf{Diff}}(-, M)$. As the sets $\{\mathrm{Hom}_{\mathsf{Diff}}(\mathbb{R}^n, M)\}_{m \in M} = \amalg_n M_n$ are enough to uniquely determine the smooth manifold M, this induces an embedding $\mathsf{Diff} \to \mathcal{SmoothSet}$.

Diffeological Spaces This does not necessarily make smooth sets particularly useful to work with: a general smooth set can still undergo many failings that render it "un-spatial". There is, for instance, nothing restricting the set of plots of surfaces in X from being bigger than the entire set of set-functions from X_2 to X_0, nor is there any notion of what the maps Xf must be like. Hence, we restrict this definition: a ***diffeological space*** is a set X_0 along with a functor $X \in \mathrm{Sh}(\mathsf{CartSp})$ with $X(\mathbb{R}^0) = X_0$ (i.e., a smooth set, presented as a set), such that X is a subsheaf of $\mathrm{Hom}_{\mathsf{Set}}(-, X_0)$. This ensures that the set of plots of $\mathbb{R}^n$ in X_0 is at *most* the set of set-functions from $\mathbb{R}^n$ to X_0, and that a smooth function $f : \mathbb{R}^m \to \mathbb{R}^n$ sends n-plots to m-plots by precomposition: an n-plot $p \in X_n$, identifiable with a set-function $p : \mathbb{R}^n \to X_0$, is sent to an m-plot $pf \in X_m$, identified with the set function $pf : \mathbb{R}^m \to \mathbb{R}^n \to X_0$. The embedding of Diff into $\mathcal{S}mooth\mathcal{S}et$ clearly restricts to an embedding of Diff into the category of diffeological spaces, denoted DiffSp. DiffSp is no longer a topos, but is a *quasi-topos*: while it is locally cartesian closed and finitely (co)complete, it does not have a subobject classifier.

Despite this limitation, diffeological spaces allow us to work with any kind of set of points that has some sort of smooth structure, be it an infinite-dimensional moduli space or any other sort of space that cannot strictly be described by a smooth manifold. The immediate example is the internal hom $[X, Y]_n = \mathrm{Hom}_{\mathsf{DiffSp}}(X \times \underline{\mathbb{R}}^n, Y)$ of diffeological spaces, which could for instance be smooth manifolds; the set of 0-plots (points) is given by $\mathrm{Hom}_{\mathsf{DiffSp}}(X \times \underline{\mathbb{R}}^0, Y) = \mathrm{Hom}_{\mathsf{DiffSp}}(X, Y)$, the set of 1-plots is $\mathrm{Hom}_{\mathsf{DiffSp}}(X \times \underline{\mathbb{R}}, Y)$, when so on, and when X and Y are smooth manifolds $[X, Y]$ provides a diffeological structure on $C^\infty(X, Y)$ [1].

[1] We may also equip this space with a smooth structure using the theory of Fréchet manifolds, which are locally homeomorphic to (complete, Hausdorff, locally convex, metrizable) topological vector spaces rather than $\mathbb{R}^n$'s, but the structure coincides with the diffeological one anyway.

Chapter 8

Noncommutative Geometry

Chapter 9

Structured Spaces

9.1 Spectral Schemes

9.1.1 Ring Spectra

The basic building block of algebraic geometry is the *affine scheme*, which is a geometric space induced by a ring (which we will always assume to be commutative and unital). As a set, an affine scheme is the spectrum, or set of prime ideals, of a ring. We shall review some of the properties which make prime ideals worthy of this honor.

Prime Ideals An ideal $\mathfrak{p}$ of a ring A is said to be prime if

1. It does not contain 1 (and is therefore a proper ideal).
2. If for $a, b \in A$ we have $ab \in \mathfrak{p}$, then $a \in \mathfrak{p}$ or $b \in \mathfrak{p}$.

The analogy is to prime integers in $\mathbb{Z}$: the product mn of integers divides some prime p if and only if at least one of m, n already divided p.

Given a morphism $\varphi : A \to B$, the preimage of a prime ideal $\mathfrak{q}$ in B must itself be a prime ideal $\mathfrak{p} \subset A$: if $\varphi(ab) = \varphi(a)\varphi(b) \in \mathfrak{q}$, then either $\varphi(a)$ in $\mathfrak{q}$ or $\varphi(b) \in \mathfrak{q}$, and hence either $a \in \mathfrak{p}$ or $b \in \mathfrak{q}$. This implies that any morphism of rings contravariantly induces a morphism between sets of their prime ideals. Referring to the set of prime ideals of a ring A as its ***spectrum*** Spec A, this gives us a functor Spec : $\mathsf{CRing}^{\mathrm{op}} \to \mathsf{Set}$.

Prime ideals also behave nicely with respect to quotients and localizations. Specifically,

- Given an ideal $I \subset A$, prime ideals of A/I are in bijection with prime ideals of A containing I.

- Given a multiplicative system $S \subset A$, the prime ideals of the localization $S^{-1}A$ are in bijection with prime ideals of A that do not touch S.

- For any prime $\mathfrak{p} \in A$, the set $A - \mathfrak{p}$ is a multiplicative system touching all primes that are not contained in $\mathfrak{p}$, so (the image of) $\mathfrak{p}$ is the unique largest prime ideal of the ring $A_{\mathfrak{p}} := (A - \mathfrak{p})^{-1}A$.

We call a prime ideal contained in no larger proper ideals a ***maximal ideal***[1], and any ring with a unique maximal ideal a ***local ring***. Hence, $A_{\mathfrak{p}}$, which has as its unique maximal ideal $\mathfrak{m} = \mathfrak{p}$, is a local ring.

In general, the quotient of a ring by a prime ideal will be an integral domain, and the quotient of a ring by a maximal ideal will be a field (these properties in fact *define* prime and maximal ideals); in particular, a local ring $(A, \mathfrak{m})$ naturally induces a field $A/\mathfrak{m}$. The field $A_{\mathfrak{p}}/\mathfrak{p}$ is known as the ***residue field*** of A at $\mathfrak{p}$.

The Zariski Topology Given a ring A, we may consider the function V sending an ideal I of A to the set of prime ideals of A that contain I, as well as its complement $D(I) = V(I)^c$. We have $V((0)) = \operatorname{Spec} A$, as all prime ideals contain (0), and $V(A) = \varnothing$, as (by definition) no prime set contains A. Pairwise multiplication of ideals I, J corresponds to taking the union of $V(I)$ and $V(J)$, and arbitrary summation corresponds to arbitrary intersection; it follows that we may take the $V(I)$ for all ideals I of A to be the closed sets of a topology on $\operatorname{Spec} A$. This is known as the ***Zariski topology***.

9.1.2 Ringed Spaces

Given a topological space X, we define a sheaf of rings on X in the usual sense: as a functor $\mathcal{F} : \mathrm{Op}(X)^{\mathrm{op}} \to \mathsf{CRing}$, satisfying the sheaf conditions:

[1] It is conventional to use the Fraktur letters $\mathfrak{p}, \mathfrak{q}, \ldots$ for prime ideals, and $\mathfrak{m}, \mathfrak{n}, \ldots$ for maximal ideals.

- (Locality) If for an open cover $\{U_\lambda \subseteq U\}_{\lambda\in\Lambda}$ and sections $x, y \in \mathcal{F}(U)$, if the restrictions of x and y to each U_λ agree, then $x = y$.
- (Gluing) If for an open cover $\{U_\lambda \subseteq U\}_{\lambda\in\Lambda}$ we have sections $x_\lambda \in \mathcal{F}(U_\lambda)$ whose restrictions agree on all intersections $U_\lambda \cap U_{\lambda'}$, there's a unique section $x \in \mathcal{F}(U)$ restricting to all x_λ.

A ***ringed space*** is a space X equipped with a sheaf of rings $\mathcal{O}_X$ known as the ***structure sheaf***. A morphism of ringed spaces $(X, \mathcal{O}_X) \to (Y, \mathcal{O}_Y)$ is a pair $(f, f^\sharp)$, where $f : X \to Y$ is a continuous map, and $f^\sharp$ is a morphism of sheaves from $\mathcal{O}_Y$ to the sheaf $f_*\mathcal{O}_X$ on Y defined by $(f_*\mathcal{O}_X)(V) = \mathcal{O}_X(f^{-1}(V))$, known as the direct image or ***pushforward*** of $\mathcal{O}_X$ by f [2].

The Zariski topology has a convenient base: take the *distinguished open sets* $D_f = D((f)) = \{\mathfrak{p} \in \operatorname{Spec} A \mid f \notin \mathfrak{p}\}$. We will define a sheaf of rings $\mathcal{O}_X$ on the topological space $X = \operatorname{Spec} A$ by defining it on this base, simply letting $\mathcal{O}_X(D_f) = A_f$, the localization of A at the multiplicative system $\{1, f, f^2, \ldots\}$. Taking stalks yields $\mathcal{O}_{X,\mathfrak{p}} = A_\mathfrak{p}$, and taking global sections yields $\mathcal{O}_X(A) = \mathcal{O}_X(D_1) = A_1 = A$.

For $X = \operatorname{Spec} A$, the ringed space $(X, \mathcal{O}_X)$ has the special property that all its stalks are local rings. A ringed space with such a property is known as a ***locally ringed space***, and a morphism $(X, \mathcal{O}_X) \to (Y, \mathcal{O}_Y)$ of locally ringed spaces is a ringed space morphism $(f, f^\sharp)$ with the additional property that the induced ring map $f^\sharp_x : \mathcal{O}_{Y,f(x)} \to \mathcal{O}_{X,x}$ sends the maximal ideal of $\mathcal{O}_{Y,f(x)}$ to the maximal ideal of $\mathcal{O}_{X,x}$. Hence, we have categories of ringed and locally ringed spaces. It is straightforward to show that taking spectra is functorial, i.e. that the set map $\psi : \operatorname{Spec} B \to \operatorname{Spec} A$ generated by a ring homomorphism $\varphi : A \to B$ is not only a continuous map w.r.t. the topologies, but a map of locally ringed spaces.

[2]Abusing terminology, we may often refer to the ringed space $(X, \mathcal{O}_X)$ simply as X, with the understanding that the symbol $\mathcal{O}$ refers to the structure sheaf of the relevant space. Similarly, we will often refer to morphisms of ringed spaces $(f, f^\sharp) : (X, \mathcal{O}_X) \to (Y, \mathcal{O}_Y)$ simply as f, understanding that the sharp $\sharp$ denotes the corresponding map of sheaves. When necessary, we will use $\operatorname{sp} X$ to refer to the topological space underlying X.

9.1.3 Schemes

We are now in a position to introduce these fundamental building blocks, as well as some of their basic properties.

Affine Schemes An ***affine scheme*** is a locally ringed space $(X, \mathcal{O}_X)$ which is isomorphic as a locally ringed space to the spectrum of some ring. Hence, affine schemes form a full subcategory AffSch of the category of locally ringed spaces. We know that there is a functor $\mathrm{Spec} : \mathsf{CRing}^{\mathrm{op}} \to \mathsf{AffSch}$, as well as a global sections functor $\Gamma : \mathsf{AffSch}^{\mathrm{op}} \to \mathsf{CRing}$ sending an affine scheme $X = \mathrm{Spec}\, A$ to $\mathcal{O}_X(X) = A$, and a morphism $f : X \to Y = \mathrm{Spec}\, B$ to the morphism $f^\sharp(Y) : B \to A$. These two functors, in fact, form halves of a contravariant *equivalence of categories*, $\mathsf{AffSch} \cong \mathsf{CRing}^{\mathrm{op}}$.

So the initial object $\mathbb{Z}$ of CRing forms the terminal object $\mathrm{Spec}\, \mathbb{Z}$ of AffSch, the product of rings yields the coproduct of affine schemes, and so on.

Schemes In the same way that a topological manifold is defined to be a space locally homeomorphic to some fixed $\mathbb{R}^n$, a scheme is defined to be a *locally affine* locally ringed space. Specifically, a ***scheme*** is a locally ringed space X such that every point $x \in X$ admits a neighborhood $U \ni x$ such that $(U, \mathcal{O}_U := \mathcal{O}_X|_U)$ is an affine scheme. The category AffSch is a subcategory of this larger category Sch of schemes (again full in locally ringed spaces), and is reflective: the "affineification" functor $\mathsf{Sch} \to \mathsf{AffSch}, X \mapsto \mathrm{Spec}\, \mathcal{O}_X(X)$ is left adjoint to the inclusion $\mathsf{AffSch} \hookrightarrow \mathsf{Sch}$, with morphisms from a scheme X to an affine scheme $Y = \mathrm{Spec}\, A$ factoring uniquely through $\mathrm{Spec}\, \mathcal{O}_X(X)$.

We often refer to schemes *over* a base scheme S, or elements of the slice category Sch/S, which are called S-schemes. Morphisms between S-schemes are, as usual, morphisms between schemes that commute with the fixed morphisms into S. A common case is $S = \mathrm{Spec}\, k$, for a field k; the only prime ideal of a field is (0), so $\mathrm{Spec}\, k$ is a single point, trivializing the topological part of a scheme X over $\mathrm{Spec}\, k$. The information such a scheme contains is a morphism $k \to \mathcal{O}_X(X)$, or a k-algebra structure on the ring of global sections of X, and a morphism between k-schemes must only preserve this structure.

Properties of Schemes Like the average topological space, the average scheme isn't very pretty! Hence, as with topological spaces, we have a great deal of niceness conditions we desire our schemes and morphisms to have. For the following, let $(X, \mathcal{O}_X)$ be a scheme with affine open cover $\{U_i = \operatorname{Spec} A_i\}_{i \in I}$, $(Y, \mathcal{O}_Y)$ be a scheme with affine open cover $\{V_j = \operatorname{Spec} B_j\}_{j \in J}$, and $f : X \to Y$ a morphism of schemes.

- X is ***connected*** if sp X is connected, and ***irreducible*** if sp X is irreducible.
- X is ***locally noetherian*** if each R_i can be taken to be a noetherian ring, and ***noetherian*** if furthermore I can be taken to be finite.
- X is ***integral*** if $\mathcal{O}_X(U)$ is an integral domain for every open U, and ***reduced*** if all stalks $\mathcal{O}_{X,x}$ have no nilpotent elements.
- f is ***locally of finite type*** if we can take the V_j to be such that each $f^{-1}(V_i)$ is affine and can be covered by open affines each of which is the spectrum of some B_i-algebra, finitely generated as a B_i-module. If we only need one open affine for each $f^{-1}(V_i)$, f is ***finite***.
- If the diagonal morphism $\Delta_f : X \to X \times_Y X$ yields a closed inclusion $i : \Delta_f(X) \to \operatorname{sp} X$ with $i^\sharp$ a surjective map, f is ***separated***. X itself is separated if the terminal morphism $X \to \operatorname{Spec} \mathbb{Z}$ is separated.

9.2 Fractured ∞-Topoi

9.3 Pregeometry

9.4 Derived Algebraic Geometry

Chapter 10

The Phenomenology of Space

10.1 Transcendental Aesthetic

10.2 Relativity

10.3 Compactification

10.4 Emergent Spacetime

10.5 Isbell Duality

Part IV

Ontology

Chapter 11

All About $\otimes$

11.1 String Diagrams

11.2 Hilbert Spaces

11.3 States and Measurements

11.4 Categorical Quantum Mechanics

Chapter 12

Ontology in Topoi

12.1 Topos Quantum Theory

12.2 Behavior Types

Chapter 13

Process Theories

13.1 Graphical Calculi

13.2 Lenses

13.3 Profunctor Optics

13.4 The Topos of Dynamical Systems

Appendix A

Homotopy and Cohomology

A.1 Homotopy Theory

A.1.1 Homotopy Equivalence

Given two continuous functions $f, g : X \rightrightarrows Y$ between topological spaces, we may ask whether there is a "continuous transformation" of f into g. For instance, we may wonder whether two different loops on a torus (continuous functions $\gamma : [0,1] \to T^2$ with $\gamma(0) = \gamma(1)$) can be morphed into one another continuously, i.e. without breaking one of the loops. Such a transformation between two paths, say f and g, would look like a *family* of paths $F_t(x)$, where $s \in [0,1]$, such that $F_0(x) = f(x)$ and $F_1(x) = g(x)$. The right definition is as follows: A ***homotopy*** between two continuous maps $f, g : X \rightrightarrows Y$ is a continuous map $F : X \times [0,1] \to Y$ such that $F(x,0) = f(x)$ and $F(x,1) = g(x)$. If there is a homotopy from f to g, the two maps are said to be ***homotopic***, written as $f \simeq g$. We think of the second argument t as moving along the continuous family, and the first argument x as selecting a point in F_t.

Homotopy is an equivalence relation on the set of continuous maps $X \to Y$, and composition is compatible with this relation.

Proof. Every map f is homotopic to itself, by letting $F(x,t) = f(x)$. If F is a homotopy from f to g, then $F'(x,t) = F(x, 1-t)$ is a homotopy from g to f. Finally, if F is a homotopy from f to g and G a homotopy from g to h, defining $H(x,t) = F(x,2t)$ for $0 \leq t \leq 1/2$ and $G(x, 2t-1)$ for $1/2 \leq t \leq 2$ yields a homotopy from f to h. So the relation whereby $f \sim g$ if $f \simeq g$ is reflexive, symmetric, and transitive, and hence an equivalence relation on $\mathsf{Top}(X, Y)$. Given two homotopies $f \simeq g : X \to Y$ and $h \simeq k : Y \to Z$, we may extend the homotopy $f \simeq g$ to a homotopy $h \cong f \simeq h \cong g$ that leaves h fixed but moves f to g, and likewise obtain a chain of homotopic maps $h \circ f \simeq h \circ g \simeq k \circ f \simeq k \circ g$. Therefore, we can define $[h] \circ [f]$ by taking the homotopy class of the composition of any representative of $[h]$ with any representative of $[f]$. ■

We may define a new category whose objects are those of Top, but whose morphisms are *homotopy classes* of morphisms in Top. This category, which is famously *not* concrete, is known as hTop.

We may sometimes want to restrict the set of homotopies between two maps $f, g : X \rightrightarrows Y$, requiring that all morphisms in our continuous family $F(x, -)$ preserve all points p in a subspace $X_0 \subseteq X$; such a homotopy is known as a homotopy relative to X_0. This is also an equivalence relation, the proof being more or less unchanged. The prototypical example is when $X = I, X_0 = \{0, 1\}$, and f, g are paths $I \to Y$; in this case, f is homotopic to g relative to the endpoints $\{0, 1\}$ when $F(x, 0) = f(x)$, $F(x, 1) = g(x)$, and $F(s, t) = f(s) = g(s)$ for all $s \in \{0, 1\}$.

A ***pointed space*** is a topological space X equipped with a specified element $x \in X$ known as the ***basepoint***. A ***basepoint-preserving map*** f between pointed spaces (X, x) and (Y, y) is a continuous map $X \to Y$ sending x to y. When working in the category Top_* of pointed spaces and basepoint-preserving maps, we often denote all basepoints as $*$, lazily stating that $f(*) = *$ and so on. Homotopies in this category must necessarily be relative to the basepoint.

Quotienting the hom-sets in Top_* by the equivalence relation of basepoint-preserving homotopy yields the homotopy category hTop_* of pointed topological spaces. The product in Top_* is the product in Top, with the basepoint being the product of the two basepoints. The coproduct is not the disjoint union, however, since there would be no canonical basepoint; Top_* remedies this in the most obvious possible way, by identifying the basepoints of the two spaces with a single point. This forms the ***wedge product*** $X \vee Y$. Denoting the basepoints of X and Y by x_0 and y_0, there is a canonical inclusion $X \vee Y \hookrightarrow X \times Y$ sending $x \in X \subseteq X \vee Y$ to (x, y_0) and $y \in Y \subseteq X \vee Y$ to (x_0, y). Identifying this subspace of $X \times Y$ with a point yields the smash product $X \wedge Y = X \times Y / X \vee Y$.

A.1.2 Categories of Topological Spaces

Since $\mathrm{Hom}_{Top_*}(-,-)$ is a bifunctor, we can immediately form four important endofunctors on Top_*. Letting S_1 have an arbitrary basepoint 0, and defining I to be the interval $[0,1]$ with the basepoint 0, these are:

- The loop space functor $\Omega = \mathrm{Hom}_{\mathsf{Top}_*}(S^1, -)$
- The path space functor $P = \mathrm{Hom}_{\mathsf{Top}_*}(I, -)$
- The reduced suspension functor $\Sigma = S^1 \wedge -$
- The reduced cylinder functor $C = I \wedge -$

The action of Ω and P on functions are canonically defined. The action of Σ and C on functions comes from the universal property of quotient spaces: if $A_0 \subseteq A$ and $B_0 \subseteq B$, then $f : A \to B$ extends to a unique map. Since $X \vee Y$ is sent to $X \vee f(Y) \subseteq X \vee Z$, this lets us define Σf and Cf for $X = S^1, I$. The action of the functor $X \wedge -$ on a map $f : Y \to Z$ is to send the image of (x, y) in $X \wedge Y$, which we can denote $x \wedge y$, to $x \wedge f(y) \in X \wedge Z$.

There are many nice properties of these functors which hold for most conceivable examples but fail to hold in general; for instance, the smash product is "usually" associative up to natural isomorphism, but fails to be so in general: as detailed in [May and Sigurdsson, 2006], $(\mathbb{Q} \wedge \mathbb{Q}) \wedge \mathbb{N}$ is not homeomorphic to $\mathbb{Q} \wedge (\mathbb{Q} \wedge \mathbb{N})$. As such, we may want to move to a more nicely behaved subcategory of Top_*, of which there are many. To specify certain subcategories, we need additional topological definitions. A space X is ***weak Hausdorff*** if, for all compact Hausdorff spaces Y and continuous functions $f : Y \to X$, the image of f is closed in X. X is a ***k-space*** if any subset $X_0 \subset X$ all of whose preimages are closed is itself closed. X is ***compactly generated*** if it is both weak Hausdorff and a k-space.

Topological manifolds, metric spaces, and compact Hausdorff spaces are all both compactly generated and Hausdorff, and are therefore contained in all of the following full subcategories of Top:

- kTop, the category of k-spaces
- wHaus, the category of weak Hausdorff spaces
- $\mathsf{CG} = \mathsf{kTop} \cap \mathsf{wHaus}$, the category of compactly generated spaces
- CGHaus, the category of compactly generated Hausdorff spaces

All of these have pointed, homotopy, and pointed homotopy variants. Letting i_k denote the inclusion functor $\mathsf{kTop} \to \mathsf{Top}$ and i_{wH} the inclusion functor $\mathsf{wHaus} \to \mathsf{Top}$, we have a triplet of adjunctions:

$$\begin{tikzcd}
\mathsf{kTop} \arrow[r, bend left, "i_k"] \arrow[r, phantom, "\perp"] & \mathsf{Top} \arrow[l, bend left, dashed, "k"] \arrow[r, bend left, dashed, "wH"] \arrow[r, phantom, "\perp"] & \mathsf{wHaus} \arrow[l, bend left, "i_{wH}"] \arrow[r, bend left, "k"] \arrow[r, phantom, "\top"] & \mathsf{CG} \arrow[l, bend left, "i_{CG}"]
\end{tikzcd}$$

The right adjoint k to i_k is known as k-ification, and the left adjoint wH to i_{wH} as weak Hausdorffification; k-ification turns a weak Hausdorff space into a compactly generated space, and, as a functor $\mathsf{wHaus} \to \mathsf{CG}$, is itself left adjoint to the inclusion functor $\mathsf{CG} \to \mathsf{wHaus}$. wHaus is complete, and right adjoints preserve limits, allowing us to construct limits in CG by constructing them in wHaus and then k-ifying. We will implicitly work in CG, letting $X \times Y$ denote the k-ification of the product in wHaus, and Y^X the k-ification of the space of maps from X to Y [1].

In CG, there is an adjunction $- \times Z \dashv (-)^Z$ for all Z, such that maps $X \times Z \to Y$ can be identified in a natural way with maps $X \to Y^Z$. In particular, $\mathrm{Hom}_{\mathsf{CG}}(X \times I, Y) \cong \mathrm{Hom}_{\mathsf{CG}}(X, PY)$ and $\mathrm{Hom}_{\mathsf{CG}}(X \times S^1, Y) \cong \mathrm{Hom}_{\mathsf{CG}}(X, \Omega Y)$. In the based version, CG_*, this becomes $- \wedge Z \dashv (-)^Z$, yielding the adjunctions $C \dashv P$ and $\Sigma \dashv \Omega$. (The exponential here ranges over basepoint-preserving maps, and its basepoint is the map that sends all points in the domain to the basepoint of the codomain). These adjunctions are preserved upon passing to homotopy classes. We will write $[X, Y]$ for $\mathrm{Hom}_{\mathsf{hCG}_*}(X, Y)$, leaving the basepoints implicit.

A.1.3 Homotopy Groups

Given two loops $\gamma_0, \gamma_1 : (S^1, *) \to (X, *)$, the ***composite loop*** $\gamma_0 * \gamma_1$ is defined by $(\gamma_0 * \gamma_1)(t) = \gamma_0(2t)$ if $0 \leq t \leq 1/2$, and $\gamma_1(2t - 1)$ if $1/2 \leq t \leq 1$. Under the operation of composition of loops, $[S^1, X]$ has the structure of a group.

Proof. The proof that $*$ respects homotopy equivalence is similar to that of $\circ$ respecting homotopy equivalence. We define the identity element on $[S^1, X]$ to be the constant loop $e(t) = *$,

[1]This space is equipped with the compact-open topology, whose subbase contains, for all $X_0 \subseteq X$, $Y_0 \subseteq Y$, the set of all functions $f : X \to Y$ with $f(X_0) \subseteq Y_0$.

and define the inverse of a loop $\gamma : S^1 \to X$ by the loop $\gamma^{-1}(t) = \gamma(1-t)$. To see that $[\gamma^{-1} * \gamma] = [e]$, use the homotopy $F(s,t) = \gamma_s(t) * \gamma_s(t)^{-1}$, where $\gamma_s(t) = \gamma(t)$ for $t \leq s$ and $\gamma(s)$ for $t \geq s$. This implies that $[\gamma * \gamma^{-1}] = [\left(\gamma^{-1}\right)^{-1} * \gamma^{-1}] = [e]$ as well, so $([S_1, X], *)$ has a multiplication, inverses, and a two-sided identity. ■

The ***fundamental group*** of a pointed space $(X, *)$ is defined as $\pi_1(X, *) := [S^1, X]$, with the group structure defined above. We will generally omit the $*$, just writing $\pi_1(X)$. The ***higher homotopy groups*** of a pointed space X are defined as $\pi_n(X) := [S^1, \Omega^{n-1}X] = \pi_1(\Omega^{n-1}X)$, $n \geq 1$. Since $S^n = \Sigma S^{n-1}$, we have $\pi_n(X) = [S^1, \Omega^n X] \cong [\Sigma^n S^1, X] \cong [S^n, X]$. This alternative definition allows us to interpret the nth homotopy group of a space X as the homotopically distinct ways of mapping the n-sphere into X in a basepoint-preserving manner, as well as to clearly demonstrate the functoriality of π_n; Every based map $f : X \to Y$ induces a map $\pi_n(X) \to \pi_n(Y)$ given by sending a loop $\ell : S^1 \to X$ to the loop $f \circ \ell : S^1 \to Y$. We can also define a *zeroth* homotopy group $\pi_0(X)$; this is just the set of path-connected components of X, and doesn't necessarily have a group structure.

As a consequence of the functoriality of homotopy groups, homeomorphic spaces have isomorphic fundamental groups. In fact, the motivation behind the introduction algebraic topology was the development of algebraic tools to figure out when two groups are homeomorphic.

A based map $f : X \to Y$ that induces isomorphisms $\pi_n(X) \cong \pi_n(Y)$ is known as a ***weak equivalence***. Two spaces X, Y are ***weakly equivalent***, written as $X \simeq Y$, when there is a weak equivalence between them. Homeomorphisms are weak equivalences, but the converse is not true in general; this means that, while two spaces X, Y with differing homotopy groups cannot be homeomorphic, verifying that all homotopy groups are the same isn't enough to verify that X and Y are homeomorphic.

Homotopy groups will serve as one of our primary methods of classifying topological spaces, and weak equivalence will serve as an important notion of equality in this classification. Another notion of equivalence is similar to that of categories: two spaces X and Y are ***homotopy equivalent*** if there are continuous $f : X \to Y$ and $g : Y \to X$ such that fg and gf are homotopic to the identity maps on Y and X, respectively. This is also a weaker property than homeomorphism.

A.1.4 CW Complexes

The vast majority of spaces that come to mind when one thinks of a topological space all share a common trait: they can be pieced together using points and n-disks in a systematic manner. The circle S^1, for instance, is constructed by attaching $D^1 \cong [0,1]$, to a single point at both ends. Attaching two copies of D^2 to the circle along their boundaries yields a sphere. A torus can be constructed in a similar manner with one point, two 1-disks, and one 2-disk, as shown below.

We can make this construction pattern rigorous. The general process is as follows:

1. Start with a set of points X^0.

2. Form an n-skeleton X^n from X^{n-1} by attaching a collection of open n-disks e^n_α via maps specifying where their boundary goes, $\varphi_n : S^{n-1} \to X^{n-1}$. We can say that X^n is the quotient space $X^{n-1} \coprod_\alpha D^n_\alpha$ of X^{n-1} under the identifications $x \sim \varphi_\alpha(x)$ for $x \in \partial D^n_\alpha$; as a set, $X^n = X^{n-1} \coprod_\alpha e^n_\alpha$.

3. Either stop at a finite stage (in which case X is finite-dimensional, and its dimension is n), or take the infinite union $X = \bigcup_n X^n$ and give it the weak topology, where A is open/closed in X iff $A \cap X^n$ is open/closed in X^n for all n.

Spaces constructed in this way are called ***CW complexes***, a.k.a. cell complexes. Some examples:

- A 1-dimensional CW complex is a ***graph***. (It's actually a multigraph, but we call it a graph).

- S^n is constructed with the cells e^0, a single point, and e^n, the disk D^n attached by the constant map $S^{n-1} \to e_0$. By part 2 of the construction, we can see that $S^n = D^n/\partial D^n$.

A ***subcomplex*** of a CW complex X is a closed subspace $A \subset X$ that's a union of cells in X; the closedness implies that the characteristic map of each of these cells has image contained in A, making A itself a CW complex. A pair (X, A) of a CW complex X and a subcomplex A is called a ***CW pair***. Since each skeleton X^n of a subcomplex X is a closed subspace of X, (X, X^n) is a CW pair.

CW complexes are especially well behaved; they are all compactly generated Hausdorff, locally contractible, and paracompact; the full subcategory CW of Top consisting of the CW complexes is closed under topological products, wedge sums, and smash products. Homotopy equivalence between CW complexes is equivalent to weak equivalence, and every topological space is weakly equivalent to a CW complex.

A.2 Homological Algebra

In this section, we'll add an increasing amount of structures to an arbitrary Ab-category, culminating in the definition of an *abelian category*. Such categories allow us to define homology and cohomology, and are very useful in the study of algebraic topology. R-Mod is the prototypical example of an abelian category, and in a sense is the universal example: the *Freyd-Mitchell embedding theorem* allows us to embed any category C, by means of a full and faithful functor, into some R-Mod. As such, we'll think of the elements of abelian categories as being R-modules, allowing us to work with elements rather than arrow-theoretic language.

A.2.1 Abelian Categories

In an Ab-category C, every hom-set is an abelian group, and composition is a bilinear operation $\circ_{XYZ} : \mathrm{Hom}_{\mathsf{C}}(X,Y) \times \mathrm{Hom}_{\mathsf{C}}(Y,Z) \to \mathrm{Hom}_{\mathsf{C}}(X,Z)$. An Ab-functor $F : \mathsf{C} \to \mathsf{D}$ between Ab-categories is a functor such that each mapping $\mathrm{Hom}_{\mathsf{C}}(X,Y) \to \mathrm{Hom}_{\mathsf{D}}(FX,FY)$ is a morphism in Ab, i.e. an abelian group homomorphism. Since Ab is a concrete category whose morphisms $1 = \mathbb{Z} \to G$ are in bijection with elements of G, the definition of an Ab-natural transformation simplifies to a family of homomorphisms $FX \to GX$ satisfying the usual commutativity condition.

Additive Categories In an Ab-category C, the finite product is, if it exists, equivalent to the coproduct. To see this, suppose for objects $X, Y \in \mathsf{C}$ we have a product $X \times Y$ with projections p_X and p_Y. Then, the pair of maps $(\mathrm{id}_X, 0_{XY})$ induces a morphism $i_X : X \to X \times Y$ such that $p_X i_X = \mathrm{id}_X$ and $p_Y i_X = 0_{XY}$; likewise, the pair of maps $(0_{YX}, \mathrm{id}_Y)$ induces a morphism $i_Y : Y \to X \times Y$. Take an object Z with morphisms $f : X \to Z$ and $g : Y \to Z$, and let $\varphi : X \times Y \to Z = f p_X + g p_Y$, such that $\varphi i_X = f p_X i_X + g p_Y i_X = f + g 0_{XY} = f$ and likewise $\varphi i_Y = g$. This construction satisfies the universal property of the coproduct, so $X \times Y$ is both a product and a coproduct. We call it the ***biproduct***, and denote it $\oplus$.

In an arbitrary category C, a ***zero object*** 0 is, if it exists, an object that is both initial and final. It has the special property that it defines a unique morphism, a ***zero morphism***, between any two objects X and Y: this morphism, denoted 0_{XY}, is given by the composition $X \to 0 \to Y$. We interpret the object 0 as carrying no information, and therefore zero morphisms destroy all information. An arbitrary Ab-category C has zero morphisms in a literal sense: they're the identities of the hom-groups. If C has a zero object 0, then $\mathrm{Hom}_{\mathsf{C}}(0, X)$, necessarily being the trivial group, generates these zero morphisms in the manner described above. An Ab-category with a zero object and finite biproducts is known as an ***additive category***.

Kernels In the Ab-category R-Mod, the zero object is simply the zero module. Once we have a zero object, we can take a morphism $f : X \to Y$ and define its ***kernel*** to be the equalizer of f with 0_{XY}, and its ***cokernel*** to be the coequalizer of f with 0_{XY}. Specifically, the kernel is an object K along with a morphism $\varphi : K \to X$ such that $f\varphi = 0_{KY}$, and any other K' with a ψ satisfying $f\psi = 0_{K'Y}$ has a unique $\rho : K' \to K$ such that $\psi = \varphi\rho$. In pictures,

$$
\begin{tikzcd}
 & X \arrow[dr, "f"] & \\
 & K \arrow[u, "\varphi"'] \arrow[r, "0_{KY}"] & Y \\
K' \arrow[uur, "\psi"] \arrow[ur, dashed, "\rho"] \arrow[urr, "0_{K'Y}"'] & &
\end{tikzcd}
$$

In Grp and Ab, K ends up being (isomorphic to) the set of all $x \in X$ that are mapped to 0 by f, with φ the inclusion map from K to X, recovering the normal definition of kernel. (While this case works out very nicely, as do cokernels, it must be emphasized that (co)kernels have not just objects but *morphisms* as well). The cokernel is an object Q along with a morphism $\varphi : Y \to Q$ such that $\varphi f = 0_{XQ}$, and any other Q' with a ψ satisfying $\psi f = 0_{XQ'}$ has a unique $\rho : Q \to Q'$ such that $\psi = \rho\varphi$. Another picture:

$$\begin{tikzcd}
 & Y \arrow[d, "\varphi"] \arrow[ddr, "\psi"] & \\
X \arrow[ur, "f"] \arrow[r, "0_{XQ}"] \arrow[drr, "0_{XQ'}"'] & Q \arrow[dr, dashed, "\rho"] & \\
 & & Q'
\end{tikzcd}$$

In Ab, the cokernel ends up being $Y/Im(f)$. In summary, if $f\psi = 0$, then ψ factors uniquely through $\ker f$, but if $\psi f = 0$, then ψ factors uniquely through coker f. Zero morphisms restrict the flow of information between two objects X and Y, kernels tell you how difficult it is to silence an X with a morphism $f : X \to Y$, and cokernels tell you how difficult it is to censor Y. The ***image*** of a morphism φ is defined by $\ker \operatorname{coker} \varphi$, and the ***coimage*** of φ is $\operatorname{coker} \ker \varphi$.

An additive category A is ***abelian*** if it has all kernels and cokernels, any monomorphism can be presented as the kernel of some morphism, and any epimorphism can be presented as the cokernel of some morphism.

A.2.2 Chain Complexes

In an abelian category A, a ***chain complex*** $C_\bullet$ is a collection $\{C_n\}_{n\in\mathbb{Z}}$ along with morphisms $\{d_n : C_n \to C_{n-1}\}_{n\in\mathbb{Z}}$, generally represented as a diagram of the form

$$\cdots \longrightarrow C_{n+1} \xrightarrow{d_{n+1}} C_n \xrightarrow{d_n} C_{n-1} \longrightarrow \cdots$$

We require that $d_n \circ d_{n+1} = 0$ for all n. This implies that $\ker d_n \subseteq \operatorname{im} d_{n+1}$ for all n; if these two submodules of C_n are equal for all n, then the chain complex $C_\bullet$ is said to be ***exact***. Dually, a ***cochain complex*** $C^\bullet$ is a collection of objects $\{C^n\}_{n\in\mathbb{Z}}$ and morphisms $\{d^n : C^{n-1} \to C^n\}$ such that $d^{n+1} \circ d^n = 0$. In specific instantiations of such complexes there may be a specific reason for going in one direction or the other. In the abstract sense, though, flipping the indices is really all we have to do; for this reason, chain and cochain complexes are more or less equivalent, and a chain complex $(C_\bullet, d_\bullet)$ generates a cochain complex $(C^{-\bullet}, d^{-\bullet})$.

Homology An arbitrary chain complex $C_\bullet$ may or may not be exact; the extent to which it fails to be exact at an index n is equivalent to the extent to which $\operatorname{im} d_{n+1}$ fails to be as large as $\ker d_n$. It will always be a submodule, though, so we can record this failure of exactness by taking the quotient module $\ker d_n / \mathrm{im} d_{n+1}$. The ***homology*** of the chain complex $(C_\bullet, d_\bullet)$ is defined by

$$H_n(C_\bullet) = \ker d_n / \mathrm{im} d_{n+1}$$

and the ***cohomology*** of a cochain complex $(C^\bullet, d^\bullet)$ is given by

$$H^n(C^\bullet) = \ker d_n / \mathrm{im} d_{n-1}$$

In R-Mod, elements of $\mathrm{im} d_{n+1}$ are known as the ***boundaries*** of C_n, and elements of $\ker d_n$ are known as the ***cycles*** of C_n; $H_n(C_\bullet)$ is then simply the submodule of cycles modulo the relation that identifies two cycles that differ only by a boundary.

The Category of Chain Complexes A ***morphism of chain complexes*** $C_\bullet \to D_\bullet$ is a family $u_\bullet$ of morphisms in A such that

$$\begin{tikzcd}
\cdots \arrow[r] & C_{n+1} \arrow[r, "d_{n+1}"] \arrow[d, "u_{n+1}"] & C_n \arrow[r, "d_n"] \arrow[d, "u_n"] & C_{n-1} \arrow[r] \arrow[d, "u_{n-1}"] & \cdots \\
\cdots \arrow[r] & D_{n+1} \arrow[r, "d'_{n+1}"] & D_n \arrow[r, "d'_n"] & D_{n-1} \arrow[r] & \cdots
\end{tikzcd}$$

is a commutative diagram. The set of all chain complexes on A, along with chain maps between chain complexes, forms a category $\mathsf{Ch}(\mathsf{A})$. This is itself an abelian category, with all kernels, cokernels, sums of morphisms, etc. being computed pointwise. Given a chain map $f : C_\bullet \to D_\bullet$ in $\mathsf{Ch}(R\text{-}\mathsf{Mod})$, we note that if $d_i(g) = 0$ for $g \in C_i$, then $d'_i f_i(g) = f_{i-1} d_i(g) = 0$, and that if $g = d_{i+1}(h)$, then $f_i(g) = f_i d_{i+1}(h) = d'_{i+1} f_{i+1}(h)$; chain maps send boundaries to boundaries and cycles to cycles, and hence induce well-defined maps $H_i(C_\bullet) \to H_i(D_\bullet)$. In this way, the map $H_i : \mathsf{Ch}(R\text{-}\mathsf{Mod}) \to R\text{-}\mathsf{Mod}, C_\bullet \mapsto H_i(C_\bullet)$ acts functorially; this holds for an arbitrary abelian category A. Two chain complexes are ***quasi-isomorphic*** if all of their homology objects are isomorphic; this provides a weaker notion of equivalence than isomorphism.

A chain complex is ***bounded*** if all but finitely many of the C_n are 0. If C_n is non-zero solely when $n \in [a, b]$, we say that $C_\bullet$ has ***amplitude*** in $[a, b]$. $C_\bullet$ is ***bounded above*** if there's a b such that $C_n = 0$ for all $n > b$, and ***bounded below*** if there's an a such that $C_n = 0$ for all $n < a$. Keeping in line with the identification $C_n = C^{-n}$, a cochain complex is bounded above/below iff its associated chain complex is bounded below/above. These allow us to form full subcategories of Ch(A): the categories of bounded, bounded above, bounded below, and non-negative chain complexes are denoted $\mathrm{Ch}(\mathsf{A})_b$, $\mathrm{Ch}(\mathsf{A})_-$, $\mathrm{Ch}(\mathsf{A})_+$, and $\mathrm{Ch}(\mathsf{A})_{\geq 0}$, respectively.

Chain Homotopies A chain complex $C_\bullet$ is ***split*** if there are maps $s_n : C_n \to C_{n+1}$ such that $d = dsd$. It is ***split exact*** if it is also exact; equivalently, it is split exact if and only if $ds + sd$ is the identity map. If we have a chain map $f : C_\bullet \to D_\bullet$, f is called ***null homotopic*** if there are maps $s_n : C_n \to D_{n+1}$ such that $f = ds + sd$. Two chain maps $f, g : C_\bullet \rightrightarrows D_\bullet$ are ***chain homotopic*** if their difference $f - g$ is null homotopic, i.e. there are maps $s_n : C_n \to D_{n+1}$ such that $f - g = ds + sd$. A diagram:

$$\begin{tikzcd}
\cdots \arrow[r] & C_{n+1} \arrow[r, "d"] \arrow[d, bend right, "f"'] \arrow[d, dashed, "f-g" description] \arrow[d, bend left, "g"] & C_n \arrow[r, "d"] \arrow[dl, "s"] \arrow[d, bend right, "f"'] \arrow[d, dashed, "f-g" description] \arrow[d, bend left, "g"] & C_{n-1} \arrow[r] \arrow[dl, "s"] \arrow[d, bend right, "f"'] \arrow[d, dashed, "f-g" description] \arrow[d, bend left, "g"] & \cdots \\
\cdots \arrow[r] & D_{n+1} \arrow[r, "d"] & D_n \arrow[r, "d"] & D_{n-1} \arrow[r] & \cdots
\end{tikzcd}$$

The maps $\{s_n\}$ are collectively called a ***chain homotopy***. We will regard the notion of chain homotopy as an extension of the notion of a homotopy between maps between topological spaces. Correspondingly, we call two chain complexes $C_\bullet$ and $D_\bullet$ ***chain homotopy equivalent*** if there are maps $f : C_\bullet \to D_\bullet$ and $g : D_\bullet$ to $C_\bullet$ such that gf and fg are equivalent to the identities on $D_\bullet$ and $C_\bullet$, respectively.

A.2.3 Resolutions

Let $F : A \to B$ be an Ab-functor between abelian categories A, B. If, for all exact sequences in A of the form $0 \to X \to Y \to Z \to 0$, F yields an exact sequence $0 \to FX \to FY \to FZ \to 0$, F is known as a ***exact functor***. If just $0 \to FX \to FY \to FZ$ is exact, F is known as ***left exact***, and if $FX \to FY \to FZ \to 0$ is exact, F is known as ***right exact***.

For a fixed $M \in \mathsf{A}$, the covariant representable functor $\mathrm{Hom}_{\mathsf{A}}(M, -)$ is left exact. To see this, let $0 \to X \xrightarrow{f} Y \xrightarrow{g} Z \to 0$ be exact. As in R-Mod, f must be monic and g must be epic. Take the map $f_* := \mathrm{Hom}_{\mathsf{A}}(X, f)$ sending $\varphi : M \to X$ to $f\varphi : M \to Y$. If $f\varphi = 0_{MY}$, then since f is monic, φ must be 0_{MX}. So f_* is monic, and likewise $g_* f_*(\varphi) = gf\varphi = 0_{XZ}\varphi = 0_{MZ}$, so $g_* f_* = 0_{\mathrm{Hom}_{\mathsf{A}}(M,X), \mathrm{Hom}_{\mathsf{A}}(M,Z)}$. Finally, if $\varphi : M \to Y$ satisfies $g_*(\varphi) = 0$, then, since $\mathrm{im}\varphi$

is a subobject of $\mathrm{im} f$, φ factors through f as $\varphi = f\psi = f_*(\psi)$ for some $\psi : M \to X$. So $0 \to \mathrm{Hom}_{\mathsf{A}}(M, X) \to \mathrm{Hom}_{\mathsf{A}}(M, Y) \to \mathrm{Hom}_{\mathsf{A}}(M, Z)$ is exact. Hence, $\mathrm{Hom}_{\mathsf{A}}(M, -)$ is a left exact functor.

Projective Objects It is not in general true that the final arrow $\mathrm{Hom}_{\mathsf{A}}(M, Y) \to \mathrm{Hom}_{\mathsf{A}}(M, Z)$ is an epimorphism, so that we could extend the left exact sequence to an exact sequence. For this to be true, we require the following (equivalent) universal lifting property on M: given any surjection $g : Y \to Z$ in A, and any map $\varphi : M \to Z$, there is a (not necessarily unique) map $\psi : M \to Y$ such that $\varphi = f\psi$. If M had this property, it would follow immediately that $\mathrm{Hom}_{\mathsf{A}}(M, Y) \to \mathrm{Hom}_{\mathsf{A}}(M, Z)$ is an epimorphism, and hence that $\mathrm{Hom}_{\mathsf{A}}(M, -)$ is an exact functor. If M satisfies this universal lifting property, or equivalently if $\mathrm{Hom}_{\mathsf{A}}(M, -)$ is an exact functor, we call M a ***projective object***. For instance, free modules are projective. For some nice rings R, including $\mathbb{Z}$, fields, and division rings, the projective R-modules are the free modules, but this isn't always the case. In general, an R-module is projective if and only if it's a direct summand of a free R-module.

Injectives The dual notion is that of an ***injective object***, or an object $M \in \mathsf{A}$ such that every monomorphism $f : X \to Y$ and map $\varphi : X \to M$ yields at least one $\psi : Y \to M$ such that $f\psi = \varphi$. The contravariant functor $\mathrm{Hom}_{\mathsf{A}}(-, M)$ is right exact, since it is $\mathrm{Hom}_{\mathsf{A}^{op}}(M, -)$ which, A^{op} being abelian, sends exact sequences in A^{op} to left exact sequences in Ab, and hence exact sequences in A to right exact sequences in Ab). $\mathrm{Hom}_{\mathsf{A}}(-, M)$ is exact if and only if M is injective. Injective modules are harder to characterize then projective modules, but if $\mathsf{A} = R\text{-}\mathsf{Mod}$ for R a principal ideal domain, then M is injective if and only if for every $r \neq 0 \in r$ and $m \in M$, $m = rm'$ for some $m' \in M$, so that we can "divide" elements of M by nonzero elements of R. For instance, $\mathbb{Q}$ is injective as a $\mathbb{Z}$-module.

It is in general true that left adjoints are right exact and right adjoints are left exact, since left adjoints preserve colimits, and hence cokernels, and right adjoints preserve kernels. In the case $\mathsf{A} = R\text{-}\mathsf{Mod}$, this observation is another way to show that $\mathrm{Hom}_R(M, -)$ is left exact, and its left adjoint $M \otimes_R -$ is right exact.

Resolutions For some nice rings R, including $\mathbb{Z}$, fields, and division rings, the projective R-modules are the free modules, but this isn't always the case. In general, an R-module is projective if and only if it's a direct summand of a free R-module. R-Mod has enough projectives: given an R-module A, take the free R-module on the set of elements of A, $\pi(A) := FJA$. The counit of the $F \dashv J$ adjunction gives us a natural map $\pi(A) \to A$ (that sends a sequence of elements of A to its sum) which is a surjection.

An abelian category A has ***enough projectives*** if for every $M \in \mathsf{A}$ there is an epimorphism from a projective object P to M, and ***enough injectives*** if there is a monomorphism from X to an injective object I. A ***left resolution*** of M is a complex $X_\bullet$ along with a map $\epsilon : X_0 \to M$ such that the following sequence

$$\dots \longrightarrow X_2 \xrightarrow{d_2} X_1 \xrightarrow{d_1} X_0 \xrightarrow{\epsilon} M \longrightarrow 0$$

is exact. If furthermore all X_i are projective objects, then $X_\bullet$ is known as a ***projective resolution*** of M. Dually, a ***right resolution*** of M is a cochain complex $X^\bullet$ along with a map $\epsilon : M \to X^0$ such that the sequence

$$0 \longrightarrow M \xrightarrow{\epsilon} X^0 \xrightarrow{d^1} X^1 \xrightarrow{d^2} X^2 \longrightarrow \dots$$

is exact. If all X^i are injective, $X^\bullet$ is known as a ***injective resolution***.

In an abelian category A with enough projectives (injectives), *every* object $M \in \mathsf{A}$ has a projective (injective) resolution.

Proof. Choosing a projection $\epsilon_0 : P_0 \to M$, we recursively choose a projective P_n and an epimorphism $\epsilon_n : P_n \to M_{n-1}$, set $M_n = \ker \epsilon_n$, and let $d_n : P_n \to P_{n-1}$ be the composition $P_n \to M_{n-1} \to P_{n-1}$. See:

$$\begin{tikzcd}
 & 0 & & 0 \arrow[d] & & 0 & & & \\
0 \arrow[r] & M_4 \arrow[u] \arrow[r] & P_4 \arrow[r] \arrow[dr, "d"] & M_3 \arrow[r] \arrow[d] & 0 \arrow[r] & M_0 \arrow[r] \arrow[u] & P_0 \arrow[dr, "\epsilon_0"] & & \\
 & P_5 \arrow[u] \arrow[ur, "d"] & & P_3 \arrow[d] \arrow[dr, "d"] & & P_1 \arrow[u] \arrow[ur, "d"] & & M \arrow[r] & 0 \\
\cdots \arrow[r] \arrow[ur] & M_5 \arrow[u] \arrow[r] & 0 \arrow[r] & M_2 \arrow[r] \arrow[d] & P_2 \arrow[r] \arrow[ur, "d"] & M_1 \arrow[r] \arrow[u] & 0 & & \\
 & 0 \arrow[u] & & 0 & & 0 \arrow[u] & & &
\end{tikzcd}$$

Using our $\pi(A) \to A$ projection as ϵ_0, we see that M_0 consists of all sequences in $\pi(A)$ that sum to 0 (and comes with an injection into P_0), P_1 is $\pi(M_0)$, coming with a canonical ϵ_1, and so on. The kernel of each d is the image of the next, by design, so this is a projective resolution of M.

The proof for injective objects is dual to the above proof. ■

Maps between objects M, N naturally induce chain maps between projective resolutions. Letting $P_\bullet \xrightarrow{\epsilon} M$, $Q_\bullet \xrightarrow{\eta} N$ be projective resolutions of M and N, and f a morphism $M \to N$, there is a chain map $\alpha : P_\bullet \to Q_\bullet$ that lifts f in the sense that the following diagram commutes:

$$\begin{tikzcd}
\cdots \arrow[r] & P_2 \arrow[r] \arrow[d, "\alpha_2"'] & P_1 \arrow[r] \arrow[d, "\alpha_1"'] & P_0 \arrow[r, "\epsilon"'] \arrow[d, "\alpha_0"'] & M \arrow[r] \arrow[d, "f"'] & 0 \\
\cdots \arrow[r] & Q_2 \arrow[r] & Q_1 \arrow[r] & Q_0 \arrow[r, "\eta"] & N \arrow[r] & 0
\end{tikzcd}$$

This chain map is unique up to chain homotopy equivalence.

The dual phenomenon is observed with injective objects: an injective resolution $N \xrightarrow{\theta} I^\bullet$ is naturally lifted *by* f to an injective resolution $M \xrightarrow{\zeta} E^\bullet$ in a way that makes the following diagram commute:

$$\begin{tikzcd}
0 \arrow[r] & M \arrow[r, "\zeta"'] \arrow[d, "f"'] & E^0 \arrow[r] \arrow[d, "\alpha^0"'] & E^1 \arrow[r] \arrow[d, "\alpha^1"'] & E^2 \arrow[r] \arrow[d, "\alpha^2"'] & \cdots \\
0 \arrow[r] & N \arrow[r, "\theta"] & I^0 \arrow[r] & I^1 \arrow[r] & I^2 \arrow[r] & \cdots
\end{tikzcd}$$

A.2.4 Derived Functors

Left Derived Functors Fix a right exact functor F, and take an R-module M. Given a projective resolution $P_\bullet$ of M, $FP_1 \to FP_0 \to FM \to 0$ is an exact sequence, but the rest of $FP_\bullet$ isn't necessarily exact. The ith homology of $FP_\bullet$ is known as the ith ***left derived functor*** of F, $L_iF(M) := H_i(FP_\bullet)$. The homology at the zeroth position is given by $L_0F(M) = FM$, so the ith derived functor of F can be seen as the ith "homological extension" of F, with the zeroth extension obviously being F itself. The module $L_iF(M)$ is independent of the projective resolution we choose for M: any two different projective resolutions $P_\bullet, Q_\bullet$ will yield a pair of chain maps $f : P_\bullet \to Q_\bullet$, $g : Q_\bullet \to P_\bullet$ each lifting the identity map id_M, implying that $h = gf$ is a map $P_\bullet \to P_\bullet$ lifting id_M from $P_\bullet$ to itself. Since $\mathrm{id}_{P_\bullet}$ also serves this role, and h is unique up to chain homotopy, h and $\mathrm{id}_{P_\bullet}$ must be chain homotopic, and hence induce equivalent maps on homology, implying that the transformation induced by using $Q_\bullet$ instead of $P_\bullet$ – which is a natural transformation – has an inverse, and hence a natural isomorphism.

Example. Our canonical example of a right exact functor on R-Mod is $- \otimes_R N$; its corresponding left derived functors are known as the ***Tor*** functors, defined by

$$\mathrm{Tor}_i^R(M, N) := L_i(- \otimes_R N)(M)$$

$\mathrm{Hom}_R(-, N)$ is also right exact, and we define the ***Ext*** functors by

$$\mathrm{Ext}_R^i(M, N) := L_i(\mathrm{Hom}_R(-, N))(M)$$

Right Derived Functors Given a *left* exact functor F and an R-module M with an (again, arbitrary) injective resolution $I^\bullet$, we can define the ***right derived functor*** $R^iF(M)$ to be the ith cohomology of $FI^\bullet$, $R^iF(M) := H^i(FI^\bullet)$. When $F = \mathrm{Hom}_R(M, -)$, we again arrive at $\mathrm{Ext}_R^i(M, N) := R^i(\mathrm{Hom}_R(M, -))(N)$. Namely, it doesn't matter if we compute the Ext functor via a left or right derived functor, and in the same vein we can show that $L_i(- \otimes_R N)(M) \cong L_i(M \otimes_R -)(N) \cong \mathrm{Tor}_i^R(M, N)$; further exposition can be found in [Weibel, 1995].

A table of correspondences:

Left derived functor L_iF	Right derived functor R^iG
Right exact functor F	Left exact functor G
Projective resolution $P_\bullet \to A$	Injective resolution $A \to I_\bullet$
$L_iF(A) = H_i(F(P))$	$R^iG(A) = H^i(G(P))$
Tor functor	Ext functor

For computational purposes, it's useful to note that Tor_i^R preserves filtered colimits – colimits over what are essentially directed preorders – and in particular directed limits (which are, confusingly, actually colimits) in both variables. In the case of $\mathsf{Ab} = \mathbb{Z}\text{-}\mathsf{Mod}$, since every abelian group G is the direct limit of its finitely generated subgroups, we only need to know a few values of $\mathrm{Tor}_i^{\mathbb{Z}}$, perhaps computed directly via selecting convenient projective resolutions, in order to compute a wide variety of Tor groups.

Example. For an arbitrary abelian group G, we may calculate $\mathrm{Tor}_i^{\mathbb{Z}}(\mathbb{Z}_n, G)$ by selecting the projective resolution $0 \to \mathbb{Z} \overset{\times n}{\to} \mathbb{Z} \to \mathbb{Z}_n \to 0$, which upon tensoring with G becomes $0 \to G \overset{\times n}{\to} G \to 0$. The homology of this complex at the 0th position is G/nG, and the homology at the first position is the n-torsion subgroup ${}_nG = \{g \in G \mid ng = 0\}$. So $\mathrm{Tor}_0^{\mathbb{Z}}(\mathbb{Z}_n, G) = G/nG$, and $\mathrm{Tor}_1^{\mathbb{Z}}(\mathbb{Z}_n, G) = {}_nG$. (The ability of Tor to compute torsion subgroups is where Tor gets its name). In fact, since every abelian group G can be written as the direct limit of its finitely generated subgroups, each of which is either some $\mathbb{Z}^n$ or some $\mathbb{Z}_n$, this approach can be used to show that $\mathrm{Tor}_i^{\mathbb{Z}}(G, H)$ vanishes for $i \geq 2$.

In contrast, Ext is named after its ability to compute extensions of R-modules. An ***extension*** of M by N is an exact sequence $0 \to N \to X \to M \to 0$, and such an extension ***splits*** if $X \cong M \oplus N$. If $\mathrm{Ext}_R^1(M, N)$ vanishes, then every extension of M by N splits; Ext^1 therefore tells us what obstruction prevents a given extension of M by N from splitting.

A.3 Cohomology Theories

A.3.1 Spectra

Generalized Cohomology Theories A ***generalized homology theory*** is a covariant functor h_n from the space $(\mathsf{Top}, \mathsf{Top})_*$ of pairs of pointed topological spaces, or inclusions $* \in A \subseteq X$ and inclusion-preserving continuous pointed maps, to $\mathsf{Ab}^{\mathbb{Z}}$, satisfying the following conditions:

1. h_* sends homotopic maps to equalities: writing f_* for $h_*(f)$, if $f \simeq g : (X, A) \to (Y, B)$, then $f_* = g_*$.
2. h_* preserves coproducts, sending arbitrary joins to direct sums.
3. Writing $h_*(X)$ for $h_*(X, *)$ and h_n for the nth element of an object of $\mathsf{Ab}^{\mathbb{Z}}$, there are exact sequences
$$\ldots \to h_{n+1}(X, A) \to h_n(A) \to h_n(X) \to h_n(X, A) \to h_{n-1}(A) \to \ldots$$
4. For a pair (X, A) and $U \subset A$ whose closure is contained in the interior of A, the inclusion $(X - U, A - U)$ induces an isomorphism on homology.

A ***generalized cohomology theory*** is a contravariant functor $h^* : (\mathsf{Top}, \mathsf{Top})_*^{\mathrm{op}} \to \mathsf{Ab}^{\mathbb{Z}}$ satisfying the duals of the above conditions (since products and coproducts coincide in Ab, we only need to turn the chain complex in condition 3 into a cochain complex). We may reduce a generalized (co)homology theory by restricting it to the category of (pointed, connected) CW complexes, CW, identifying a complex X with the pair $(X, *)$. The reduction will be denoted as $\tilde{h}_*$ for homology, $\tilde{h}^*$ for cohomology. We have for any reduced cohomology theory a Mayer-Vietoris sequence: given a homotopy pushout square

$$\begin{tikzcd} X \arrow[r] \arrow[d] & Y \arrow[d] \\ Z \arrow[r] \arrow[ur, phantom, "\ulcorner"] & Y +^h_X Z \end{tikzcd}$$

there is a natural long exact sequence

$$\ldots \to \tilde{h}^{n-1}(X) \to \tilde{h}^n(Y +^h_X Z) \to \tilde{h}^n(Y) \oplus \tilde{h}^n(Z) \to \tilde{h}^n(X) \to \tilde{h}^{n+1}(Y +^h_X Z) \to \ldots$$

By the Brown representability theorem, every reduced cohomology theory $\tilde{h}^* : \mathsf{CW}^{\mathrm{op}} \to \mathsf{Ab}^{\mathbb{Z}}$ yields a family $\{E_n\}_{n \in \mathbb{Z}}$ of CW complexes such that E_n represents $\tilde{h}^n$.

Spectra Recall that, given a pointed topological space X, we may construct the (reduced) suspension $\Sigma X = S^1 \wedge X$, with the smash product $\wedge$ defined on pointed topological spaces X, Y by taking the product $X \times Y$ and identifying the inclusions $X, Y \to X \times Y$ with the same point; that is, $X \wedge Y = X \times Y / X \vee Y$. Equivalently, we may construct ΣX as a homotopy pushout diagram:

$$\begin{tikzcd} X \arrow[r] \arrow[d] & CX \arrow[d] \\ CX \arrow[r] \arrow[ur, phantom, "\ulcorner"] & \Sigma X \end{tikzcd}$$

This defines a functor left adjoint to the loop space functor Ω. This adjunction descends to homotopy: the set of all homotopy classes of maps from ΣX to Y, denoted $[\Sigma X, Y]$, is in natural bijection with $[X, \Omega Y]$.

Given a reduced cohomology theory $\tilde{h}^*$ and its representation $\tilde{h}^n(X) \cong [X, E_n]$, we may apply the Mayer-Vietoris sequence to the homotopy pushout $\Sigma X = CX +^h_X CX$: since CX is contractible, this results in natural isomorphisms $\tilde{h}^n(X) \cong \tilde{h}^{n+1}(\Sigma X)$. Hence, $[X, E_n] \cong [\Sigma X, E_{n+1}] \cong [X, \Omega E_{n+1}]$; these isomorphisms are natural in X, implying by the Yoneda lemma that they arise from a homotopy equivalence $E_n \to \Omega E_{n+1}$. By adjunction, this gives a sequence of homotopy classes of maps $\Sigma E_n \to E_{n+1}$, though these aren't necessarily homotopy equivalences. This leads us to the definition of a spectrum.

A ***spectrum*** is a sequence $\{E_n\}_{n \in \mathbb{N}}$ of pointed, connected CW complexes, along with *structure maps* $\Sigma E_n \to E_{n+1}$. Every cohomology theory defines a spectrum in the manner outlined above (use the Brown representability theorem to obtain a representing space E_n for each $\tilde{h}^n$, apply Mayer-Vietoris and the $\Sigma \dashv \Omega$ adjunction to get natural isomorphisms $[X, E_n] \cong [X, \Omega E_{n+1}]$, apply Yoneda's lemma and the same adjunction to get a spectrum). In addition, every space $X \in \mathsf{CW}$ defines a ***suspension spectrum*** $\Sigma^\infty X$ whose nth space is $\Sigma^n X$, and where the inclusions are identities.

The Stable Homotopy Category The fundamental result motivating the theory of spectra is the *Freudenthal suspension theorem*: for an n-connected pointed CW complex X, the map $\pi_k(X) \to \pi_k(\Omega\Sigma X) \cong \pi_{k+1}(\Sigma X)$ induced by the unit map $X \to \Omega\Sigma X$ is an isomorphism for $k \leq 2n$. Hence, the reduced suspension will be at least $(n+1)$-connected, and there will be an m such that the maps $\pi_{m+k}(\Sigma^m X) \to \pi_{m+k+1}(\Sigma^{m+1} X)$ are isomorphisms for all k. We say that the homotopy groups $\pi_{n+k}(\Sigma^n X)$ *stabilize*. Correspondingly, we define the kth homotopy group of a spectrum E by

$$\pi_k(E) = \varinjlim_n \pi_{n+k}(E_n)$$

where the maps $\pi_{n+k}(E_n) \to \pi_{n+1+k}(E_{n+1})$ are given by the unit maps $\pi_{n+k}(E_n) \to \pi_{n+1+k}(\Sigma E_n)$ followed by the images of the structure maps $\Sigma E_n \to E_{n+1}$.

Define the ***sphere spectrum*** $\mathbb{S}$ to be $\Sigma^\infty S^0$, so that $\mathbb{S}_n = S^n$. The homotopy groups of $\mathbb{S}$ are known as the *stable homotopy groups of spheres*, denoted by π_n^S; their calculation is the subject of immense study, and are very well known. While one can deduce that $\pi_0^S = \mathbb{Z}$ from the fact that $\pi_n(S^n) = \mathbb{Z}$ for $n \geq 1$, the higher stable homotopy groups are a bit trickier to deduce: we tabulate a few below, from [Ravenel, 2003].

Stable homotopy groups of spheres π_n^S up to $n = 10$.

n	0	1	2	3	4	5	6	7	8	9	10
π_n^S	$\mathbb{Z}$	$\mathbb{Z}_2$	$\mathbb{Z}_2$	$\mathbb{Z}_{24}$	0	0	$\mathbb{Z}_2$	$\mathbb{Z}_{240}$	$\mathbb{Z}_2 \oplus \mathbb{Z}_2$	$\mathbb{Z}_2 \oplus \mathbb{Z}_2 \oplus \mathbb{Z}_2$	$\mathbb{Z}_6$

Spectra study what is "eventually true" about homotopy groups. This motivates the next few definitions. Given a spectrum E, a subspectrum E' of E is a sequence of subcomplexes $E'_n \subseteq E_n$ with $\Sigma E'_n$ a subcomplex of E'_{n+1}. If for every cell X in an E_m there is an n such that X is a cell of some E'_n, we call E' a cofinal subspectrum. Thus, every cell of E is eventually contained within the cofinal subspectrum. Clearly, the intersection of cofinal subspectra is a cofinal subspectrum, and cofinality is a transitive relation: if E'' is cofinal in E' is cofinal in E, then E'' is cofinal in E.

We must define a notion of a morphism between spectra in several steps.

1. First, consider collections of maps $\{f_n : E_n \to F_n\}_{n\in\mathbb{N}}$ such that Σf_n and f_{n+1} commute with the inclusions $\Sigma E_n \to E_{n+1}, \Sigma F_n \to F_{n+1}$. Call these *functions*.
2. Second, we remember our motto of "eventual truth", and consider functions from a cofinal subspectrum $E' \subset E$ to F, calling two functions $f : E' \to F, g : E'' \to F$ equivalent if they coincide on some cofinal subspectrum $E'' \subseteq E' \cap E''$.
   Functions from a cofinal subspectrum $E' \subset E$ will *eventually* be defined on all cells in E, and restricting to equivalence classes allows us to consider two functions which are eventually equal identical to one another. Call these *maps*.

3. Third, given a CW complex X and a spectrum E, define the spectrum $X \wedge E$ by $(X \wedge E)_n = X \wedge E_n$, with obvious structure maps. Call two maps $f, g : E \to F$ *homotopic* if there's a map $h : E \wedge [0,1] \to F$ such that the inclusions from $E \wedge \{0\}$ and $E \wedge \{1\}$ yield f and g respectively when composed with h. Call a homotopy class of maps $E \to F$ a *morphism*.

Spectra and their maps define the category of spectra Sp, while spectra and their morphisms define the ***stable homotopy category*** SHC. The theory of this category bears an incredible resemblance to homological algebra, as described in [Weibel, 1995], 10.9. In particular, we can attach an additive structure compatible with the monoidal closed structure induced by the smash product $\wedge$.

Constructing Cohomology Theories As we have seen, we can associate to each reduced cohomology theory $\tilde{h}^* : \mathsf{CW}^{\mathrm{op}} \to \mathsf{Ab}^{\mathbb{Z}}$ a spectrum $E = \{E_n, f_n : \Sigma E_n \to E_{n+1}\}_{n \in \mathbb{N}}$ such that $\tilde{h}^n(X) \cong [X, E_n]$. The spectra obtained in this manner are ***Ω-spectra***, characterized by the property that the maps $E_n \to \Omega E_{n+1}$ adjunct to the structure maps $\Sigma E_n \to E_{n+1}$ are weak homotopy equivalences.

Conversely, suppose we have a spectrum E in the stable homotopy category SHC, whose internal hom we will also denote $[-,-]$. We may obtain a reduced cohomology theory E^* from E by the formula

$$E^n(X) = [\Sigma^\infty X, \Sigma^n E]$$

where $\Sigma^n E$ is the obvious spectrum: $(\Sigma^n E)_m = \Sigma^n(E_m)$ and the structure maps $(\Sigma^n f)_m : \Sigma(\Sigma^n E)_m \to (\Sigma^n E)_{m+1}$ are given by $(\Sigma^n f)_m = \Sigma^n(f_m)$; we may obtain a reduced homology theory E_* as $E_n(X) = \pi_n(X \wedge E)$. (A proof that these do indeed define (co)homology theories is given in [Switzer, 2017], 8.33).

For instance, starting with an $X \in \mathsf{CW}$, we may construct a cohomology theory $(\Sigma^\infty X)^n(Y) = [\Sigma^\infty Y, \Sigma^n \Sigma^\infty X]$; for $X = S^0$, with $\Sigma^\infty X = \mathbb{S}$ the sphere spectrum, this yields the ***stable cohomotopy*** cohomology theory.

Higher Algebra Denote by $\mathcal{CAlg}$ the ∞-category of $\mathbb{E}_\infty$-rings and their morphisms (as commutative monoids). The stable homotopy groups $\pi_k(E)$ of an $\mathbb{E}_\infty$-ring E carry an abelian group structure; since the π_k are functorial, they translate the commutative monoid structure on E into a commutative monoid structure on each abelian group $\pi_k(E)$, forming a *ring*. Every commutative ring R with underlying abelian group $\underline{R}$ can be obtained in this way by equipping $\Sigma^\infty K(\underline{R}, 1)$ with the right commutative monoid structure, giving us an embedding $\mathsf{CRing} \to \mathcal{CAlg}$. For $\mathbb{Z}$, this gives the sphere spectrum $\mathbb{S}$, with $\mathbb{E}_\infty$-ring structure given by the isomorphism $S^m \wedge S^n \cong S^{m+n}$. This is the sense in which $\mathbb{E}_\infty$-rings are higher homotopical generalizations of rings.

A.3.2 Singular Cohomology

Take a topological space X. Let $\mathrm{Hom}(\Delta^n, X)$ be the set of maps from the space

$$\Delta^n = \{(x_0, \ldots, x_n) \in \mathbb{R}^{n+1} \mid x_0 + \ldots + x_n = 1, x_0, \ldots, x_n \geq 0\}$$

known as the n-***simplex***, to X. The images of maps $\alpha, \beta, \ldots$ in this set are known as ***singular*** n-***simplices***, and denoted $\alpha|[v_0, \ldots, v_n]$, where each vertex v_i is the image of the vertex e_i of Δ^n. We write $\alpha|[v_0, \ldots, \widehat{v}_i, \ldots, v_n]$ for the singular $(n-1)$-simplex obtained by projecting the regular n-simplex onto the face opposing the ith vertex and sending that to X. Let $C_n(X)$ be the free abelian group on $\mathrm{Hom}(\Delta^n, X)$, whose elements are known as n-***chains***, and $\partial_n : C_n(X) \to C_{n-1}(X)$ the linear map defined on bases as

$$\partial_n(\alpha) = \sum_{i=0}^{n} (-1)^i \alpha|[v_0, \ldots, \widehat{v}_i, \ldots, v_n]$$

known as the ***boundary operator***. For instance, ∂_1 sends a singular 1-simplex, or a path in X, to the 0-chain consisting of its end minus its beginning. It's easy to check that $\partial_{n-1}\partial_n = 0$, so $(C_\bullet, \partial_\bullet)$ forms a chain complex of abelian groups. Its homology groups are known as the ***singular homology groups*** of X.

A map $f : X \to Y$ generates a map $f_\sharp : C_n(X) \to C_n(Y)$ sending $\alpha : \Delta^n \to X$ to $f\alpha : \Delta^n \to Y$. $f_\sharp \partial_n^{(X)} = \partial_n^{(Y)} f_\sharp$, so this map is a chain map, and hence extends to a map $f_* : H_n(X) \to H_n(Y)$ evidencing H_n as a functor $\mathsf{Top} \to \mathsf{Ab}$; homotopic maps induce the same map, so H_n is in fact a map $\mathsf{hTop} \to \mathsf{Ab}$.

Given a group G, let $C^n(X)$ be the set of all homomorphisms $C_n(X) \to G$, known as n-***cochains***, which is itself an abelian group. We may precompose any morphism with ∂_{n+1} to obtain a homomorphism $\delta^{n+1} : C^n(X) \to C^{n+1}(X), \varphi \mapsto \alpha\partial_{n+1}$ known as the ***coboundary operator***. Since $\delta^n\delta^{n-1}(\alpha)(\varphi) = \varphi\partial_{n-1}\partial_n = 0$, $(C^\bullet, \delta^\bullet)$ is a cochain complex, whose cohomology groups $H^n(X; G)$ are known as X's ***singular cohomology groups with coefficients in*** G. The failure of $H^n(X; G)$ to be equivalent to $\mathrm{Hom}_{\mathsf{Ab}}(H_n(X), G)$ is given by the ***universal coefficient theorem for homology***, which states that the sequence

$$0 \longrightarrow \mathrm{Ext}(H_{n-1}(X), G) \longrightarrow H^n(X; G) \longrightarrow \mathrm{Hom}_{\mathsf{Ab}}(H_n(X), G)$$

is split exact; this is a purely algebraic fact, but evidences $H^n(-; G)$ as a functor $\mathsf{hTop} \to \mathsf{Ab}$ as well, and is often useful in computing cohomology groups in cases where Ext is easy to calculate. Methods for doing actual calculations, such as specific instantiations of the Mayer-Vietoris sequence, are given in [Hatcher, 2005].

Eilenberg-MacLane Spaces For G an abelian group, the ***Eilenberg-MacLane*** space $K(G, n)$ is the CW complex, unique up to weak equivalence, such that $\pi_i(K(G, n)) = G$ when $i = n$ and 0 otherwise. It is expedient to give a few examples: $K(\mathbb{Z}, 1) \simeq S^1$, $\mathbb{RP}^\infty \simeq K(\mathbb{Z}_2, 1)$ and $\mathbb{CP}^\infty \simeq K(\mathbb{Z}, 2)$. Since $\pi_n(\Omega X) = \pi_{n+1}(X)$, there are isomorphisms $K(G, n) \cong \Omega K(G, n+1)$. By adjunction, we have an Ω-spectrum $(HG)_n = K(G, n)$ known as the Eilenberg-MacLane spectrum of G. This spectrum represents singular cohomology with coefficients in G, $H^n(-; G) \cong [X, K(G, n)]$.

A.3.3 Lie Algebra Cohomology

A ***Lie algebra*** is an R-module $\mathfrak{g}$ equipped with an R-bilinear, antisymmetric bracket $[-, -]$ satisfying the ***Jacobi identity***,

$$[X, [Y, Z]] + [Y, [Z, X]] + [Z, [X, Y]] = 0$$

The Lie algebra generalizes the idea of the commutator of a product, and we correspondingly say that two elements $X, Y \in \mathfrak{g}$ ***commute*** if $[X, Y] = 0$. Every associative R-algebra can be made into a Lie algebra with the ***commutator*** bracket $[X, Y] = XY - YX$; this defines a functor R-$\mathsf{Alg} \to R$-Lie, where R-Lie is the category of R-Lie algebras whose morphisms are ***Lie algebra homomorphisms***, or R-module homomorphisms satisfying $\phi([X, Y]) = [\phi(X), \phi(Y)]$. Given a Lie algebra $\mathfrak{g}$, we can consider the R-module $\mathrm{End}(\mathfrak{g})$ of endomorphisms of $\mathfrak{g}$, which when equipped with composition as a product and the commutator bracket, is a Lie algebra itself. The homomorphism $\mathfrak{g} \to \mathrm{End}(\mathfrak{g})$ sending X to $[X,-]$ is a Lie algebra homomorphism: the Jacobi identity implies that $[[X,Y],-] = [X,-] \circ [Y,-] - [Y,-] \circ [X,-] = [[X,-],[Y,-]]$. The map $\mathfrak{g} \to \mathrm{End}(\mathfrak{g}), X \mapsto [X,-]$ is generally written as ad, with $X \mapsto \mathrm{ad}_X$, and is known as the ***adjoint representation***.

An ***ideal*** of $\mathfrak{g}$ is a subalgebra $\mathfrak{h}$, or submodule closed under the bracket, such that $[X,H] \in \mathfrak{h}$ for all $H \in \mathfrak{h}$. We can quotient $\mathfrak{g}$ by any of its subalgebras, but for the bracket to keep existing on the quotient module $\mathfrak{g}/\mathfrak{h}$, we require that $[X,0] = 0$, which is satisfied when $\mathfrak{h}$ is an ideal.

A $\mathfrak{g}$-***module*** is an R-module M equipped with an R-bilinear action of $\mathfrak{g}$, written $X, f \mapsto Xf$. We do not require that $XYf = YXf$, but instead that $([X,Y])f = X(Yf) - Y(Xf)$. Every Lie algebra $\mathfrak{g}$ is a $\mathfrak{g}$-module in a natural way, since $XYZ - YXZ = [[X,Y],Z] - [[Y,X],Z] = [[X,Y],Z]$. Defining a $\mathfrak{g}$-module homomorphism to be an R-module homomorphism φ such that $\varphi(Xf) = X\varphi(f)$, we have a category $\mathfrak{g}\text{-}\mathsf{Mod}$ of $\mathfrak{g}$-modules.

The ***tensor algebra*** $T(\mathfrak{g})$ of an R-Lie algebra $\mathfrak{g}$ is the R-module

$$T(\mathfrak{g}) := \bigoplus_{i=0}^{\infty} \otimes_R^n \mathfrak{g} = R \oplus \mathfrak{g} \oplus (\mathfrak{g} \otimes_R \mathfrak{g}) \oplus (\mathfrak{g} \otimes_R \mathfrak{g} \otimes_R \mathfrak{g}) \oplus \dots$$

with product given by $\otimes_R$. We can write elements of this algebra as finite sums of the form $r + X + X_1 \otimes_R X_2 + \dots$. In general, this process turns any R-module M into an R-algebra $T(M)$, and does so in a functorial manner (in the obvious way). The functor T is left adjoint to the forgetful functor $R\text{-}\mathsf{Alg} \to R\text{-}\mathsf{Mod}$, and the unit of this adjunction is the obvious inclusion $i : M \to T(M), m \mapsto m$. We interpret this as the "most general" way to turn an arbitrary module into an associative, unital algebra. If we wish to make any identifications of elements on M, we often do so by moving to $T(M)$ and quotienting by a suitable ideal. For instance, given a k-vector space V and a quadratic form $q : V \to k$, we can endow V with a multiplication in which $v \cdot v = q(v)$ by taking the quotient $T(V)/I$, where I is the ideal consisting of all elements of the form $v \otimes v - q(v)$. This yields the ***Clifford algebra*** $C\ell(V,q)$. For instance, if we want to endow $\mathbb{R}$ (as a 1-dimensional $\mathbb{R}$-vector space spanned by e_1) with a multiplication such that $e_1 \cdot e_1 = -1$, we take $T(\mathbb{R})/(e_1 \otimes e_1 + 1)$, which is isomorphic to $\mathbb{C}$.

In the case of Lie algebras, we want to make $\mathfrak{g}$ an associative, unital algebra, while still preserving its bracket. The tensor algebra lets us do even better, turning the abstract bracket into an actual commutator. Specifically, we take $T(\mathfrak{g})$, and quotient it out by the ideal generated by the relation $i([X,Y]) = i(X) \otimes i(Y) - i(Y) \otimes i(X)$. Concretely, we identify the rank 2 tensor $X \otimes Y - Y \otimes X$ with the single element $[X,Y]$. This is known as the ***universal enveloping algebra*** $U\mathfrak{g}$, and satisfies the following universal property: any R-module homomorphism $\varphi : \mathfrak{g} \to M$ such that $\varphi([X,Y]) = XY - YX$ extends to a unique R-algebra map $U\mathfrak{g} \to M$. This is indicative of an adjunction: denoting Lie for the functor $R\text{-}\mathsf{Alg} \to R\text{-}\mathsf{Lie}$ that equips the algebra M with its Lie algebra structure $[m,n] = mn - nm$, we have $U \dashv$ Lie. In particular, we obtain for any module an equivalence $\mathrm{Hom}_{R\text{-}\mathsf{Lie}}(\mathfrak{g}, \mathrm{Lie}(\mathrm{End}(M))) \cong \mathrm{Hom}_{R\text{-}\mathsf{Alg}}(U\mathfrak{g}, \mathrm{End}(M))$. A Lie algebra homomorphism $\mathfrak{g} \to \mathrm{Lie}(\mathrm{End}(M))$ is precisely a $\mathfrak{g}$-module structure on M, and an R-algebra homomorphism $U\mathfrak{g} \to \mathrm{End}(M)$ is precisely a $U\mathfrak{g}$-module structure (where we treat $U\mathfrak{g}$ as a normal ring) on M. This adjunction therefore yields an equivalence of categories between $\mathfrak{g}$-Mod and $U\mathfrak{g}$-Mod. Since $U\mathfrak{g}$-Mod is in particular an R-Mod, this gives us a *lot* of structure on $\mathfrak{g}$-Mod. In particular, we note that $\mathfrak{g}$-Mod has enough injectives, enough projectives, and is an abelian category. When $\mathfrak{g}$ is a free R-module with basis $X_1, \ldots, X_n$, as is often the case, the ***Poincaré-Birkhoff-Witt theorem*** allows us to characterize $U\mathfrak{g}$ as spanned by all elements of the form $\bigotimes_{j=1}^{n} \otimes_R^{n_j} i(X_j)$, where $n_j \geq 0$.

Since $\mathfrak{g}$-Mod is abelian and has enough injectives and projectives, we may set up a Lie algebra (co)homology theory by finding a (left) right exact functor and taking its derived functors. There are two natural functors:

Given a $\mathfrak{g}$-module M, the ***invariant submodule*** $M^{\mathfrak{g}}$ is given by the set of all $m \in M$ for which $Xm = 0$ for all $X \in \mathfrak{g}$. The ***coinvariant submodule*** $M_{\mathfrak{g}}$ is the set of all orbits, $M/\mathfrak{g}M$.

$-^{\mathfrak{g}}$ acts as a functor by restriction: if $m \in M^{\mathfrak{g}}$ and $\varphi : M \to N$, then $X\varphi(m) = \varphi(Xm) = 0$ for all X, so $\varphi(M^{\mathfrak{g}}) \subseteq N^{\mathfrak{g}}$. The action of any $X \in \mathfrak{g}$ on $M_{\mathfrak{g}}$ is trivial: XYm is identified with YXm is identified with $[X,Y]m = XYm - YXm$, so $Xm = 0$ for all X. Any $\mathfrak{g}$-module on which $Xm = 0$ for all X and all m is known as a ***trivial*** $\mathfrak{g}$-module; there's an obvious functor $\iota : R\text{-}\mathsf{Mod} \to \mathfrak{g}\text{-}\mathsf{Mod}$ sending the R-module M to the trivial $\mathfrak{g}$-module M.

$M^{\mathfrak{g}}$ is obviously the largest trivial submodule of the $\mathfrak{g}$-module M, and $M_{\mathfrak{g}}$ is (less obviously) the largest trivial quotient $\mathfrak{g}$-module of M. It follows that there is a triplet of adjunctions $-_{\mathfrak{g}} \dashv \iota \dashv -^{\mathfrak{g}}$, and therefore that $-^{\mathfrak{g}}$ is left exact, whereas $-_{\mathfrak{g}}$ is right exact. We define the ***homology*** of $\mathfrak{g}$ with coefficients in a $\mathfrak{g}$-module M as $H_i(\mathfrak{g}, M) := L_i(-_{\mathfrak{g}})(M)$, and the ***cohomology*** of $\mathfrak{g}$ as

$H^i(\mathfrak{g}, M) := R^i(-^{\mathfrak{g}})(M)$. In $U\mathfrak{g}$-Mod, $M_{\mathfrak{g}}$ is simply $R \otimes_{U\mathfrak{g}} M$, whereas $M^{\mathfrak{g}}$ is $\mathrm{Hom}_{U\mathfrak{g}}(R, M)$. So, we may equivalently define $H_i(\mathfrak{g}, M) = \mathrm{Tor}_i^{U\mathfrak{g}}(R, M)$ and $H^i(\mathfrak{g}, M) = \mathrm{Ext}^i_{U\mathfrak{g}}(R, M)$.

Lie algebra cohomology can also be constructed directly by means of a chain complex known as the ***Chevalley-Eilenberg complex***. Given a free R-Lie algebra, let $\Lambda^n\mathfrak{g}$ denote the nth antisymmetric power of $\mathfrak{g}$, or $\otimes_R^n \mathfrak{g}$ modulo the relation $X \otimes X = 0$ for all $X \in \mathfrak{g}$. Elements of $\Lambda^n\mathfrak{g}$, denoted by $X_1 \wedge \ldots \wedge X_n$, take a negative sign when any two terms are switched, since $X \otimes X = 0 \implies (X+Y) \otimes (X+Y) = X \otimes X + X \otimes Y + Y \otimes X + Y \otimes Y = 0 \implies X \otimes Y = -Y \otimes X$. $\Lambda^n\mathfrak{g}$ is a free R-module, as is $U\mathfrak{g}$, so we may define an $\mathbb{N}$-indexed family of free modules $C^n(\mathfrak{g}, M) := \mathrm{Hom}_R(\Lambda^n\mathfrak{g}, M) = \Lambda^n\mathfrak{g}^* \otimes M$ whose elements are known as n-cochains; we'll set up a cochain complex structure on $\{C^n(\mathfrak{g})\}$ that allows us to directly calculate Lie algebra cohomology.

First, define the ***augmentation map*** $\epsilon : U\mathfrak{g} \to R$ to be the R-algebra homomorphism associated to the zero map $\mathfrak{g} \to \mathrm{Lie}(R)$ by the $U \dashv$ Lie adjunction. This map sends the inclusion of $\mathfrak{g}$ in $U\mathfrak{g}$ to zero, and therefore sends all elements of $U\mathfrak{g}$ that aren't elements of R to 0. The kernel of ϵ, therefore, is the ideal $(i(X_1), \ldots, i(X_n))$ of $U\mathfrak{g}$, where $X_1, \ldots, X_n$ are the generators of $\mathfrak{g}$. This ideal is known as the ***augmentation ideal*** $\mathfrak{I}$.

An n-cochain $f : \Lambda^n\mathfrak{g} \to M$ is sent to an $(n+1)$-cochain by the differential d given by

$$(df)(X_1, \ldots, X_{n+1}) = \sum_{i=1}^{n+1} (-1)^{i+1} X_i f(X_1, \ldots, \widehat{X}_i, \ldots, X_{n+1})$$

$$+ \sum_{j=1}^{n+1} \sum_{i=1}^{j-1} (-1)^{i+j} f([X_i, X_j], X_1, \ldots, \widehat{X}_i, \ldots, \widehat{X}_j, \ldots, X_{n+1})$$

The zeroth differential $d : \mathrm{Hom}_R(R, M) \to \mathrm{Hom}_R(\mathfrak{g}, M)$ sends $f : R \to M$ to $df : \mathfrak{g} \to M$, $X \mapsto Xf(1)$, and the first differential $d : \mathrm{Hom}_R(\mathfrak{g}, M) \to \mathrm{Hom}_R(\mathfrak{f}^2\mathfrak{g}, M)$ sends f to $df(X, Y) = Xf(Y) - Yf(X) - f([X, Y])$. We see that, for $f \in C^0(\mathfrak{g}, M)$, $(d^2 f)(X, Y) = XYf(1) - YXf(1) - [X, Y]f(1) = 0$, and in general $d^2 = 0$, making $C^\bullet(\mathfrak{g}, M) = \ldots \xleftarrow{d} C^2(\mathfrak{g}, M) \xleftarrow{d} C^1(\mathfrak{g}, M) \xleftarrow{d} C^0(\mathfrak{g}, M) \leftarrow 0$ a cochain complex. The cohomology of this complex agrees with the groups $H^i(\mathfrak{g}, M)$ defined above; for instance, the zeroth cohomology group is the set of all f such that $Xf(1) = 0$ for all X, which is equivalent to $M^{\mathfrak{g}}$.

The first cohomology group $H^1(\mathfrak{g}, M)$ is the module of all R-linear maps $D : \mathfrak{g} \to M$ such that $D([X, Y]) = XD(Y) - YD(X)$, known as ***derivations***, modulo the submodule of all ***inner***

derivations satisfying $D([X,Y]) = [X,Y]m$ for some $m \in M$. $H^2(\mathfrak{g}, M)$, meanwhile, is the set of all equivalence classes of ***Lie algebra extensions*** of $\mathfrak{g}$ by M, or short exact sequences $0 \to M \to \mathfrak{h} \to \mathfrak{g} \to 0$.

A.3.4 Hochschild and Cyclic Cohomology

Hochschild Cohomology Given a k-algebra R (which is not necessarily commutative) and an R-bimodule M, we define a chain complex $C_\bullet$ whose nth element is $M \otimes_k R^{\otimes_k n}$ (we will write $\otimes_k$ as $\otimes$ for convenience), and a series of partial differentials $\partial_0, \ldots, \partial_n : M \otimes R^{\otimes n} \to M \otimes R^{\otimes n-1}$ as follows:

$$\partial_0(m \otimes r_1 \otimes \ldots \otimes r_n) = mr_1 \otimes r_2 \otimes \ldots r_n$$

$$\partial_i(m \otimes r_1 \otimes \ldots \otimes r_n) = m \otimes r_1 \otimes \ldots \otimes r_i r_{i+1} \otimes \ldots \otimes r_n, \;\; i \in \{1, \ldots, n-1\}$$

$$\partial_n(m \otimes r_1 \otimes \ldots \otimes r_n) = r_n m \otimes r_1 \otimes \ldots \otimes r_{n-1}$$

We define the differential $d_n : C_n \to C_{n-1}$ as $d_n = \sum_{i=0}^{n} (-1)^i \partial_i$. From the identity $\partial_i \partial_j = \partial_{j-1} \partial_i$ for $i < j$, it is easy to show that $d^2 = 0$, and hence that $(C_\bullet, d_\bullet)$ is a chain complex. The homology groups of the complex associated to the pair (R, M) are known as the ***Hochschild homology*** $HH_n(R, M)$.

The category of R-bimodules is equivalent to the category of (left or right) modules over $R^e = R \otimes R^{\mathrm{op}}$, where R^{op} has the same addition and opposite multiplication of R; this homology is equivalently $\mathrm{Tor}_*^{R^e}(R, M)$. Analogously, we construct $\mathrm{Ext}_{R^e}^*(R, M)$ by replacing $C_n = M \otimes R^{\otimes n}$ with $C^n = \mathrm{Hom}_k(R^{\otimes n}, M)$, or k-multilinear maps $R^n \to M$. The partial differentials $\partial^i : \mathrm{Hom}_k(R^{\otimes n}, M) \to \mathrm{Hom}_k(R^{\otimes(n+1)}, M)$ are given by:

$$(\partial^0 f)(r_1, \ldots, r_{n+1}) = r_1 f(r_2, \ldots, r_{n+1})$$

$$(\partial^i f)(r_1, \ldots, r_{n+1}) = f(r_1, \ldots, r_{i-1} r_i, \ldots, r_{n+1}), \;\; i \in \{1, \ldots, n-1\}$$

$$(\partial^n f)(r_1, \ldots, r_{n+1}) = f(r_1, \ldots, r_n) r_{n+1}$$

We again define $d^n = \sum_{i=0}^{n} \partial^i$, and call the cohomology of the cochain complex $(C^\bullet, d^\bullet)$ the ***Hochschild cohomology*** $HH^n(R, M)$.

As with Lie algebra cohomology, we can concretely characterize the first few (co)homology groups.

$$HH_0(R,M) = M/\mathrm{im}(\partial_0 - \partial_1) = M/\{rm - mr \mid r \in R, m \in M\} = \frac{M}{[R,M]}$$

so the zeroth homology group is given by identifying the left and right R-actions on M, while $HH^0(R,M)$ is given by the kernel of $(\partial^0 - \partial^1) : \mathrm{Hom}_k(R^{\otimes 0} \to M) \cong M \to \mathrm{Hom}_k(R,M)$, so that

$$HH^0(R,M) = \{m \in M \mid rm = mr, \forall r \in R\} = Z(M)$$

First (co)homology The kernel of $d_1 : M \otimes R \to M$ is given by $Z(M)$ as well, though the image of d_2 is given by $\{mr_1 \otimes r_2 - m \otimes r_1 r_2 + r_2 m \otimes r_1 \mid m \in M, r_1, r_2 \in R\}$; hence, the first homology is given by not only identifying rm with mr, but $m \otimes r_1 r_2$ with $r_1 m \otimes r_2 + r_2 m \otimes r_1$ as well.

To make this clearer, we need a definition. Define the multiplication map $R \otimes_k R \to R$ by sending $r_1 \otimes_k r_2$ to $r_1 r_2$ for $r \in R$; this is well-defined by abstract properties of the tensor product, and extends to all of $R \otimes_k R$ by linearity. Let I be the kernel of this map, an ideal of $R \otimes_k R$, and define the R-module of ***Kähler differentials*** $\Omega_{R/k}$ as I/I^2. We equip this module with a map $d : R \to \Omega_{R/k}, r \mapsto 1 \otimes r - r \otimes 1$, which is a derivation in the sense that

$$d(r_1 r_2) = 1 \otimes r_1 r_2 - r_1 r_2 \otimes 1 = r_1 \otimes r_2 - r_1 r_2 \otimes 1 + 1 \otimes r_1 r_2 - r_1 \otimes r_2 = r_1\, d(r_2) + d(r_1) r_2$$

Tensoring an R-module M with the R-module $\Omega_{R/k}$ allows us to identify $m \otimes_R d(r_1 r_2)$ with $r_1 m \otimes_R d(r_2) + m \otimes d(r_1) r_2$; if R is commutative, we can move the latter r_2 to the first argument, evidencing the map $HH_1(R,M) \to M \otimes_R \Omega_{R/k}, m \otimes_k r \mapsto m \otimes_R d(r)$ as an isomorphism.

The kernel of $d^2 : \mathrm{Hom}_k(R,M) \to \mathrm{Hom}_k(R \otimes R, M)$ is given by the set of maps $f : R \to M$ for which $f(r_1 r_2) = r_1 f(r_2) + f(r_1) r_2$, hence a derivation as well. We write the k-vector space of k-derivations as $\mathrm{Der}_k(R,M)$. The image of $d^1 : M \to \mathrm{Hom}_k(R,M)$ is given by those morphisms $R \to M$ of the form $r \mapsto rm - mr$. These are known as the principal derivations $\mathrm{PDer}_k(R,M)$.

Hence,

$$HH^1(R, M) \cong \mathrm{Der}_k(R, M)/\mathrm{PDer}_k(R, M)$$

A.4 K-Theory

A.4.1 Bundles

Principal Bundles A principal bundle is a topological space equipped with an action of a topological group G that splits the space into a set of continuously connected fibers over a base space of orbits. The setup is as follows: a principal bundle over a group G, or a principal G-bundle, consists of a continuous surjection π from the *total space* E to the *base space* B and a continuous left action of G on E, all fibers $\pi^{-1}(b)$ of which are isomorphic. We require that this action restricts to a free, transitive group action on each fiber, and that the action is *locally trivial*: there is an open covering $\{U_\lambda\}_{\lambda\in\Lambda}$ of X such that $\pi^{-1}(U_\lambda)$ is homeomorphic to $U_\lambda \times G$ for each λ. Such an open covering, along with homeomorphisms $U_\lambda \times G \to \pi^{-1}(U_\lambda)$, is known as a ***local trivialization***.

A map of principal bundles $\pi : E \to B$, $\pi' : E' \to B$ is a continuous map $E \to E'$ commuting with the projections π, π', as well as the action of G. This gives us a category $\mathsf{Bun}_G(B)$ of principal G-bundles over B. A map $f : B' \to B$ induces a functor given by pullback (in Top): we send a bundle $\pi : E \to B$ to its pullback along f, giving us a map $f^*E = E \times_B B' \to B'$; this clearly preserves fibers, allowing us to lift the G-action on $E \times_B B'$ from the G-action on E. It is quick to show that the ***pullback bundle*** $f^*E \to B$ is locally trivial, and hence that $f : B' \to B$ defines a functor $f^* : \mathsf{Bun}_G(B) \to \mathsf{Bun}_G(B')$. Furthermore, homotopic maps $f \simeq g : B' \to B$ yield isomorphic elements of $\mathsf{Bun}_G(B')$. Therefore, fixing a bundle $E \to B$ and varying the maps $f : B' \to B$ yields a map from homotopy classes of maps $f : B' \to B$ to isomorphism classes of bundles over B'. We can phrase this as a functor $\underline{\mathsf{Bun}_G} : \mathsf{hTop}^{\mathrm{op}} \to \mathsf{Set}$.

It can be shown by the Brown representability theorem that the functor $\underline{\mathsf{Bun}_G}$, when restricted to CW complexes, is in fact representable: there is a space BG such that isomorphism classes of principal G-bundles over any CW complex X are in bijection with homotopy classes of maps $X \to BG$. BG is known as the ***classifying space*** for the group G.

One (functorial!) way of constructing the corresponding total space, denoted EG, is given by Milnor [Milnor, 1956]: let E_nG be the join of n copies of G, or the space consisting of formal sums $\sum c_i g_i$, where the $c_i \in [0,1]$ satisfy $\sum_i c_i = 1$, and $g_i \in G$ for each i. We can equip this with the continuous G-action $g\sum c_i g_i = \sum c_i g g_i$. Let $n \to \infty$, and define $BG = EG/G$, with the corresponding projection $EG \to BG$ being the bundle we want.

There is a nice way to characterize BG in terms of G: by noting that $EG \to BG$ is a fibration, we have a long exact sequence on homotopy groups [May, 1999] given by

$$\ldots \to \pi_n(G) \to \pi_n(EG) \to \pi_n(BG) \to \pi_{n-1}(G) \to \ldots \to \pi_0(EG) \to 0$$

EG is weakly contractible, i.e. has $\pi_n(EG) = 0$ for all n, and this sequence therefore gives us isomorphisms $\pi_n(G) \cong \pi_{n+1}(BG)$ for all n. Hence, BG is weakly equivalent to the ***delooping*** of G, or the space X such that $\Omega X \simeq G$. The function of the classifying space functor, therefore, is to bump up all of its argument's homotopy groups. This allows us to identify some basic classifying spaces: for instance, the discrete topological group $\mathbb{Z}$ has $\pi_n(\mathbb{Z}) = \mathbb{Z}$ for $n = 0$, and 0 for $n \geq 1$. $B\mathbb{Z}$ is therefore weakly equivalent to $K(\mathbb{Z}, 1)$, i.e. the circle S^1. BS^1, in turn, is a $K(\mathbb{Z}, 2)$, and therefore weakly equivalent to $\mathbb{CP}^\infty$.

Fiber and Vector Bundles The notion of a fiber bundle is given by stripping the group structure from a principal bundle: a fiber bundle is simply a continuous surjection $E \xrightarrow{\pi} B$ of topological spaces such that all fibers $\pi^{-1}(b)$, which we will denote E_b, are isomorphic to a single *typical fiber* F, along with a local trivialization, or an open cover $\{U_\lambda\}_{\lambda\in\Lambda}$ and homeomorphisms $\phi_\lambda : U_\lambda \times F \to \pi^{-1}(U_\lambda)$.

A (k-)***vector bundle*** is a continuous surjection $E \xrightarrow{\pi} B$ in which each fiber has the structure of a k-vector space V (which is not necessarily the same for each fiber) and the local trivialization is compatible with the local trivialization: for all $b \ni U_\lambda$, the map sending $v \in V$ to $\phi(b, v)$ is an isomorphism $V \cong E_b$. A morphism of vector bundles E_1, E_2 with a common base B is a morphism $E_1 \to E_2$ commuting with the projections. k-vector bundles over a fixed space B and their morphisms form a category $\mathsf{VB}_k(B)$. If all fibers of a vector bundle are isomorphic to k^n, we say that the bundle has ***rank*** n. In the case $n = 1$, the bundle is called a ***line bundle***. We shall denote the k-vector bundle on X given by the projection $\pi_2 : k^n \times X \to X$ as T^n_k.

Consider the functor VB^n_k sending a space X to the set $\mathsf{VB}^n_k(X)$ of k-vector bundles of dimension n over some space X. It is known (see, e.g., [Weibel, 2013, Husemoller, 1975]) that if X

is paracompact, then VB_k^n is representable by the infinite Grassmannian $\mathrm{Gr}_n(k)$, which is the space of all n-dimensional subspaces of k^∞. Namely, $\mathsf{VB}_k^n(X) \cong [X, \mathrm{Gr}_n(k)]$. If $\mathrm{Gr}_n(k)$ is also some $K(G,n)$, we can chain isomorphisms to obtain $\mathsf{VB}_k^n(X) \cong H^n(X;G)$.

In the case $n = 1$, for instance, $\mathrm{Gr}_n(k) = k\mathbb{P}^\infty$. This gives us the following isomorphisms:

$$\mathsf{VB}_{\mathbb{R}}^1(X) \cong [X, \mathbb{RP}^\infty] \cong H^1(X;\mathbb{Z}_2)$$

$$\mathsf{VB}_{\mathbb{C}}^1(X) \cong [X, \mathbb{CP}^\infty] \cong H^2(X;\mathbb{Z})$$

So we may send a complex line bundle $E \xrightarrow{\pi} X$ to an element of the second singular cohomology class of X; this element is known as the ***first Chern class*** $c_1(X)$. If $E \xrightarrow{\pi} X$ is a real line bundle, it is represented by an element of the first mod 2 cohomology class of X, known as the ***first Stiefel-Whitney class*** $w_1(X)$.

Appendix B

Some Physics

B.1 Functional Analysis

B.1.1 Banach Spaces

In the theory of finite dimensional vector spaces, everything goes right. More specifically, every such space V satisfies the following:

- The double dual of V, V^{**}, is canonically isomorphic to V itself.
- All norms on V are equivalent, and induce the same topology.
- With this topology, any linear map from V is continuous.
- An endomorphism on V is injective iff it is surjective.
- The unit ball in V (under any norm) is compact.

The theory of infinite dimensional vector spaces, however, is far more dangerous: *none* of these statements hold, nor can they be easily fixed. In such an infinite dimensional vector space W, the following properties are satisfied:

- As cardinals, $\dim W^{**} > \dim W^{*} > \dim W$, these inequalities being strict.
- W generally has many different topologies of interest.
- Linear maps from W aren't necessarily continuous.
- There are non-surjective injections $W \to W$.
- The unit ball is never compact.

In nature, infinite dimensional vector spaces tend to occur as spaces of functions, hence the name functional analysis. There is a hierarchy of classes of infinite-dimensional vector spaces, with each level of the hierarchy introducing a new structure, or a new condition to be fulfilled by a structure provided at the lower tier. At the bottom rung are simply k-vector spaces, where we assume k is either $\mathbb{R}$ or $\mathbb{C}$.

Normed Spaces The first thing we can do with a vector space V is put a ***norm*** on it. This is a function $||\cdot|| : V \to [0,\infty)$ which satisfies the following properties:

1. *Homogeneity*: $||cv|| = |c|\,||v||$, for $c \in k$.
2. *Triangle inequality*: $||v+w|| \leq ||v|| + ||w||$
3. *Definiteness*: $||v|| = 0$ iff $v = \vec{0}$.

Equipped with such a norm, V becomes a ***normed space***. This norm induces a topology on V whose basis consists of open sets

$$B_r(v) = \{w \in V \mid ||v-w|| < r\}$$

for all $r \in [0,\infty)$ and all $v \in V$. Given two normed vector spaces V, W, we may ask which linear maps $A : V \to W$, also known as ***operators***, preserve the norm, in the sense that $||Av||_W \leq c||v||_V$ for all $v \in V$, for some fixed $c \geq 0$. Such an operator is known as a ***bounded operator***. It's well known that an operator is bounded if and only if it is continuous: in this sense, the structure on V that a norm provides is equivalent to the structure that the topology induced by the norm itself provides. The smallest such c satisfying $||Av||_W \leq c||v||_V$ is given by $\sup_{||f||_V \leq 1} ||Av||_W$, and is known as the ***operator norm*** $||A||$. With this norm, the space $\mathcal{B}(V,W)$ of bounded operators $V \to W$, with its natural vector space structure, becomes a normed space itself.

An important family of normed spaces can be constructed as follows: take a measure space $(\Omega, \mathcal{F}, \mu)$ and consider the vector space of measurable functions $\Omega \to k$, $k \in \{\mathbb{R}, \mathbb{C}\}$. Define the p-norm of a function f to be

$$||f||_p := \left(\int_\Omega |f|^p\right)^{1/p}$$

for $1 \leq p < \infty$. The space of functions f for which $||f||_p < \infty$ is a vector space $\mathcal{L}^p(\Omega, \mu)$, but it isn't a normed space, since functions which are 0 almost everywhere have norm zero. The set of all such functions forms a linear subspace of $\mathcal{L}^p(\Omega, \mu)$, though, and quotienting out by it yields a proper normed space $L^p(\Omega, \mu)$, known as an L^p ***space***, whose elements aren't strictly measurable functions $\Omega \to k$, but equivalence classes of measurable functions which differ by sets of measure zero [Rudin, 1973]. As $p \to \infty$, $||f||_p$ converges to the essential supremum of $|f|$, since raising $|f|$ to a power $p > 1$ makes a greater change when $|f|$ is large, with the size of p exaggerating this change. This allows us to define $||f||_\infty$ to be the essential supremum of $|f|$ over Ω, and thereby obtain the space $L^\infty(\Omega, \mu)$.

In the special case when μ is the counting measure, which sends a finite $S \subseteq \Omega$ to $|S|$ and an infinite S to ∞, the set $L^p(\mathbb{N}, \mu)$ is known as the ℓ^p ***space***; its elements are sequences $\{c_0, c_1, \ldots\}$ and the norm of a sequence $c = \{c_n\}_{n\in\mathbb{N}}$ is just $\left(\sum_{n=0}^{\infty} |c_n|^p\right)^{1/p}$ when $1 \leq p < \infty$, and $\sup c$ when $p = \infty$.

Inner Product Spaces Given a (k-)vector space V, an ***inner product*** on V is a mapping $\langle \cdot, \cdot \rangle : V \times V \to k$ which is

1. *Conjugate-symmetric*: $\langle v, w \rangle = \overline{\langle w, v \rangle}$
2. *Positive definite*: $\langle v, v \rangle \geq 0$, and $\langle v, v \rangle = 0$ iff $v = \vec{0}$.
3. *Sesquilinear*: Linear in the first argument, and conjugate linear in the second argument.

A vector space equipped with an inner product is known as a ***inner product space***. The norm ***induced*** by the inner product $\langle \cdot, \cdot \rangle$ is given by $||v|| = \sqrt{\langle v, v \rangle}$; it is straightforward to check that this is indeed a norm, and therefore that inner product spaces are a subset of normed spaces. This norm satisfies the ***polarization identity***

$$||f+g||^2 - ||f-g||^2 = 4\,\mathrm{Re}\left(\langle f, g \rangle\right)$$

as well as the ***parallelogram law***

$$||f+g||^2 + ||f-g||^2 = 2\left(||f||^2 + ||g||^2\right)$$

In fact, an arbitrary norm on a vector space is induced by an inner product if and only if it satisfies the parallelogram law [Haase, 2014].

Example. As in the finite dimensional case, two vectors v, w on an inner product space are ***orthogonal*** if $\langle v, w\rangle = 0$. For instance, consider the k-vector space of continuous functions $[0,1] \to k = \mathbb{C}$, with inner product $\langle f, g\rangle = \int_0^1 f\overline{g}\,dx$. For $f_n = e^{2\pi inx}$, $n \in \mathbb{Z}$, we have

$$\langle f_m, f_n\rangle = \int_0^1 e^{2\pi i(m-n)x}\,dx$$

which when $m = n$ is 1 and when $m \neq n$ is $\frac{1}{2\pi i(m-n)}\left(e^{2\pi i(m-n)} - 1\right) = 0$. So, in fact, $\{f_n\}$ is not only a set of pairwise orthogonal vectors, but an orthonormal set. It is not an orthonormal basis, since an arbitrary $f \in C[0,1]$ cannot be expressed as a finite linear combination of the f_n, but (since this is just a Fourier transform) we know that we can specify coefficients $c_n = \langle f, f_n\rangle$ such that the sum $\sum_{i\in\mathbb{Z}} c_n f_n$ *converges* to f under the norm induced by the inner product. Such a "basis" in which every element of the vector space can be expressed as the limit of a countable sum is known as a ***Schauder basis***.

Banach Spaces A ***Banach space*** is a normed space $(V, ||\cdot||)$ which is complete with respect to its norm, having for each Cauchy sequence $\{v_n\}_{n\in\mathbb{N}}$ a vector v such that $\lim_{n\to\infty} ||v_n - v|| = 0$. This completeness condition ensures that V has "no holes", so that all sequences that should converge (Cauchy sequences) do converge. An incomplete normed space V can be made complete in the following manner: take the set of all Cauchy sequences $\{v_n\}_{n\in\mathbb{N}}$ in V, and, given $v = \{v_n\}, w = \{w_n\}$, define a "metric" on Cauchy sequences by $D(v, w) = \lim_{n\to\infty} ||v_n - w_n||$. If V isn't already complete, this isn't an actual metric: let v and w be the same sequence except at the first element to get $D(v, w) = 0$ with $v \neq w$. To fix this, we declare v and w to be equivalent to be equal if $\lim_{n\to\infty} ||v_n - w_n|| = 0$. This is an equivalence relation by the triangle inequality, and quotienting the set of Cauchy sequences out by it makes D a proper metric on what is now a complete space, which we denote by $\widehat{V}$. Of course, if V is already complete, we can identify Cauchy sequences with the vector they converge to, so $\widehat{V}$ can be identified with V. If not, then V naturally *embeds* into V, this embedding being given by sending a $v \in V$ to the equivalence class of the Cauchy sequence $(v, v, v, \ldots)$. In this way, every normed vector space V naturally embeds into the Banach space $\widehat{V}$ known as the ***completion*** of V.

$L^p(\Omega, \mu)$ is always a Banach space, a fact often known as the Riesz-Fischer theorem. For V an arbitrary Banach space, the set $\mathcal{B}(V) := \mathcal{B}(V, V)$ of bounded operators on V is, when equipped with the operator norm, a Banach space. This space can be equipped with an associative multiplication given by composition. A Banach space equipped with an associative algebra structure is known as a ***Banach algebra***; we also require that $||AB|| \leq ||A||\,||B||$, but this holds trivially for the Banach algebra $\mathcal{B}(V)$. In addition, the normed vector space $V^* := \mathcal{B}(V, k)$ is also a Banach space, known as the ***dual*** of V.

Example. Banach spaces often appear in the study of differential equations and dynamical systems, since they allow us to use linear algebra in sufficiently nice topological spaces. For instance, let X be a Banach space, and f a continuous map $X \to X$. f is called a ***contracting*** map if there's a $\lambda < 1$ s.t. $d(f(x), f(y)) \leq \lambda d(x, y)$, where $d(x, y) := ||x - y||$. f and its positive iterates $f^2, f^3, \ldots$ form what is known as a discrete-time topological dynamical system. Of course, $d(f^n(x), f^n(y)) \to 0$ as $n \to \infty$; every $\{f^n(x)\}_{x \in \mathbb{N}}$ is, in fact, a Cauchy sequence, so, given that X is complete by virtue of being a Banach space, there's a unique limit p to which all points converge, known as the ***fixed point***.

We verify that it's a Cauchy sequence as follows: for $n \geq m$,

$$d(f^m(x), f^n(x)) \leq d(f^m(x), f^{m+1}(x)) + \ldots + d(f^{n-1}(x), f^n(x))$$

$$\leq (\lambda^m + \lambda^{m+1} + \ldots + \lambda^{n-1}) d(f(x), x) \leq \lambda^m (1 + \lambda + \lambda^2 + \ldots) d(f(x), x) = \frac{\lambda^m}{1 - \lambda} d(f(x), x) \overset{m \to \infty}{\longrightarrow} 0$$

In particular, as $m \to \infty$, $d(p, f^n(x)) \to 0$, implying that $f(p) = p$. As $n \to \infty$, we get $d(f^m(x), p) \leq \frac{\lambda^n}{1-\lambda} d(f(x), x)$. We say that two sequences of points $\{x_n\}_{n \in \mathbb{N}}$ and $\{y_n\}_{n \in \mathbb{N}}$ ***converge exponentially*** to each other if $d(x_n, y_n) < c\lambda^n$ for some $c > 0$, $\lambda < 1$. In the case that $\{y_n\}_{n \in \mathbb{N}}$ is a constant sequence $y_n = y$, we just say that $\{x_n\}_{n \in \mathbb{N}}$ converges exponentially to y. We therefore have the ***Contraction Mapping Principle***: under the action of iterates of a contracting map f on a complete metric space X, all points converge with exponential speed to the unique fixed point of f.

B.1.2 Hilbert Spaces

Hilbert spaces combine the theories of inner product and Banach spaces. In particular, a ***Hilbert space*** is an inner product space $\mathcal{H}$ which is Banach with respect to the norm induced by its inner product or, equivalently, a Banach space whose norm satisfies the parallelogram law. Among the L^p spaces, this is only satisfied for $p = 2$, in which case the norm on $L^2(\Omega, \mu)$ is induced by the inner product $\langle f, g\rangle = \int_\Omega f\overline{g}\,d\mu$. When a Hilbert space $\mathcal{H}$ has an orthonormal Schauder (countable) basis, it's called ***separable***. This is essentially a size restriction on $\mathcal{H}$; every Hilbert space has a possibly uncountable orthonormal basis (assuming the AC), but we will assume that our Hilbert spaces are separable to avoid size issues. We'll also assume that $k = \mathbb{C}$ unless otherwise specified.

C^*-Algebras For the purposes of quantum mechanics, we're not interested in Hilbert spaces per se, but in algebras of operators on Hilbert spaces. A Banach algebra of the form $\mathcal{B}(\mathcal{H})$ has a natural involution operation given by taking adjoints: the ***adjoint*** of an operator $A \in \mathcal{B}(\mathcal{H})$ is an operator $A^\dagger$ satisfying $\langle Av, w\rangle = \langle v, A^\dagger w\rangle$ for all $v, w \in \mathcal{H}$. (In the real or complex finite dimensional case, this simply corresponds to taking the transpose or conjugate transpose, respectively). The fact that adjoints always exist is a consequence of the ***Riesz representation theorem***, which states that any $\varphi \in \mathcal{H}^*$ can be represented as $\langle -, v\rangle$ for some $v \in \mathcal{H}$; if we set $\varphi = \langle A-, w\rangle$ for a fixed $w \in \mathcal{H}$, this theorem gives us a v such that $\varphi = \langle -, v\rangle$, and therefore an identification $\langle Ax, w\rangle = \langle x, v\rangle$. This v depends linearly and continuously on w, and hence can be represented as $A^\dagger w$, giving us the adjoint $A^\dagger$. We can check that this really does define an involution on $\mathcal{B}(\mathcal{H})$: $\langle A^{\dagger\dagger}v, w\rangle = \overline{\langle w, A^{\dagger\dagger}v\rangle} = \overline{\langle A^\dagger w, v\rangle} = \langle v, A^\dagger w\rangle = \langle Av, w\rangle$, so $A^{\dagger\dagger} = A$. Furthermore, $\langle ABv, w\rangle = \langle Bv, A^\dagger w\rangle = \langle v, B^\dagger A^\dagger w\rangle$, so $(AB)^\dagger = B^\dagger A^\dagger$. It can also be verified that $||A^\dagger A|| = ||A^\dagger||\,||A||$, and this property, along with the previous two, makes $\mathcal{B}(\mathcal{H})$ a ***C^*-algebra*** when equipped with the involution $(\cdot)^\dagger$. In general, a C^*-algebra is a Banach algebra with an involution satisfying $(AB)^\dagger = B^\dagger A^\dagger$ and $||A^\dagger A|| = ||A^\dagger||\,||A||$; the ***Gelfand-Naimark theorem*** allows us to identify any C^*-algebra as a subalgebra of some $\mathcal{B}(\mathcal{H})$.

Observables and Projections Three especially important subsets of $\mathcal{B}(\mathcal{H})$ must be distinguished: first are the ***self-adjoint*** operators, which satisfy $A^\dagger = A$. (Physicists often call a self-adjoint operator a ***Hermitian operator***, or an ***observable***). For $\mathcal{H} = \mathbb{R}^n$, these are the symmetric matrices $A = A^T$, and for $\mathcal{H} = \mathbb{C}^n$, these are the conjugate symmetric/Hermitian matrices $A = A^H$. The eigenvalues of a self-adjoint operator, i.e. those $\lambda \in \mathbb{C}$ such that $Av = \lambda v$ for some v known as λ's eigenvector, can easily be seen to be real even if $\mathcal{H}$ is complex:

$$\lambda||v||^2 = \langle \lambda v, v\rangle = \langle Av, v\rangle = \langle v, Av\rangle = \langle v, \lambda v\rangle = \overline{\lambda}||v||^2$$

In addition, the eigenvectors v_1, v_2 of a self-adjoint A are orthogonal given that they have different eigenvalues $\lambda_1 \neq \lambda_2$:

$$\lambda_1\langle v_1, v_2\rangle = \langle Av_1, v_2\rangle = \langle v_1, Av_2\rangle = \lambda_2\langle v_1, v_2\rangle$$

so $(\lambda_1 - \lambda_2)\langle v_1, v_2\rangle = 0$, implying that $\langle v_1, v_2\rangle = 0$. We denote the set of all self-adjoint operators as $\mathcal{O}(\mathcal{H}) \subset \mathcal{B}(\mathcal{H})$; it isn't closed as an algebra, since $(AB)^\dagger = B^\dagger A^\dagger = BA$ isn't necessarily equal to AB, but it is closed under the ***commutator*** $i[A,B] = i(AB - BA)$, with $(i[A,B])^\dagger = (-i)(B^\dagger A^\dagger - A^\dagger B^\dagger) = i[A,B]$.

The second subset of $\mathcal{B}(\mathcal{H})$ consists of the ***positive operators***, for which $\langle v, Av\rangle$ is real and non-negative for all $v \in \mathcal{H}$. Obviously, $\langle v, Av\rangle = \overline{\langle Av, v\rangle}$, suggesting that positive operators are self-adjoint. Given two self-adjoint operators A_1, A_2, we write $A_1 \geq A_2$ if $A_1 - A_2$ is positive; this forms a partial order on $\mathcal{O}(\mathcal{H})$.

Finally, there are the ***projection operators***, those operators $P \in \mathcal{B}(\mathcal{H})$ which satisfy $P^2 = P$. These operators are necessarily self-adjoint and positive, satisfying $I \geq P \geq 0$, where I is the identity operator $Iv = v$ and 0 is the zero operator $0v = \vec{0}$. In fact, projections can be characterized as orthogonal projections onto some linear subspace of $\mathcal{H}$. For instance, every $v \in \mathcal{H}$ induces a projection operator $P_v w = v\langle w, v\rangle / \langle v, v\rangle$. Given an operator A, define its ***range*** to be $R(A) = A\mathcal{H}$ and its ***null space*** to be $N(A) = \{v \in \mathcal{H} \mid Av = \vec{0}\}$. Given a family $\{P_\alpha\}$ of projections, we can then define the ***meet*** $\wedge_\alpha P_\alpha$ to be the smallest closed subspace of $\mathcal{H}$ containing $\bigcap_\alpha R(P_\alpha)$, and the ***join*** $\vee_\alpha P_\alpha$ to be the smallest closed subspace containing $\bigcup_\alpha R(P_\alpha)$. Denoting by $\mathcal{P}(\mathcal{H})$ the subset of $\mathcal{O}(\mathcal{H})$ containing the projections, these are the inf and sup operations with respect to the partial order on $\mathcal{P}(\mathcal{H})$ inherited from $\mathcal{O}(\mathcal{H})$.

Diagonalizability Since we've assumed $\mathcal{H}$ to be separable, we can fix a countable orthonormal basis $(e_1, e_2, \ldots)$, and represent any $v \in \mathcal{H}$ as the converging sum $\sum_{i=1}^{\infty} v_i e_i$, where $v_i = \langle v_i, e_i\rangle$. This allows us to write $\langle v, w\rangle = \langle \sum_i v_i e_i, \sum_j w_j e_j\rangle = \sum_i v_i w_i$, and to express an operator A in terms of a "matrix" $A_{ij} = \langle e_i, Ae_j \rangle$. With this notation, the usual formulas for finite-dimensional vector spaces can be extended: $(Av)_i = \sum_{ij} A_{ij} v_j$, $(AB)_{ij} = \sum_k A_{ik} B_{kj}$, and so on. In particular, A is ***diagonal*** when $A_{ij} = 0$ for $i \neq j$, and ***diagonalizable*** when there is an orthonormal countable basis in which A is diagonal. Two self-adjoint operators A, B are ***mutually diagonalizable*** when there is a single basis in which they're both diagonal; this happens when $[A, B] = 0$.

The ***trace*** of an operator A is given by $\operatorname{Tr} A = \sum_i A_{ii} = \sum_i \langle e_i, Ae_i \rangle$; this value is independent of the basis chosen, being a property of the operator A itself. This sum may not always converge, but when it does, A is said to be of ***trace class***. For instance, we can take the trace of a projection operator of the form P_v:

$$\operatorname{Tr} P_v = \sum_i \langle e_i, P_v e_i \rangle = \sum_i \left\langle e_i, v \frac{\langle e_i, v \rangle}{\langle v, v \rangle} \right\rangle = \frac{1}{\langle v, v \rangle} \sum_i |\langle e_i, v \rangle|^2 = \frac{1}{\sum_i |v_i|^2} \sum_i |v_i|^2 = 1$$

On the set of trace class operators in $\mathcal{B}(\mathcal{H})$, denoted $\mathcal{T}(\mathcal{H})$, the trace generates a norm: take an operator A and define the self-adjoint operator $A^\dagger A$, which has real eigenvalues $\sigma_1, \sigma_2, \ldots$. The ***trace norm*** of A, denoted variously as $||A||_*$ or $\operatorname{Tr} \sqrt{A^\dagger A}$, is then $\sum_i \sqrt{\sigma_i}$. With this norm, $\mathcal{T}(\mathcal{H})$ is a Banach space; in fact, its dual can be identified with $\mathcal{B}(\mathcal{H})$ itself. We say that $\mathcal{T}(\mathcal{H})$ is the ***predual*** of $\mathcal{B}(\mathcal{H})$, and write $\mathcal{T}(\mathcal{H}) = \mathcal{B}(\mathcal{H})_*$.

Any operator $\rho \in \mathcal{T}(\mathcal{H})$ with trace 1 is said to be a ***state***; the projection operators P_v are special among these, and are called ***pure states***. It is a consequence of the Hilbert-Schmidt theorem that an arbitrary state A can be decomposed into a sum of finitely many pure states as $A = \sum_{i=1}^N c_i P_{v_i}$, where the $\{v_i\}$ are orthonormal and $\sum_{i=1}^N c_i = 1$.

Bra-Ket Notation Let $\mathcal{H}$ be a complex Hilbert space. Dirac's ***bra-ket notation*** prescribes that we write an element ψ of $\mathcal{H}$ as $|\psi\rangle$, calling them ***kets***, and elements ϕ of $\mathcal{H}^*$ as $\langle \phi | := \langle \phi, - \rangle$, calling them ***bras***. Note that physicists tend to write the inner product as being conjugate linear in the *first* argument, rather than the second, which is why we've used $\langle \phi, - \rangle$ instead of the $\langle -, \phi \rangle$ above. We'll continue to use this convention for this section. The inner product of two kets $|\phi\rangle$, $|\psi\rangle$ is written as $\langle \phi | \psi \rangle$. The correspondence between $\mathcal{H}$ and $\mathcal{H}^*$ given by the Riesz representation theorem sends a $c|\psi\rangle$ to $\bar{c}\langle \phi|$, and a term of the form $A|\psi\rangle$ to $\langle \phi | A^\dagger$, where we define the action of an operator A on a bra $\langle\phi|$ as

$$(\langle\phi|A)\,|\psi\rangle = \langle\phi|\,(A|\psi\rangle)$$

We generally require our bras and kets to be ***normalized***, requiring that $\langle\psi|\psi\rangle = 1$; an arbitrary element of $\mathcal{H}$ can be normalized by dividing it by its norm. In contrast to the inner product, bra-ket notation allows us to express the ***outer product*** of a bra $\langle\phi|$ with a ket $|\psi\rangle$, which is the operator $|\psi\rangle\langle\phi|$ that acts on a $|\xi\rangle$ as $(|\psi\rangle\langle\phi|)(|\xi\rangle) = |\psi\rangle\langle\phi|\xi\rangle$. We may also speak of the outer product of two bras $\langle\phi_1|\langle\phi_2|$ or kets $|\psi_1\rangle|\psi_2\rangle$, which is just defined to be the tensor product in $\mathcal{H}\otimes\mathcal{H}$. We'll rewrite a few of our above formulas in this notation: $A_{ij} = \langle e_i|A|e_j\rangle$, $|v\rangle = \sum_i |e_i\rangle\langle e_i|v\rangle$, $\operatorname{Tr} A = \sum_i \langle e_i|A|e_i\rangle$, and $P_v = |v\rangle\langle v|$ (note that $|v\rangle$ is assumed to be normalized). Note that since $v = \sum_i |e_i\rangle\langle e_i|v\rangle = (\sum_i |e_i\rangle\langle e_i|)\,v$, we can write $\sum_i |e_i\rangle\langle e_i| = I$. This is known as a ***resolution of the unity***, and can be inserted anywhere: for instance, $\langle v|w\rangle = \sum_i \langle v|e_i\rangle\langle e_i|w\rangle$. Commonly, $\mathcal{H} = L^2(M,\mathbb{C})$, where M is a Riemannian manifold with metric g and the inner product is $\langle\psi|\phi\rangle = \int_M \overline{\psi}(x)\phi(x)\,\omega$, where ω is the volume form on M.

Resolution of the identity works when we replace the $\{e_i\}$ with an arbitrary orthonormal basis, for instance the eigen*kets* of a self-adjoint operator A, when they form a complete set. In physical systems, we often use the case of $A = H$, where H is an operator representing the Hamiltonian, whose eigenvalues are thought of as the allowed energy levels of the system. The eigenvalue equation $H|\psi\rangle = E|\psi\rangle$ is known as the ***time-independent Schrödinger equation***. The Hamiltonian H generates a one-parameter group of operators $U_t := e^{-\frac{i}{\hbar}tH}$, where $e^A := I + A + \frac{1}{2}A^2 + \ldots$ satisfies the usual properties of the exponential, and $\hbar$ is a positive constant. We have $U_tU_s = e^{-\frac{i}{\hbar}(t+s)H} = U_{t+s}$ and $U_t^\dagger = e^{\frac{i}{\hbar}tH} = U_{-t}$, so that $U_tU_t^\dagger = U_t^\dagger U_t = I$; operators whose adjoints are their inverses are known as ***unitary***, and the group $\{U_t\}_{t\in\mathbb{R}}$ is known as the ***unitary group*** generated by H. For instance, the unitary group generated by $i\hbar\frac{d}{dx}$ on $L^2(\mathbb{R})$ is given by

$$U_x f(x') = e^{-\frac{i}{\hbar}xi\hbar\frac{d}{dt}} f(x') = e^{x\frac{d}{dx}} f(x') = \left(f + xf' + \frac{x^2}{2}f'' + \ldots\right)(x') = f(x'-x)$$

where we've identified the penultimate step as a Taylor expansion. The operator $i\hbar\frac{d}{dx}$ is the quantum analog of momentum, and we correspondingly say that momentum is the generator of translation.

B.1.3 Von Neumann Algebras

A C^*-algebra $\mathcal{M}$ given by the subset of some $\mathcal{B}(\mathcal{H})$ is a ***von Neumann algebra*** if it has a predual. Defining the ***commutant*** of an arbitrary unital C^*-subalgebra $\mathcal{M} \subseteq \mathcal{B}(\mathcal{H})$ to be

$$\mathcal{M}' := \{A \in \mathcal{B}(\mathcal{H}) \mid AB = BA \text{ for all } B \in \mathcal{M}\}$$

von Neumann's ***double commutant theorem*** states that $\mathcal{M}$ is a von Neumann algebra if and only if $\mathcal{M} = \mathcal{M}''$. Note that if a von Neumann algebra $\mathcal{M}$ is contained in its commutator $\mathcal{M}'$, it must be abelian; if it is in fact equal to its commutator, we call it ***maximally abelian***. On the other hand, a commutator might be called "maximally noncommutative" if $\mathcal{M}$ and $\mathcal{M}'$ are as disjoint as possible, having only in common scalar multiples of the identity. Such a von Neumann algebra for which $\mathcal{M} \cap \mathcal{M}' = \{zI \mid z \in \mathbb{C}\}$ is known as a ***factor***. The most obvious example is $\mathcal{B}(\mathcal{H})$ itself.

Any von Neumann algebra can be reconstructed from its set of projections $\mathcal{P}(\mathcal{M})$, as $\mathcal{M} = \mathcal{P}(\mathcal{M})''$. In this way, we can study $\mathcal{M}$ simply by studying its projections which, as noted previously, form a lattice with meets and joins. We may put an equivalence relation on $\mathcal{P}(\mathcal{M})$, whereby $A \sim B$ if there's an $X \in \mathcal{M}$ satisfying $X^\dagger X = A$ and $XX^\dagger = B$. This generates a partial ordering on $\mathcal{P}(\mathcal{M})$, whereby $A \preceq B$ if there is some A' with $R(A') \subseteq R(B')$ and $A \sim A'$. We can "approximate" arbitrary operators $P \in \mathcal{P}(\mathcal{H})$ from the perspective of an arbitrary von Neumann algebra $\mathcal{M}$ by taking its ***outer*** $\mathcal{M}$***-support***, or the smallest operator in $\mathcal{M}$ greater than or equal to P:

$$\delta^o(P)_{\mathcal{M}} = \bigwedge\{Q \in \mathcal{P}(\mathcal{M}) \mid Q \succeq P\}$$

We may also take its ***inner*** $\mathcal{M}$***-support***, or the largest operator in $\mathcal{M}$ less than or equal to P:

$$\delta^i(P)_{\mathcal{M}} = \bigvee\{Q \in \mathcal{P}(\mathcal{M}) \mid Q \preceq P\}$$

Every factor $\mathcal{M}$ admits a function $d : \mathcal{P}(\mathcal{M}) \to [0, \infty]$, known as the ***dimension function***, which satisfies the following properties [Rédei and Summers, 2007]:

1. $d(A) = 0$ iff $A = 0$.
2. $d(A) < \infty$ iff $B \preceq A$ and $B \sim A$ imply $B = A$.
3. $d(A) \leq d(B)$ iff $A \preceq B$.

4. $d(A) = d(B)$ iff $A \sim B$.
5. $d(A + B) = d(A \wedge B) + d(A \vee B)$.
6. $d(A + B) = d(A) + d(B)$ if $A \perp B$.

Up to multiplication by a constant, there is exactly one function satisfying these properties; we can normalize it as we wish, but we cannot change what its range looks like. The nature of the ranges of the dimension functions of factors allows us to classify them into several different types. We list the classifications corresponding to different (normalized) values of range(d):

Discrete

- $\{0, 1, \ldots, n-1, n\}$: I_n
- $\{0, 1, \ldots, \infty\}$: I_∞

Continuous

- $[0, 1]$: II_1
- $[0, \infty]$: II_∞

Purely Infinite

- $\{0, \infty\}$: III

More concretely, we can characterize these in the following way: $\mathcal{M}$ is a type I factor if there is a non-zero projection A such that there is no other projection B with $B \leq A$ and $B \sim A$ but $B \neq A$. Such a projection is known as a ***minimal projection***. Any type I factor is isomorphic to some $\mathcal{B}(\mathcal{H})$, which is a type I_n factor if $\dim \mathcal{H} = n < \infty$ and type I_∞ otherwise. $\mathcal{M}$ is type II if there are no minimal projections, but there are projections A such that $B \leq A$ and $B \sim A$ imply $B = A$ (known as ***finite*** projections); if there's an infinite projection, $\mathcal{M}$ is type II_∞, else II_1. If $\mathcal{M}$ is neither type I or type II, i.e. contains *no* finite projections, then it is type III. There is a further refinement, due to A. Connes, of type III factors into type III_λ, $\lambda \in [0, 1]$, though we will not discuss it.

In a von Neumann algebra $\mathcal{M}$, there's a natural embedding $\mathcal{M}_* \to \mathcal{M}^*$ given by taking a $\phi \in \mathcal{M}_*$ and defining its action on $\mathcal{M}$ as $\phi(A) = A(\Phi)$. If a $\phi \in \mathcal{M}^*$ can be obtained in this way, and it is a state, it is known as a ***normal state***. Normal states can additionally be characterized by the following continuity property: for any countable family $\{P_n\}$ of mutually orthogonal projections in $\mathcal{M}$, $\phi\left(\bigvee P_n\right) = \sum \phi(P_n)$. On the von Neumann algebra $\mathcal{B}(\mathcal{H})$, every normal state ϕ acts on operators A as $\phi(A) = \mathrm{Tr}(\Phi A)$ for some unique state $\Phi \in \mathcal{T}(\mathcal{H}) \subset \mathcal{B}(\mathcal{H})$; in this context, Φ is known as the ***density operator*** corresponding to ϕ.

Gelfand Representations Given an abelian von Neumann algebra $\mathcal{M}$, denote by $\Sigma_{\mathcal{M}}$ the set of $\mathbb{C}$-algebra homomorphisms $\lambda : \mathcal{M} \to \mathbb{C}$ such that $\lambda(I) = 1$, known as its ***Gelfand spectrum***. With the weak-* topology, $\Sigma_{\mathcal{M}}$ is a compact Hausdorff space. The ***Gelfand representation theorem*** states that $\mathcal{M}$ is isomorphic as a C^*-algebra to the C^*-algebra of continuous complex functions on $\Sigma_{\mathcal{M}}$; this construction, which is functorial, is in fact half of a contravariant equivalence between the categories of unital C^*-algebras and compact Hausdorff spaces. The isomorphism sends an operator $A \in \mathcal{M}$ to a continuous function $\overline{A} : \Sigma_{\mathcal{M}} \to \mathbb{C}$, $\overline{A}(\lambda) = \lambda(A)$, known as its ***Gelfand transform***; if $A = A^\dagger$, then $\overline{A} = \overline{A}^\dagger$, implying that self-adjoint operators are transformed into real functions.

Of particular interest is the image of projections $P \in \mathcal{M}$ under the Gelfand transform, $\overline{P}(\lambda) = \lambda(P)$. Since $\lambda(P)^2 = \lambda(P^2) = \lambda(P)$ for any $\lambda \in \Sigma_{\mathcal{M}}$, the range of $\overline{P}$ must be $\{0, 1\}$. The function λ judges a projection P either *true* or *false*, and the transformed projection $\overline{P}$ judges a function λ as λ judges $\overline{P}$. We denote by S_P the set of $\lambda \in \Sigma_{\mathcal{M}}$ on which $\overline{P}$ is 1; since $\overline{P}$ is continuous, S_P is closed, being $\overline{P}^{-1}(\{1\})$, and open, being the complement of $\overline{P}^{-1}(\{0\})$, making it a clopen subset of $\Sigma_{\mathcal{M}}$.

B.1.4 Quantum Probability Theory

The structure of classical probability theory is as follows: fix a set X and a σ-algebra $\mathcal{F}$ on X. A function $\mu : \mathcal{F} \to \mathbb{R}$ such that $\mu(\varnothing) = 0, \mu(X) = 1$, and $\mu\left(\cup_{i=1}^{\infty} A_i\right) = \sum_{i=1}^{\infty} \mu(A_i)$ for pairwise disjoint A_i is known as a ***probability measure*** on the measurable space $(X, \mathcal{F})$, and the measure space $(X, \mathcal{F}, \mu)$ is known as a ***probability space***. The sets $F \in \mathcal{F}$ are known as ***events***, $\mu(F)$ as the ***probability*** of F, $\mathcal{F}$ as the ***event space***, and X as the ***sample space***. A function $f : X = (X, \mathcal{F}, \mu) \to \mathbb{R}^n$, where $\mathbb{R}^n$ is given its Borel σ-algebra, is ***$\mathcal{F}$-measurable*** if preimages of measurable sets are measurable; every measurable function $f : (X, \mathcal{F}) \to (Y, \mathcal{G})$ defines a

pushforward measure μ_f on Y by $\mu_f(G) = f(\mu^{-1}(G))$ for $G \in \mathcal{G}$.

For quantum probability, we fix a von Neumann algebra $\mathcal{M}$ and identify elements of its predual $\mathcal{M}_*$ with unit vectors $|\psi\rangle$, on which an operator A acts as $|\psi\rangle \mapsto \langle A\rangle_\psi$, which is the expectation value $\langle\psi|A|\psi\rangle$. The embedding of $|\psi\rangle$ in $\mathcal{M}^*$, i.e. a normal state, is defined by $\psi(A) := \langle A\rangle_\psi = \mathrm{Tr}(P_\psi A)$. With this in mind, we define a ***quantum probability space*** to be a von Neumann algebra $\mathcal{M}$ equipped with a specified normal state $|\psi\rangle$.

If $\mathcal{M}$ is abelian, then, since von Neumann algebras are C^*-algebras, $\mathcal{M}$ is isomorphic to the C^*-algebra $C(X)$ of continuous functions from a compact Hausdorff space X to $\mathbb{C}$ by the Gelfand representation. With this isomorphism, it can be shown that the normal states on $\mathcal{M}$ are in bijection with the Radon measures on X. The correspondence goes in the other direction as well: any probability space $(X, \mathcal{F}, p)$ gives a von Neumann algebra $L^\infty(X, \mathcal{F}, p)$, which is interpreted as acting on $L^2(X, \mathcal{F}, p)$. $L^\infty(X, \mathcal{F}, p)$ is spanned by the family of projections $\{\chi_F : F \in \mathcal{F}\}$, whose span is in fact dense in $L^\infty(X, \mathcal{F}, p)$. The measure p defines a state ϕ on $L^\infty(X, \mathcal{F}, p)$ by $\phi(f) = \int_X f\, dp$ which, since p is countably additive, makes ϕ normal.

So, an abelian quantum probability space gives us a classical probability space, and a classical probability space gives us an abelian quantum probability space. We can therefore loosely state that the study of abelian von Neumann algebras is equivalent to classical probability theory. Noncommutative von Neumann algebras, then, must give us *noncommutative* measure theory.

Every normal state ϕ on $\mathcal{M}$ determines a probability measure μ on $\mathcal{P}(\mathcal{M})$ (with the discrete σ-algebra), with $\phi\left(\bigvee P_n\right) = \sum \phi(P_n)$ for countable disjoint n being the equivalent of σ-additivity. We have $p(I) = 1$ and $p(0) = 0$, with all projections being placed between those two dependent on their position in the complete lattice $\mathcal{P}(\mathcal{M})$. ***Gleason's theorem*** says that the converse is true when $\mathcal{M} = \mathcal{B}(\mathcal{H})$: every map $p : \mathcal{P}(\mathcal{H}) \to [0, 1]$ satisfying the equivalent of σ-additivity extends uniquely to a normal state on $\mathcal{B}(\mathcal{H})$. In fact, every finitely σ-additive map extends to a (not necessarily normal) state. This theorem can be extended to any von Neumann algebra with no direct summand of type I_2.

B.2 Quantum Mechanics

The section on functional analysis is based on [Haase, 2014,Rudin, 1973], and the natural segue into quantum probability theory relies on *many* sources, including [Takhtadzhian, 2008,Meyer, 2006,Holevo, 2003,Rédei and Summers, 2007], each of which tells a small part of a large story. In addition to the sources used in our discussion of functional analysis and quantum probability theory, we use [Sakurai et al., 2014] as a source for quantum mechanics.

B.2.1 Classical Mechanics

We'll sketch out the basics, using [Landau and Lifshitz, 2013] as our primary source for classical mechanics in its traditional, analytic sense; [Arnold, 2013] concerns the porting of this theory over to manifolds, which will later allow us to discuss general relativity and more abstract models of mechanics such as those encountered in synthetic differential geometry.

Equations of Motion Suppose we have a system consisting of N particles in a three-dimensional space. Each particle has an x, y, and z component, and we require $3N$ ***degrees of freedom*** to express the state of this system at any given moment. Generalizing this, suppose the quantities $q_1, \ldots, q_s$ completely define a system: these q_i are ***generalized coordinates***, and their time derivatives $\dot{q}_i$ are their ***generalized derivatives***. Heuristically, if all coordinates $q = \{q_i\}$ and velocities $\dot{q}$ are given, the accelerations $\ddot{q}$ are uniquely determined.

The most general formulation of classical mechanics is given by the ***principle of least action***, which states that (a) there is a function $L(q, \dot{q}, t)$ (known as the ***Lagrangian*** of a system's generalized coordinates at a given time (of which q and $\dot{q}$ are themselves functions), and that q and $\dot{q}$ are specified so as to extremize the ***action***

$$S = \int_{t_1}^{t_2} L(q, \dot{q}, t)\, dt$$

To play around with this, we'll need some concepts from the calculus of variations. For a functional $F[f]$, the ***functional derivative*** is given by

$$\frac{\delta F}{\delta f} = \lim_{\epsilon \to 0} \frac{F[f + \epsilon\eta] - F[f]}{\epsilon}$$

For instance, the functional derivative of the action is given by

$$\frac{\delta S}{\delta q(t_0)} = \lim_{\epsilon \to 0} \frac{1}{\epsilon} \int_{t_1}^{t_2} L(q + \epsilon\eta, \dot{q} + \epsilon\dot{\eta}, t) - L(q, \dot{q}, t)\, dt$$

$$= \lim_{\epsilon \to 0} \frac{1}{\epsilon} \int_{t_1}^{t_2} \epsilon\eta \frac{\partial}{\partial q} L(q, \dot{q}, t) + \epsilon\dot{\eta} \frac{\partial}{\partial \dot{q}} L(q, \dot{q}, t) + O(\epsilon^2)\, dt$$

If we set boundary conditions on what $q(t_1), \dot{q}(t_1), q(t_2)$, and $\dot{q}(t_2)$ are, we must also set $\eta(t_1) = \eta(t_2) = 0$, so as not to alter these conditions. Then, applying integration by parts, we get

$$\frac{\delta S}{\delta q} = \int_{t_1}^{t_2} \eta \left[\frac{\partial L}{\partial q} - \frac{d}{dt}\left(\frac{\partial L}{\partial \dot{q}} \right) \right] dt$$

Since the principle of least action implies that q is selected so as to extremize the action, we must be at a peak (or trough) of the action, and $\delta S/\delta q$ must be zero, regardless of what η is; the expression in the brackets must therefore be zero. Therefore, any q obeying the principle of least action must also obey the equation

$$\frac{d}{dt}\left(\frac{\partial L}{\partial \dot{q}} \right) - \frac{\partial L}{\partial q} = 0$$

which is known as the ***Euler-Lagrange equation***.

Lagrangians and Hamiltonians In a vacuum, we can assume by symmetry that we're in a reference frame where space is homogeneous and isotropic (the same regardless of orientation); such a reference frame is called an ***inertial frame***. In an inertial frame, the Lagrangian can't refer explicitly to the radius vector, the time, or the direction of the velocity, implying that the Lagrangian for a *free* particle is solely a function of $\vec{v} \cdot \vec{v} = v^2$. Plugging this finding into the Euler-Lagrange equations, we see that $\dfrac{d}{dt}\dfrac{\partial L}{\partial \vec{v}} = 0$, so $\partial L/\partial \vec{v}$ is constant; since this is a function of $\vec{v}$ only, it follows that $\vec{v}$ is constant, and therefore that free motion in an inertial frame has a constant velocity: this is known as the ***law of inertia***. Heuristically, two inertial frames, perhaps moving at different velocities, are equivalent in all mechanical respects: this is known as ***Galileo's relativity principle***.

For a *system* of particles which interact with each other, but which are isolated from exterior forces (a ***closed*** system), we subtract from the kinetic energy term $T = \sum \frac{1}{2} m_i v_i^2$ a potential energy term U that depends on the locations r_i of the particles, giving us

$$L = \sum \frac{1}{2} m_i v_i^2 - U(r_1, \ldots, r_n)$$

Solving the Euler-Lagrange equations gives us

$$m_i \frac{dv_i}{dt} = -\frac{\partial U}{\partial r_i}$$

Such equations of motion are called ***Newton's equations***, and the term on the LHS, $m\dot{v}_i$, is known as the ***force***. Note that, since the equations of motion depend entirely on derivatives of the Lagrangian, the potential is effectively only defined up to a constant; we generally choose this constant such that the potential goes to zero as the particles get infinitely far away from one another.

Given a Lagrangian L, we may define the ***conjugate momentum*** to a coordinate q_i to be $p_i := \frac{\partial L}{\partial \dot{q}_i}$. For instance, when $L = \frac{1}{2} m\dot{q}^2 - U(q)$, $p = m\dot{q}$. If the kinetic energy T is a function of $\dot{q}$ alone and the potential energy a function of q alone, then $\sum_i p_i \dot{q}_i = 2T$, and the quantity $H = \sum_i p_i \dot{q}_i - L$ yields $T + U$, the total energy of the system. This quantity, which is in general conserved, is known as the ***Hamiltonian***. While we express the Lagrangian as a function of q, $\dot{q}$, and t, we conventionally express the Hamiltonian as a function of p, q, and t. By matching different expressions for the total differential dH of the Hamiltonian,

$$dH = \frac{\partial H}{\partial p} dp + \frac{\partial H}{\partial q} dq + \frac{\partial H}{\partial t} dt = d(p\dot{q} - L)$$

we can obtain ***Hamilton's equations***,

$$\frac{dp}{dt} = -\frac{\partial H}{\partial q} \qquad \frac{dq}{dt} = \frac{\partial H}{\partial p}$$

B.2.2 Measurements

As per Dirac, "a measurement always causes the system to jump into an eigenstate of the dynamical variable that is being measured." To illustrate, say an operator A with some corresponding physical variable (e.g., position) has eigenkets $\{|a_i\rangle\}$, where a_i refers to an actual value of the variable. A normalized ket $|\alpha\rangle$ is represented in this basis as $|\alpha\rangle = \sum c_i |a_i\rangle$, where $c_i = \langle a_i | \alpha \rangle$. When we make a measurement of the variable corresponding to A, $|\alpha\rangle$ jumps into *one* of the $|a_i\rangle$, and the probability of a specific ket $|a_i\rangle$ is $|\langle a_i|\alpha\rangle|^2$. Since $|\alpha\rangle$ is normalized, we know that $\sum_i |\langle a_i|\alpha\rangle|^2 = 1$, so the probabilities sum to 1. The ***expectation value*** of A in the state $|\alpha\rangle$, denoted as $\langle A\rangle_\alpha$ (α is often suppressed, especially when it is some ground state), can be calculated as

$$\langle A\rangle_\alpha = \sum_{a_i} a_i P(a_i) = \sum_{a_i} a_i |\langle a_i|\alpha\rangle|^2 = \sum_{a_i}\sum_{a_j}\langle\alpha|a_j\rangle\langle a_j|A|a_i\rangle\langle a_i|\alpha\rangle = \langle\alpha|A|\alpha\rangle$$

Define the ***commutator*** of two observables A, B as

$$[A, B] = AB - BA$$

and the ***anticommutator*** of A and B as

$$\{A, B\} = AB + BA$$

The observables A, B are said to be ***compatible*** when $[A, B] = 0$, and ***incompatible*** otherwise.

Suppose A's eigenvalues are nondegenerate and generate a basis, in which the matrix representation of A is diagonal. If B is compatible with A, B is diagonal in A's basis as well. Why? $\langle a_i|[A, B]|a_j\rangle = \langle a_i|0|a_j\rangle = 0 = (a_i - a_j)\langle a_i|B|a_j\rangle$, which by nondegeneracy implies that $\langle a_i|B|a_j\rangle = 0$ unless $i = j$. So, really, the eigenkets of A are the eigenkets of B, though they may have different eigenvalues: they are said to be ***simultaneous eigenkets***, and are sometimes denoted by $|a_i, b_i\rangle$. We may also use a ***collective index***, $|K_i\rangle = |a_i, b_i\rangle$. Due to the simultaneity of the eigenkets, measurements of A do not interfere with measurements of B, and vice-versa; this can be extended to larger sets of pairwise compatible operators. Of course, if A and B are incompatible, then simultaneous eigenkets generally do *not* exist and successive measurements *do* interfere with each other.

To represent our uncertainty in the result of a measurement, we adopt the statistical notion of variance, calling it ***dispersion***: defining $\Delta A = A - \langle A\rangle$, the dispersion, also known as the variance or mean square deviation, is given by the expectation of $(\Delta A)^2$, or

$$\langle(\Delta A)^2\rangle = \langle(A^2 - 2A\langle A\rangle + \langle A\rangle^2)\rangle = \langle A^2\rangle - \langle A\rangle^2$$

It is more convenient to denote this by σ_A^2.

For observables A, ΔA is also Hermitian, since the expectation is a real number (implicitly multiplied by the identity) and thus equal to its own adjoint. We can use the fact that an operator can be defined by its action on all possible kets to lift certain identities on vectors in Hilbert spaces to corresponding identities on their operators: for instance, if we assume that operators A and B are Hermitian, we can lift the Cauchy-Schwarz identity $\langle\alpha|\alpha\rangle\langle\beta|\beta\rangle \geq |\langle\alpha|\beta\rangle|^2$ to a corresponding identity $\langle A^2\rangle\langle B^2\rangle \geq |\langle AB\rangle|^2$. Since the dispersion operators of observables are Hermitian, this implies that $\sigma_A^2\sigma_B^2 \geq |\langle\Delta A\Delta B\rangle|^2$ for any observables A, B. In fact, expanding this yields:

$$\sigma_A^2\sigma_B^2 \geq |\langle\Delta A\Delta B\rangle|^2 = \left|\frac{1}{2}\langle[A,B]\rangle + \frac{1}{2}\langle\{\Delta A, \Delta B\rangle\right|^2 = \frac{1}{4}|\langle[A,B]\rangle|^2 + \frac{1}{4}|\langle\{\Delta A, \Delta B\}\rangle|^2$$

giving us the important inequality

$$\sigma_A\sigma_B \geq \frac{1}{2}|\langle[A,B]\rangle|$$

(Note that "σ_A" is notational trickery, since σ_A^2 itself is not a square, but the expectation value of a square; however, as σ_A^2 corresponds to variance, σ_A corresponds to the standard deviation of A).

B.2.3 Position, Momentum, and Time

We've been dealing with finite-dimensional spaces so far, where spectra are finite and everything converges. Now we'll move to infinite-dimensional spaces, replacing Kronecker deltas by Dirac deltas and sums by integrals: for instance, $\langle a_i|a_j\rangle = \delta_{ij}$ becomes $\langle a_i|a_j\rangle = \delta(i-j)$, and $\sum_i |a_i\rangle\langle a_i| = 1$ becomes $\int |a_i\rangle\langle a_i|\, di = 1$.

Consider a position operator x on one dimension, whose eigenkets $x|x_i\rangle = x_i|x_i\rangle$ form a complete set. An arbitrary physical state $|\alpha\rangle$ can be expanded as $|\alpha\rangle = \int_{-\infty}^{\infty} |x_i\rangle\langle x_i|\alpha\rangle\, dx_i$. Suppose we centered a detector of length ℓ at position x_0: when the detector registers a particle, the state changes:

$$|\alpha\rangle = \int_{-\infty}^{\infty} |x_i\rangle\langle x_i|\alpha\rangle\, dx_i \longrightarrow \int_{x_0-\ell/2}^{x_0+\ell/2} |x_i\rangle\langle x_i|\alpha\rangle\, dx_i$$

The probability of the particle being detected in this range is given by

$$\int_{x_0-\ell/2}^{x_0+\ell/2} |\langle x_i|\alpha\rangle|^2\, dx_i$$

Of course, as $\ell \to \infty$, this probability goes to 1 as long as $|\alpha\rangle$ is normalized.

To consider three dimensions x, y, z, we must be assured that measurement in one dimension does not affect the other two, so $[x, y] = [x, z] = [y, z] = 0$. Defining $\vec{x}$ as a collective index for x, y, z, such that $|\vec{x}\rangle$ is a simultaneous eigenket for the observables x, y, z, consider the ***infinitesimal translation*** operator $\mathcal{J}(d\vec{x})|\vec{x}\rangle = |\vec{x} + d\vec{x}\rangle$. What properties should we expect such an operator to have? It should preserve normalized eigenkets, implying that $\langle\alpha|\mathcal{J}^\dagger(d\vec{x})\mathcal{J}(d\vec{x})|\alpha\rangle = \langle\alpha|\alpha\rangle = 1$ and therefore that $\mathcal{J}(d\vec{x})$ is unitary. We should also have $\mathcal{J}(d\vec{x}_1)\mathcal{J}(d\vec{x}_2) = \mathcal{J}(d\vec{x}_1 + d\vec{x}_2)$ and $\mathcal{J}(-d\vec{x}) = \mathcal{J}^{-1}(d\vec{x})$. Finally, as $d\vec{x}$ goes to zero, $\mathcal{J}(d\vec{x})$ should go to the identity operator: $\lim_{d\vec{x}\to 0} \mathcal{J}(d\vec{x}) = 1$.

If we take $\mathcal{J}(d\vec{x}) = 1 - i\vec{K} \cdot d\vec{x}$ for some hermitian $\vec{K} = (K_x, K_y, K_z)$, all these properties are satisfied (up to $O((d\vec{x})^2)$, which is good enough, since $d\vec{x}$ is infinitesimal). Accepting this to be the correct form for $\mathcal{J}(d\vec{x})$, we note that $[\vec{x}, \mathcal{J}(d\vec{x})] = d\vec{x}$ and therefore that $[x_i, K_j] = i\delta_{ij}$. This $\vec{K}$ seems to generate translations, so it must be in some way related to momentum. Since $\vec{K} \cdot d\vec{x}$ is dimensionless, $\vec{K}$ has units L^{-1}. We can define it as $\vec{p}$ divided by some constant with the dimension of action, L^2MT^{-1}. Calling this constant $\hbar$, we rewrite $\mathcal{J}(d\vec{x}) = 1 - i\vec{p} \cdot d\vec{x}/\hbar$, assuring that momentum really is the generator of translation. Our commutation relation becomes $[x_i, p_j] = i\hbar\delta_{ij}$, and we can now state the ***Heisenberg uncertainty principle*** as a special case of the more general relation above:

$$\sigma_x \sigma_{p_x} \geq \frac{\hbar}{2}$$

Note: $[p_i, p_j] = 0$, and we can use $\vec{p} = (p_x, p_y, p_z)$ to create a simultaneous momentum eigenket $|\vec{p}\rangle$. This forms one of the three ***canonical commutation relations*** of quantum mechanics:

$$[x_i, x_j] = 0 \qquad [p_i, p_j = 0] \qquad [x_i, p_j] = i\hbar\delta_{ij}$$

Time Evolution Suppose a state $|\alpha\rangle$ is pictured at some time t_0. We write this state as $|\alpha, t_0\rangle$, and its evolution to an arbitrary time t we write $|\alpha, t_0; t\rangle$. We want a time evolution operator $\mathcal{U}(t, t_0)|\alpha, t_0\rangle = |\alpha, t_0; t\rangle$ with the same conditions as the above infinitesimal position operator. We again make the choice $\mathcal{U}(t_0 + dt, t_0) = 1 - i\Omega dt$ for some Hermitian Ω. In classical mechanics, the Hamiltonian H is the generator of time evolution, and we correspondingly define

$\Omega = H/\hbar$, giving us $\mathcal{U}(t_0 + dt, t_0) = 1 - iH\,dt/\hbar$. We find that

$$\mathcal{U}(t + dt, t_0) - \mathcal{U}(t, t_0) = -i(H/\hbar)\,dt\,\mathcal{U}(t, t_0)$$

and therefore that

$$ih\frac{\partial}{\partial t}\mathcal{U}(t, t_0) = HU(t, t_0)$$

Multiplying both sides by a state ket $|\alpha\rangle$ immediately leads to the ***time-dependent Schrodinger equation***,

$$i\hbar\frac{\partial}{\partial t}|\alpha, t_0; t\rangle = H|\alpha, t_0; t\rangle$$

Defining the exponential of an operator A by the Taylor series for the usual exponential, $e^A = 1 + A + A^2/2 + A^3/6 + \ldots$, the solution to this equation is the same as it is for a normal differential equation:

$$\mathcal{U}(t, t_0) = e^{-\frac{i}{\hbar}H(t-t_0)}$$

when H is not a function of time,

$$\mathcal{U}(t, t_0) = e^{-\frac{i}{\hbar}\int_{t_0}^{t} H(t')\,dt'}$$

when H is a function of time but $[H(t_1), H(t_2)] = 0$, and

$$\mathcal{U}(t, t_0) = 1 + \sum_{n=1}^{\infty}\left(-\frac{i}{\hbar}\right)^n \int_{t_0}^{t}\int_{t_0}^{t_1}\ldots\int_{t_0}^{t_{n-1}} H(t_1)H(t_2)\ldots H(t_n)\,dt_n\,dt_{n-1}\ldots dt_1$$

when H is a function of time and $[H(t_1), H(t_2)] \neq 0$. We'll generally deal only with the first case.

Suppose that H is time-independent and generates a complete basis $\{|a_i\rangle\}$, with $H|a_i\rangle = E_{a_i}|a_i\rangle$. Setting $t_0 = 0$ and expanding the time evolution operator in terms of $|a_i\rangle\langle a_i|$, we find that

$$e^{-\frac{i}{\hbar}Ht} = \sum_i\sum_j |a_j\rangle\langle a_j|e^{-\frac{i}{\hbar}Ht}|a_i\rangle\langle a_i| = \sum_i |a_i\rangle e^{-\frac{i}{\hbar}E_{a_i}t}\langle a_i|$$

For an arbitrary ket $|\alpha\rangle = \sum_i |a_i\rangle\langle a_i|\alpha\rangle = \sum_i c_{a_i}|a_i\rangle$, we have

$$|a; t\rangle = e^{-\frac{i}{\hbar}Ht}|\alpha\rangle = \sum_i c_{a_i}e^{-\frac{i}{\hbar}E_{a_i}t}|a_i\rangle$$

So the coefficent $c_{a_i}(t)$ is given by $c_{a_i}(t) = c_{a_i}e^{-\frac{i}{\hbar}E_{a_i}t}$.

How does the expectation value of an observable change over time? Observe:

$$\langle B\rangle_{a_i} = \langle a_i, t|B|a_i, t\rangle = \langle a_i|\mathcal{U}^\dagger(t,0)B\mathcal{U}(t,0)|a_i\rangle = \langle a_i|e^{\frac{i}{\hbar}E_{a_i}t}Be^{-\frac{i}{\hbar}E_{a_i}t}|a_i\rangle = \langle a_i|B|a_i\rangle$$

implying that the expectation values of observables taken with respect to energy eigenstates does *not* change over time. Energy eigenstates are correspondingly known as ***stationary states***. In general, this does not hold true for expectation values taken with respect to superpositions of energy eigenstates, which are correspondingly known as ***nonstationary states***.

The above exposition is an example of the ***Schrodinger picture*** of quantum dynamics, in which state kets are postulated to change over time while observables stay constant. We can view this in another way, though: state kets are constant, while observables change. This is known as the ***Heisenberg picture***, and relies on the following mathematical equality: consider two state kets $|\beta\rangle$ and $|\alpha\rangle$ and an observable U. Since observables are unitary, $\langle\beta|\alpha\rangle = \langle\beta|U^\dagger U|\alpha\rangle$. For an operator X, consider the action of a unitary transformation $X \mapsto U^\dagger XU$ on $\langle\beta|U|\alpha\rangle$. We have

$$\langle\beta|X|\alpha\rangle \mapsto \langle\beta|U^\dagger XU|\alpha\rangle$$

But we can view this in two equivalent ways:

$$\left(\langle\beta|U^\dagger\right) X \left(U|\alpha\rangle\right) \;=\; \langle\beta|\left(U^\dagger XU\right)|\alpha\rangle$$

So either the bras and kets change as $|\alpha\rangle \mapsto U|\alpha\rangle$, or the operator itself changes as $X \mapsto U^\dagger XU$. These two pictures have different physical interpretations, but are entirely equivalent; in the case that $U = \mathcal{U}$, the time evolution operator, we recover the Schrodinger-Heisenberg distinction.

B.3 Relativity

As in classical mechanics, to talk about nature we need a reference frame, or coordinate system. We would like moving bodies not acted upon by external forces to move at constant velocities; a reference frame in which this holds is known as an inertial reference frame. We can have multiple reference frames, each attached to a distinguished point serving as the origin; if one is inertial, and the other moves uniformly relative to the first, the other is inertial.

Galileo's principle of relativity states that laws of nature are identical in all inertial reference frames. This principle, however, was formulated with the idea of instantaneous transmission of physical signals in mind; in experiment, we find that this doesn't happen, and that the maximum velocity of propagation is a finite constant known as the speed of light, $c \approx 3 \times 10^8\ m/s$. Einstein's principle of relativity states that physical laws are invariant under choice of inertial reference frame; in particular, they all measure the same c. Theories of mechanics built upon this principle are called relativistic.

B.3.1 Intervals

In special relativity, the primitive objects of study are events, or points in spacetime ($\mathbb{R}^4$). Suppose two events happen with spacetime coordinates in a reference frame K given by (x_1, y_1, z_1, t_1) and (x_2, y_2, z_2, t_2), respectively, corresponding to the emission and receiving of a light-speed signal, respectively. The signal covers a distance $c(t_2 - t_1)$ which is equal to $\sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2 + (z_2$ so we can write

$$(x_2 - x_1)^2 + (y_2 - y_1)^2 + (z_2 - z_1)^2 - c^2(t_2 - t_1)^2 = 0$$

In a system K' where the coordinates of the two events are (x'_1, y'_1, z'_1, t'_1) and (x'_2, y'_2, z'_2, t'_2), respectively, the velocity c^2 is still the same due to the principle of invariance, so we have

$$(x'_2 - x'_1)^2 + (y'_2 - y'_1)^2 + (z'_2 - z'_1)^2 - c^2(t'_2 - t'_1)^2 = 0$$

In general, in a reference frame K where two events have coordinates (x_1, y_1, z_1, t_1) and (x_2, y_2, z_2, t_2), the interval between those two coordinates is given by

$$s_{12}^2 = c^2(t_2 - t_1)^2 - (x_2 - x_1)^2 - (y_2 - y_1)^2 - (z_2 - z_1)^2$$

We've deduced that if the interval is zero in any one reference frame, it's zero in all reference frames. If two events are infinitely close to each other, the interval ds between them is given by

$$ds^2 = c^2\, dt^2 - dx^2 - dy^2 - dz^2$$

If we measure the same interval in two different reference frames K and K' to get ds and ds', it follows from the facts that (1) if $ds = 0$ then $ds' = 0$ and (2) ds and ds' are infinitesimals of the same order, that ds and ds' are proportional to each other: $ds = a\,ds'$. Since space and time are homogeneous and isotropic, the constant of proportionality cannot depend on the coordinates or the time, nor can it depend on the direction of the relative velocity. Therefore, $ds' = a\,ds$, with the *same* constant of proportionality. It follows that $ds = a^2\,ds$, so $a^2 = 1$ and $a = \pm 1$. a obviously can't be -1, since moving between *three* reference frames would give us $ds = -ds$, so we must have $a = 1$. Therefore, $ds = ds'$ and $s = s'$. The interval between two events is independent of the frame of reference.

The Light Cone Suppose we have two events in spacetime, viewed from a reference frame K, and you, a massive object (no offense) want to get from one to the other by traveling along a straight line. Were we to attach a reference frame K' to you, putting you at the origin, we'd find that both events have the same space coordinates in K'. Introducing the notation $t_{12} = t_2 - t_1$ and $l_{12}^2 = (x_2 - x_1)^2 + (y_2 - y_1)^2 + (z_2 - z_1)^2$, the intervals in K and K' are $s_{12}^2 = c^2 t_{12}^2 - l_{12}^2$ and $s_{12}'^2 = c^2 t_{12}'^2 - l_{12}'^2$. Since $l_{12}'^2 = 0$ and $s_{12}^2 = s_{12}'^2$, we have $s_{12}^2 = c^2 t_{12}^2 - l_{12}^2 = c^2 t_{12}'^2 > 0$. So you can get from one to the other if $s_{12}^2 > 0$. We call such an interval ***timelike***, since all that's keeping you from traveling along it is time. If we want the two events to happen at the same time, we require $s_{12}^2 < 0$, and call the interval ***spacelike***, since you'd have to *teleport* through space to get from one to the other. Because of the invariance of intervals, the spacelike/timelike divide is an absolute division, independent of reference frames; at any point p in a coordinate system there is a cone defined by $x^2 + y^2 + z^2 - c^2 t^2 = 0$ known as the ***light cone***, any point outside of which is absolutely remote relative to p, and any point inside which is either in the absolute past or absolute future relative to p, where $t < 0$ and $t > 0$, respectively.

Proper Time Suppose that we're at the center of an inertial reference frame K, we have two clocks C and C', and we chuck C' away at an arbitrary velocity. During an infinitesimal period of time dt as measured by our clock C, C' will travel a distance $\sqrt{dx^2 + dy^2 + dz^2}$. Because of the invariance of intervals, $ds^2 = c^2 dt^2 - dx^2 - dy^2 - dz^2 = c^2 dt'^2$, so

$$dt' = \frac{ds}{c} = dt\sqrt{1 - \frac{dx^2 + dy^2 + dz^2}{c^2\,dt^2}} = dt\sqrt{1 - \frac{v^2}{c^2}}$$

Integrating this expression, we see that over a time interval $t_2 - t_1$ measured by C, C' experiences a time interval

$$t_2' - t_1' = \int_{t_1}^{t_2} \sqrt{1 - \frac{v^2}{c^2}}\, dt$$

Since this interval is less than $t_2 - t_1$, C' is seen as lagging. Paradoxically, however, from C''s reference frame, C is lagging!

The proper time for an object is the time read by a clock moving along with that object, which is the integral $\int_a^b \frac{ds}{c}$ taken along the world line of the clock. For two points separated by a timelike interval, this integral has the maximum value when taken along the straight world line joining these two points.

B.3.2 Lorentz Transformations

We want to translate the set of coordinates (x, y, z, t) in the reference frame K to another set of coordinates (x', y', z', t') in a reference frame K'. Supposing K' moves along K's x axis at a velocity V, in classical mechanics we'd set $x' = x + Vt, y' = y, z' = z, t' = t$, which is known as the Galilean transformation, but this fails to leave intervals invariant, making it unacceptable for relativistic mechanics.

Setting $\tau = ict$, such that $s^2 = x^2 + y^2 + z^2 + \tau^2$, and changing coordinates to (x, y, z, τ), what we're looking for is precisely an isometry of this space. It's then either a parallel displacement or a rotation. Displacement doesn't matter, since it only changes the origin, so we want a rotation: every rotation can be broken up into six rotations in the $xy, zy, xz, \tau x, \tau y, \tau z$ planes. We don't care about $xy, zy, xz, \tau y$, or τz rotations, so this must be a τx rotation, changing coordinates as $x = -\tau' \sin\psi$, $\tau = \tau' \cos\psi$. From this it follows that $\tan\psi = iV/c$, so simple algebra leads us to the change of coordinates

$$x = \frac{x' + Vt'}{\sqrt{1 - \frac{V^2}{c^2}}} \quad y = y' \quad z = z' \quad t = \frac{t' + V\frac{x'}{c^2}}{\sqrt{1 - \frac{V^2}{c^2}}}$$

This transformation is known as the Lorentz transformation. Clearly, it yields the Galilean transformation as $c \to \infty$. As a consequence, suppose a rod moving along the x axis at velocity

V relative to us measures its own length as $\Delta x'$: we will then measure its length as

$$\Delta x = \frac{\Delta x'}{\sqrt{1 - \frac{V^2}{c^2}}}$$

In other words, the faster it goes, the shorter it appears to us. This is known as Lorentz contraction.

By considering such a transformation for infinitesimal dx, dt, we can find formulas for the transformation of velocities: under the same conditions as above, we have

$$v_x = \frac{v'_x + V}{1 + v'_x \frac{V}{c^2}} \quad v_y = \frac{v'_y \sqrt{1 - \frac{V^2}{c^2}}}{1 + v'_y \frac{V}{c^2}} \quad v_z = \frac{v'_z \sqrt{1 - \frac{V^2}{c^2}}}{1 + v'_z \frac{V}{c^2}}$$

Again, as $c \to \infty$, we get the classical transformation, in which $v_x = v'_x + V$.

We generally denote the factor $\frac{1}{\sqrt{1 - \frac{V^2}{c^2}}}$ as γ, the Lorentz factor. So, for instance, we can restate Lorentz contraction and time dilation as $\Delta x = \gamma \Delta x'$ and $\Delta t = \gamma \Delta t'$, respectively.

B.3.3 Four-vectors

We'll set $c = 1$ from now on; if you want, you can figure out where it's been hidden via dimensional analysis. In the four dimensional spacetime manifold in which relativistic mechanics take place, Minkowski space, vectors have three space components and one time component, and are known as four-vectors. The inner product on this space is given by

$$a \cdot b = a^0 b^0 - a^1 b^1 - a^2 b^2 - a^3 b^3$$

We can write this neatly by introducing a metric tensor η_{ij} on this manifold, given by

$$\eta_{ij} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix}$$

So $a \cdot b = \eta_{ij} a^i b^j$. We can restate several of the above developments in sleeker ways: the infinitesimal interval (or line element) is given by $ds^2 = -\eta_{ij} dx^i dx^j$, the path length and proper time are given by

$$\Delta s = \int \sqrt{-\eta_{ij}\frac{dx^i}{d\lambda}\frac{dx^j}{d\lambda}}\,d\lambda \qquad \Delta\tau = \int \sqrt{\eta_{ij}\frac{dx^i}{d\lambda}\frac{dx^j}{d\lambda}}\,d\lambda$$

Recall the Einstein summation notation: (i) when the same index appears in both a raised and a lowered position, we implicitly sum over it, e.g. $v_i w^i = \sum_{i=1}^4 v_i w^i$ (ii) we use the metric to raise and lower indices at will, e.g. $v^i = \eta^{ij} v_j$, and (iii) putting indices in square (curly) brackets indicates that we wish to take their commutator (anticommutator), e.g. $v_{[i,}w_{j]} = v_i w_j - v_j w_i$. By rewriting everything in terms of tensors, we can express relationships without invoking any sort of reference frame; doing this makes an equation, relationship, or theory covariant (which has nothing to do with covariance/contravariance of tensors).

The velocity of a particle x^i, parametrized by its proper time, is given by $v^i = \partial_\tau x^i$; since $d\tau^2 = \eta_{ij} dx^i dx^j$, we have $\eta_{ij} v^i v^j = 1$, the interpretation being that we're *always* traveling at the same speed through spacetime (light-speed, really; examining units, the 1 yields a hidden c), and that moving faster through space just means moving slower through time. The momentum of a particle is given by $p^i = \gamma m v^i$, and the energy is γm. The force on a particle is given by $f^i = \partial_\tau u^i$.

B.3.4 General Relativity

General relativity is far more subtle, though a significant portion of the legwork was performed in the previous discussion of Riemannian geometry. We postulate that gravitational force on an observer is equivalent to the "pseudo"-force experienced by an observer in an accelerating reference, a postulate known as the *equivalence principle*. Our sources include [Wald, 2007, Carroll, 2019, Misner et al., 1973]. The differential geometry book [Kühnel, 2015] discusses general relativity as well, focusing in particular on "Einstein manifolds", or Riemannian manifolds whose metrics are solutions to the vacuum Einstein field equations.

Pseudo-Riemannian Manifolds We begin by recapping some constructions on a pseudo-Riemannian manifold (M, g). The *Levi-Civita connection* ∇_i is the unique connection on M that preserves g and has vanishing torsion tensor, and its difference from the ordinary derivative ∂_i

is given by the *Christoffel symbols*,

$$\Gamma^i{}_{jk} := \frac{1}{2} g^{i\ell} \left(\partial_k g_{\ell j} + \partial_j g_{\ell k} - \partial_\ell g_{jk} \right)$$

Having written these down, we can express the action of ∇_i on a vector v^j as

$$\nabla_i v^j = \partial_i v^j + \Gamma^j{}_{ik} v^k$$

In local coordinates, the Christoffel equations give us second-order differential equations for the position x^i of a "particle" traveling on a geodesic, known as the *geodesic equations*:

$$\frac{d^2 x^i}{dt^2} + \Gamma^i{}_{jk} \frac{dx^j}{dt} \frac{dx^k}{dt} = 0$$

(Compare this with the result that the geodesics in a flat space are straight lines, i.e. $\ddot{x} = 0$). For any given initial position x^i and velocity $\frac{dx^i}{dt}$, the theory of ordinary differential equations tells us that a unique solution exists to the geodesic equations.

Given an infinitesimal square with sides v^i and w^i, parallel transport of a vector x^i around the square generally fails to leave x^i unaltered. The difference, as a vector, is linear in v^i, w^i, and x^i, and hence is given by $y^\ell = R^\ell{}_{ijk} v^j w^k x^i$ for some tensor $R^\ell{}_{ijk}$ known as the *Riemann curvature tensor*. In terms of the Christoffel symbols, this tensor can be given as

$$R^\ell{}_{ijk} = \partial_j \Gamma^\ell{}_{ki} - \partial_k \Gamma^\ell{}_{ji} + \Gamma^\ell{}_{jm} \Gamma^m{}_{ki} - \Gamma^\ell{}_{km} \Gamma^m{}_{ji}$$

Contracting it yields the *Ricci curvature* R_{ij} and *scalar curvature* R:

$$R_{ij} = R^\ell{}_{i\ell j} \qquad\qquad R = R^i{}_i$$

We define the *Einstein tensor* G_{ij} by

$$G_{ij} = R_{ij} - \frac{1}{2} R g_{ij}$$

A metric g_{ij} which solves the equations $G_{ij} = 0$ is one which distributes the curvature of M "most evenly" [Kühnel, 2015]. A key property of the Einstein tensor is its vanishing divergence: $\nabla^i G_{ij} = 0$.

The Stress-Energy Tensor General relativity historically has its roots in an attempt to generalize the Poisson equation, a field-theoretic version of Newtonian gravity. Given a mass density ρ and a gravitational field $\vec{g}$ expressed as the gradient of a scalar potential ϕ, Gauss's law reads $\nabla \cdot \vec{g} = -4\pi G\rho$, where G is a gravitational constant. Plugging in $\vec{g} = -\nabla\phi$, we obtain *Poisson's equation*,

$$\nabla^2\phi = 4\pi G\rho$$

To generalize this to the framework of special relativity, we first need to figure out how to replace ρ with something that respects mass-energy equivalence and transforms like a tensor. The solution is a symmetric tensor T_{ij} known as the ***stress-energy tensor***. An observer with velocity v^i will measure a mass-energy per unit volume of $T_{ij}v^iv^j$. Given an x^j orthogonal to v^μ, the component $-T_{ij}x^jv^i$ is interpreted as the momentum density of the matter in the x^j direction. A y^k also orthogonal to v^i can be plugged in along with x^j, and $T_{ij}x^iy^j$ is interpreted as the x^i-y^j component of the stress tensor for a point in an arbitrary material body. To summarize, the stress-energy-momentum tensor T_{ij} gives us *stress* when we plug in two position vectors, *momentum* when we plug in a position vector and an orthogonal velocity vector, and *energy* when we plug in one velocity vector twice. Conservation of energy implies that the stress-energy tensor has vanishing divergence: $\nabla^iT_{ij} = 0$.

The Einstein Field Equations We've identified the mass density ρ with the mass-energy density $T_{ij}v^iv^j$. Now we have to replace $\nabla^2\phi$ with a tensorial quantity as well; it should have at most second-order derivatives of the metric, and it should be divergence-free.

A first guess is given by the observation that the differential acceleration of two nearby particles with separation vector x is given by $-(x \cdot \nabla)\nabla\phi$. However, since their world lines will be geodesics, and a fortiori curves on our spacetime manifold, we know that this same acceleration is given by $-R^\ell{}_{jik}v^jv^kx^i$. So let's make the correspondence $R^\ell{}_{jik}v^jv^k = \partial_i\partial^\ell\phi$, and therefore $\partial^2\phi = R^\ell{}_{j\ell k} = R_{jk}$, and conclude that the correct covariant generalization of the Poisson equation is given by $R_{ij}v^iv^j = 4\pi GT_{ij}v^iv^j$, or more concisely $R_{ij} = 4\pi GT_{ij}$.

This was, in fact, one of Einstein's guesses. It is wrong. It is in general true that $\nabla^iG_{ij} = \nabla^i(R_{ij} - \frac{1}{2}Rg_{ij}) = 0$, and hence the divergence of R_{ij} is given by $\nabla^i\frac{1}{2}Rg_{ij} = \frac{1}{2}\nabla_jR$. Hence, divergence-freeness of R_{ij} implies that $\nabla_iR = 0$, i.e. that R and hence $T = T^i_i$ are constant throughout the universe! The correct solution to the problem is contained within the problem itself: we replace R_{ij} with $\frac{1}{2}G_{ij}$, which we already know to be divergence-free. This yields the ***Einstein field equations***:

$$G_{ij} = 8\pi G T_{ij}$$

Comparing units, we see that there's a hidden c^{-4} on the right-hand side; it is convenient to define ***Einstein's constant*** by $\kappa = 8\pi G/c^4$ and simply write $G_{ij} = \kappa T_{ij}$.

The Lagrangian Formulation In Lagrangian mechanics, we associate to a physical system a function of time $L(t)$ known as the Lagrangian, which governs the dynamics of the system; the Lagrangian is allowed to operate on the positions and velocities of the particles, e.g. as $L(t) = L(q(t), \dot{q}(t)) = \frac{1}{2}m\dot{q}(t)^2 - mgq(t)$. In a field-theoretic context, such as general relativity, we may also consider the Lagrangian as a function of fields ϕ and their first derivatives $\partial_\mu\phi$, e.g. as $L(t) = \int \frac{1}{2}\partial_\mu\phi\partial^\mu\phi - \frac{1}{2}m^2\phi^2\, d^3x$. In this case, we refer to the term which is integrated over space to get the Lagrangian as the Lagrangian density $\mathcal{L}$. Integrating the Lagrangian over time yields the action, $S = \int L\, dt$; the principle of least action states that the positions/fields involved in the Lagrangian are chosen so as to minimize the variation of the action under an arbitrary variation in said positions/fields $\delta S = 0$.

A covariant formulation of Lagrangian mechanics requires us to replace ∂_μ with the covariant derivative ∇_μ, so as to make all terms appearing in the Lagrangian tensorial; further, if we wish to work on an n-dimensional Riemannian manifold (M, g), we must integrate the scalar Lagrangian density $\mathcal{L}$ with respect to the volume form $\sqrt{|g|}\, d^n x$, where $d^n x := dx_1 \wedge \ldots \wedge dx_n$ and $|g|$ is the determinant of the metric tensor.

In a vacuum, the Einstein-Hilbert action of general relativity is given by the Lagrangian density $\mathcal{L}_V = R/2\kappa$:

$$S_V = \int \frac{R}{2\kappa}\sqrt{|g|}\, d^4x$$

Upon variation of the metric, this yields

$$\delta S_V = \int \left(R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu}\right)\delta g^{\mu\nu}\sqrt{|g|}\, d^4x$$

(A detailed derivation is given in [Carroll, 2019]). Since this must be zero for all variations of the metric, we obtain $R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu} = G_{\mu\nu} = 0$, Einstein's equations for a vacuum.

To add mass-energy fields, we add an arbitrary density $\mathcal{L}_M$ to the Lagrangian density, which by the linearity of integration splits the action S into $S_V + S_M$, the sum of the vacuum and mass-energy actions. Working in reverse, we define the stress-energy tensor as

$$T_{\mu\nu} = -\frac{2}{\sqrt{|g|}}\frac{\delta S_M}{\delta g^{\mu\nu}}$$

guaranteeing that the principle of least action reduces to Einstein's equation, $G_{\mu\nu} = \kappa T_{\mu\nu}$.

B.4 Quantum Field Theory

This section discusses the Lorentz covariant generalization of quantum mechanics to fields known as quantum field theory. Our sources for vanilla quantum field theory are [Peskin, 2018, Lancaster and Blundell, 2014, Ticciati et al., 1999]; the two-volume series [Deligne et al.,] delivers mathematical rigor to the field. Being especially confusing, we have tried to root our discussion of spinors in representation theory, for which the books [Weinberg, 1995, Bleecker, 2005] are useful.

B.4.1 Representations of the Lorentz Group

Recall that the distance between two points x^μ, y^μ of Minkowski space X is given by

$$(\eta_{\mu\nu}x^\mu y^\nu)^{1/2} = \sqrt{(x^0-y^0)^2-(x^1-y^1)^2-(x^2-y^2)^2-(x^3-y^3)^2}$$

An ***isometry*** of Minkowski space is a continuous map $X \to X$ preserving the distance between points; the set of all such isometries is a Lie group known as the ***Poincaré group***. It is ten-dimensional, with 4 dimensions dedicated to translations, three to rotations (x-y, x-z, y-z), and three to boosts, or rotations involving the time dimension (t-x, t-y, t-z).

Discarding the translations gives us a six-dimensional Lie group known as the ***Lorentz group*** $L = O(1,3)$; its objects are all linear maps, and hence can be written as matrices $\Lambda^\mu{}_\nu$ satisfying

$$\eta_{\mu\nu}\Lambda^\mu{}_\sigma\Lambda^\nu{}_\rho x^\sigma y^\rho = \eta_{\mu\nu}x^\mu y^\nu$$

In matrix notation, such a Λ satisfies $x^T\eta y = (\Lambda x)^T\eta(\Lambda y)$ for all x, y, and hence $\Lambda^T\eta\Lambda = \eta$.

It follows that $\det(\Lambda^T \eta \Lambda) = -(\det \Lambda)^2 = \det \eta = -1$, so that $\det \Lambda \in \pm 1$. Also, letting $e^0 = (1,0,0,0)$, we have

$$1 = (e^0)^T \eta (e^0) = (\Lambda e^0)^T \eta (\Lambda e^0) = (\Lambda^0{}_0)^2 - (\Lambda^1{}_0)^2 - (\Lambda^2{}_0)^2 - (\Lambda^3{}_0)^2$$

so that $(\Lambda^0{}_0)^2 \geq 1$, implying that either $\Lambda^0{}_0 \geq 1$ or $\Lambda^0{}_0 \leq 1$. It follows that L is composed of four connected components, each consisting of all transformations Λ with a specified determinant and sign of $\Lambda^0{}_0$. We write these components as

$$L_+^\uparrow = \{\Lambda \in L \mid \det \Lambda = 1, \Lambda^0{}_0 \geq 1\} \qquad L_-^\downarrow = \{\Lambda \in L \mid \det \Lambda = -1, \Lambda^0{}_0 \leq 1\}$$

and likewise for $L_-^\uparrow, L_+^\downarrow$. $L_+^\uparrow$, which contains the identity, is often known as the ***restricted*** or ***proper orthochronous*** Lorentz group, $SO^+(1,3)$. Defining the space inversion and time reversal operators $P = \mathrm{diag}(+1,-1,-1,-1)$ and $T = \mathrm{diag}(-1,+1,+1,+1)$ gives the structure of the Klein four-group $\{I_4, P, T, PT\}$ to these four connected components.

Since the exponentiation operator e^- from a Lie algebra $\mathfrak{g}$ to its Lie group G is continuous, and hence has an image contained in one connected component, $\mathfrak{g}$ depends solely on the special component of G containing the identity. Thus, the Lie algebras of $L = \mathrm{O}(1,3), \mathrm{SO}(1,3)$, and $L_+^\uparrow = \mathrm{SO}^+(1,3)$ are all the same. This algebra is generally written as $\mathfrak{so}(1,3)$.

Fix a Lie group G and Lie algebra $\mathfrak{g}$. A ***Lie group representation*** of G is a smooth homomorphism $\Pi : G \to \mathrm{GL}\,(n;\mathbb{C})$ for some n. A ***Lie algebra representation*** of $\mathfrak{g}$ is a Lie algebra homomorphism $\pi : \mathfrak{g} \to \mathfrak{gl}\,(n;\mathbb{C}) \cong \mathrm{End}(\mathbb{C}^n)$. Since the Lie algebra of a Lie group is the tangent space to its identity, the pushforward of any Lie group representation defines a homomorphism between Lie algebras; this homomorphism preserves brackets, so that Lie group representations induce Lie algebra representations. If $\mathfrak{g}$ is the Lie algebra of G, it isn't true in general that (Lie algebra) representations of $\mathfrak{g}$ come from (Lie group) representations of G, but, if G is connected, we may find a group G_1 fitting into a short exact sequence of groups

$$1 \longrightarrow \pi_1(G) \longrightarrow G_1 \xrightarrow{\varphi} G \longrightarrow 1$$

known as the ***universal covering group*** of G. Representations of $\mathfrak{g}$ are in bijection with representations of G_1 rather than G.

Define the 2×2 Hermitian ***Pauli matrices*** as

$$\sigma^0 = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad \sigma^1 = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \quad \sigma^2 = \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix} \quad \sigma^3 = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$$

(Generally, σ^0 is omitted, giving us three Pauli matrices). These obviously span the space $\mathrm{H}(2, \mathbb{C})$ of 2×2 Hermitian matrices, and in fact we have a pair of isomorphisms $\underset{\sim}{-}, \tilde{-} : \mathbb{R}^4 \to \mathrm{H}(2, \mathbb{C})$ defined by $\underset{\sim}{x} = \delta_{\mu\nu} x^\mu \sigma^\nu$, $\tilde{x} = \eta_{\mu\nu} x^\mu \sigma^\nu$. We can computationally verify that $\det \underset{\sim}{x} = \det \tilde{x} = x \cdot x$, and $\tilde{x}\underset{\sim}{x} = \underset{\sim}{x}\tilde{x} = (x \cdot x) I_2$. It follows that, for an arbitrary determinant 1 complex matrix A, the linear map $\varphi(A)(x) = (\underset{\sim}{-})^{-1}(A\underset{\sim}{x}A^\dagger)$ defines a homomorphism $\varphi : \mathrm{SL}\,(2; \mathbb{C}) \to L$; in fact, we can show that it is a surjection $\mathrm{SL}\,(2; \mathbb{C}) \to L_+^\uparrow$ with kernel $\varphi^{-1}(I_4) = \{\pm I_2\} \cong \mathbb{Z}/2\mathbb{Z}$.

Topologically, $L_+^\uparrow$ is equivalent to $\mathbb{R}^3 \times \mathrm{SO}(3)$, and therefore $\pi_1(L_+^\uparrow) = \pi_1(\mathrm{SO}(3)) = \mathbb{Z}/2\mathbb{Z}$. It follows that the homomorphism $\varphi : \mathrm{SL}\,(2; \mathbb{C}) \to L_+^\uparrow$ fits into a short exact sequence

$$1 \to \mathbb{Z}/2\mathbb{Z} \to \mathrm{SL}\,(2; \mathbb{C}) \to L_+^\uparrow \to 1$$

evidencing $\mathrm{SL}\,(2; \mathbb{C})$ as the universal covering group of $L_+^\uparrow$.

Given a Lie group or algebra representation M, a subspace V of $\mathbb{C}^n$ mapped into itself by all $\Pi(g)$ is known as ***invariant***; $\{\vec{0}\}$ and $\mathbb{C}^n$ are trivially invariant, but any representation with no nontrivial invariant subspaces is known as ***irreducible***. Every representation of $\mathrm{SL}\,(2; \mathbb{C})$ decomposes as the direct sum of irreducible representations, i.e. $\Pi(g) = \Pi_1(g) \oplus \Pi_2(g) \oplus \ldots \oplus \Pi_k(g)$ with each $\Pi_j(g)$ an $n_j \times n_j$ matrix, where $\sum_{j=1}^k n_j = n$. We define a pair of representations $\Pi^{(1/2,0)}, \Pi^{(0,1/2)} : \mathrm{SL}\,(2; \mathbb{C}) \to \mathrm{GL}\,(2; \mathbb{C})$ given by

$$\Pi^{(1/2,0)}(A) = A \qquad \Pi^{(0,1/2)}(A) = (A^\dagger)^{-1}$$

For $\mu, \nu \in \{0, 1/2, 1, 3/2, \ldots\}$, we define $\Pi^{(\mu,\nu)} : \mathrm{SL}\,(2; \mathbb{C}) \to \mathrm{GL}\,(4^{\mu+\nu}; \mathbb{C})$ by

$$\Pi^{(\mu,\nu)}(A) = \left(\bigotimes_{i=1}^{2\mu} \Pi^{(1/2,0)}(A) \right) \otimes \left(\bigotimes_{i=1}^{2\nu} \Pi^{(0,1/2)}(A) \right)$$

The $\Pi^{(\mu,\nu)}$ are the irreducible representations of $\mathrm{SL}\,(2; \mathbb{C})$. Every irreducible representation of the Lorentz algebra can be recovered as the pushforward of some $\Pi^{(\mu,\nu)}$, which we denote $\pi^{(\mu,\nu)}$. Under an infinitesimal Lorentz transformation, or an element $g \in \mathfrak{so}(1,3)$, an n-component complex field $\Phi = (\phi_1, \ldots, \phi_n)$ described by a Lorentz covariant theory must experience an infinitesimal change described by a matrix $M(g) \in \mathfrak{gl}(n;\mathbb{C})$, where M is a representation of $\mathfrak{so}(1,3)$ and thus decomposes as $M = \bigoplus_{i=1}^{k} \pi^{(\mu_i,\nu_i)}$. The largest $\mu_i + \nu_i$ is known as the ***spin*** of Φ.

Spinors The Lorentz algebra is a 6-dimensional vector space, with three rotation dimensions and three boost dimensions. It is spanned by the set $J^{\mu\nu} = i(x^\mu \partial^\nu - x^\nu \partial^\mu)$ of tangent vectors (since $J^{\mu\nu} = -J^{\nu\mu}$, there really are only six), and satisfy the commutation relations

$$[J^{\mu\nu}, J^{\rho\sigma}] = i(\eta^{\nu\rho} J^{\mu\sigma} - \eta^{\mu\rho} J^{\nu\sigma} - \eta^{\nu\rho} J^{\mu\rho} + \eta^{\mu\sigma} J^{\nu\rho})$$

Any set of six $n \times n$ matrices $S^{\mu\nu}$ satisfying the same commutation relations (in particular, $[S^{\mu\nu}, S^{\nu\mu}] = 0$, so that $S^{\nu\mu} = -S^{\mu\nu}$) defines a Lie algebra homomorphism $\mathfrak{so}(1,3) \to \mathfrak{gl}(n;\mathbb{C})$, and hence a representation of the Lorentz algebra.

Any set of four $n \times n$ matrices γ^μ γ^μ such that $\gamma^\mu \gamma^\nu + \gamma^\nu \gamma^\mu = 2\eta^{\mu\nu} I_n$ yields a set of matrices $S^{\mu\nu} = \frac{i}{4}[\gamma^\mu, \gamma^\nu]$ satisfying these relations. One such set of ***gamma matrices*** is given in block diagonal form by

$$\gamma_0 = \begin{bmatrix} 0 & I_2 \\ I_2 & 0 \end{bmatrix} \quad \gamma_i = \begin{bmatrix} 0 & \sigma^i \\ -\sigma^i & 0 \end{bmatrix}$$

This yields matrices

$$S^{0i} = -\frac{i}{2} \begin{bmatrix} \sigma^i & 0 \\ 0 & -\sigma^i \end{bmatrix} \quad S^{ij} = \frac{1}{2}\varepsilon^{ijk} \begin{bmatrix} \sigma^k & 0 \\ 0 & \sigma^k \end{bmatrix}$$

and, for a family of scalars $c_{\mu\nu}$, gives the representation $c_{\mu\nu} J^{\mu\nu} \mapsto c_{\mu\nu} S^{\mu\nu}$, known as the ***chiral representation***. This representation decomposes as $\pi^{(1/2,0)} \oplus \pi^{(0,1/2)}$; complex 2-dimensional vector fields transforming according to $\pi^{(1/2,0)}$ and $\pi^{(0,1/2)}$ are known as the ***left-handed*** and ***right-handed Weyl spinors***, whereas a 4-dimensional complex vector field transforming according to $\pi^{(1/2,0)} \oplus \pi^{(0,1/2)}$ is known as a ***Dirac spinor***.

The most notable property of Dirac spinors is their behavior under rotations: consider for instance the action of an infinitesimal θ degree rotation in the xy-plane on a Dirac spinor, which we obtain by exponentiating its representation:

$$\begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & \theta i & 0 \\ 0 & -\theta i & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \mapsto \begin{bmatrix} e^{i\theta/2} & 0 & 0 & 0 \\ 0 & e^{-i\theta/2} & 0 & 0 \\ 0 & 0 & e^{i\theta/2} & 0 \\ 0 & 0 & 0 & e^{-i\theta/2} \end{bmatrix}$$

Under a full $360^\circ = 2\pi$ revolution, a Dirac spinor doesn't return to its original state, but picks up a minus sign; it takes a $720^\circ = 4\pi$ rotation to return the spinor to its original state. Dirac spinor fields are spin $1/2$ fields, as opposed to scalar fields, which transform under the trivial representation of the Lorentz algebra and are hence spin 0. In general, a spin $n > 0$ field requires a $2\pi/n$ degree rotation to return to its original state; spin 0 fields are invariant under any rotation.

B.4.2 Spin Structures

Spin structures on Riemannian manifolds offer a way to abstractly study spinorial structures such as those introduced in B.4.1.

The Spin Group Recall from there that, given a connected but not simply connected Lie group G, Lie algebra representations of the Lie algebra $\mathfrak{g}$ of G are *not* in bijection with Lie group representations of G, but instead in bijection with Lie group representations of the ***universal covering group*** of G, or the simply connected Lie group G_1 fitting into a short exact sequence of groups

$$1 \longrightarrow \pi_1(G) \longrightarrow G_1 \longrightarrow G \longrightarrow 1$$

Take $G = \mathrm{SO}(n)$, the special orthogonal group in n dimensions. This Lie group consists of all orthogonal matrices $A \in \mathbb{R}^{n\times n}$ such that $A^T A = AA^T = I_n$ and $\det A = 1$. Its Lie algebra $\mathfrak{so}(n)$ consists of the skew-symmetric matrices $A = -A^T$, with Lie bracket given by the commutator $[A, B] = AB - BA$. $\pi_1(\mathrm{SO}(n)) = \mathbb{Z}_2$, so $\mathrm{SO}(n)$ is not simply connected[1]; its universal covering group, or equivalently its double cover, is known as the ***spin group*** $\mathrm{Spin}(n)$.

[1]Consider the loop of θ-degree rotations around a fixed axis, $\theta = [0, 2\pi]$.

Spin(n) is constructed as follows: let V be the vector space $\mathbb{R}^n$ equipped with quadratic form $q : \mathbb{R}^n \to \mathbb{R}, q(v) = \langle v, v\rangle$, and let TV be the *tensor algebra* on V, or the vector space

$$\begin{aligned} TV &= \mathbb{R} \oplus V \oplus (V \otimes V) \oplus (V \otimes V \otimes V) \oplus \dots \\ &= \bigoplus_{k=0}^{\infty} V^{\otimes k} \end{aligned}$$

Elements of this algebra look like formal sums $c + v_0 + v_1 \otimes v_2 + \dots$. Define the *Clifford algebra* $C\ell(V, q)$ to be the quotient algebra given by identifying any term of the form $v \otimes v$ with the real number $q(v)$. This is naturally a graded vector space, with $C\ell^k(V, q)$ being the vector space of tensors of k vectors (after reduction)[2].

We define the ***pin group*** Pin(n) to be the set of all formal sums of strings of the form $v_1 v_2 \dots v_n$, where $q(v_1) = q(v_2) = \dots = q(v_n) = 1$. Using the canonical basis $\{e_1, \dots, e_n\}$ of $\mathbb{R}^n$, we may express an element of Pin(n) as a formal sum of even-length strings of the e_i. The ***spin group*** Spin(n) is the set of all formal sums of *even* strings in Pin(n). This group satisfies an anticommutation relation: $2 = (e_i + e_j)(e_i + e_j) = e_i e_i + e_i e_j + e_j e_i + e_j e_j = 2 + e_i e_j + e_j e_i$, implying that $e_i e_j = -e_j e_i$.

Frame Bundles Given a smooth manifold M, the tangent bundle $TM = \amalg_{x \in M} T_x M$, equipped with its canonical smooth structure, yields a vector bundle $TM \to M$ sending a point x and tangent vector v^i to x to x alone. Define a ***frame*** at a point $x \in M$ to be an ordered linear basis of $T_x M$, and let F_x denote the set of all frames at x. The bundle $FM = \amalg_{x \in M} F_x$ has a natural $\mathrm{GL}(n; \mathbb{R})$ action; as invertible matrices freely and transitively send frames to frames, $\pi : FM \to M$ is a principal $\mathrm{GL}(n; \mathbb{R})$-bundle.

If M has a Riemannian metric g_{ij}, then we may consider the set of orthonormal frames at x, or ordered bases $(v^i_1, \dots, v^i_n)$ such that $v^i_k v^j_\ell g_{ij} = \delta_{k\ell}$. This yields the orthogonal frame bundle $\pi_O : F_O M \to M$, which is a principal n-bundle. If furthermore M is orientable, such that we may distinguish between positively and negatively oriented orthonormal frames, then we may form the special orthogonal frame bundle $\pi_{SO} : F_{SO} M \to M$, which has as its fiber at x the set

[2] Note that the multiplication inherited from TV does *not* respect this grading.

of all positively oriented orthonormal frames at x; this forms a principal $\mathrm{SO}(n)$-bundle.

Let $\rho : \mathrm{Spin}(n) \to \mathrm{SO}(n)$ denote the double covering defining $\mathrm{Spin}(n)$, and let M be an oriented Riemannian manifold, with notation as above. A ***spin structure*** on M is a principal $\mathrm{Spin}(n)$-bundle $\pi_{Sp} : F_{Sp}M \to M$ equipped with a 2-fold covering $\phi : F_{Sp}M \to F_{SO}M$ of bundles, such that $\pi_{SO} \circ \phi = \pi_{Sp}$ and, for all $x \in \mathrm{Spin}(n)$, $\phi(x \cdot f) = \rho(x) \cdot \phi(f)$. Such a structure may not necessarily exist; if one does, M is said to be a ***spin manifold***[3].

Spin Representations Suppose we have a principal G-bundle $\pi : E \to M$, where M is a smooth manifold, and a continuous homomorphism $\rho : G \to \mathrm{Diff}(M')$, where $\mathrm{Diff}(M')$ is the group of diffeomorphisms on the smooth manifold M' with the C^∞-topology. G then has a free action on $E \times M$, given by $g \cdot (e, x) = (g \cdot e, \rho(g)(x))$. Taking the quotient of $E \times M$ by its G-orbits yields a fiber bundle $\pi \circ \pi_1 : E \times_\rho M' \to M$ with typical fiber M', known as the ***associated bundle*** to ρ.

In particular, if we have a continuous representation ρ of G on a k-vector space V, or a continuous homomorphism $\rho : G \to \mathrm{GL}(V) \subset \mathrm{Diff}(V)$, the fiber bundle $E \times_\rho V \to M$ has typical fiber V, yielding a k-vector bundle. If G has a canonical representation, such as $\mathrm{SO}(n)$'s representation as determinant-1 orthogonal matrices over $\mathbb{R}^n$, we can turn principal G-bundles into vector bundles in a canonical way.

We initially motivated the construction of $\mathrm{Spin}(n)$ as the universal covering group of G, the Lie group whose representations are in bijection with Lie algebra representations of $\mathfrak{so}(n)$; these representations aren't in bijection with those of $\mathrm{SO}(n)$ due to the latter's not being simply connected. Hence, there are representations of $\mathfrak{so}(n)$ which are associated to representations of $\mathrm{Spin}(n)$ but *not* to representations of $\mathrm{SO}(n)$. A representation $\rho : \mathrm{Spin}(n) \to \mathrm{GL}(V)$ is such a representation when $-1 \notin \ker \rho$, and is known as a ***spin representation***. By virtue of our construction, we can generalize the constructions of the spin and special orthogonal groups to any vector space V equipped with a quadratic form q; in particular, starting with an $\mathbb{R}$-vector space V, we may construct its complexification $V_{\mathbb{C}} := V \otimes_{\mathbb{R}} \mathbb{C}$, a $\mathbb{C}$-vector space, yielding the

[3] M admits a spin structure iff the second Stiefel-Whitney class $w_2(M) \in H^2(M, \mathbb{Z}_2)$ vanishes, in which case spin structures on M are in correspondence with elements of $H^1(M, \mathbb{Z}_2)$.

Lie groups $\mathrm{Spin}(n, \mathbb{C})$ and $\mathrm{SO}(n, \mathbb{C})$. If we instead keep $V = \mathbb{R}^n$ and define the quadratic form by $q = \mathrm{diag}(1, \ldots, 1, -1, \ldots, -1)$, with a copies of 1 and $b = n - a$ copies of -1, we obtain the Lie groups $\mathrm{Spin}(a, b)$ and $\mathrm{SO}(a, b)$.

Spinor Bundles Recall from B.4.1 the definition of the Pauli matrices:

$$\sigma^0 = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad \sigma^1 = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \quad \sigma^2 = \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix} \quad \sigma^3 = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$$

We will use these to construct a complex representation of $\mathrm{Spin}(n, \mathbb{C})$. Suppose that $n = 2k$ for some $k \in \mathbb{N}$. There is then an algebra isomorphism of $C\ell(\mathbb{C}^n, q)$ with $\mathbb{C}^{2^k \times 2^k}$ sending the basis e_j, $j \in \{1, \ldots, n\}$, to:

$$(j \text{ even}) \quad (\sigma^0)^{\otimes(k-j/2)} \otimes (i\sigma^1) \otimes (\sigma^2)^{\otimes(j/2-1)}$$

$$(j \text{ odd}) \quad (\sigma^0)^{\otimes(k-(j+1)/2)} \otimes (i\sigma^3) \otimes (\sigma^2)^{\otimes((j-1)/2)}$$

(This construction is given in [Friedrich, 2000]). If $n = 2k + 1$ is odd instead of even, a similar isomorphism yields $C\ell(\mathbb{C}^n, q) \cong \mathbb{C}^{2^k \times 2^k} \oplus \mathbb{C}^{2^k \times 2^k}$. Hence, if we define the vector space $\Delta_n = \mathbb{C}^{2^k}$, for $n = 2k, 2k + 1$, we obtain $C\ell(\mathbb{C}^n, q) \cong \mathrm{End}(\Delta_n)$ for n even, $\mathrm{End}(\Delta_n) \oplus \mathrm{End}(\Delta_n)$ for n odd.

For n even, this isomorphism yields a faithful (and therefore spin) representation of $\mathrm{Spin}(n, \mathbb{C})$ over the vector space $\mathbb{C}^{2^{n/2}}$; as there is an inclusion $\mathrm{Spin}(n) \to \mathrm{Spin}(n, \mathbb{C})$, this yields a representation on $\mathrm{Spin}(n)$ as well, which we will denote by κ_n. By our previous argument, given a spin manifold M of even dimension, we can construct a complex vector bundle $F_{Sp}M \times_{\kappa_n} \Delta_n \to M$ known as the ***spinor bundle*** SM. A section of this bundle is known as a ***spinor field***; we may now call the elements of Δ_n themselves ***spinors***.

The Dirac Operator The bases $e_1, \ldots, e_n$ of $\mathrm{Spin}(n)$ will be sent to a set of fixed $2^{n/2} \times 2^{n/2}$ complex matrices by the above isomorphism, which matrices which we denote $\gamma^1, \ldots, \gamma^n$; these generalize the gamma matrices[4]. For convenience, we define an "extra" gamma matrix $\gamma^{n+1} = i^{n(n+1)/2}\gamma^1 \dots \gamma^n$.

Recall the definition of the Christoffel symbols on the Riemannian manifold (M, g) as

$$\Gamma^i{}_{jk} := \frac{1}{2} g^{i\ell} \left(\partial_k g_{\ell j} + \partial_j g_{\ell k} - \partial_\ell g_{jk} \right)$$

and the definition of the Levi-Civita connection as

$$\nabla_i v^j = \partial_i v^j + \Gamma^j{}_{ik} v^k$$

The Levi-Civita connection naturally lifts to a connection 1-form ω on $F_{SO}M$, and therefore to a covariant derivative $\nabla_i : \Gamma(SM) \to \Gamma(T^*M \otimes SM)$ on the spinor bundle. (See [Lawson and Michelsohn, 2016], II.4, for details). Given a choice of coordinates $\{e^1, \dots, e^n\}$, we define the ***Dirac operator*** $D : \Gamma(S) \to \Gamma(S)$ by $D\psi = \sum_i \gamma^i \nabla_i \psi$; using Feynman's slash notation, we may write this as $D = \not\nabla$.

B.4.3 Quantization of Classical Field Theories

The setup for studying classical field theories in Minkowski space with metric $\eta^{\mu\nu} = \mathrm{diag}(+1, -1, -1, -1)$ is as follows:

1. Obtain the *Lagrangian density* of the theory, typically by subtracting its potential energy term from its kinetic energy term. As an example, we will work with the Lagrangian of a free scalar theory,
$$\mathcal{L} = \frac{1}{2}\partial_\mu \phi \partial^\mu \phi - \frac{1}{2} m^2 \phi^2$$
2. Plug $\mathcal{L}$ into the *Euler-Lagrange equations*,
$$\partial_\mu \left(\frac{\partial \mathcal{L}}{\partial \left(\partial_\mu \phi \right)} \right) - \frac{\partial \mathcal{L}}{\partial \phi} = 0$$
and then simplify to obtain constraints on the fields (equations of motion). We iterate this process over each free variable ϕ in our theory; thankfully, the free scalar theory only has

[4]The isomorphism $C\ell(\mathbb{C}^n, q) \cong \mathrm{End}(\Delta^n)$ defined above is one of many possible isomorphisms, each of which defines a different set of gamma matrices.

one free variable. Plugging the above $\mathcal{L}$ into the Euler-Lagrange equations, we calculate

$$\frac{\partial(\partial_\mu\phi\partial^\mu\phi)}{\partial(\partial_\mu\phi)} = \frac{\partial(\eta^{\mu\nu}\partial_\mu\phi\partial_\nu\phi)}{\partial(\partial_\mu\phi)} = \eta^{\mu\nu}\frac{\partial(\partial_\mu\phi)}{\partial(\partial_\mu\phi)}\partial_\nu\phi + \eta^{\mu\nu}\partial_\mu\phi\frac{\partial(\partial_\nu\phi)}{\partial(\partial_\mu\phi)}$$
$$= \partial^\mu\phi + (\partial^\nu\phi)\delta^\mu_\nu = 2\partial^\mu\phi$$

and hence obtain the equation

$$\partial_\mu\partial^\mu\phi + m^2\phi = 0$$

Writing $\partial_\mu\partial^\mu$ as ∂^2, this becomes the ***Klein-Gordon equation***

$$(\partial^2 + m^2)\phi = 0$$

3. Look for solutions to the equations of motion. In the case of the free scalar theory, whose single equation of motion is the Klein-Gordon equation, the equations of motion look like plane-waves,

$$\phi(x_\mu) = e^{-ip^\mu x_\mu}$$

with $p_\mu = (\omega, \vec{k})$ consisting of an angular frequency $p_0 = \omega$ and wavenumber $(p_1, p_2, p_3) = \vec{k}$ such that $p_\mu p^\mu = \omega^2 - k^2 = m^2$.

4. If we want more information, we may calculate the *Hamiltonian density* of the theory. In a theory with n free variables $\phi_1, \ldots, \phi_n$, this is first done by associating to each ϕ_i a *conjugate momentum*

$$\Pi^\mu_i = \frac{\partial\mathcal{L}}{\partial(\partial_\mu\phi_i)}$$

and then deriving the Hamiltonian as

$$\mathcal{H} = \left(\sum_{i=1}^{n}\Pi^0_i\partial_0\phi_i\right) - \mathcal{L}$$

For our free scalar theory, we have

$$\Pi^\mu = \frac{\partial\mathcal{L}}{\partial(\partial_\mu\phi)} = \partial^\mu\phi$$

and hence

$$\mathcal{H} = \partial^0\phi\partial_0\phi - \mathcal{L} = \frac{1}{2}\partial^0\phi\partial_0\phi - \frac{1}{2}\sum_{i=1}^{3}\partial_i\phi\partial^i\phi + \frac{1}{2}m^2\phi^2$$

$$= \frac{1}{2}\left[\left(\frac{\partial \phi}{\partial t}\right)^2 + (\nabla\phi)^2 + m^2\phi^2\right]$$

5. Alternatively, we can define the *stress-energy tensor* $T^{\mu}{}_{\nu}$ of the theory, given by

$$T^{\mu}{}_{\nu} = \frac{\partial \mathcal{L}}{\partial\left(\partial_{\mu}\phi\right)}\partial_{\nu}\phi - \mathcal{L}\delta^{\mu}_{\nu}$$

This gives rise to four conserved quantities,

$$P^i = \int T^{0i}\, d^3x$$

For the Klein-Gordon Lagrangian, we obtain a stress energy tensor of

$$T^{\mu}{}_{\nu} = \partial^{\mu}\phi\partial_{\nu}\phi - \frac{1}{2}\delta^{\mu}_{\nu}\left(\partial_{\rho}\phi\partial^{\rho}\phi - m^2\phi^2\right)$$

For $\mu = \nu = 0$, we reclaim the Hamiltonian, and for $\mu = 0, \nu \neq 0$, we obtain

$$T^{0i} = \sum_{j=1}^{3}\eta^{ij}T^{0}{}_{j} = -T^{0}{}_{i} = \partial^0\phi\partial_i\phi$$

This gives us a set of tools for the analysis of classical fields.

Another example is given by classical electromagnetism. Setting $c = 1$, define the electromagnetic four-potential A_{μ} to have as its timelike component the electric potential ϕ and as its spacelike components the magnetic vector potential $\vec{A}$. The exterior derivative of this one-form is known as the ***electromagnetic tensor*** $F_{\mu\nu}$, and as a matrix looks like

$$\begin{bmatrix} 0 & E_x & E_y & E_z \\ -E_x & 0 & -B_z & B_y \\ -E_y & B_z & 0 & -B_x \\ -E_z & -B_y & B_x & 0 \end{bmatrix}$$

The Lagrangian of classical electromagnetism is given by

$$\mathcal{L} = \overbrace{-\frac{1}{4\mu_0}F^{\mu\nu}F_{\mu\nu}}^{\text{field}} - \overbrace{A_{\mu}J^{\mu}}^{\text{source}}$$

where $J^\mu = (\rho, \vec{j})$ is a four-current. With some effort, we may show that the Euler-Lagrange equations read

$$\partial_\mu F^{\mu\nu} = \mu_0 J^\nu$$

For $\nu = 0$ this reduces to $\nabla \cdot \vec{E} = \mu_0 \rho = \rho/\varepsilon_0$, Gauss's law. For $\nu = 1, 2, 3$, we obtain $\nabla \times \vec{B} = \mu_0 \vec{j} + \frac{\partial \vec{E}}{\partial t}$, or Ampere's law.

Canonical Quantization To quantize a classical field theory with position variables $\phi_1, \ldots, \phi_n$ and conjugate momenta $\Pi_1^\mu, \ldots, \Pi_n^\mu$, we turn the position and momentum variables into operators $\widehat{\phi}_1, \ldots, \widehat{\phi}_n, \widehat{\Pi}_1^\mu, \ldots, \widehat{\Pi}_n^\mu$, and impose the ***equal-time commutation relations***

$$[\widehat{\phi}_i(t, \vec{x}), \widehat{\Pi}_j^0(t, \vec{y})] = i\delta^{(3)}(\vec{x} - \vec{y})\delta_{ij}$$

with all commutators among $\widehat{\phi}$s and among $\widehat{\Pi}$s being zero. The Hamiltonian $\mathcal{H}$, being a function of ϕ and Π, becomes an operator $\widehat{\mathcal{H}}$ as well, as does $H = \int \mathcal{H}\, d^3x$.

Fundamentally, quantizing $\widehat{H}$ gives it a quantized spectrum. In the case where we have one variable ϕ with no self-interactions (i.e., the Euler-Lagrange equations are linear in ϕ), we have a lowest-energy ***vacuum state*** $|0\rangle$ to which we can add a "particle" with momentum $\vec{p}$ via the ***creation operator*** $\widehat{a}_{\vec{p}}^\dagger$, and remove a particle with momentum $\vec{q}$ via the ***annihilation operator*** $\widehat{a}_{\vec{q}}$.

Additional variables will define additional pairs of annihilation and creation operators, generally denoted $(\widehat{b}_{\vec{p}}^\dagger, \widehat{b}_{\vec{q}}), (\widehat{c}_{\vec{p}}^\dagger, \widehat{c}_{\vec{q}})$, and so on. We may reconstruct $\widehat{\phi}$ from the annihilation and creation operators by means of a ***mode expansion*** which, in the case of the Klein-Gordon field, is given by

$$\widehat{\phi}(t, \vec{x}) = \int \frac{d\vec{p}}{(2\pi)^{3/2}} \frac{1}{\sqrt{2E_{\vec{p}}}} \left(\widehat{a}_{\vec{p}} e^{-ip\cdot x} + \widehat{a}_{\vec{p}}^\dagger e^{ip\cdot x} \right)$$

where $p \cdot x = (t, \vec{p}) \cdot (t, \vec{x}) = t^2 - \vec{p} \cdot \vec{x}$, and $E_{\vec{p}} = \sqrt{\vec{p}^2 + m^2}$. We interpret $\widehat{\phi}(x)$ as creating a particle at position x. We define the state $|\vec{p}\rangle$ consisting of one particle with momentum $\vec{p}$ by $|\vec{p}\rangle = \widehat{a}_{\vec{p}}^\dagger |0\rangle$, so that $\langle \vec{p} | \vec{q} \rangle = \delta^{(3)}(\vec{p} - \vec{q})$.

In general, though, our theory will not be free from self-interactions, so we have to replace the vacuum state $|0\rangle$ with a more mysterious ground state $|\Omega\rangle$. While acting on $|0\rangle$ with $\widehat{a}_{\vec{p}}^\dagger$ yields a state with a single particle of momentum $\vec{p}$, acting on $|\Omega\rangle$ with $\widehat{a}_{\vec{p}}^\dagger$ guarantees nothing but a superposition of particles whose momenta sum to $\vec{p}$.

The dynamics of a quantum field theory can be analyzed via its ***correlation functions***, numbers of the form

$$\langle \Omega | \widehat{\phi}(x_1) \ldots \widehat{\phi}(x_n) \phi(y_1)^\dagger \ldots \widehat{\phi}(y_n)^\dagger | \Omega \rangle$$

which express the probability for particles created at positions $y_1, \ldots, y_n$ to travel to positions $x_1, \ldots, x_n$. To evaluate these, we need some additional machinery.

Green's Functions Given a linear differential operator L, e.g. $Lx(t) = m\frac{d^2}{dt^2}x(t) + cx(t)$, we define the ***Green's function*** of L to be a function $G(t,u)$ such that $LG(t,u) = \delta(t-u)$. Given a differential equation $Lx(t) = f(t)$, we may use G to solve for x as

$$x(t) = \int G(t,u) f(u)\, du$$

noting that

$$Lx(t) = L\left(\int G(t,u) f(u)\, du\right) = \int LG(t,u) f(u)\, du = \int \delta(t-u) f(u)\, du = f(t)$$

Normal and Time Ordering When we have a series of scalar fields $\widehat{\phi}(x_1), \widehat{\phi}(x_n)$ being multiplied, we define the ***time-ordering symbol*** T by $T\widehat{\phi}(x_1) \cdot \ldots \cdot \widehat{\phi}(x_n) = \widehat{\phi}(x_{i_1}) \cdot \ldots \cdot \widehat{\phi}(x_{i_n})$, where the x_{i_j} are such that $x^0_{i_j} \leq x^0_{i_k}$ iff $j \geq k$; T simply orders the scalar fields from latest to earliest in time. Similarly, the ***normal ordering symbol*** N puts all creation operators on the left, e.g. as $N\widehat{a}_{\vec{p}}\widehat{a}^\dagger_{\vec{q}}\widehat{a}_{\vec{r}} = \widehat{a}^\dagger_{\vec{q}}\widehat{a}_{\vec{p}}\widehat{a}_{\vec{r}}$ (note that $\widehat{a}_{\vec{p}}$ and $\widehat{a}_{\vec{r}}$ commute, so it doesn't matter what order they're placed in). We define the ***contraction*** of two operators as

$$\widehat{A}\widehat{B} = \langle 0 | T\widehat{A}\widehat{B} | 0 \rangle$$

So, for instance,

$$\widehat{A}\widehat{B}\widehat{C}\widehat{D}\widehat{E}\widehat{F} = \widehat{A}\widehat{E}\langle 0 | T\widehat{B}\widehat{D} | 0 \rangle \langle 0 | T\widehat{C}\widehat{F} | 0 \rangle$$

Wick's theorem states that applying T to a given string of operators is equivalent to applying N to that string plus all of its possible contractions. For instance,

$$T\widehat{A}\widehat{B}\widehat{C}\widehat{D} = N\widehat{A}\widehat{B}\widehat{C}\widehat{D} + \langle 0 | T\widehat{A}\widehat{B} | 0 \rangle N\widehat{C}\widehat{D} + \langle 0 | T\widehat{A}\widehat{C} | 0 \rangle N\widehat{B}\widehat{D} + \ldots + \langle 0 | T\widehat{A}\widehat{B} | 0 \rangle \langle 0 | T\widehat{C}\widehat{D} | 0 \rangle + \ldots$$

where we first list the term with zero contractions, then those with one contraction, then with two. As a particular case, this allows us to evaluate terms of the form $\langle 0|T\widehat{A}\widehat{B}\widehat{C}\ldots|0\rangle$: since

Since taking $\langle 0|N\widehat{A}\widehat{B}\ldots|0\rangle$ always yields zero, we see that this simplifies to the sum of all terms which contract *all* elements.

Propagators We define the ***Feynman propagator*** by

$$G(x,y) = \langle \Omega|T\widehat{\phi}(x)\widehat{\phi}^\dagger(y)|\Omega\rangle$$

The interpretation of this is as follows: starting from the ground state $|\Omega\rangle$, create a particle at spacetime point y, wait a while, and then attempt to annihilate it at spacetime point x; the extent to which the state no longer resembles $|\Omega\rangle$ is given by taking its product against $\langle\Omega|$. When we're in a free theory with $|\Omega\rangle = |0\rangle$, $G(x,y)$ is known as the ***free propagator***

$$\Delta(x,y) = \langle 0|T\widehat{\phi}(x)\widehat{\phi}^\dagger(y)|0\rangle$$

Perturbation Expansions To see this machinery in action, we need a non-free, interacting field theory. One such theory is given by the "ϕ^4" theory, with Lagrangian

$$\mathcal{L} = \frac{1}{2}\partial_\mu\phi\partial^\mu\phi - \frac{1}{2}m^2\phi^2 - \frac{\lambda}{4!}\phi^4$$

This is similar to the Klein-Gordon Lagrangian, except for the ϕ^4 term which induces a non-linear Euler-Lagrange equation

$$(\partial^2 + m^2)\phi = -\frac{\lambda}{3!}\phi^3$$

The quantized Hamiltonian $\widehat{\mathcal{H}}$ is similar to that of the Klein-Gordon Hamiltonian, but with an extra "interaction" term $\frac{\lambda}{4!}\widehat{\phi}^4$. We correspondingly decompose $\widehat{\mathcal{H}}$ as $\widehat{\mathcal{H}}_0 + \widehat{\mathcal{H}}'$, where $\widehat{\mathcal{H}}_0$ is the Klein-Gordon Hamiltonian and $\widehat{\mathcal{H}}'$ is this interaction term. When λ is small, we can approximate the evolution of an arbitrary operator $\widehat{O}$ as $\widehat{O}_I(t) = e^{i\widehat{H}_0 t}\widehat{O}e^{-i\widehat{H}_0 t}$, where the subscript I denotes that we're working in the "interaction picture". We define the *S-**matrix*** by

$$\widehat{S} = T\left[e^{-i\int_{-\infty}^{\infty}\widehat{\mathcal{H}}_I\,d^4x}\right]$$

Since this is generally insoluble, we expand in powers of $-i\int_{-\infty}^{\infty}\widehat{\mathcal{H}}_I\,d^4x$:

$$\begin{aligned}\widehat{S} &= T\left[1 - i\int \widehat{\mathcal{H}}_I(x)\,d^4x + \frac{(-i)^2}{2}\int \widehat{\mathcal{H}}_I(x)\widehat{\mathcal{H}}_I(y)\,d^4x\,d^4y + \dots\right] \\ &= T\left[1 + \sum_{n=1}^{\infty}\frac{(-i)^n}{n!}\int\prod_{m=1}^{n}\widehat{\mathcal{H}}_I(x_m)\,d^4x_m\right]\end{aligned}$$

We can analyze the probability that a particle with momentum $\vec{p}$ turns into a particle with momentum $\vec{q}$ by plugging the two probabilities into the S-matrix: for instance, in the ϕ^4 theory, we obtain

$$\langle\vec{q}|\widehat{S}|\vec{p}\rangle \propto \langle 0|\widehat{a}_{\vec{q}}\widehat{S}\widehat{a}^{\dagger}_{\vec{p}}|0\rangle =$$

$$T\left[\langle 0|\widehat{a}_{\vec{q}}\widehat{a}^{\dagger}_{\vec{p}}|0\rangle + (-i)\left(\frac{\lambda}{4!}\right)\int\langle 0|\widehat{a}_{\vec{q}}\widehat{\phi}(x)^4\widehat{a}^{\dagger}_{\vec{p}}|0\rangle\,d^4x + \frac{(-i)^2}{2}\left(\frac{\lambda}{4!}\right)^2\int\langle 0|\widehat{a}_{\vec{q}}\widehat{\phi}(x)^4\phi(y)^4\widehat{a}^{\dagger}_{\vec{p}}|0\rangle\,d^4x\,d^4y + \dots\right]$$

$$= \langle 0|\widehat{a}_{\vec{q}}\widehat{a}^{\dagger}_{\vec{p}}|0\rangle + \sum_{n=1}^{\infty}\frac{(-i)^n}{n!}\left(\frac{\lambda}{4!}\right)^n\int\langle 0|T\left[\widehat{a}_{\vec{q}}\left(\prod_{m=1}^{n}\widehat{\phi}(x_m)^4\right)\widehat{a}^{\dagger}_{\vec{p}}\right]|0\rangle\prod_{m=1}^{n}d^4x_m$$

Thus, the higher-order corrections to $\langle 0|\widehat{a}_{\vec{q}}\widehat{S}\widehat{a}^{\dagger}_{\vec{p}}|0\rangle$ arise in powers proportional to λ.

Let's analyze the first-order correction, given by

$$\frac{-i\lambda}{4!}\int\langle 0|T\,\widehat{a}_{\vec{q}}\widehat{\phi}(x)\widehat{\phi}(x)\widehat{\phi}(x)\widehat{\phi}(x)\widehat{a}^{\dagger}_{\vec{p}}|0\rangle\,d^4x$$

As stated above, the integrand can be reduced to the sum of all total contractions over its six members. Given $2n$ operators, there are $\frac{(2n)!}{2^n n!}$ distinguishable ways to contract all operators (i.e., form n pairs); $2n = 6$ here, there are 15 terms to consider. In each of these, either the annihilation and creation operators have been contracted with one another, or they have not. The cases in which they have number $\frac{4!}{2^2\cdot 2!} = 3$, and the cases in which they have not, so that each one is contracted with a $\widehat{\phi}(x)$, number $15 - 3 = 12$. The three terms are of the form $\langle 0|\widehat{a}_{\vec{q}}\widehat{a}^{\dagger}_{\vec{p}}|0\rangle = \delta^{(3)}(\vec{q}-\vec{p})$, and we may also calculate $\langle 0|\widehat{\phi}(x)\widehat{a}^{\dagger}_{\vec{p}}|0\rangle = \frac{1}{(2\pi)^{3/2}}\frac{1}{\sqrt{2E_{\vec{p}}}}e^{-ip\cdot x}$, $\langle 0|\widehat{a}_{\vec{q}}\widehat{\phi}(x)|0\rangle = \frac{1}{(2\pi)^{3/2}}\frac{1}{\sqrt{2E_{\vec{q}}}}e^{iq\cdot x}$.

B.4.4 The Dirac Field

While Dirac spinors are four-component vectors, they will be treated analogously to the scalars seen in previous field theories: we will generally not give them indices. Consequently, four-component vectors of four-component vectors, or 4×4 matrices, will have one index. To refer to the space-like components, or the in the case of matrices the latter three components, though, we may still use vector notation (or, in the case of ∂, $\nabla = (\partial^1, \partial^2, \partial^3)$). For a four-component object x_μ, we write the contraction $\gamma^\mu x_\mu$ as $\not{x}$; note that $\not{x}^2 = \gamma^\mu \gamma^\nu x_\mu x_\nu = \frac{1}{2}(\gamma^\mu \gamma^\nu + \gamma^\nu \gamma^\mu) x_\mu x_\nu$ (because we are summing over all μ, ν) $= \eta^{\mu\nu} x_\mu x_\nu = x^2$.

A ***Dirac field*** is a Dirac spinor field ψ with Lagrangian

$$\overline{\psi}(i\not{\partial} - m)\psi = 0$$

where $\overline{\psi} = \psi^\dagger \gamma^0$, $m = mI_4$, and ∂_μ acts on ψ coordinate-wise. The Euler-Lagrange equation for ψ yields the ***Dirac equation***

$$(i\not{\partial} - m)\psi = 0$$

where $m = mI_4$. It follows that $(-i\not{\partial} - m)(i\not{\partial} - m)\psi = (\not{\partial}^2 + m^2)\psi = (\partial^2 + m^2)\psi = 0$, so that the Dirac equation implies the Klein-Gordon equation in each coordinate. The Hamiltonian is given by $\mathcal{H} = -\overline{\psi}(i\vec{\gamma} \cdot \nabla - m)\psi$, so the conjugate momentum of ψ is given by $\Pi^\mu_\psi = i\overline{\psi}\gamma^\mu$ and the conjugate momentum of $\overline{\psi}$ is given by $\Pi^\mu_{\overline{\psi}} = 0$.

Splitting ψ into left-handed and right-handed Weyl spinor fields as $\psi = (\psi_L, \psi_R)$, or equivalently by separating it into eigenvalues of the ***chirality operator*** $\gamma^5 = i\gamma^0\gamma^1\gamma^2\gamma^3 = \begin{bmatrix} -I_2 & 0 \\ 0 & I_2 \end{bmatrix}$, we see that the mass operator leaves ψ_L and ψ_R in their place, whereas gamma operators switch them. In general, this causes the two fields to interact with one another, but when $m = 0$, they do not, and the Dirac equation splits into two separate equations known as the ***Weyl equations***:

$$i(\partial_0 - \vec{\sigma} \cdot \nabla)\psi_L = 0 \qquad i(\partial_0 + \vec{\sigma} \cdot \nabla)\psi_R = 0$$

The solutions to the Dirac equation are given by waves of the form $\psi(x) = \begin{bmatrix} \xi\sqrt{p\cdot\sigma} \\ \xi\sqrt{p\cdot\sigma} \end{bmatrix} e^{-ip\cdot x}$ for positive energy, and $\psi(x) = \begin{bmatrix} \eta\sqrt{p\cdot\sigma} \\ -\eta\sqrt{p\cdot\sigma} \end{bmatrix} e^{ip\cdot x}$ for negative energy. The ξ and η forming the spinors $u(p) = \begin{bmatrix} \xi\sqrt{p\cdot\sigma} \\ \xi\sqrt{p\cdot\sigma} \end{bmatrix}$ and $v(p) = \begin{bmatrix} \eta\sqrt{p\cdot\sigma} \\ -\eta\sqrt{p\cdot\sigma} \end{bmatrix}$ are arbitrary, so we choose to normalize, setting $\xi^\dagger\xi = \eta^\dagger\eta = 1$. We write u^i for ξ^i, $i = 1, 2$, and likewise for v^i. We can write down some useful properties of the u^i and v^i: $u^\dagger(p)u(p) = v^\dagger(p)v(p) = 2E_{\vec{p}}$, $\sum_j u^j(p)\overline{u}^j(p) = \gamma \cdot p + m$, $\sum_j v^j(p)\overline{v}^j(p) = \gamma \cdot p - m$.

Quantization To quantize the Dirac field, we can *not* impose the equal-time commutation relation $[\psi(x), i\psi^\dagger(y)] = i\delta^{(4)}(x-y)$. The particles described by any field with half-integer spin are *fermions*, meaning that interchanging the position of any two fermions adds a negative sign to the state of the field. In particular, any state with two fermions occupying the same position in spacetime must be zero. This is in contrast to particles described by integer spin fields, such as the spin 0 Klein-Gordon equation, which can be stacked on top of one another indefinitely; these particles are known as *bosons*. Hence, we impose equal-time *anti*commutation relations,

$$\{\widehat{\psi}_j(\vec{x}), i\widehat{\psi}_j^\dagger(\vec{y})\} = i\delta^{(3)}(\vec{x}-\vec{y})$$

where j indexes the components of each spinor.

The mode expansions for $\widehat{\psi}$ and $\widehat{\overline{\psi}}$ can be given as

$$\widehat{\psi}(x) = \int \frac{d^3p}{(2\pi)^{3/2}} \frac{1}{\sqrt{2E_{\vec{p}}}} \sum_{j=1}^{2} u^j(p)\widehat{a}_{j\vec{p}}e^{-ip\cdot x} + v^j(p)\widehat{b}^\dagger_{j\vec{p}}e^{ip\cdot x}$$

$$\widehat{\overline{\psi}}(x) = \int \frac{d^3p}{(2\pi)^{3/2}} \frac{1}{\sqrt{2E_{\vec{p}}}} \sum_{j=1}^{2} \overline{u}^j(p)\widehat{a}^\dagger_{j\vec{p}}e^{ip\cdot x} + \overline{v}^j(p)\widehat{b}_{j\vec{p}}e^{-ip\cdot x}$$

The interpretation is that $\widehat{a}^\dagger_{s\vec{p}}$ creates a fermion with momentum $\vec{p}$ and handedness given by j, whereas $\widehat{b}^\dagger_{s\vec{p}}$ creates an *anti*fermion.

Quantum Electrodynamics The Dirac equation obviously has a global $\mathrm{U}(1)$ symmetry, since the Dirac Lagrangian $\mathcal{L}$ remains invariant under phase shifts $\psi \mapsto \psi e^{i\alpha}$, $\alpha \in \mathbb{R}$. We're going to outline a procedure by which we can turn global symmetries of Lagrangians into local symmetries, and then analyze the Dirac Lagrangian with local $\mathrm{U}(1)$ invariance.

In general, given a principal G-bundle $E \xrightarrow{\pi} X$ with specified connection one-form ω, we write v^V and v^H for the restrictions of an arbitrary vector field v to its vertical and horizontal components, which satisfy $\partial i_*(v^V) = \omega(v^H) = 0$ and $v^V + v^H = v$. Given a (possibly $\mathfrak{g}$-valued) k-form η on E, we define $\eta^H(v_1, \ldots, v_k) := \eta(v_1^H, \ldots, v_k^H)$ and likewise for η^V. The ***exterior***

covariant derivative on the bundle with connection $(E \xrightarrow{\pi} X, \omega)$ is given by $D\eta := (d\eta)^H$. The ***curvature*** of the connection form ω is given by $\Omega := D\omega$. Cartan's structure equation states that $\Omega = d\omega + \frac{1}{2}[\omega, \omega]$, where $[\omega, \omega](v, w) = [\omega(v), \omega(w)] - [\omega(w), \omega(v)] = 2[\omega(v), \omega(w)]$. It follows that $d\Omega = d(d\omega + \frac{1}{2}[\omega, \omega]) = \frac{1}{2}d[\omega, \omega] = \frac{1}{2}([d\omega, \omega] - [\omega, d\omega]) = [d\omega, \omega]$. Since $[[\omega, \omega], \omega] = 0$, we can write $d\omega = [\Omega, \omega]$. A locally U(1) invariant version of the Dirac equation, in which E is spinors and X is spacetime, has $d\psi = \partial_\mu \psi = (d\psi)^H + (d\psi)^V = D\psi + (d\psi)^V$. Hence, the gauge covariant derivative $D_\mu \psi$ differs from ∂_μ by a one-form: we will write $D_\mu \psi = \partial_\mu \psi + iqA_\mu \psi$, where q is a constant and A_μ is known as the ***gauge field***, transforming under a shift $\psi \mapsto \psi e^{i\alpha}$ as $A_\mu \mapsto A_\mu - \frac{1}{q}\partial_\mu \alpha$.

To make the Dirac equation as we know it *locally* U(1) invariant, we will simply make the derivative covariant, replacing ∂_μ with $D_\mu = \partial_\mu + iqA_\mu$. This gives us a U(1) gauge theory $\mathcal{L} = \overline{\psi}(i\not{D} - m)\psi = \overline{\psi}(i\not{\partial} - m)\psi - q\overline{\psi}\not{A}\psi$. In order to use this to model electromagnetism, we simply *add* the Lagrangian of classical electromagnetism, obtaining a Lagrangian

$$\mathcal{L} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} + \overline{\psi}(i\not{D} - m)\psi$$

Note that $F_{\mu\nu} = dA_\mu$, so that this is a restriction on the gauge field itself. Hence, A_μ serves two purposes: it both enforces local U(1) invariance and serves as an electromagnetic current.

$\mathcal{L} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} + \overline{\psi}(i\not{D} - m)\psi$ is the Lagrangian of ***quantum electrodynamics***. The current density is recovered from the interacting part as $J^\mu = \overline{\psi}\gamma^\mu\psi$. ψ creates fermions (electrons), $\overline{\psi}$ creates antifermions (positrons), and A_μ is a massless boson (photon) field interacting with electrons via the interaction term $\mathcal{L}_I = -q\overline{\psi}\not{A}\psi$. S-matrix terms see photons interacting with pairs of electrons and fermions, creating many of the same interactions seen in the previously encountered Yukawa interaction theory.

Bibliography

[Abramsky and Coecke, 2009] Abramsky, S. and Coecke, B. (2009). Categorical quantum mechanics. *Handbook of quantum logic and quantum structures*, 2:261--325.

[Aluffi, 2009] Aluffi, P. (2009). *Algebra: chapter 0*, volume 104. American Mathematical Soc.

[Arnold, 2013] Arnold, V. I. (2013). *Mathematical methods of classical mechanics*, volume 60. Springer Science & Business Media.

[Bergner, 2007] Bergner, J. (2007). A model category structure on the category of simplicial categories. *Transactions of the American Mathematical Society*, 359(5):2043--2058.

[Bleecker, 2005] Bleecker, D. (2005). *Gauge theory and variational principles*. Courier Corporation.

[Carroll, 2019] Carroll, S. M. (2019). *Spacetime and geometry*. Cambridge University Press.

[Coecke, 2010] Coecke, B. (2010). Quantum picturalism. *Contemporary physics*, 51(1):59--83.

[Deligne et al.,] Deligne, P., Etingof, P. I., and Freed, D. S. *Quantum fields and strings: a course for mathematicians*, volume 1.

[Fearns, 2002] Fearns, J. D. (2002). A physical quantum model in a smooth topos. *arXiv preprint quant-ph/0202079*.

[Fong and Spivak, 2018] Fong, B. and Spivak, D. I. (2018). Seven sketches in compositionality: An invitation to applied category theory. *arXiv preprint arXiv:1803.05316*.

[Fong and Spivak, 2019] Fong, B. and Spivak, D. I. (2019). Supplying bells and whistles in symmetric monoidal categories. *arXiv preprint arXiv:1908.02633*.

[Friedrich, 2000] Friedrich, T. (2000). *Dirac operators in Riemannian geometry*, volume 25. American Mathematical Soc.

[Guts and Grinkevich, 1996] Guts, A. K. and Grinkevich, E. B. (1996). Toposes in general theory of relativity. *arXiv preprint gr-qc/9610073*.

[Guts and Zvyagintsev, 2000] Guts, A. K. and Zvyagintsev, A. A. (2000). Interpretation of intuitionistic solution of the vacuum einstein equations in smooth topos. *arXiv preprint gr-qc/0001076*.

[Haase, 2014] Haase, M. (2014). *Functional analysis: an elementary introduction*, volume 156. American Mathematical Society.

[Hatcher, 2005] Hatcher, A. (2005). *Algebraic topology*.

[Holevo, 2003] Holevo, A. S. (2003). *Statistical structure of quantum theory*, volume 67. Springer Science & Business Media.

[Husemoller, 1975] Husemoller, D. (1975). *Fiber bundles*. Springer.

[Johnstone, 2002] Johnstone, P. T. (2002). Sketches of an elephant (2 vols.), volume 43 of oxford logic guides.

[Kelly and Kelly, 1982] Kelly, G. M. and Kelly, M. (1982). *Basic concepts of enriched category theory*, volume 64. CUP Archive.

[Kock, 1986] Kock, A. (1986). Lie group valued integration in well-adapted toposes. *Bulletin of the Australian Mathematical Society*, 34(3):395--410.

[Kühnel, 2015] Kühnel, W. (2015). *Differential geometry*, volume 77. American Mathematical Soc.

[Lancaster and Blundell, 2014] Lancaster, T. and Blundell, S. J. (2014). *Quantum field theory for the gifted amateur*. OUP Oxford.

[Landau and Lifshitz, 2013] Landau, L. D. and Lifshitz, E. M. (2013). *Mechanics and electrodynamics*. Elsevier.

[Lawson and Michelsohn, 2016] Lawson, H. B. and Michelsohn, M.-L. (2016). *Spin Geometry (PMS-38), Volume 38*. Princeton university press.

[Lurie, 2006] Lurie, J. (2006). Stable infinity categories. *arXiv preprint math/0608228*.

[Lurie, 2012] Lurie, J. (2012). Higher algebra.

[Mac Lane, 2013] Mac Lane, S. (2013). *Categories for the working mathematician*, volume 5. Springer Science & Business Media.

[MacLane and Moerdijk, 2012] MacLane, S. and Moerdijk, I. (2012). *Sheaves in geometry and logic: A first introduction to topos theory*. Springer Science & Business Media.

[May, 1999] May, J. P. (1999). *A concise course in algebraic topology*. University of Chicago press.

[May and Sigurdsson, 2006] May, J. P. and Sigurdsson, J. (2006). *Parametrized homotopy theory*. Number 132. American Mathematical Soc.

[Meyer, 2006] Meyer, P. A. (2006). *Quantum probability for probabilists*. Springer.

[Milnor, 1956] Milnor, J. (1956). Construction of universal bundles, ii. *Annals of Mathematics*, pages 430--436.

[Misner et al., 1973] Misner, C. W., Thorne, K. S., Wheeler, J. A., et al. (1973). *Gravitation*. Macmillan.

[Munkres, 2018] Munkres, J. R. (2018). *Elements of algebraic topology*. CRC Press.

[nLab authors, 2020] nLab authors (2020). geometry of physics – categories and toposes. `http://ncatlab.org/nlab/show/geometry%20of%20physics%20--%20categories%20and%20toposes`. Revision 122.

[Peskin, 2018] Peskin, M. E. (2018). *An introduction to quantum field theory*. CRC Press.

[Ravenel, 2003] Ravenel, D. C. (2003). *Complex cobordism and stable homotopy groups of spheres*. American Mathematical Soc.

[Rédei and Summers, 2007] Rédei, M. and Summers, S. J. (2007). Quantum probability theory. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 38(2):390--417.

[Riehl, 2014] Riehl, E. (2014). *Categorical homotopy theory*, volume 24. Cambridge University Press.

[Riehl, 2017] Riehl, E. (2017). *Category theory in context*. Courier Dover Publications.

[Riehl and Verity, 2016a] Riehl, E. and Verity, D. (2016a). Homotopy coherent adjunctions and the formal theory of monads. *Advances in Mathematics*, 286:802--888.

[Riehl and Verity, 2016b] Riehl, E. and Verity, D. (2016b). Infinity category theory from scratch. *arXiv preprint arXiv:1608.05314*.

[Rudin, 1973] Rudin, W. (1973). Functional analysis.

[Sakurai et al., 2014] Sakurai, J. J., Napolitano, J., et al. (2014). *Modern quantum mechanics*, volume 185. Pearson Harlow.

[Switzer, 2017] Switzer, R. M. (2017). *Algebraic topology-homotopy and homology*. Springer.

[Takhtadzhian, 2008] Takhtadzhian, L. A. (2008). *Quantum mechanics for mathematicians*, volume 95. American Mathematical Soc.

[Ticciati et al., 1999] Ticciati, R. et al. (1999). *Quantum field theory for mathematicians*, volume 72. Cambridge University Press.

[Wald, 2007] Wald, R. M. (2007). *General relativity*. University of Chicago Press (Chicago, 1984).

[Weibel, 1995] Weibel, C. A. (1995). *An introduction to homological algebra*. Number 38. Cambridge university press.

[Weibel, 2013] Weibel, C. A. (2013). *The K-book: An introduction to algebraic K-theory*, volume 145. American Mathematical Society Providence, RI.

[Weinberg, 1995] Weinberg, S. (1995). *The quantum theory of fields*, volume 1. Cambridge university press.