Sheaves in Topology

Master's Course
Summer Semester 2025

Julian Holstein *
University of Hamburg
Department of Mathematics

*Please email comments and corrections to julian.holstein@uni-hamburg.de

Contents

These are the lecture notes as of April 15, 2025.

An up to date version of these notes can be found at `http://www.math.uni-hamburg.de/home/holstein/lehre/STnotes.pdf`.

1. Introduction

1.1. Overview

In this course we study the basic theory of sheaves with a view to applications in topology.

– presheaves and sheaves, stalks and sheafifiaciton, pushforward and pullback functors. , sheaf cohomology.

This will require some background in category theory and homological algebra, in particular the notion of derived functors, that I will review very very briefly.

Here is an outline of the course as it is planned at the moment. There may well be changes.

1. Basic definitions, examples and constructions. Presheaves, sheaves, stalks, sheafificaiton, pushforward, inverse image.

2. A very brief introduction to homological algebra. Derived functors, the derived category.

3. Cohomology as derived global sections. Injective, flasque and soft sheaves, de Rham and Cech cohomology.

4. Computations. Cohomology and pushforward with compact support; Mayer-Vietoris, base change; Projection formula.

5. Local systems. Cohomology with local coefficients, Riemann-Hilbert, constructible sheaves.

6. If time permits: Advanced topics.

This is an advanced graduate course, the main pre-requisites is a course and on advanced algebra (language of functors and homological algebra). A course on algebraic topology (including cohomology) is extremely useful, but can be taken at the same time.

The course is not complete in the sense that I reserve the right to leave out some details and use non-trivial results from the literature.

You can influence the pace and focus of the course somewhat by making requests, asking questions or telling me to slow down or speed up.

2. Basic theory of sheaves

2.1. Definitions and Examples

Let X be a topological space and $\mathrm{Op}(X)$ the category (poset) of open sets. The category has the open subsets of X as objects and a unique morphism $U \to V$, written $U \subset V$ if U is a subset of V and no other morphisms.

Definition 2.1. A *presheaf* on X with values in a category $\mathscr{C}$ is a functor $\mathscr{F} : \mathrm{Op}(X)^{\mathrm{op}} \to \mathscr{C}$
We call $\mathscr{F}(U)$ the *sections* of $\mathscr{F}$ on U.
A *morphism of presheaves* $\mathscr{F} \to \mathscr{G}$ is just a natural transformation.

We can unravel these abstract definitions: A presheaf on X provides an object $\mathscr{F}(U)$ of $\mathscr{C}$ for any open set in X and a restriction map $r_{UV} : \mathscr{F}(U) \to \mathscr{F}(V)$ for any inclusion $V \to U$ that is compatible with compsition: $r_{UW} = r_{UV} \circ r_{VW}$. A morphism $f : \mathscr{F} \to \mathscr{G}$ is a map $f_U : \mathscr{F}(U) \to \mathscr{G}(U)$ for every U such that $f_V \circ r^{\mathscr{F}}_{UV} = r^{\mathscr{G}}_{UV} \circ f_V$.

We will be mostly interested in the case that $\mathscr{C}$ is the category of abelian groups or more generally R-modules for some commutative ring R. We will always assume that $\mathscr{C}$ has all small limits and that it is a concrete category equipped with a forgetful functor to sets, i.e. we can characterise $\mathscr{F}(U)$ by its elements.

For a section $s \in \mathscr{F}(U)$ we also write $s|_V$ for $r_{UV}(s) \in \mathscr{F}(V)$.

Example 2.2. 1. On any X the functor sending any open set U to $\mathbb{Z}$ is a presheaf with values in abelian groups called the *constant presheaf.*

2. Ony any X the functor sending any open U to the set $\mathscr{C}^0(U, \mathbb{R})$ of continuous functions on U is a presheaf.

Definition 2.3. A collection $\{U_i\}_{i\in I}$ in $\mathrm{Op}(X)$ such that $\cup U_i = U$ is called a *cover*.
A presheaf $\mathscr{F}$ is called a *sheaf* if for any cover U_i of an open U and for any collection of sections $s_i \in \mathscr{F}(U_i)$ such that $\forall i, j \in I$

$$s_i|_{U_i \cap U_j} = s_j|_{U_i \cap U_j}$$

there exists a unique section $s \in \mathscr{F}(U)$ such that $s_i = s|_{U_i}$ for all $i \in I$.

The uniqueness of the section means that sections of a sheaf are determined by their restrictions, they are *locally determined.* A presheaf satisfying this condition is sometimes called *separated.*

The existence of the section means that sheaves can be *glued* from consistent local data.

We can write the sheaf condition somewhat compactly as a limit:

Lemma 2.4. *A presheaf $\mathscr{F}$ on X is a sheaf if and only if for any cover $\{U_i\}_{i\in I}$ of any open $U \subset X$ we have*

$$\mathscr{F}(U) = \mathrm{eq}\left(\prod_{i\in I}\mathscr{F}(U_i) \rightrightarrows \prod_{i,j\in I} F(U_i \cap U_j)\right)$$

Proof. Unravelling this limit returns the definition in words. □

From either definition we can read off two useful facts:

1. For any sheaf $\mathscr{F}(\amalg_i U_i) = \prod_i \mathscr{F}(U_i)$ as the U_i form a cover and all intersections are by definition empty.
2. For any sheaf $\mathscr{F}(\emptyset) = *$, the final object of the category $\mathscr{C}$. This is a special case of the previous point, we can cover the empty set by the empty set and read off that $\mathscr{F}(\emptyset)$ is the limit over the empty category, i.e. the final object!

Example 2.5. The constant presheaf on a topological space is typically *not* a sheaf. Assume X has two disjoint open subsets U, V and consider the constant sheaf with value $\mathbb{Z}$. Then for a sheaf $\mathscr{F}$ we have $\mathscr{F}(U \cup V) = \mathscr{F}(U) \times \mathscr{F}(V)$, but the constant sheaf takes value $\mathbb{Z} \neq \mathbb{Z} \times \mathbb{Z}$.

Example 2.6. Let Y be a topological space, for example $Y = \mathbb{R}$. Let X be an arbitrary toplogical space. Define $\mathscr{C}(U)$ to be the set of continuous maps $U \to Y$. Then $\mathscr{C}$ is a sheaf.

Let U_i be a cover of U. Then U is the colimit of the U_i, to be precise $U = \mathrm{coeq}(\amalg_i U_i \leftleftarrows U_i \cap U_j)$, which we write $\operatorname{colim} U_i$ by abuse of notation to simplify things. But then $\mathscr{C}$ is a sheaf because

$$\mathscr{C}(\operatorname{colim} U_i) := \mathrm{Hom}(\operatorname{colim} U_i, Y) = \lim \mathrm{Hom}(U_i, Y) = \lim \mathscr{C}(U_i)$$

by the fundamental property of limits and homs.

Alternatively, one can unravel the definitions.

In the case $Y = \mathbb{R}$ we call this the sheaf of real-valued (continuous) functons on X. I.e. the presheaf of real-valued continuous functions on X is a sheaf.

Example 2.7. In the previous example let Y have the discrete topology, for example $Y = \mathbb{Z}$. Then we have constructed the sheaf of locally constant functions on X with values in Y. We call it the *constant sheaf* and denote it by $\underline{Y}$. This is not to be confused with the constant presheaf. To be precise, the value on a set U is $\mathbb{Z}^{c(U)}$ where $c(U)$ is the number of connected components of U.

Example 2.8. Let E be a vector bundle of rank n on a topological space X, i.e. a space E with a surjection $p : E \to X$ such that X has a cover U_i and each $p^{-1}(U_i)$ is homeomorphic to $U_i \times \mathbb{R}^n$.

Then $\mathscr{E}$ defined by $\mathscr{E}(U) = \{s : U \to p^{-1}(U) \mid p \circ s = \mathbf{1}_U\}$ is a sheaf, the *sheaf of sections* of E. If $E = X \times \mathbb{R}$ is the trivial rank one vector bundle its sheaf of sections is the sheaf of $\mathbb{R}$-valued functions.

Example 2.9. More generally for any continuous map $p : Y \to X$ we may define the sheaf of sections $\mathscr{S}$ that sends any $U \subset X$ to the set of maps $s : U \to Y$ satisfying $ps = \mathbf{1}_U$. By definition $\mathscr{S}(U) = \mathscr{C}(U) \times_{\mathrm{Hom}(U,X)} \{\iota_U\}$ where ι_U is the inclusion $U \subset X$ and thus for a cover we have

$$\begin{aligned}\mathscr{S}(\operatorname*{colim}_i U_i) &\cong \mathscr{C}(\operatorname*{colim}_i U_i) \times_{\mathrm{Hom}(\mathrm{colim}_i U_i, X)} \{\iota_U\} \\ &\cong \left(\lim_i \mathscr{C}(U_i)\right) \times_{\lim \mathrm{Hom}(U_i,X)} \{\iota_{U_i}\} \\ &\cong \lim_i \left(\mathscr{C}(U_i) \times_{\mathrm{Hom}(U_i,X)} \{\iota_{U_i}\}\right) \cong \lim_i \mathscr{S}(U_i)\end{aligned}$$

as limits commute with limits, in particular the pullback commutes with the equalizer of products in the sheaf condition.

Example 2.10. As sheaves are defined locally we may make local modifications: If E is a smooth vector bundle on a smooth manifold the presheaf of smooth sections of E is a sheaf: As the presheaf of smooth sections is contained in the sheaf of continuous sections we can always glue compatible smooth sections to a unique continuous section. But this continuous section must be smooth as it restricts to a smooth section on each open in our cover.

Similarly we may define the sheaf of locally constant functions or holomorphic functions as a subsheaf of the sheaf of all continuous functions into $\mathbb{C}$.

Here and in future a *subsheaf* $\mathscr{F}$ of a sheaf $\mathscr{G}$ is just a sheaf on the same space such that $\mathscr{F}(U) \subset \mathscr{G}(U)$ for all U.

Example 2.11. Let $X = *$. Then a $\mathscr{C}$-valued sheaf on X is exactly an object of $\mathbb{C}$.

Let $*$ be a terminal object in $\mathscr{C}$. Then the constant presheaf with value $*$ is a sheaf.

Example 2.12. Let R be a commutative ring and M an R-module. We let Spec(R) be the set of all prime ideals of R and define a topology a follows. Let for each $f \in R$ $D_f \subset \operatorname{Spec} R$ be the set of prime ideals not containing f. This is a basis of open sets for a topology on $\operatorname{Spec} R$ called the *Zariski topology*. Define a presheaf $\tilde{M}$ as follows:

1. on the D_f by $\tilde{M}(D_f) = M_f$, the localisation of M at f, i.e. the R-module of formal quotients $\{\frac{m}{f^j} \mid m \in M, j \in \mathbb{N}\}$.
2. on an arbitrary $U = \cup_f D_f$ we define $\tilde{M}(U) = \lim \tilde{M}(D_f)$.

Then one can show with some commutative algebra that this is sheaf on $\operatorname{Spec} R$. In particular R itself gives rise to a sheaf on $\operatorname{Spec} R$ called the *structure sheaf* with the property that every $\tilde{M}(U)$ is a module over $\tilde{R}(U)$. We say $\tilde{M}$ is a *quasi-coherent sheaf* an the afine scheme $\operatorname{Spec} R$ and these (and their generalizations to general schemes) play a huge role in algebraic geometry, but our focus will lie elsewhere.

Definition 2.13. A topological space X equipped with a sheaf of rings $\mathscr{R}$ is called a *ringed space*. A *sheaf of $\mathscr{R}$-modules* is a sheaf $\mathscr{M}$ of abelian groups on X such that $\mathscr{M}(U)$ is a (left) $\mathscr{R}(U)$-module for every open set U in X. A morphism of sheaves of $\mathscr{R}$-modules is a morphism of sheaves $\mathscr{F} \to \mathscr{G}$ such that each $\mathscr{F}(U) \to \mathscr{G}(U)$ is $\mathscr{R}(U)$-linear.

We will probably only look at sheaves of commutative rings, but there is no reason not to define things in general.

Definition 2.14. Given a topological space X and a category $\mathscr{C}$ we define the category $\mathsf{PSh}(X, \mathscr{C})$ as the category of presheaves on X.

We denote by $\mathsf{Sh}(X, \mathscr{C})$ the full subcategory of sheaves.

We will be particularly interested in sheaves with values in the category of R-modules for some commutative ring R.

We write $\mathsf{Sh}(X, R)$ for $\mathsf{Sh}(X, R\text{-Mod})$ for a commutative ring R and $\mathsf{Sh}(X)$ for $\mathsf{Sh}(X, \mathbb{Z}) = \mathsf{Sh}(X, \mathsf{Ab})$ for the category of sheaves of abelian groups. If $(X, \mathscr{R})$ is a ringed space we write $\mathsf{Sh}(X, \mathscr{R})$ for the category of sheaves of $\mathscr{R}$-modules.

2.2. Stalks and sheafification

As sheaves are local we may look at them at a point. We begin by looking at presheaves at points. To simplyif things we look at sheaves with values in an abelian category $\mathscr{A}$, for example abelian groups. Bt everything will be true in greater generality, for sheaves of sets one needs minor modifications of the proofs.

Definition 2.15. The *stalk* $\mathscr{F}_x$ of a presheaf $\mathscr{F}$ on X at a point $x \in X$ is defined as $\operatorname{colim}_{x \in U} \mathscr{F}(U)$ where the colimit is taken in the category $\mathscr{A}$ over all open sets containing x.

Given $s \in \mathscr{F}(U)$ we denote by $s|_x$ its image in $\mathscr{F}_x$, called the *germ* of s.

Explicitly, objects of $\mathscr{F}_x$ are pairs (U, s) with $x \in U \subset X$ open and $s \in \mathscr{F}(U)$ up to the equivalence $(U, s) \sim (W, t)$ if there is $V \subset U \cap W$ with $s|_V = t|_V$.

This is an example for a filtered colimit, which is sometimes (confusingly!) called a direct limit. See the section in the appendix if you are unfamiliar with these kinds of colimits.

Example 2.16. The constant presheaf with value R has stalk $R = \operatorname{colim} R$.

The constant sheaf $\underline{R}$ also has stalk R. The connected open neighbourhoods of a point P are final in all open neighbourhoods, thus we can compute the stalk on connected open sets, see Lemma A.35. But on a connected open set $\underline{R}(U) = R$.

Example 2.17. The presheaf of continuous functions $\mathscr{C}$ on a manifold M has as stalk at the point p the set (in fact, ring) of germs of functions at p.

Any morphism $f : \mathscr{F} \to \mathscr{G}$ induces a morphism of stalks $f_x : \mathscr{F}_x \to \mathscr{G}_x$ by sending the germ represented by (U, s) to the germ represented by $(U, f(s))$.

Lemma 2.18. *Two morphisms $f, g : \mathscr{F} \to \mathscr{G}$ of sheaves agree if they agree on stalks.*

Proof. For any U we have a commutative diagram

$$\begin{array}{ccc} \mathscr{F}(U) & \longrightarrow & \mathscr{G}(U) \\ \big\downarrow & & \big\downarrow \\ \prod_{x\in U}\mathscr{F}_x & \longrightarrow & \prod_{x\in U}\mathscr{G}_x \end{array} \tag{2.1}$$

and the vertical maps are injections: Assume given $s \in \mathscr{G}(U)$ with $s_x = 0$ for all $x \in U$. This means for any x there is some U_x on which s vanishes. But the $\{U_x\}$ form a cover of U and by the uniqueness part of the sheaf condition s must be 0.

As the maps induced by f, g in the bottom row agree, they must also agree in the top row. □

Lemma 2.19. *A morphism $f : \mathscr{F} \to \mathscr{G}$ of sheaves is an isomorphism if and only if all induced morphisms on stalks are isomorphisms.*

Proof. The only if direction is clear.

So let f be such that f_x is an isomorphism for all $x \in X$. We will show that for all U we have an isomorphism $f_U : \mathscr{F}(U) \to \mathscr{G}(U)$, then $U \mapsto f_U^{-1}$ is an inverse morphism in the category of sheaves.

To show f is injective assume $f(s) = 0$ for all $s \in U$. In particular $f(s)_x = 0$ for all x, thus by injectivitiy $s_x = 0$, so there is some U_x with $s|_{U_x} = 0$. By the uniqueness property of sheaves this means $s|_U = 0$ as in Diagram 2.1.

To show surjectivity assume we have $t \in \mathscr{G}(U)$. By surjectivity on stalks at the point x there is some U_x and $s^x \in \mathscr{F}(U_x)$ such that $(f(s^x), U_x)$ represents t_x. Shrinking U_x if necessary we may even assume $f(s^x) = t|_{U_x}$.

We want to glue the s^x into a section of $\mathscr{F}(U)$. The U_x cover U, so we have to check overlaps. Let $U_{xy} = U_x \cap U_y$ be nonempty. Then $s^x|_{U_{xy}}$ and $s^y|_{U_{xy}}$ are sent to $t_{U_{xy}}$ by assumption. By the injectivity we have already established we have $s^x|_{U_{xy}} = s^y|_{U_{xy}}$. Thus by the sheaf property of $\mathscr{F}$ we can glue to obtain $s \in \mathscr{F}(U)$. As $f(s)$ agrees with t on all stalks we see that s maps to t by Diagram 2.1. □

The constant presheaf seemed like a reasonable construction and we did then construct something we called the constant sheaf. Could we have obtained the constant sheaf directly from the constant presheaf?

Definition 2.20. The *sheafification* of a presheaf $\mathscr{F}$ is defined as follows.

$$\mathscr{F}^{\mathsf{sh}}(U) := \{(f_p \in \mathscr{F}_p)_{p\in U} \mid f_p \text{ are compatible}\}$$

where compatibility means that for any $q \in U$ there is an open $q \in V \subset U$ and a section $s \in \mathscr{F}(V)$ with $f_p = s_p$ for $p \in V$. The restriction maps are the natural restriction maps.

Here the product is taken in the category $\mathscr{A}$ and the compatibility condition is expressible as an equaliser, so if $\mathscr{F}$ takes values in $\mathscr{A}$ so does $\mathscr{F}^{\mathsf{sh}}(U)$.

Theorem 2.21. *Given a presheaf $\mathscr{F}$ on X there is a natural map $u : \mathscr{F} \to \mathscr{F}^{\mathsf{sh}}$ such that any presheaf morphism $f : \mathscr{F} \to \mathscr{G}$ for a sheaf $\mathscr{G}$ factors uniquely through u.*

Proof. Let $\mathscr{F} \in \mathsf{PSh}(X)$. We first note that $\mathscr{F}^{\mathsf{sh}}$ is indeed a sheaf. Given any cover we have (U_i) and compatible sections $s_i \in \mathscr{F}^{\mathsf{sh}}(U_i)$ we define s by $((s_i)_x) \mid x \in U_i)$, i.e. we have to specify an element of the stalk $\mathscr{F}_x$ for any $x \in U$, and just choose any $x \in U_i$ in our cover and choose the germ $(s_i)_x$. By definition of the stalks this is well-defined. Thus we have existence of sections. But the construction is also unique as $s|_{U_i} = s_i$ implies $s_x = (s_i)_x$.

We now consider the map of presheaves $u : \mathscr{F} \to \mathscr{F}^{\mathsf{sh}}$ given on U by $s \in \mathscr{F}(U) \mapsto (s_x)_{x \in U} \in \mathscr{F}^{\mathsf{sh}}(U)$.

Let $\mathscr{G}$ be a sheaf and $f : \mathscr{F} \to \mathscr{G}$ a map of presheaves. We define $\mathscr{F}^{\mathsf{sh}}(U) \to \mathscr{G}(U)$ for any open U as follows. Take $s = (s_x)_{x \in U} \in \mathscr{F}^{\mathsf{sh}}(U)$. By definition there is a cover $\{U_i\}$ of U and sections $s_i \in \mathscr{F}(U_i)$ such that for all x we have $s_x = (s_i)_x$ for a suitable i. We consider $f(s_i) \in \mathscr{G}(U_i)$. By the sheaf property of $\mathscr{G}$ they glue to a section of $\mathscr{G}(U)$ that we call $f(s)$. (Note that $s_i|_{U_i \cap U_j} = s_j|_{U_i \cap U_j}$ as they agree on stalks.) This defines $f^{\#} : \mathscr{F}^{\mathsf{sh}} \to G$. This morphism is unique as morphisms of sheaves are determined on stalks by Lemma 2.18. □

Example 2.22. Let $\mathscr{F}$ be the constant presheaf with value R. Then $\mathscr{F}^{\mathsf{sh}}(U)$ is given by functions from U to R which locally come from a section of $\mathscr{F}(U) = R$, i.e. they are locally constant functions. Thus $\mathscr{F}^{\mathsf{sh}} = \underline{R}$, the constant sheaf is the sheafification of the constant presheaf.

Corollary 2.23. *We have $u_x : \mathscr{F}_x \cong (\mathscr{F}^{\mathsf{sh}})_x$ for any $x \in X$*

Proof. The morphism is from Theorem 2.21, the result follows by unravelling the definition of $(\mathscr{F}^{\mathsf{sh}})_x$. □

Corollary 2.24. *If $\mathscr{F}$ is a sheaf $\mathscr{F}$ is uniquely isomorphic to $\mathscr{F}^{\mathsf{sh}}$.*

Proof. We have a map $\mathscr{F} \to \mathscr{F}^{\mathsf{sh}}$ by Theorem 2.21. By Lemma 2.19 it suffices to compare stalks, so the result follows from Corollary 2.23. □

Corollary 2.25. *Sheafification provides a functor left adjoint to the inclusion $\iota : \mathsf{Sh}(X, \mathscr{A}) \to \mathsf{PSh}(X, \mathscr{A})$ of presheaves into sheaves, i.e. $\mathrm{Hom}_{\mathsf{Sh}(X,\mathscr{A})}(\mathscr{F}^{\mathsf{sh}}, \mathscr{G}) \cong \mathrm{Hom}_{\mathsf{PSh}(X,\mathscr{A})}(\mathscr{F}, \mathscr{G})$ for a sheaf $\mathscr{G}$ and presheaf $\mathscr{F}$ on X.*

Proof. Given $f : \mathscr{F} \to \mathscr{G}$ a map of presheaves we obtain a map $f^{\mathsf{sh}} : \mathscr{F}^{\mathsf{sh}} \to \mathscr{G}^{\mathsf{sh}}$ by applying Theorem 2.21 to $\mathscr{F} \to \mathscr{G} \to \mathscr{G}^{\mathsf{sh}}$. Uniqueness ensures that this is functorial.

Theorem 2.21 provides the isomorphism of hom spaces for the adjunction. The map $u : \mathscr{F} \to \iota(\mathscr{F}^{\mathsf{sh}})$ is the unit and the identity map is the counit of this adjunction. □

Remark 2.26. There are different ways of considering sheafification. We may view the sheafification of a presheaf as the sheaf of sections of a certain space associated to the presehaf, the espace étalé, which is the union of all stalks of $\mathscr{F}$, equipped with a topology such that the natural projection map to X is a local homeomorphism.

This is just a different flavour of the construction we chose, but there are generally different constructions. Grothendieck's plus construction associates to any presheaf a separated presehaf and to any separated presheaf a sheaf, doing it twice is sheafification.

We could have of course also just defined sheafification as a left adjoint. We could have then shown existence by constructing it explicitly, or by some general machinery like an adjoint functor theorem. The main ingredient is checking that the inclusion of presheaves into sheaves preserves limits (see below for (co)limits of (pre)sheaves).

2.3. Limits and colimits

Recall that a category is called (co)complete if it has all (co)limits.

Theorem 2.27. *Let X be a topological spaces. If $\mathscr{C}$ is complete then so are* $\mathsf{PSh}(X, \mathscr{C})$ *and* $\mathsf{Sh}(X, \mathscr{C})$*. Limits of presheaves and sheaves are computed objectwise.*

If $\mathscr{C}$ is cocomplete then so are $\mathsf{PSh}(X, \mathscr{C})$ *and* $\mathsf{Sh}(X, \mathscr{C})$*. Colimits of presheaves are computed objectwise while the colimit of a diagram of sheaves is the sheafifiaciton of the (objectwise) colimit of the underlying diagram of presheaves.*

In particular the stalk of a colimit of sheaves is the colimit of the stalks.

Proof. We first observe that limits and colimits in the category of presheaves are determined objectwise. If you are less familiar with (co)limits it's a good exercise to check this for yourself.

By the adjunction $(-)^{\mathsf{sh}} \leftrightarrows \iota$ of Lemma 2.25 sheafification preserves colimits, thus with Corollary 2.24 we have

$$\operatorname*{colim}_j \mathscr{F}_j = \operatorname*{colim}_j (\iota \mathscr{F}_j)^{\mathsf{sh}} = (\operatorname*{colim}_j \iota \mathscr{F}_j)^{\mathsf{sh}}.$$

By Corollary 2.23 the statement about stalks follows.

To compute the limit of sheaves not that the objectwise limit of a diagram of sheaves is again a sheaf: The sheaf condition may be formulated as a limit and limits commute with limits. In other words, we may compute that for a cover $\{U_j\}$ of U and our diagram $\mathscr{F}_i$ of sheaves we have

$$\begin{aligned} \lim_i \mathscr{F}_i(U) &\cong \lim_i \lim_j \mathscr{F}_i(U_j) \\ &\cong \lim_j \lim_i \mathscr{F}_i(U_j) \end{aligned}$$

where we used that the $\mathscr{F}_i$ are sheaves and then that limits commute with limits (by what it means to be a limit). So the objectwise limit is a sheaf and satisfies the universal property of being a limit of presheaves, but then it also satisfies the weaker universal property of being a limit of sheaves.

Note that the fact that limits of sheaves exist and are given by the limit of presheaves also follows from the (non-trivial) category-theoretic statement that any inclusion with a left adjoint creates limits. □

We now consider sheaves with values in a fixed abelian category $\mathscr{A}$, for example R-modules for a fixed commutative ring R.

Then in particular a kernel of a map of sheaves is determined pointwise. We say that a map of sheaves is *injective* if its kernel is the 0 sheaf, i.e. it is injective on each open.

We say $f : \mathscr{F} \to \mathscr{G}$ is *surjective* if the cokernel is the 0 sheaf, which is the case if and only if all the maps $f_x : \mathscr{F}_x \to \mathscr{G}_x$ on stalks are surjective. In particular the map does not have to be surjective on each open. The condition is also called locally surjective to emphasize this point.

Remark 2.28. In fact these are precisely monomorphisms and epimorphisms in the category of sheaves and arguably these are the better terms to use. But enough people use the words injections and surjections.

Example 2.29. The need to sheafify the cokernel may look like a formal inconvenience, but it has a mathematical meaning. Let X be a complex manifold (like $\mathbb{C} \setminus \{0\}$) and $\mathscr{O}$ the sheaf of holomorphic functions.

Consider for example the inclusion of sheaves $\underline{\mathbb{Z}} \xrightarrow{2\pi i} \mathscr{O}$. This is the kernel of the exponential map from $\mathscr{O} \to \mathscr{O}^\times$ whose image as a presheaf we denote by $\mathscr{F}$. Then $\mathscr{F}$ is the presheaf of functions admitting a logarithm. We obtain a short exact sequence of presheaves

$$0 \to \underline{\mathbb{Z}} \to \mathscr{O} \to \mathscr{F} \to 0$$

which is just a compact way of saying $\mathscr{O} \to \mathscr{F}$ is an epimorphism with kernel $\mathbb{Z}$.

However, the presheaf cokernel $\mathscr{F}$ is not a sheaf. Having a logarithm is not a local property so if we try to glue locally defined functions which admit logarithms into a global function, the result will not in general have a logarithm.

The sheafification of $\mathscr{F}$ is $\mathscr{O}^\times$, the sheaf of invertible holomorphic functions. It is clear this is a sheaf so it suffices to check that $\mathscr{O}^\times$ is the stalkwise cokernel of the map $\underline{\mathbb{Z}} \to \mathscr{O}$. The sheaf of locally constant funnctions is the kernel of the exponentiation map, s we need to check surjectivity. Let (s, U) be a nonzero holomorphic function on some open U containing y. Shrinking U if necessary we may assume $s(y) \in B_{\frac{1}{2}|f(x)|}(f(x))$ and we have a well-defined logarithm.

The proof of the following lemma contains a brief reminder what an abelian category is.

Lemma 2.30. *The category* $\mathsf{Sh}(X, \mathscr{A})$ *of sheaves with values in the abelian category* $\mathscr{A}$ *is itself abelian.*

Proof. $\mathsf{Sh}(X)$ clearly has hom spaces which are abelian groups, it has a zero object given by the constant sheaf taking the value zero and we have seen it has finite limits and colimits in Theorem 2.27 as $\mathscr{A}$ has finite limits and colimits. We also observe that finite coproducts are equal to finite products. The presheaf finite product and coproduct agree, and this shows the finite coproduct is already a sheaf and thus equal to its own sheafifiaciton by Corollary 2.24 which is the coproduct of sheaves.

It remains to show that the natural map from the image of a map f (defined as $\ker\operatorname{coker}(f)$) to the coimage (defined as $\operatorname{coker}\ker(f)$) is an isomorphism. But this may be checked on stalks by Theorem 2.27 and Lemma 2.31 below, and on stalks it follows from the result in $\mathscr{A}$. □

Lemma 2.31. *Let $(\mathscr{F}_i)_{i\in I}$ be a finite diagram of sheaves on X. Then* $(\lim \mathscr{F}_i)_x \cong \lim_i(\mathscr{F}_i)_x$ *for all $x \in X$.*

Proof. By definition the stalk is a filtered colimit and colimits commute with finite limits in categories sufficiently like Set, see Theorem A.37.

But one can also prove this in a more elementary way. Every finite limit is an equalizer of maps between finite products by a variation of Lemma A.38. In an abelian category the finite products are finite coproducts and commute with stalks, and the equalizer may be replaced by a kernel. Thus it suffices to show that given a map of sheaves $f : \mathscr{F} \to \mathscr{G}$ we have $\ker(f)_x = \ker(\mathscr{F}_x \to \mathscr{G}_x)$ and this follows by unravelling definitions: Elements of the left hand side are germs (U, s) with $f(s) = 0$ and elements of the right hand side are germs (V, t) with $f(t|_{V'} = 0$ for some $x \in V' \subset V$. Up to equivalence of germs these sets agree. □

Note that infinite limits cannot usually be computed stalkwise.

2.4. Functors of sheaves

Given a continuous map $f : X \to Y$ of topological spaces we would like to transport sheaves along f.

Definition 2.32. Let $f : X \to Y$ be continuous and let $\mathscr{F}$ be a sheaf on X. Then we define the *pushforward sheaf* $f_*\mathscr{F}(U) = \mathscr{F}(f^{-1}U)$ on Y.

Lemma 2.33. *The pushforward sheaf is indeed a sheaf.*

Proof. This follows as the preimage of a cover is a cover. □

Example 2.34. Let X be any topological space and $p : X \to *$ the only map to the one element space. Then for any $\mathscr{F}$ in $\mathsf{Sh}(X, \mathscr{C})$ the object $p_*\mathscr{F} = \mathscr{F}(X)$ in $\mathsf{Sh}(*, \mathscr{C}) = \mathscr{C}$ is also written as $\Gamma(X, \mathscr{F})$, the *global sections* of $\mathscr{F}$.

A. Basic category theory

I will give a rapid fire overview of category theory. The focus is on definitions and examples, with a few results thrown in, but no proofs (those can be found in any standard reference, e.g. Mac Lane's "Categories for the working mathematician").

If you have met a few concepts here and there this should be nice refresher putting everything we need together in a systematic way

If you are comfortable with categories up to limits and adjunctions you can skip this. The least standard part is probably Section A.2.3 on filtered colimits.

A.1. Basics

A.1.1. Categories and Functors

Definition A.1. A *category* $\mathscr{C}$ consists of the following data:

- a class of *objects* $\mathrm{Ob}(\mathscr{C})$,
- for every pair of objects $X, Y \in \mathrm{Ob}(\mathscr{C})$ a class of *morphisms* $\mathrm{Hom}_{\mathscr{C}}(X, Y)$ (also called arrows),
- for every object X a distinguished morphism $\mathbf{1}_X \in \mathrm{Hom}_{\mathscr{C}}(X, X)$, the *identity*
- for every three objects $X, Y, Z \in \mathrm{Ob}(\mathscr{C})$ a *composition* $\circ : \mathrm{Hom}_{\mathscr{C}}(Y, Z) \times \mathrm{Hom}_{\mathscr{C}}(X, Y) \to \mathrm{Hom}_{\mathscr{C}}(X, Z)$,

such that

- compositon is associative: $(f \circ g) \circ h = f \circ (g \circ h)$,
- the identity is an identity for composition: $\mathbf{1}_Y \circ f = f = f \circ \mathbf{1}_X$ for $f \in \mathrm{Hom}_{\mathscr{C}}(X, Y)$.

Given f in $\mathrm{Hom}_{\mathscr{C}}(X, Y)$ we call X the *source* and Y the *target* of f.

Example A.2.

1. Sets and functions form a category we denote by Set. (Since we want to consider the category of all sets and want to avoid paradoxa we referred to a class of objects in our definition.)

2. Topological spaces and continuous maps form a category Top. It is easy to consider the subcategory of CW complexes or path connected spaces etc.

3. There is also a category Top_* whose objects are pointed topological spaces (X, x_0) and whose morphisms are base-point preserving maps, i.e. $f : (X, x_0) \to (Y, y_0)$ is given by $f : X \to Y$ with $f(x_0) = y_0$.

   This is an example of an *undercategory*: Given any category $\mathscr{C}$ with an object C there is a category whose objects are arrows $f : C \to D$ in $\mathscr{C}$, and whose morphisms are maps $g : D \to D'$ making the obvious triangle commute: $g \circ f = f' : C \to D1$. Top_* is the category of topological spaces under the one point space.

4. In algebra we find many further categories: Groups and homomorphisms form the category Group, vector spaces over k and linear maps form Vect_k, abelian groups, rings, fields, etc. all form categories

5. There is a category with one object and one morphism (the identity of the object). In general a category is called *discrete* if the identities are the only morphisms. Every set I can be considered as a discrete category $\mathbf{I}$ with $\mathrm{Ob}(\mathbf{I}) = I$.

6. For every category $\mathscr{C}$ there is an *opposite category* $\mathscr{C}^{op}$ with the same objects, $\mathrm{Hom}_{\mathscr{C}^{op}}(A, B) = \mathrm{Hom}_{\mathscr{C}}(B, A)$ and $f \circ_{\mathscr{C}^{op}} g := g \circ_{\mathscr{C}} f$. Thus we obtain the opposite category $\mathscr{C}^{op}$ from $\mathscr{C}$ by turning around all arrows.

We will often abuse notation and write $C \in \mathscr{C}$ as a shortcut for "C is an object of $\mathscr{C}$".

Definition A.3. A morphism $f : C \to D$ is called *isomorphism*, if there is $g : D \to C$ such that $g \circ f = \mathbf{1}_C$ and $f \circ g = \mathbf{1}_D$.

Homeomorphisms and (group/ring/vector space) isomorphisms are examples.
In all categories we consider isomorphic object as equivalent and (almost) interchangeable.

Remark A.4. If the objects and morphisms of a category form sets we call it a *small category*. If there may be a class of objects but the morphisms between any two pair of objects form a set we say the category is *locally small*.

Many categories we are interested in, like Top, Set and Group are not small, but locally small.

Example A.5. A small category in which there is at most one morphism between any two objects and in which any isomorphism is an identity is called a *partial order*. Then the composition is uniquely determined by the morphisms (as there is only one function into a set with one element).

An example is the category $\mathbb{N}$ whose objects are the natural numbers and where there is a morphism $i \to j$ if and only if $i \leq j$.

An important motivation for the study of category theory is the observation that mathematical objects are often better understood through the morphisms between them. The same principle holds for categories.

Definition A.6. A *functor* F between two categories $\mathscr{C}$ and $\mathscr{D}$ consists of the following data:

- a map that associates to any $X \in \mathrm{Ob}(\mathscr{C})$ an object $F(X) \in \mathrm{Ob}(\mathscr{D})$.
- for each pair of objects $X, Y \in \mathrm{Ob}(\mathscr{C})$ a map from $\mathrm{Hom}_{\mathscr{C}}(X, Y)$ to $\mathrm{Hom}_{\mathscr{D}}(F(X), F(Y))$ which we write as $f \mapsto F(f)$,

such that

- F is compatible with composition: $F(f \circ g) = F(f) \circ F(g)$,
- F preserves the identities: $F(\mathbf{1}_X) = \mathbf{1}_{F(X)}$.

Example A.7.

1. For every category $\mathscr{C}$ there is an identity functor $\mathbf{1}_{\mathscr{C}}$ that does nothing on objects and morphisms.
2. Let $\mathscr{C}$ and $\mathscr{D}$ be categories and D an object of $\mathscr{D}$. Then there is a constant functor $c_D : \mathscr{C} \to \mathscr{D}$ that sends every object of $\mathscr{C}$ to D and any morphism of $\mathscr{C}$ to $\mathbf{1}_D$.
3. A family of topological spaces $(X_i)_{i \in I}$ is nothing but a functor from I, considered as a discrete category, to Top.
4. From every category whose objects have an underlying set e.g. $\mathsf{Top}, \mathsf{Group}, \mathsf{Vect}_k$) there is a *forgetful functor* to Set, that forgets all additional structure.
5. Algebraic Topology is in no small part the study of functors from topological spaces to algebraic categories.

   The homotopy groups are functors $\pi_n : \mathsf{Top}_* \to \mathsf{Group}$ associating to any pointed topological space (X, x_0) the homotopy group $\pi_n(X, x_0)$ and to any map $f : X \to Y$ the induced map f_*.

   Similary homology groups are functors $H_n : \mathsf{Top} \to \mathsf{Ab}$.

   Cohomology groups are functors $H^n : \mathsf{Top}^{\mathrm{op}} \to \mathsf{Ab}$. Note that these functors turns around the direction of arrows, which is why we write it as a functor from the opposite category. We also call such functors *contravariant*.

It is easy to see that functors can be composed, so there is a *category of categories* whose objects are (small) categories and whose morphisms are functors.

A.1.2. Natural Transformations

Remarkably, there are not just maps between categories (the functors) but also maps between maps betwen categories.

Definition A.8. Let $F, G : \mathscr{C} \to \mathscr{D}$ be two functors. A *natural transformation* α from F to G consists of maps $\alpha_C : FC \to GC$ for every $C \in \mathscr{C}$ such that for every map $f : C \to C'$ in $\mathscr{C}$ there is a commutative diagram:

$$\begin{array}{ccc} FC & \xrightarrow{Ff} & FC' \\ \downarrow{\scriptstyle \alpha_C} & & \downarrow{\scriptstyle \alpha_{C'}} \\ GC & \xrightarrow{Gf} & GC' \end{array}$$

Remark A.9. You might think that it is easier to write $\alpha_{C'} \circ Ff = Gf \circ \alpha_C$ instead of drawing the commutative diagram.

The commutative diagram has the advantage that it keeps track of all the objects as well as the morphisms between them. More importantly, in category theory, algebraic topology and homological algebra there is often a plethora of maps whose compositions we want to compare, and it is much easier to keep track if one arrange them all in a beautiful diagram.

Example A.10.

1. There is a functor $D : \mathsf{Vect}_k \to \mathsf{Vect}_k$ that takes every vector space to its double dual $V \mapsto (V^*)^*$. Then for every vector space there is a map $\iota : V \to DV$ that sends $v \in V$ to the functional $\alpha \mapsto \alpha(v)$. This map is natural, meaning it is compatible with linear maps. In other words, ι is a natural transformation from the identity functor $\mathbf{1}_{\mathsf{Vect}}$ to the double dual D.

2. For any functor $F : \mathscr{C} \to \mathscr{D}$ there is the identity natural transformation $\mathbf{1}_F$ defined by $(\mathbf{1}_F)_C = \mathbf{1}_{FC}$ for every $C \in \mathscr{C}$.

3. Fix two categories I and $\mathscr{C}$, where we may think of I as being somehow small.

   We will consider a functor $F : I \to \mathscr{C}$ as a *diagram* in $\mathscr{C}$, given by objects $F(i)$ together with arrows $F(f) : F(i) \to F(j)$ for every morphism $f : i \to j$ in I.

   Any object C of $\mathscr{C}$ determines a constant functor $c_C : I \to \mathscr{C}$ that sends any i to C and any $f : i \to j$ to $\mathbf{1}_C$.

   Then natural transformation from c to another functor $F : I \to \mathscr{C}$ is given by maps $\alpha_i : C \to F(i)$ for every $i \in I$ such that $F(f) \circ \alpha_i = \alpha_j$ for every $f : i \to j$.

   We call a natural transformation from a constant diagram to F a *cone* over F. We think of C as the tip of the cone, and there are arrows going to all the vertices of the diagram, making all the triangles commute.

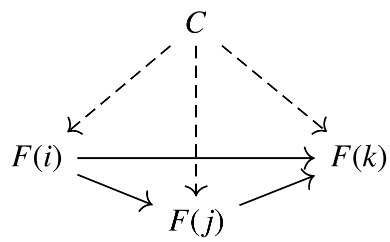

4. For every $n \geq 1$ the Hurewicz homomorphism $h_n : \pi_n(X, *) \to H_n(X, \mathbb{Z})$ from homotopy to homologoy of path connected spaces is a natural transformation. (To be precise it is a natural transformation from π_n to the composition of homology with the functor forgetting basepoints. If $n = 1$ we also have to compose with the inclusion functor from abelian groups to all groups.)

5. For every topological space X we have a functor which takes the underlying set of X and equips it with the discrete topology, write this as X^δ. Then the identity map from X^δ to X is continuous. In fact it is a natural transformation from the discretization functor to the identity functor $X^\delta \mapsto X$.

Natural transformations may be composed and form the morphism in the *category of functors* $\mathrm{Fun}(\mathscr{C}, \mathscr{D})$ between two categories.

Definition A.11. A natural tranformation α such that all α_C are isomorphisms is an isomorphism in the category of functors and is called a *natural isomorphism*.

A.1.3. Equivalences

Definition A.12. Two categories are *equivalent* if there are functor $F : \mathscr{C} \to \mathscr{D}$ and $G : \mathscr{D} \to \mathscr{C}$ such that $F \circ G$ is naturally isomorphic to $\mathbf{1}_{\mathscr{D}}$ and $G \circ F$ is naturally isomorphic to $\mathbf{1}_{\mathscr{C}}$.

We can give a more concrete description, for which we need some definitions.

Definition A.13. functor $F : \mathscr{C} \to \mathscr{D}$ is *full* if it induces surjections on all hom sets, i.e. every $g : FC \to FC'$ in $\mathscr{D}$ is $F(f)$ for some $f : C \to C'$.

The functor F is *faithful* if it induces injections on all hom sets, i.e. $F(f) = F(f')$ only if $f = f'$.

F is *fully faithful* if it is both full and faithful.

F is *essentially surjective* if every object in $\mathscr{D}$ is isomorphic to some object FC in the image of F.

Then one can prove that $F : \mathscr{C} \to \mathscr{D}$ is an equivalence of categories if and only if it is fully faithful and essentially surjective. (The "if" direction needs the axiom of choice.)

Example A.14. 1. Let k be a field. There is an equivalenc of categories from finite-dimensional k-vector spaces to its opposite category, given by $V \mapsto V^*$ on objects.

2. Let Mat be the category whose objects are non-negative integers and whose morphisms from m to n are $(m \times n)$-matrices. Composition is given by matrix multiplication.

   Then there is a natural functor from Mat to the category of finite-dimensional $\mathbb{R}$-vector spaces, given by $n \mapsto \mathbb{R}^n$ on objects. This is an equivalence of categories.

A.1.4. Opposite categories

We recall the following Example A.2.6:

Definition A.15. Let $\mathscr{C}$ be any category. Then its *opposite category* $\mathscr{C}^{\mathrm{op}}$ is defined to have the same objects as $\mathscr{C}$ but $\mathrm{Hom}_{\mathscr{C}^{\mathrm{op}}}(C, D) := \mathrm{Hom}_{\mathscr{C}}(D, C)$ and $f \circ_{\mathscr{C}^{\mathrm{op}}} g := g \circ_{\mathscr{C}} f$.

In words $\mathscr{C}$ is obtained by turning around all the arrows in $\mathscr{C}$.
Clearly any functor $F : \mathscr{C} \to \mathscr{D}$ induces an opposite functor $F^{\mathrm{op}} : \mathscr{C}^{\mathrm{op}} \to \mathscr{D}^{\mathrm{op}}$.
Many natural functors, like cohomology, turn around the order of arrows, i.e. cohomology is a functor $\mathsf{Top}^{\mathrm{op}} \to \mathsf{Ab}$.

Definition A.16. We call a functor $\mathscr{C}^{\mathrm{op}} \to \mathscr{D}$ a *contravariant functor* from $\mathscr{C} \to \mathscr{D}$.

By using the opposite of categories and functors, we can dualize all the definitions and results in category theory.
Moreover, whenever we prove a statement about a category $\mathscr{C}$ then the *dual statement* holds for its opposite category.
This is a very powerful idea, which we will come back to soon.

A.1.5. The hom functor

Forming the hom sets in a category is actually functorial. Let us explain what this means.
Let $\mathscr{C}$ be a locally small category, i.e. the morphisms between any two objects form a set (rather than a proper class). Let C be an object of $\mathscr{C}$.

Definition A.17. The *hom-functor*, denoted $h_C : \mathscr{C} \to \mathsf{Set}$, sends any object D to $\mathrm{Hom}_{\mathscr{C}}(C, D)$ and any morphism $f : D \to D'$ to the map $f_* : \mathrm{Hom}_{\mathscr{C}}(C, D)$ to $\mathrm{Hom}_{\mathscr{C}}(C, D')$ defined by $g \mapsto f \circ g$.

We can of course also put the object C in the second place of Hom. Then our functor will be contravariant and turn around the order of arrows. We obtain $h^C : \mathscr{C}^{\mathrm{op}} \to \mathsf{Set}$ which is defined by $D \mapsto \mathrm{Hom}_{\mathscr{C}}(D, C)$ and $f \mapsto f^*$, where $f^*(g) = g \circ f$.
For another level of abstraction, $h_{(-)}$ defines a functor from $\mathscr{C}^{\mathrm{op}}$ to the category of functors $\mathrm{Fun}(\mathscr{C}, \mathsf{Set})$. This is a fully faithful functor that is called the *Yoneda embedding*. Any functor naturally isomorphic to h_C is called *representable*.

Example A.18. The forgetful funtor U : Group $\to$ Set is representable by the group of integers.

Unravelling our definition this means that there for every group G there is an isomorphism $\mathrm{Hom}_{\mathsf{Group}}(\mathbb{Z}, G) \cong U(G)$, and these isomorphisms are compatible with group homomorphisms.

But this just says that the set of morphisms from $\mathbb{Z}$ to G is exactly the set of elements of G, the isomorphism is given by sending $f : \mathbb{Z} \to G$ to $f(1) \in G$.

Remark A.19. A key result in category theory is the *Yoneda lemma*. It states that natural transformations from h^C to some other functor $F : \mathscr{C} \to$ Set are in natural bijection with $F(C)$. It's not hard, but very consequential. (Although we won't need it.)

A.2. Universal constructions

A.2.1. Limits

Category theory allows us to unify many constructions in mathematics, in particular those characterised by *universal properties*.

Definition A.20. Let I be a small category and $\mathscr{C}$ any category. A *diagram of shape I* in $\mathscr{C}$ is just a functor $D : I \to \mathscr{C}$.

A *cone over D* is an object C in $\mathscr{C}$ together a natural transformation from the constant diagram C to D.

Explicitly a cone consists of C with maps $\gamma_i : C \to D(i)$ for all objects i in I such that for any $a : i \to j$ we have $D(a) \circ \gamma_i = \gamma_j$.

A map of cones $(C, \gamma) \to (E, \epsilon)$ is a map $f : C \to E$ compatible with the maps, i.e. $\epsilon_i \circ f = \gamma_i$.

We will often write F_i for the objects $F(i)$ for $i \in I$.

Definition A.21. A *limit* of the diagram $F : I \to \mathscr{C}$ is a cone (L, α_i) over F that is universal in the sense that any cone (C, γ_i) maps uniquely to (L, α_i).

In other words, L and α have the property that whenever we have C in the following diagram there is exactly one dashed arrow $C \to L$ making the diagram commute.

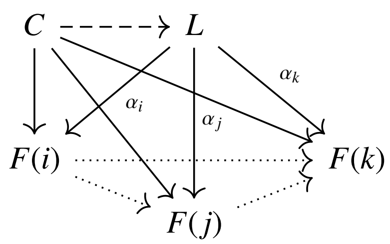

This universal property (like all universal property) ensures that if there are two limits L and L' there is a unique isomorphism between them: As L is a limit there is a unique map

$g : L' \to L$ and as L' is a limit there is a unique map $g' : L \to L'$. As $g'g$ and $\mathbf{1}_{L'}$ are both maps of cones from L' to itself they must agree and g' and g are inverse.

We thus also speak of *the limit* and denote it by $\lim_I F$ or $\lim F_i$.

Remark A.22. Note that the limit need not exist! If we can form arbitrary (small) limits in a category $\mathscr{C}$ we say that $\mathscr{C}$ *has all small limits*.

Let us make this more concrete.

Definition A.23. Let I a set considered as a discrete category. The limit of $F : I \to \mathscr{C}$ is called the product of the $F(i)$, often written $\prod_{i \in I} F_i$.

Thus $\prod_i F_i$ has the property that there are natural maps $\pi_j : \prod_i F_i \to F_j$ for all j (called *projection*) and whenever we are given maps $\beta_j : C \to F_j$ for all j we obtain a map $\beta : C \to \prod_i F_i$ such that $\beta_j = \pi_j \circ \beta$.

This recovers the familiar product of sets, topological spaces, abelian groups etc.

We consider a special case:

Definition A.24. Let I be the empty set considered as a discrete category without objects! The limit of the unique functor $I \to \mathscr{C}$ is called the *terminal* object of $\mathscr{C}$, often written $*$. It has the property that for every $C \in \mathscr{C}$ there is a unique morphism $C \to *$.

The terminal object in Set is the set with 1 Element.

Definition A.25. Let I be the category with two objects and two arrows in the same direction $\bullet \rightrightarrows \bullet$. The limit of $F : I \to \mathscr{C}$ is called *equalizer*.

Definition A.26. Let I be the category with three objects $\bullet \to \bullet \leftarrow \bullet$. The limit of $F : I \to \mathscr{C}$ is called *pullback*.

Example A.27.

1. The terminal object in Groups is the group with 1 element.
2. The terminal object in Top is the topological space with 1 point.
3. In the diagram $\bullet \to \bullet \leftarrow \bullet$ that defines pull-backs the middle object is terminal.
4. If a pull-back diagram in Set or Top takes the form $* \to Y \xleftarrow{f} X$ then the pull-back is the fiber of f (equipped with the subspace topology in the case of Top).
5. If a pull-back diagram takes the form $X \to * \leftarrow Y$, i.e. the middle object goes to the terminal object of $\mathscr{C}$, then the limit is the product $X \times Y$.
6. In the category Groups there is a unique map from $*$ to any group H and the pullback of the diagram $* \to H \xleftarrow{f} G$ is nothing but the kernel of f.
7. The equalizer of two maps $f, g : A \to B$ in Set is exactly the subset of A given by all elements a with $f(a) = g(a)$, this explains the name.

A.2.2. Colimits

We now apply the idea of *dualizing* categorical notions by turning around all the arrows to the previous section.

So we change the orientation of all the arrows in the definition of a limit. This gives the dual notion of a limit, called the colimit.

Definition A.28. A *colimit* of the diagram $F : I \to \mathscr{C}$, denoted by $\mathrm{colim}_I F$, is an object D of $\mathscr{C}$ together with a natural transformation $\alpha : F \Rightarrow c_D$ that is *universal*, in the sense that any natural transformation from F to a constant functor c_C factors uniquely through c_D.

The corresponding diagram looks like this:

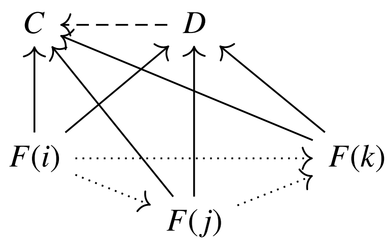

Remark A.29. To make the duality of limit and colimit more precise we can observe that (D, α) is a colimit of the diagram $F : I \to \mathscr{C}$ exactly if $(D, \alpha^{\mathrm{op}})$ is a limit of the diagram $F^{\mathrm{op}} : I^{\mathrm{op}} \to \mathscr{C}^{\mathrm{op}}$. Here $\alpha^{\mathrm{op}} : c_D^{\mathrm{op}} \Rightarrow F^{\mathrm{op}}$ is the natural transformation corresponding to $\alpha : F \Rightarrow c_D$ under the correspondence of morphisms in $\mathscr{C}$ and $\mathscr{C}^{\mathrm{op}}$.

Definition A.30. The colimit over a discrete category is called the *coproduct* or sum.

The colimit of the empty diagram is called the *initial object*.

The colimit of the diagram $\bullet \leftarrow \bullet \rightarrow \bullet$ is called *pushout*.

The colimit of a diagram of shape $\bullet \leftleftarrows \bullet$ is called *coequalizer*.

Example A.31.

1. In Set and Top ithe coproduct is given by the disjoint union.
2. In Group the coproduct is given by the free product of groups.
3. In Vect the product and coproduct of two vector spaces V and W agree, both are given by $V \oplus W$. (This holds for all finite products and coproducts in Vect, but it is no longer true for infinite products and coproducts!)
4. The initial object in Set is given by the empty set.
5. The group with one object is both initial and terminal.
6. The pushout of the diagram $0 \leftarrow V \rightarrow W$ of vector spaces is the quotient space W/V.
7. The coequalizer of two maps $f, g : A \to B$ in Set is given by the quotient of B by the relation generated by $f(a) \sim g(a)$ for all $a \in A$.

From the definition of limit and colimits it is not hard to obtain the following extremely useful result:

Lemma A.32. *Let $F : I \to \mathscr{C}$ and $G : J \to \mathscr{C}$ be diagrams. Then we have natural isomorphisms*

$$\mathrm{Hom}_{\mathscr{C}}(C, \lim_I F_i) \cong \lim_I \mathrm{Hom}_{\mathscr{C}}(C, F_i)$$

and

$$\mathrm{Hom}_{\mathscr{C}}(\operatorname{colim}_J G_i, C) \cong \lim_J \mathrm{Hom}_{\mathscr{C}}(G_j, C)$$

A.2.3. Filtered colimits

A special kind of colimit is given by the following.

A category I is *filtered* if any finite diagram in I has a cone. Equivalently I is filtered when it is not empty, for every two objects i, i' there exists an object k with two arrows $i \to k$ and $i' \to k$; for any two parallel arrows $u, v : i \rightrightarrows j$ there is an object k and morphism $f : j \to k$ with $fu = fv$.

A *filtered diagram* is a diagram $I \to \mathscr{C}$ with I filtered.

Definition A.33. A colimit over a filtered diagram is a *filtered colimit*

Example A.34. 1. The category $(\mathbb{N}, \leq)$ with objects the natural numbers and a single morphism $a \to b$ whenevere $a \leq b$ is filtered. A colimit indexed by $(\mathbb{N}, \leq)$ is also called a sequential colimit. Increasing unions are a typical example: $\mathbb{R} = \operatorname{colim}_{a \in \mathbb{N}}(-a, a)$ as sets or topological spaces.

2. The set of all neighbourhoods of a point x in a topological space X is a filtered category under inclusion.

Such examples where there is at most one morphism between two objects are also called posets.

A functor $F : I \to J$ is called *cofinal* if

1. For any object j in J there is i in I with a morphism $j \to F(i)$
2. For any two arrows $j \to F(i)$ and $j \to F(i')$ there is a zig-zag of arrows $i \xleftarrow{f_1} \cdots \xrightarrow{f_n} i1$ making the natural diagram commute:

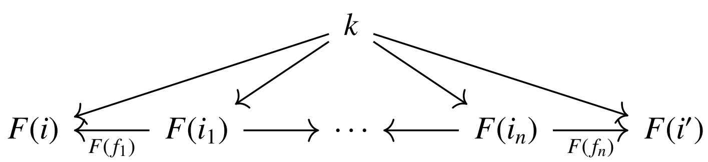

Note that the second condition is automatic if J is filtered.

Lemma A.35. *Let $F : I \to J$ be a final functor and $G : J \to \mathscr{C}$ a diagram. Then if* $\operatorname{colim}_I GF$ *exists then* $\operatorname{colim}_J G$ *also exists and agrees with* $\operatorname{colim}_I GF$.

Example A.36. The inclusion of all prime numbers into $(\mathbb{N}, \leq)$ is final.

The inclusion of connected open neighbourhoods in all neighbourhoods of a point in a topological set is final.

The key result about filtered colimits is the following:

Theorem A.37. *In the category* Set *and A-Mod for any ring A finite limits commute with filtered colimits.*

A.2.4. Existence of (co)limits

We say a category $\mathscr{C}$ *has all small limits* or *is complete* if every diagram $I \to \mathscr{C}$ has a limit. Similarly we say $\mathscr{C}$ has all small colimits or is cocomplete if every diagram $I \to \mathscr{C}$ has a colimit.

This may seem extremely difficult to check, but in fact one can build any limit from just two types of limit:

Recall that an equalizer is a limit for a diagram of the shape $\bullet \rightrightarrows \bullet$ and a product is a diagram whose shape is a discrete category.

We say a category $\mathscr{C}$ has all equalizers if any equalizer diagram has a limit, and similarly for products (and other shapes of diagrams).

Lemma A.38. *A category $\mathscr{C}$ has all limits if and only if it has all products and equalizers. It has all colimits if and only if it has all coproducts and coequalizers.*

A.2.5. Adjunctions

It is rare that categories are equivalent, but a weaker notion is extremely fruitful.

Definition A.39. We say $F : \mathscr{C} \to \mathscr{D}$ is *left adjoint* to $G : \mathscr{D} \to \mathscr{C}$, in symbols $F \dashv G$ if for all $C \in \mathscr{C}$ and $D \in \mathscr{D}$ there are natural isomorphisms

$$\phi_{C,D} : \operatorname{Hom}_{\mathscr{C}}(C, GD) \cong \operatorname{Hom}_{\mathscr{D}}(FC, D)$$

Here naturality means that for every map $C \to C'$ in $\mathscr{C}$ the natural diagram commutes:

$$\begin{array}{ccc} \operatorname{Hom}_{\mathscr{C}}(C', GD) & \xrightarrow{\phi_{C',D}} & \operatorname{Hom}_{\mathscr{D}}(FC', D) \\ \downarrow{\scriptstyle f^*} & & \downarrow{\scriptstyle Ff^*} \\ \operatorname{Hom}_{\mathscr{C}}(C, GD) & \xrightarrow{\phi_{C,D}} & \operatorname{Hom}_{\mathscr{D}}(FC, D) \end{array}$$

and a similar diagram commutes for $g : D \to D'$ in $\mathscr{D}$.

If $\mathscr{C}$ and $\mathscr{D}$ are locally small we can also phrase naturality as saying that the two functors $\operatorname{Hom}_{\mathscr{C}}(-, G(-))$ and $\operatorname{Hom}_{\mathscr{D}}(F(-), -)$ from $\mathscr{C}^{\mathrm{op}} \times \mathscr{D}$ to Set are naturally isomorphic.

Example A.40. 1. Throughout algebra there are adjunctions between free and forgetful functors. For example the forgetful functor $U : \mathsf{Group} \to \mathsf{Set}$ has a left adjoint given by taking a set X to the free group with set of X as set of generators.

2. The forgetful functor $\mathsf{Top} \to \mathsf{Set}$ has a left adjoint given by equipping any set with the discrete topology. It also has a right adjoint given by equipping any set with the indiscrete topology.

Left and right adjoints are naturally dual: If $F : \mathscr{C} \to \mathscr{D}$ is left adjoint to G, then $F^{\mathrm{op}} : \mathscr{C}^{\mathrm{op}} \to \mathscr{D}^{\mathrm{op}}$ is right adjoint to G^{op}.

Let $F \dashv G : \mathscr{C} \rightleftarrows \mathscr{D}$ and $C \in \mathscr{C}$. By the adjunction the identity map $\mathbf{1}_{FC} : FC \to FC$ corresponds to a map $\epsilon_C : C \to GFC$. By naturality in the definition of an adjunction the ϵ assemble into a natural transformation $\epsilon : \mathbf{1}_{\mathscr{C}} \Rightarrow GF$. This is called the *unit* of the adjunction.

Simlarly there is a natural transformation $\eta : FG \Rightarrow \mathbf{1}_{\mathscr{D}}$, called the *counit* of the adjunction.

Lemma A.41. *Let $F \dashv G$. Then unit and counit satisfy the following identities of natural transformations: For every $C \in \mathscr{C}$ we have*

$$\eta_{FC} \circ F(\epsilon_C) = \mathbf{1}_{FC}$$

and for every $D \in \mathscr{D}$ we have

$$G(\eta_C) \circ \epsilon_{GD} = \mathbf{1}_{GD}.$$

Put a little differently, we have the following identities of natural transformations: $G\eta \circ \epsilon_G = \mathbf{1}_G$ and $\eta_F \circ F\epsilon = \mathbf{1}_F$.

In fact, adjoints may be equivalently characterized by the existence of unit and counit.

Remark A.42. An adjunction induces an equivalence of categories if and only if unit and counit are natural isomorphisms.

One can also show that adjoints are given by a universal property and are thus unique up to unique natural isomorphism.

Adjoints are closely related to limts:

Lemma A.43. *Let F be a left adjoint. Then F preserves colimits, i.e. whenever (D, α) is a colimit of a diagram $G : I \to \mathscr{C}$ then $(FD, F\alpha)$ is a colimit for $F \circ G : I \to \mathscr{D}$.*

Dually, if G is a right adjoint then G preserves limits.

Remark A.44. Under some assumption on the categories $\mathscr{C}$ and $\mathscr{D}$ there is even a converse to the lemma: Any functor preserving all colimits has a left adjoint. There are different theorems, depending on the precise assumptions made, but they are all called *adjoint functor theorems*.

We can even characterize limits using adjoints.

Lemma A.45. *Consider the category* $\mathrm{Fun}(I, \mathscr{C})$ *of I-shaped diagrams in $\mathscr{C}$. There is a diagonal functor $\Delta : \mathscr{C} \to$* $\mathrm{Fun}(I, \mathscr{C})$ *sending any object C to the constant functor c_C. Then taking the limit of a diagram is right adjoint to Δ, and taking the colimit is left adjoint.*