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 February 23, 2026.

Please contact me with any comments or corrections at `julian.holstein@uni-hamburg.de`.

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

Here are some useful books:

[Wei95] C. Weibel, *An introduction to homological algebra*, CUP (1994). Very thorough general reference for homological algebra, occasionally a bit dated.

[GM03] S. Gelfand, Yu. Manin, *Methods of homological algebra*, Springer (2003). Another general reference for homological algebra, more modern but watch out for typos.

[KS90] M. Kashiwara, P. Schapira, *Sheaves on Manifolds*, Springer (1990). Very dense and not always easy to use, but probably the most comprehensive book around.

[Ive86] B. Iversen, *Cohomology of sheaves*, Springer (1986). Good treatment of many of the topics of this course.

[Dim04] A. Dimca, *Sheaves in topology*, Springer (2004). Nice overview getting to some advanced topics, but light on details with very few proofs.

Many standard references in Algebraic Geometry or Topology have a useful perspective on some aspect of this course, e.g. Hartshorne *Algebraic Geometry*, Vakil *Foundations of Algebraic Geometry*, Voisin *Hodge Theory and Complex Algebraic Geometry I*, Bott & Tu *Differential forms in algebraic geometry* ...

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 was planned at the beginning. There have been 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 $\mathcal{C}$ is a functor $\mathcal{F} : \mathrm{Op}(X)^{\mathrm{op}} \to \mathcal{C}$
We call $\mathcal{F}(U)$ the *sections* of $\mathcal{F}$ on U.
A *morphism of presheaves* $\mathcal{F} \to \mathcal{G}$ is just a natural transformation.

We can unravel these abstract definitions: A presheaf on X provides an object $\mathcal{F}(U)$ of $\mathcal{C}$ for any open set in X and a restriction map $r_{UV} : \mathcal{F}(U) \to \mathcal{F}(V)$ for any inclusion $V \to U$ that is compatible with compsition: $r_{UW} = r_{UV} \circ r_{VW}$. A morphism $f : \mathcal{F} \to \mathcal{G}$ is a map $f_U : \mathcal{F}(U) \to \mathcal{G}(U)$ for every U such that $f_V \circ r^{\mathcal{F}}_{UV} = r^{\mathcal{G}}_{UV} \circ f_V$.
We will be mostly interested in the case that $\mathcal{C}$ is the category of abelian groups or more generally R-modules for some commutative ring R. We will always assume that $\mathcal{C}$ has all small limits and that it is a concrete category equipped with a forgetful functor to sets, i.e. we can characterise $\mathcal{F}(U)$ by its elements.
For a section $s \in \mathcal{F}(U)$ we also write $s|_V$ for $r_{UV}(s) \in \mathcal{F}(V)$.

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 $\mathcal{C}^0(U, \mathbb{R})$ 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 $\mathcal{F}$ is called a *sheaf* if for any cover U_i of an open U and for any collection of sections $s_i \in \mathcal{F}(U_i)$ 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 $s \in \mathcal{F}(U)$ 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 $\mathcal{F}$ 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*

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

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

From either definition we can read off two useful facts:

1. For any sheaf $\mathcal{F}(\amalg_i U_i) = \prod_i \mathcal{F}(U_i)$ as the U_i form a cover and all intersections are by definition empty.

2. For any sheaf $\mathcal{F}(\emptyset) = *$, the final object of the category $\mathcal{C}$. This is a special case of the previous point, we can cover the empty set by the empty set and read off that $\mathcal{F}(\emptyset)$ 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 $\mathcal{F}$ we have $\mathcal{F}(U \cup V) = \mathcal{F}(U) \times \mathcal{F}(V)$, 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 $\mathcal{C}(U)$ to be the set of continuous maps $U \to Y$. Then $\mathcal{C}$ 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 $\mathrm{colim}\, U_i$ by abuse of notation to simplify things. But then $\mathcal{C}$ is a sheaf because

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

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 $\mathcal{E}$ defined by $\mathcal{E}(U) = \{s : U \to p^{-1}(U) \mid p \circ s = \mathbf{1}_U\}$ 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 $\mathcal{S}$ that sends any $U \subset X$ to the set of maps $s : U \to Y$ satisfying $ps = \mathbf{1}_U$. By definition $\mathcal{S}(U) = \mathcal{C}(U) \times_{\mathrm{Hom}(U,X)} \{\iota_U\}$ where ι_U is the inclusion $U \subset X$ and thus for a cover we have

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

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* $\mathcal{F}$ of a sheaf $\mathcal{G}$ is just a sheaf on the same space such that $\mathcal{F}(U) \subset \mathcal{G}(U)$ for all U.

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

Let $*$ be a terminal object in $\mathcal{C}$. 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 $\mathcal{R}$ is called a *ringed space*. A *sheaf of $\mathcal{R}$-modules* is a sheaf $\mathcal{M}$ of abelian groups on X such that $\mathcal{M}(U)$ is a (left) $\mathcal{R}(U)$-module for every open set U in X. A morphism of sheaves of $\mathcal{R}$-modules is a morphism of sheaves $\mathcal{F} \to \mathcal{G}$ such that each $\mathcal{F}(U) \to \mathcal{G}(U)$ is $\mathcal{R}(U)$-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 $\mathcal{C}$ we define the category $\mathsf{PSh}(X, \mathcal{C})$ as the category of presheaves on X.

We denote by $\mathsf{Sh}(X, \mathcal{C})$ 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 $(X, \mathcal{R})$ is a ringed space we write $\mathsf{Sh}(X, \mathcal{R})$ for the category of sheaves of $\mathcal{R}$-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 $\mathcal{A}$, 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* $\mathcal{F}_x$ of a presheaf $\mathcal{F}$ on X at a point $x \in X$ is defined as $\operatorname{colim}_{x \in U} \mathcal{F}(U)$ where the colimit is taken in the category $\mathcal{A}$ over all open sets containing x.

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

Explicitly, objects of $\mathcal{F}_x$ are pairs (U, s) with $x \in U \subset X$ open and $s \in \mathcal{F}(U)$ 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.

Note that the stalk of a sheaf of rings is again a ring (whose underlying abelian group is the stalk of the underlying sheaf of abelian groups) by defining multiplication and addition of representatives in the obvious way: $[(U, s)] \cdot [(V, t)] = [(U \cap V, s|_{U \cap V} \cdot t|_{U \cap V})]$ etc.

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 $\mathcal{C}$ on a manifold M has as stalk at the point p the set (in fact, ring) of germs of functions at p.

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

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

Proof. For any U we have a commutative diagram

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

and the vertical maps are injections: Assume given $s \in \mathcal{G}(U)$ 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 $f : \mathcal{F} \to \mathcal{G}$ 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 $f_U : \mathcal{F}(U) \to \mathcal{G}(U)$, 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 $t \in \mathcal{G}(U)$. By surjectivity on stalks at the point x there is some U_x and $s^x \in \mathcal{F}(U_x)$ 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 $\mathcal{F}(U)$. 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 $\mathcal{F}$ we can glue to obtain $s \in \mathcal{F}(U)$. 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 $\mathcal{F}$ is defined as follows.

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

where compatibility means that for any $q \in U$ there is an open $q \in V \subset U$ and a section $s \in \mathcal{F}(V)$ 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 $\mathcal{A}$ and the compatibility condition is expressible as an equaliser, so if $\mathcal{F}$ takes values in $\mathcal{A}$ so does $\mathcal{F}^{\mathsf{sh}}(U)$.

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

Proof. Let $\mathcal{F} \in \mathsf{PSh}(X)$. We first note that $\mathcal{F}^{\mathsf{sh}}$ is indeed a sheaf. Given any cover we have (U_i) and compatible sections $s_i \in \mathcal{F}^{\mathsf{sh}}(U_i)$ we define s by $((s_i)_x) \mid x \in U_i)$, i.e. we have to specify an element of the stalk $\mathcal{F}_x$ 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 $u : \mathcal{F} \to \mathcal{F}^{\mathsf{sh}}$ given on U by $s \in \mathcal{F}(U) \mapsto (s_x)_{x \in U} \in \mathcal{F}^{\mathsf{sh}}(U)$.

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

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

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

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

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

Proof. We have a map $\mathcal{F} \to \mathcal{F}^{\mathsf{sh}}$ 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 $\iota : \mathsf{Sh}(X, \mathcal{A}) \to \mathsf{PSh}(X, \mathcal{A})$ of presheaves into sheaves, i.e. $\mathrm{Hom}_{\mathsf{Sh}(X,\mathcal{A})}(\mathcal{F}^{\mathsf{sh}}, \mathcal{G}) \cong \mathrm{Hom}_{\mathsf{PSh}(X,\mathcal{A})}(\mathcal{F}, \mathcal{G})$ for a sheaf $\mathcal{G}$ and presheaf $\mathcal{F}$ on X.*

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

Theorem 2.21 provides the isomorphism of hom spaces for the adjunction. The map $u : \mathcal{F} \to \iota(\mathcal{F}^{\mathsf{sh}})$ 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 $\mathcal{F}$, 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 $\mathcal{C}$ is complete then so are $\mathsf{PSh}(X, \mathcal{C})$ and $\mathsf{Sh}(X, \mathcal{C})$. Limits of presheaves and sheaves are computed objectwise.*

If $\mathcal{C}$ is cocomplete then so are $\mathsf{PSh}(X, \mathcal{C})$ and $\mathsf{Sh}(X, \mathcal{C})$. 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

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

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 $\mathcal{F}_i$ of sheaves we have

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

where we used that the $\mathcal{F}_i$ 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 $\mathcal{A}$, 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 $f : \mathcal{F} \to \mathcal{G}$ is *surjective* if the cokernel is the 0 sheaf, which is the case if and only if all the maps $f_x : \mathcal{F}_x \to \mathcal{G}_x$ 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 $\mathcal{O}$ the sheaf of holomorphic functions.

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

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

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

However, the presheaf cokernel $\mathcal{F}$ 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 $\mathcal{F}$ is $\mathcal{O}^\times$, the sheaf of invertible holomorphic functions. It is clear this is a sheaf so it suffices to check that $\mathcal{O}^\times$ is the stalkwise cokernel of the map $\underline{\mathbb{Z}} \to \mathcal{O}$. 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* $\mathsf{Sh}(X, \mathcal{A})$ *of sheaves with values in the abelian category* $\mathcal{A}$ *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 $\mathcal{A}$ 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 $\mathcal{A}$. □

Lemma 2.31. *Let $(\mathcal{F}_i)_{i\in I}$ be a finite diagram of sheaves on X. Then $(\lim \mathcal{F}_i)_x \cong \lim_i (\mathcal{F}_i)_x$ 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 of two maps f, g may be replaced by a kernel of $g - f$. Thus it suffices to show that given a map of sheaves $f : \mathcal{F} \to \mathcal{G}$ we have $\ker(f)_x = \ker(f_x : \mathcal{F}_x \to \mathcal{G}_x)$ 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 $\mathcal{F}$ be a sheaf on X. Then we define the *pushforward sheaf* or *direct image* $f_*\mathcal{F}(U) = \mathcal{F}(f^{-1}U)$ 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 $\mathcal{F}$ in $\mathsf{Sh}(X, \mathcal{C})$ the object $p_*\mathcal{F} = \mathcal{F}(X)$ in $\mathsf{Sh}(*, \mathcal{C}) = \mathcal{C}$ is also written as $\Gamma(X, \mathcal{F})$, the *global sections* of $\mathcal{F}$.

Example 2.35. Let $i : x \to X$ be an inclusion of a point and $M \in \mathcal{A}$. Then i_*M is the sheaf defined as $i_*M(U) = M$ if $x \in U$ and 0 otherwise. This is called the *skyscraper sheaf* at x.

Definition 2.36. Let $f : X \to Y$ be continuous and let $\mathcal{F}$ be a sheaf on X. Then we define the *pullback sheaf* $f^{-1}F$ or *inverse image* as the sheafification of the presheaf $U \mapsto \operatorname{colim}_{f(U)\subset V} \mathcal{F}(V)$.

Example 2.37. Let X be any topological space and $p : X \to *$ the only map to the one element space. For U open in X and R a ring considered as a sheaf on $*$ then $p^{-1}R(U)$ is the sheaf associated to the presheaf $U \mapsto R$, using that the index category of all V with $p(U) \subset V$ only has the element $\{*\}$. Thus $p^{-1}R = \underline{R}$.

Example 2.38. Let $i : x \to X$ be the inclusion of a point and let $\mathcal{F}$ be a sheaf on X. Then $i^{-1}\mathcal{F}$ is by definition equal to the stalk $\mathcal{F}_x$.

Example 2.39. Let $j : U \to X$ be the inclusion of an open set and $\mathcal{F}$ a sheaf on X. Then $j^{-1}\mathcal{F}(V) = \mathcal{F}(V)$ with $V \subset U \subset X$. This is a sheaf (by the sheaf condition on X) and is also denoted $\mathcal{F}|_U$ and called the restricition of $\mathcal{F}$ to U. (This is not to be confused with the restriction maps of sections of a sheaf induced by an inclusion of open sets.)

Example 2.40. A sheaf $\mathcal{F}$ on X is called *locally constant* if there is a cover of X by open sets U_i such that each $\mathcal{F}|_{U_i}$ is isomorphic as a sheaf to the constant sheaf.

Let for example $X = S^1$ and M the open Möbius strip whith projection $p : M \to S^1$. Then the sheaf $\mathcal{S}$ of sections of M, defined as locally constant maps $s : U \to M \times S^1 U$ with $ps = \mathbf{1}_U$, forms a sheaf. (As being locally constant is a local condition this is a subsheaf of the sheaf of sections from Example 2.9) It is locally constant as we can cover X by two open intervals U_1, U_2 on which the Möbius band is homeomorphic to $U_i \times \mathbb{R}$. This identifies our sheaf of sections with the locally constant functions, which is the constant sheaf.

The definition of the sheaf pullback looks unwieldy, but it is well-behaved on stalks.

Lemma 2.41. *Let $f : X \to Y$ be continuous and let $\mathcal{F}$ be a sheaf on Y. We have $(f^{-1}\mathcal{F})_x = F_{f(x)}$. In particular let $i_y : * \to Y$ be the inclusion of a point. Then $(i_y)^{-1}\mathcal{F} = \mathcal{F}_y$.*

Proof. We may take the stalk $(f^{-1}\mathcal{F})_x$ as the stalk of the underlying presheaf, thus we compute $\operatorname{colim}_{x \in U} \operatorname{colim}_{f(U) \subset V} \mathcal{F}(V)$ which is exactly $\mathcal{F}_{f(x)} = \operatorname{colim}_{f(x) \in V} \mathcal{F}(V)$ by unravelling definitions. (Any V containing $f(x)$ also contains the image of an open containing x, namely $f^{-1}V$.) □

The following fact is extremely useful.

Theorem 2.42. *Given $f : X \to Y$ there is an adjunction $f^{-1} \dashv f_* : \mathsf{Sh}(Y) \leftrightarrows \mathsf{Sh}(X)$.*

Proof. We fix $\mathcal{F} \in \mathsf{Sh}(X)$ and $\mathcal{G} \in \mathsf{Sh}(Y)$. It is possible to write down natural maps $f^{-1}f_*\mathcal{F} \to \mathcal{F}$ and $\mathcal{G} \to f_*f^{-1}\mathcal{G}$ which are the unit and counit of the adjunction, or equivalently write down natural maps between $\operatorname{Hom}(\mathcal{G}, f_*\mathcal{F})$ and $\operatorname{Hom}(f^{-1}\mathcal{G}, \mathcal{F})$. Checking the triangle equalities, respectively the fact the maps are indeed inverse is not pleasent (books like to skip this step). The following trick is from Vakil's *Foundations of Algebraic Geometry*, Exercise 2.7.B.

Define the set $\operatorname{Hom}^C(\mathcal{G}, \mathcal{F})$ as the set of all collections of maps $\phi_{UV} : \mathcal{G}(V) \to \mathcal{F}(U)$ for $f(U) \subset V$ which are compatible with restrictions.

From the point of view of the open sets $U \subset X$ these maps are represented by maps $\operatorname{colim}_{f(U) \subset V} \mathcal{G}(V) \to \mathcal{F}(U)$. Compatibility with restriction means we have a map from the diagram of all V with $f(U) \subset V$, thus we obtain a map from the colimit.

From the point of view of the open sets $V \subset Y$ these maps are represented by maps $\mathcal{G}(V) \to \mathcal{F}(f^{-1}(V))$, as for a fixed V any U with $f(U) \subset V$ is a subset of $f^{-1}(V)$. □

For example we may compute for any sheaf $\mathcal{F}$ on X that

$$\mathrm{Hom}(\underline{\mathbb{Z}}, \mathcal{F}) \cong \mathrm{Hom}(p^{-1}\mathbb{Z}, \mathcal{F}) \cong \mathrm{Hom}(\mathbb{Z}, p_*\mathcal{F}) \cong \Gamma(X, \mathcal{F})$$

for $p : X \to *$.

Definition 2.43. Given $\mathcal{F}, \mathcal{G} \in \mathsf{Sh}(X)$ we define the *sheaf of homomorphisms* $U \mapsto \mathcal{H}om_{\mathsf{Sh}(U,\mathcal{A}}(\mathcal{F}|_U, \mathcal{G}_U)$.

One can check that this is indeed a sheaf.
In particular the hom space $\mathrm{Hom}_{\mathsf{Sh}(X)}(\mathcal{F}, \mathcal{G})$ is nothing but $\Gamma(\mathcal{H}om(\mathcal{F}, \mathcal{G}))$.
Restricting Theorem 2.42 to open subset of Y shows the following:

Corollary 2.44. *For any $f : X \to Y$ and sheaves $\mathcal{F}$ on X and $\mathcal{G}$ on Y we have $f_* \mathcal{H}om_X(f^*\mathcal{G}, \mathcal{F}) \cong \mathcal{H}om_Y(\mathcal{G}, f_*\mathcal{F})$.*

Proof. Let V be an open subset of Y and apply Theorem 2.42 to the restriciton $f' : f^{-1}V \to V$ to obtain $\mathrm{Hom}_{f^{-1}V}(f'^*\mathcal{G}, \mathcal{F}) \cong \mathcal{H}om_V(\mathcal{G}, f'_*\mathcal{F})$. This verifies the corollary on each open. □

Let now $\mathcal{R}$ be a sheaf of rings on X. We let $\mathcal{F}$ and $\mathcal{G}$ be sheaves of left $\mathcal{R}$-modules on X, see Definition 2.13. We can define $\mathcal{H}om_{\mathcal{R}}(\mathcal{F}, \mathcal{G})$ as the subsheaf of $\mathcal{H}om(\mathcal{F}, \mathcal{G})$ that on each U consists of $\mathcal{R}(U)$-linear maps.

Similarly $\mathcal{R}$ be a sheaf of rings on X, $\mathcal{F}$ a sheaf of left $\mathcal{R}$-modules and $\mathcal{G}$ a sheaf of right $\mathcal{R}$-modules. (Equivalently $\mathcal{G}$ is a sheaf of left $\mathcal{R}^{\mathrm{op}}$-modules.) Then there is a presheaf of abelian groups $U \mapsto \mathcal{F}(U) \otimes_{\mathcal{R}(U)} \mathcal{G}(U)$ which we may sheafify to obtain a tensor product of sheaves.

The tensor hom adjunction of modules directly gives us a tensor hom adjunction for sheaves:

Corollary 2.45. *Let $\mathcal{R}, \mathcal{S}$ be sheaves of rings on X. Let $\mathcal{F}$ be a $\mathcal{R} \otimes \mathcal{S}^{op}$-module sheaf, $\mathcal{G}$ a sheaf of $\mathcal{S}$-modules and $\mathcal{H}$ a sheaf of $\mathbb{R}$-modules. Then there is a natural isomorphism*

$$\mathrm{Hom}_{\mathcal{R}}(\mathcal{F} \otimes_{\mathcal{S}} \mathcal{G}, \mathcal{H}) \cong \mathrm{Hom}_{\mathcal{S}}(\mathcal{G}, \mathcal{H}om_{\mathcal{R}}(\mathcal{F}, \mathcal{H}))$$

where we used that $\mathrm{Hom}_{\mathcal{R}}(\mathcal{F}, \mathcal{H})$ *has a natural $\mathcal{S}$-module structure.*

Proof. We may check on each open, using Corollary 2.25. □

3. An introduction to homological algebra

3.1. Exactness

We now work in some general abelian category. This could be R-Mod for an arbitrary unital ring R or the category $\mathsf{Sh}(X, \mathcal{A})$ of sheaves on some space X with values in some other abelian category $\mathcal{A}$.

All our functors will be additive, i.e. they preserve finite sums (which are the same as finite products). It is a key question if they preserve kernels and/or cokernels.

Definition 3.1. A *(cochain) complex* in $\mathcal{A}$ is a sequence of objects $A^i \in \mathcal{A}$ where $i \in \mathbb{Z}$ with *differentials* $d_i : A^i \to A^{i+}$ satisfying $d_{i+1} \circ d_i = 0$.

A morphism of complexes $A \to B$ is called a *chain map*, it consists of maps $f^i : A^i \to B^i$ for every i which commute with the differential.

Complexes and the morphisms between them form the category $\mathsf{Ch}(\mathcal{A})$.

The i-th *cohomology* of a complex C is $\ker(d_i)/\operatorname{Im}(d_{i-1})$.

A cochain complex C whose cohomology group $H^i(C) = 0$ is called *exact at C^i*. And if all cohomology groups vanish it is called *exact* or *acyclic*. We also call an exact cochain complex an *exact sequence*.

It is often convenient to consider a *shifted* complex $A[1]$ defined by $A[1]^i = A^{i+1}$ and $d_i^{A[1]} = -d_{i+1}^A$.

It is not hard to check that chain maps naturally induce maps on cohomology groups.

This is *cohomological grading convention*. It is often convenient to instead use *homological grading convention* where the differenital decreases degree.

We will identify objects of $\mathcal{A}$ with cochain complexes in $\mathsf{Ch}(\mathcal{A})$ concentrated in degree 0.

Definition 3.2. An exact chain complex $0 \to A \xrightarrow{f} B \xrightarrow{g} C \to 0$ in $\mathcal{A}$ is called a *short exact sequence*.

We also say B is an *extension* of C by A.

Example 3.3. For any objects A, C in $\mathcal{A}$ there is a shorte exact sequence

$$0 \to A \to A \oplus C \to C \to 0$$

called a *split* short exact sequence. One can show an exact sequence $0 \to A \xrightarrow{f} B \xrightarrow{g} C \to 0$ is split if and only if f has a left inverse or g has a right inverse.

In particular in a short exact sequence we have $\ker(f) = 0$, $\operatorname{coker}(g) = 0$ and $\ker(g) = \operatorname{Im}(f)$.

If you know homology from topology you know that the sequence of singular chains is exact at the object in degree n if there aren't any "holes" in degree n. In homological algebra you study this condition algebraically.

Lemma 3.4. *A sequence of sheaves $\mathcal{F} \xrightarrow{f} \mathcal{G} \xrightarrow{g} \mathcal{H}$ is a short exact sequence if and only if $\mathcal{F}_x \to \mathcal{G}_x \to \mathcal{H}_x$ is a short exact sequence at each x*

Proof. We have seen in the proof of Theorem 2.19 that f is injective if all f_x are injective. By definition we see that g is surjective if all g_x are surjective.

It remains to compare the image of f with the kernel of g. But as kernel and image (by definition the kernel of a cokernel) are computed stalkwise by Theorem 2.27 and Lemma 2.31 this follows. □

Example 3.5. Consider a point x_0 on the manifold $\mathbb{R}$. Then there is a short exact sequence of sheaves

$$0 \to C^0 \xrightarrow{(x-x_0)\cdot} C^0 \to \mathbb{R}_x \to 0$$

where $\mathbb{R}_x$ is the skyscraper sheaf at x.

3.2. Exact functors

Short exact sequences are thus a way to encode monomorphisms, epimorphisms and extensions. We now examine what functors do to them.

Definition 3.6. An additive functor that preserves short exact sequences is called *exact*. An additive functor that sends an exact sequence $0 \to A \to B \to C \to 0$ to an exact sequence $0 \to F(A) \to F(B) \to F(C)$ (not necessarily exact on the right!) is called *left exact*. Similarly for right exact functors.

Example 3.7. For any object M of $\mathcal{A}$ the functor $\operatorname{Hom}(M, -) : \mathcal{A} \to \mathsf{Ab}$ is left exact. The functor $\operatorname{Hom}(-, M) : \mathcal{A}^{\mathrm{op}} \to \mathsf{Ab}$ is also left exact.

Example 3.8. For any $f : X \to Y$ the functor f_* is left exact. By Theorem 2.42 we see that f_* is a right adjoint, thus it preserves all limits and in particular kernels. It follows that f_* is left exact. The lack of right exactness of f_* will occupy us for the rest of the semester.

Let us reiterate that by the argument in the example all left adjoints are right exact, and all right adjoints are left exact.

Remark 3.9. Arguing as in Lemma 2.31 a functor is left exact if and only if it preserves finite limits, and right exact if and only if it preserves finite colimits.

Example 3.10. The functor $f^{-1} : \mathsf{Sh}(Y) \to \mathsf{Sh}(X)$ is not just right exact (as it's a left adjoint) but is exact. This follows as $f^{-1}\mathcal{F}_x = \mathcal{F}_{f(x)}$ by Lemma 2.41 and we can check exactness at stalks by Lemma 3.4.

Example 3.11. Let $i : Z \to X$ be a closed inclusion. Then i_* is exact. We use again that we can check exactness at stalks, so given an exact sequence $0 \to \mathcal{F} \to \mathcal{G} \to \mathcal{H} \to 0$ we consider $(i_*\mathcal{F})_x \to (i_*\mathcal{G})_x \to (i_*\mathcal{H})_x$ at an arbitrary $x \in X$. As Z is closed we see that all stalks vanish for $x \notin Z$. If on the other hand $x \in Z$ we have $(i_*\mathcal{F})_x = \operatorname{colim}_{x\in V} \mathcal{F}(V \cap Z)$ which agrees with $\mathcal{F}_x$ as the $V \cap Z$ are exactly the open neighbourhoods of x in Z.

Example 3.12. The pushforward f_* is not exact in general. Consider the projection $p : X = \mathbb{C} \setminus \{0\} \to *$ and the short exact sequence of sheaves $0 \to \underline{\mathbb{Z}} \to \mathcal{O} \to \mathcal{O}^\times \to 0$ from Example 2.29. Pushing forward along p is taking global sections, but the function $\exp : \mathcal{O}(X) \to \mathcal{O}^\times(X)$ is not surjective as there is no global logarithm of the identity function.

3.3. Derived functors

Consider the result of applying a left exact functor to a short exact sequence $A \to B \to C$. If F is not right exact then $F(B) \to F(C)$ is not an epimorphism. So there is a cokernel. Can we compute this cokernel in terms of F and the short exact sequence? Faling that, can we find something which contains the cokernel, and then try to determine the cokernel of the new map and so on. In other words, if we do not have a short exact sequence, can we get a long exact sequence?

One useful observation is that we know that any additive functor will preserve *split* exact sequences. We can relate being split to nice properties of modules:

Definition 3.13. An object M in an abelian category is *projective* if for any epi $q : A \to B$ and any map $f : M \to B$ there is a lift $g : M \to A$ such that $q \circ g = f$. An object N in an abelian category is *injective* if for any monomorphism $i : A \to B$ and any map $f : A \to N$ there is an extension $g : B \to N$ such that $g \circ i = f$.

It is easy to see that M is projective if only if $\operatorname{Hom}(M, -)$ is an exact functor and dually N is injective if and only if $\operatorname{Hom}(-, N)$ is an exact functor.

Example 3.14. In the category of R-modules any free module is projective. In fact projectives are exactly direct summands of free modules.

In the category of abelian groups the groups $\mathbb{Q}$ and $\mathbb{Q}/\mathbb{Z}$ are injective.

Lemma 3.15. *If C is projective or A is injective then $A \to B \to C$ is split, i.e. $B \cong A \oplus C$.*

Sketch of proof. If C is projective use the identity map $C \to C$ to find a one-sided inverse of the map $B \to C$. Dually if A is injective. □

An object of $\mathcal{A}$ can be viewed as a complex concentrated in degree 0. We will now identify such objects with larger complexes consisting of nicer objects.

Definition 3.16. A *quasi-isomorphism* of complexes is a map of complexes $A \to B$ such that the induced map on cohomology is an isomorphism in every degree.

Definition 3.17. A *projective resolution* of A is a levelwise projective complex in nonpositive degrees $P^\bullet$ with a quasi-isomorphism to A.

An *injective resolution* of A is a levelwise injective complex $I^\bullet$ in nonnegative degrees with a quasi-isomorphism from A.

Definition 3.18. The *i-th left derived functor* of a right exact functor F is defined as $L_iF(A) := H_i(F(P))$ where P is a projective resolution of A.

The *i-th right derived functor* of a left exact functor G is defined as $R^iG(A) := H_i(F(I))$ where I is an injective resolution of A.

In the remainder of this section many results will have two versions, we one for left derived functors and one for right derived functors. I will only make statements for *right* derived functors, as these will be more interesting to us in this course, but it will be clear what the analogous statements for left derived functors are.

Lemma 3.19. *For any left exfact functor G we have $R^{<0}G(A) = 0$ and $R^0G(A) = G(A)$.*

Proof. The first statement follows from the definition. For the second statement we have by definition that $A = \ker(I^0 \to I^1)$ for an injective resolution. By left exactness of G and definition of R^0 we have $G(A) = \ker(GI^0 \to GI^1) = R^0G(A)$. □

Example 3.20. We define $\mathrm{Ext}^i_R(A, B)$ to be $R^i \mathrm{Hom}_R(-, B)(A)$. Consider the category of abelian groups, i.e. $R = \mathbb{Z}$. Note that an injective resolution in $\mathbb{Z}$-Mod$^{\mathrm{op}}$ is given by a projective resolution in $\mathbb{Z}$-Mod. So $\mathbb{Z} \xrightarrow{p} \mathbb{Z}$ is a suitable resolution of $\mathbb{Z}/p$ and we find $\mathrm{Ext}^*(\mathbb{Z}/p, B) = H^*(B \xrightarrow{p} B)$. So $\mathrm{Ext}^0(\mathbb{Z}/p, B) = {}_pB$, the submodule of p-torsion elements, and $\mathrm{Ext}^1(\mathbb{Z}/p, B) = B/pB$.

Definition 3.21. A category has *enough projectives* if for every object there is an epimorphism from a projective object. Dually a category has *enough injectives* if for every object there is a monomorphism to an injective object.

Example 3.22. R-Mod has enough projectives, there is always a surjection $F(M) \to M$, from the free module generated by the elements of M to M.

Lemma 3.23. *The category R-Mod has enough injectives.*

Proof. The proof is explained in [Wei95, Exercise 2.35] and the lead-up to that. □

We collect some fundamental facts for future reference which you hopefully know from previous exposure to homological algebra. Otherwise they are not hard to find, e.g. in the book of Weibel.

First we may worry that projective and injective resolutions need not be unique in general. They will however always be unique up to chain homotopy equivalence: Two chain maps $f, g : L \to M$ in $\mathsf{Ch}(\mathcal{A})$ are *chain homotopic* if there is a map $h : L \to M[1]$ such that $dh = g - f$.

Theorem 3.24 (Comparison Theorem)**.** *Let $\epsilon : M \to I^\bullet$ and $\eta : N \to J^\bullet$ be injective resolutions and $f : M \to N$ a homomorphism. Then there is a lift $\tilde{f} : I^\bullet \to J^\bullet$ of f, i.e. we have $\eta \circ \tilde{f} = f \circ \epsilon$. Moreover, $\tilde{f}$ is unique up to chain homotopy equivalence.*

Proof. See [Wei95, Theorem 2.3.7] or do it as an exercise! □

Corollary 3.25. *Injective resolutions exist in $\mathcal{A}$ if there are enough injectives in $\mathcal{A}$. These resolutions are unique up to chain homotopy equivalence.*

The second statement means given two resolutions $I^\bullet, J^\bullet$ there are maps $f : I^\bullet \to J^\bullet$ and $g : J^\bullet \to I^\bullet$ such that $g \circ f$ and $f \circ g$ are chain homotopic to the identity map of $I^\bullet$ respectively $J^\bullet$.

Corollary 3.26. *The i-ith right derived functor is a well-defined functor.*

This follows since by functoriality of cohomology with respect to chain maps two chain homotopy equivalent complexes have isomorphic cohomology.

Lemma 3.27. *[Snake Lemma] Any short exact sequence of complexes $0 \to A \to B \to C \to 0$ induces a natura, i.e. functorial,l long exact sequence in cohomology groups:*

$$\cdots \to H^k(A) \to H^k(B) \to H^k(C) \to H^{k+1}(A) \to \cdots$$

Proof. See [Wei95, Theorem 1.3.1] □

We explain how these facts give the key result on derived functors:

Corollary 3.28. *Let $F : \mathcal{A} \to \mathcal{B}$ be a left exact functor between abelian categories. A s.e.s $0 \to A \to B \to C \to 0$ in $\mathcal{A}$ gives rise to a natural long exact sequence of derived functors*

$$0 \to FA \to FB \to FC \to R^1FA \to R^1FB \to R^1FC \to R^2FA \to \cdots$$

in $\mathcal{B}$.

Sketch of proof. We need to replace all objects injectivly and lift the maps to a short exact sequence of injective complexes. The lift from Theorem 3.24 is not the right tool here, instead we resolve A by $I^\bullet$ and C by $J^\bullet$ and then we may inductively (using injectivity) construct differentials on $I^i \oplus J^i$ such that we obtain a resolution of B. This is called the Horsesehoe lemma, see [Wei95, Lemma 2.2.8]. We then have a short exact sequence of complexes which is split in every degree. We apply F to obtain a new short exact sequence and taking cohomology we finish with the Snake Lemma 3.27. □

This definition of derived functors is the most direct one, it is not the only one. It is a bit ad hoc and we have to work to show what we have is well-defined.

Injective objects are not always easy to work with, so it is good to have other ways to compute. Let F be a left exact functor between two abelian categories.

Definition 3.29. An object A is *F-acyclic* if $R^iF(A) = 0$ for all $i > 0$.

Acyclic objects can be used to compute derived functors.

Proposition 3.30. *Let A be an object in an abelian category with enough injectives and let $0 \to A \to S^0 \to S^1 \to \cdots \to S^m \to \cdots$ be a resolution of A such that each S^i is F-acyclic. Then $H^i(S^\bullet) \cong R^iF(A)$.*

Proof. The proof technique here is called *dimension shifting*. We first consider $0 \to A \to S^0 \to Q_0 \to 0$ with Q the quotient object. Then by Lemma 3.28 we have $R^iF(A) \cong R^{i-1}FQ_0$ for $i \geq 2$ while $R^1F(A) \cong FQ_0/FS^0$. With $Q_0 = \ker(S^1 \to S^2)$ and F preserving kernels we get $R^1F(A) = H^1F(S^\bullet)$.

Now Q_0 has an F-acyclic resolution $S^1 \to S^2 \to \cdots$, thus by the same argument we see $R^1F(Q_0) = H^2F(S^\bullet)$. Together with the first part we get the result for $i = 2$.

Now we proceed by induction, letting $Q_i = S^i/S^{i-1} = S^i/Q_{-1} = \ker(S^{i+1} \to S^{i+2})$ we prove $R^{i+1}F(A) = FQ_i/FS^i = H^{i+1}F(S^\bullet)$. □

3.4. The derived category

To compute derived functors we replaced objects, considered as complexes concentrated in degree 0, by quasi-isomorphic complexes. After applying the functor we have a complex which is typically no longer quasi-isomorphic to a complex concentrated in degree 0. Hence it makes sense to consider all complexes, up to quasi-isomorphisms, and try to lift functors to this new category.

As complexes are now fundamental I will drop the $-^\bullet$ from the notation.

Remark 3.31. It is non-trivial to invert quasi-isomorphisms, mainly since it is unclear what happens to morphisms. We'd have to replace them by arbitrarily long zig-zags $* \to * \leftarrow * \to * \leftarrow \cdots \to *$ where all right-to-left maps are quasi-isomorphisms. But if we do not have a set of objects but a proper class then we quickly have a proper class of morphisms to consider, which is a problem.

We first note that there is a natural complex of morphism between two complexes.

Definition 3.32. Let $\mathcal{A}$ be an abelian category and $L, M \in \mathsf{Ch}(\mathcal{A})$. The *hom complex* $\underline{\mathrm{Hom}}(L, M)$is defined by $\underline{\mathrm{Hom}}^i(L, M) = \{f^\bullet : L^\bullet \to M^{\bullet+i}\}$ and $df : a \mapsto d(fa) - (-1)^{|f|}f(da)$ where $|f|$ denotes the degree of f.

In particular a chain map is a cocycle in degree 0.

Definition 3.33. Given an abelian category $\mathcal{A}$ we define the *homotopy category $K(\mathcal{A})$* to be the category with the same objects as $\mathsf{Ch}(\mathcal{A})$ but with morphisms from L to M equal to $H^0(\mathrm{Hom}(L, M))$, i.e. the *homotopy classes* of chain maps.

There are different boundedness conditions we can put on chain complexes, and hence on $\mathsf{Ch}(\mathcal{A})$ and $K(\mathcal{A})$. Let $\mathsf{Ch}^b(\mathcal{A})$ to be the category of *bounded cochain complexes*, i.e. those A_* such that $A_n = 0$ for all but finitely many n. We also define $\mathsf{Ch}^+(\mathcal{A})$, resp. $\mathsf{Ch}^-(\mathcal{A})$, to be the categories of chain complexes that are bounded below, resp. above. $K^+(\mathcal{A})$, $K^-(\mathcal{A})$ etc. are defined similarly.

Definition 3.34. Given a category $\mathcal{A}$ and a class of morphisms S we define the *localization* of $\mathcal{A}$ at S to be a category $\mathcal{B}$ with a functor $Q : \mathcal{A} \to \mathcal{B}$ such that $Q(s)$ is an isomorphism for any $s \in S$ and which is universal with this property: Any $\mathcal{A} \to \mathcal{C}$ that sends all $s \in S$ to isomorphisms factors through Q.

Definition 3.35. We define the *derived category* $D(\mathcal{A})$ as the localization of $K(\mathcal{A})$ at the class of quasi-isomorphisms. Write $Q_{\mathcal{A}} : K(\mathcal{A}) \to D(\mathcal{A})$ for the natural functor. $D^b(\mathcal{A})$ is defined similarly from $K^b(\mathcal{A})$, etc.

Theorem 3.36. *$D(\mathcal{A})$ exists (as a locally small category).*

Not a proof. See [Wei95, Sections 10.3 and 10.4] or [Huy06, Section 2.1] for a detailed proof.

I'll just give some comments about the shape of the proof: Localization means that we throw in an inverse f^{-1} for every quasi-isomorphism f. This is a lot like Ore localization for (noncommutative) rings, if you've met that. The content is in working out suitable conditions for the existence of a localization, and checking that they are satisfied in our case. A class of morphisms is called *localising* if

- S contains the identities and is closed under composition.
- Given morphisms $A \xleftarrow{s} A' \to B$ with $s \in S$ there are morphisms $A \to B' \xleftarrow{t} B$ with $t \in S$ making the obvious diagram commute. Dually given the second pair of morphisms there exists the first one.
- Given any morphism f, g the existence of $s \in S$ with $sf = sg$ is equivalent to the existence of $t \in S$ with $ft = gt$.

These conditions hold for quasi-isomorphisms in the homotopy category, but not in the category of chain complexes. □

The second condition in the proof allows us to write any morphism in the derived category as a 2-term zig-zag or "roof" $(s, f) := A \xleftarrow{s} A' \xrightarrow{f} B$ with $s \in S$. Two morphisms (s, f) and (t, g) represent the same map if there is a common roof (r, h) with $sr = th$ and $hg = fr$.

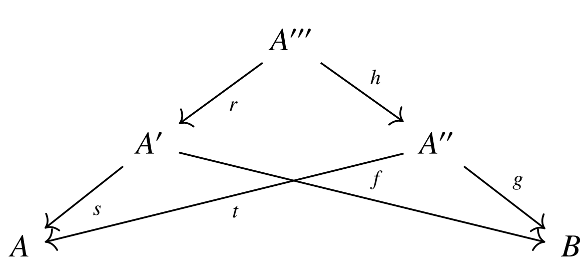

This is sometimes called a *calculus of fractions*.

Remark 3.37. Any complex with homology bounded above has a natural quasi-isomorphism from a bounded above complex. Any complex with homology bounded below has a quasi-isomorphism to a bounded below complex. Together this gives a chain of quasi-isomorphisms between a complex with bounded cohomology and a bounded complex. There will not be chain homotopies in general.

We call the subcategory of $D(\mathcal{A})$ consisting of objects with homology in bounded (resp. bounded above, resp. bounded below) degrees the bounded (above, below) derived category $D^b(\mathcal{A})$ ($D^-(\mathcal{A})$, $D^+(\mathcal{A})$).

Example 3.38. Let R be a ring. We define the *derived category* of R, $D(R)$, as the derived category of $\mathsf{Ch}(R\text{-Mod})$.

Because of its definition $D(\mathcal{A})$ is a bit hard to work with. For example, it's an additive category, but that is not obvious from the definition!

A morphism $A \xleftarrow{s} A' \to\xrightarrow{f} B$ in $D(\mathcal{A})$ is 0 if it is equivalent to a morphism homotopy equivalent to 0, and unravelling definitions this means there must exist a quasi-isomoprhism r with $f \circ r$ chain homotopic to 0.

Remark 3.39. $D(\mathcal{A})$ is not abelian but is an example of a *triangulated category*. Identifying complexes up to quasi-isomorphism is a good middle ground between the homotopy category of all complexes (which is very large) and the category $\mathcal{A}$ (which is not well-suited to cohomology and derived phenomena). One issue with the derived category is that while it admits derived functors (as we will see) it does not keep track of higher coherences. This makes it ill-suited for some applications. For example it is not possible to glue locally defined derived categories.

To keep track of this extra structure one may replace the derived category by a certain *stable ∞-category* which is obtained as the ∞-categorical localisation of $K(\mathcal{A})$ at all quasi-isomorphisms. One can do this explicitily to build a simplicial derived category using the so-called "hammock localization" or by abstract properties of ∞-categories, as explained e.g. in [Cis19].

For the purposes of this course we will need neither the details of triangulated categories nor of stable ∞-categories.

The name triangulated category refers to the following very useful construction:

Definition 3.40. Given a chain map $f : A \to B$ in $K(\mathcal{A})$ its *cone* is defined as the complex C with $C^n = A^{n+1} \oplus B^n$ and $d_n : (a, b) \mapsto (-da, db - fa)$.

By construction there are natural maps $B \to \text{cone}(f)$ and $\text{cone}(f) \to A[1]$, and we can build a sequence of morphisms $A \xrightarrow{f} B' \to \text{cone}(f) \to A[1]$ (which may of course be continued to the left and to the right).

We call any sequence $A \to B \to C \to A[1]$ in $K(\mathcal{A})$ or $D(\mathcal{A})$ that is isomorphic (in $K(\mathcal{A})$, respectively $D(\mathcal{A})$) to a sequence of the form $D \to \text{cone}\,(g) \to E \xrightarrow{g} D[1]$ an *exact triangle*. Here an isomorphism of sequences means there are isomorphisms $A \to D$, $B \to \text{cone}(g)$ and $C \to E$ making the obvious diagraom commute.

One can show that if $A \to B \to C \to A[1]$ is an exact triangle so is $B \to C \to A[1] \to B[1]$ and so forth, see [Wei95, Example 10.1.6]. Thus while exact triangles are related to short exact sequences in an abelian category there is no object singled out. (This also shows our slightly non-standard definition is equivalent to the usual ones.)

Proposition 3.41. *A chain map $f : A \to B$ is a quasi-isomorphism if and only if* $\text{cone}(f)$ *is acyclic, i.e. it has no homology.*

Proof. By definition we have a short exact sequence of complexes $B \to \text{cone}(f) \to A[1]$. In the associated long exact sequence of homology groups (Lemma 3.27) the boundary maps are the maps induced by f on cohomology, thus the result follows. □

To get a more concrete representation of hom spaces in the derived category we have the following very useful result:

Theorem 3.42. *Given a complex of injectives $I \in K^+(\mathcal{A})$ and any cochain complex A we have* $\text{Hom}_{K(A)}(A, I) \cong \text{Hom}_{D(\mathcal{A})}(A, I)$.

Proof. We first show that $\underline{\text{Hom}}_{K(\mathcal{A})}(-, I)$ sends quasi-isomorphisms to quasi-isomorphisms. This is a bit stronger than we need, but it's a nice thing to know.

We know by Proposition 3.41 that $f : A \to B$ is a quasi-isomorphism if its cone $C := \text{cone}(f)$ is acyclic. So we let $B \to C \to A[1]$ be the short exact sequence in $\mathsf{Ch}(\mathcal{A})$ associated to the map $f : A \to B$. We apply $\underline{\text{Hom}}(-, I)$ to obtain a sequence $\underline{\text{Hom}}(A[1], I) \to \underline{\text{Hom}}(C, I) \to \underline{\text{Hom}}(B, I)$.

By construction $\underline{\text{Hom}}(C, I)[1]$ is the cone of $\underline{\text{Hom}}(f, I)$, thus if $\underline{\text{Hom}}(C, I) \simeq 0$ we have that the natural map $\underline{\text{Hom}}(B, I) \to \underline{\text{Hom}}(A, I)$ is a quasi-isomorphism.

To show that $\underline{\text{Hom}}(C, I)$ is indeed acyclic if C is we build a homotopy to 0 for an arbitrary chain map $g : C \to I[i]$. This shows that $\underline{\text{Hom}}(C, I)$ has no cohomology and the desired result follows.

We build our homotopy $h : C \to I[i-1]$. As I is bounded below we may use induction and start with $h^{k_0} = 0$ for some small enough k_0. Assuming we found $h^{\leq k}$ with $g^{k-1} = dh^{k-1} - (-1)^{i-1} i h^k d : C^{k-1} \to I^{k-1+i}$ we consider $g^k - dh^k : C^k \to I^{k+i}$. Using that g is a chain map (i.e. $dg = (-1)^i gd$) this factors through C^k/C^{k-1}:

$$\begin{aligned}(g^k - dh^k)d &= (-1)^i dg^{k-1} - dh^k d \\ &= (-1)^i d(dh^{k-1} - (-1)^{i-1} h^k) - dh^k d \\ &= dh^k d - dh^k d = 0\end{aligned}$$

Since C^k/C^{k-1} injects into C^{k+1}, by injectivity of I we may extend to a map $(-1)^k h^k$ which satisfies precisely $g^k = dh^k - (-1)^k h^{k+1} d$.

We now consider the map $f \mapsto (\mathbf{1}, f)$ from $\mathrm{Hom}_{K(A)}(A, I)$ to $\mathrm{Hom}_{D(\mathcal{A})}(A, I)$. Let (s, f) be morphism $A \to B$ in $D(\mathcal{A})$, to show the theorem we have to show it is equivalent to a unique $(\mathbf{1}, g)$. By our first claim the quasi-isomorphism s induces an isomorphism of homotopy classes of maps (given by $H^0(\underline{\mathrm{Hom}}(-, I))$), thus there is g unique up to chain homotopy with $gs \sim f$ and then $(\mathbf{1}, g)$ is equivalent to (s, f). To show uniqueness we look at the equivalence criterion for fractions. It suffices to show that $(\mathbf{1}, g)$ and $(\mathbf{1}, g')$ are only equivalent if g and g' are homotopic. From the diagram we read off that $gr \simeq g'r$ for some quasi-isomorphism r. By the first claim we see that this means $g \simeq g'$.

□

The following corollary allows us to compute hom-sets in the derived category.

Corollary 3.43. *Assume $\mathcal{A}$ has enough injectives. Then for objects $A, B \in \mathcal{A}$ considered as complexes concentrated in degree 0 we have* $\mathrm{Hom}_{D(\mathcal{A})}(A, B[i]) = \mathrm{Ext}^i(A, B)$.

Proof. The two sides may be identified with the two sides of $\mathrm{Hom}_{D(\mathcal{A})}(A, I[i]) \cong \mathrm{Hom}_{K(\mathcal{A})}(A, I[i]))$ where I is an injective resolution of B. □

Remark 3.44. This result remains true with the same proof for A, B any bounded below complex as long as we define $\mathrm{Ext}^i(A, B)$ suitably, i.e. via level-wise injective resolution of B.

Corollary 3.45. *Assume $\mathcal{A}$ has enough injectives. Consider the subcategory $K^+(Inj(\mathcal{A}))$ of $K^+(\mathcal{A})$ that consists of levelwise injective complexes. Then the natural quotient map $Q_{\mathcal{A}} : K^+(Inj(\mathcal{A})) \to D^+(\mathcal{A})$ is an equivalence of categories.*

Sketch of proof. Full faithfulness follows from Theorem 3.42. To show inclusion is essentially surjective we have to injectively resolve complexes that are bounded below. There is a natural but technical proof proceding by induction and using the existence of enough injectives, details are in [GM03, III.5.25]. □

3.5. Total derived functors

We can now interpret derived functors differently: They lift functors to the derived category. But we need to change our definition a little:

Definition 3.46. Let $F : \mathcal{A} \to \mathcal{B}$ is a left exact functor. By Corollary 3.45 we can choose an equivalence of categories $\sigma : D^+(\mathcal{A}) \to K^+(\mathrm{Inj}(\mathcal{A}))$.

Then the right derived functor $RF : D^+(\mathcal{A}) \to D^+(\mathcal{B})$ is defined as $Q_{\mathcal{B}} \circ F \circ \sigma$.

We may sometimes abuse notation and also refert to the composition $RF \circ Q$ as a right derived functor. We observe that this functor preserves quasi-isomorphisms.

If we need to disambiguate we will call the R^iF the *classical derived functors* and RF the *total derived functor*. We indeed have $R^iF(A) = H^i(RF(A))$ for any $A \in \mathcal{A}$ as $\sigma(A)$ is nothing but an injective resolution of A. Note that by Theorem 3.24 all injective resolutions of A are isomorphic in $K^+(\mathrm{Inj}(\mathcal{A}))$.

What happens to exact functors? A priori we can derive them on the left or th eright, but there should be no need to do that.

Lemma 3.47. *Let $F : \mathcal{A} \to \mathcal{B}$ be an exact functor of abelian categories. Then the functor $\mathsf{Ch}(\mathcal{A}) \to \mathsf{Ch}(\mathcal{B})$ induced by F preserves quasi-isomorphism. In particular the right or left derived functor of F is quasi-isomorphic to F.*

Proof. It suffices to check that F sends acyclic complexes to acyclic complexes by Proposition 3.41.

Let L be a complex in $\mathsf{Ch}(\mathcal{A})$. We have $H^i(L) = 0$ if $\mathrm{Im}(L^{i-1} \to L^i)$ equals $\ker(L^i \to L^{i+1})$. But if F is exact it preserves kernels and images, thus $H^i(L) = 0$ implies $H^i(FL) = 0$. □

Proposition 3.48. *The derived functor of a left (or right) exact functor preserve exact triangles.*

Proof. We have to check separately that σ, F and Q preserve exact triangles. Q is an equivalence and maps standard triangles to standard triangles by definition, thus it preserves distinguished triangles and so does its inverse σ. For F we note that by construction F preserves cones and thus standard triangles, and F sends chain homotopy equivalences to chain homotopy equivalences. □

A functor of derived categories preserving exact triangles is called *exact*.

Proposition 3.49. *Given a left exact functor $F : \mathcal{A} \to \mathcal{B}$ there is a natural transformation $\epsilon : Q_{\mathcal{B}} \circ F \to RF \circ Q_{\mathcal{A}}$ between functors $K^+(\mathcal{A}) \to D^+(\mathcal{B})$. Moreover, the pair (RF, ϵ) is initial, i.e. given any exact functor $G : D^+(\mathcal{A}) \to D^+(\mathcal{B})$ and $\phi : Q_{\mathcal{B}} \circ F \to G \circ Q_{\mathcal{A}}$ there is a unique natural transformation $\rho : RF \to G$ such that $\rho \circ \epsilon = \phi$.*

Not a proof. The existence of ϵ follows from the existence of a natural transformation $\mathbf{1}_{K^+(\mathrm{Inj}(\mathcal{A}))} \to \sigma \circ Q_{\mathcal{A}}$, by Corollary 3.45 we can use the unit of the equivalence.

For the full proof of universality see [GM03, III.6.8 respectively III.6.11] or [Wei95, 10.5.6]. □

In fact, this characterisation of the derived functor via the universal property is arguably the correct definition of the derived functor.

Remark 3.50. Total derived functors may be constructed even in the absence of injective or projective resolutions. Say a class of objects $\mathcal{R} \subset \mathcal{A}$ closed under finite direct sums is *adapted* to a left exact functor $F : K^+(\mathcal{A}) \to K^+(\mathcal{B})$ if

1. F preserves acyclic complexes with entries in $\mathcal{R}$ and
2. any $A \in K^+(\mathcal{A})$ is a subobject of some R_A in $\mathcal{R}$.

Then $RF(A)$ is defined as $F(R_A)$ and this is well-defined (in particular it agrees with the definition using injectives if that is available) and it has all the desirable properties, see [GM03].

4. Sheaf cohomology

4.1. Derived global sections

Recall that the functor of *global sections* $\Gamma(X, -)$ on a space X sends any sheaf $\mathcal{F}$ to $\mathcal{F}(X)$.

If we denote by $\pi : X \to *$ the projection to the point then Γ can be identified with the pushforward π_*.

In particular by Theorem 2.42 Γ is left exact.

Definition 4.1. The *i-th cohomology* of the sheaf $\mathcal{F}$ is $R^i\Gamma(X, \mathcal{F})$, which we also write as $H^i(X, \mathcal{F})$.

In particular, as with all derived functors, $R^0\Gamma(X, \mathcal{F}) = \Gamma(X, \mathcal{F}) = \mathcal{F}(X)$.

One of the key themes of this course will be working out $R\Gamma(X, \underline{R})$ if X is a reasonably nice topological space. In particular it agrees with singular cohomology and thus we recover one of our favourite tools from algebraic topology just from considering one particular sheaf. Other sheaves will give us more tools.

In order to even define derived global sections we need to have a supply of injective sheaves. You will see an in exercise that $\mathsf{Sh}(X)$ does not have enough projectives in general. It does however have enough injectives.

Lemma 4.2. *The pushforward of an injective sheaf is injective.*

Proof. We take $f : X \to Y$ and $\mathcal{I}$ injective on X. To check $f_*\mathcal{I}$ is injective we need to show $\mathrm{Hom}(\mathcal{G}, f_*I) \to \mathrm{Hom}(\mathcal{F}, f_*\mathcal{I})$ is surjective for any injection $\mathcal{F} \to \mathcal{G}$. But by Theorem 2.42 this amounts to showing $\mathrm{Hom}(f^{-1}\mathcal{G}, I) \to \mathrm{Hom}(f^{-1}\mathcal{F}, \mathcal{I})$ is surjective for any injection $\mathcal{F} \to \mathcal{G}$. But f^{-1} is exact by Example 3.10, so we may assume $f^{-1}\mathcal{F} \to f^{-1}\mathcal{G}$ is injective and the surjectivitiy follows as $\mathcal{I}$ is injective. □

Lemma 4.3. *Let $\mathcal{A}$ be an abelian category with enough injectives and X a topological space. Then* $\mathsf{Sh}(X, \mathcal{A})$ *has enough injectives.*

Proof. We fix a sheaf $\mathcal{F}$ and consider $x \in X$. We consider $\mathcal{F}_x$ in $\mathcal{A}$ with its injective hull $\mathcal{F}_x \to E(\mathcal{F}_x)$.

By Lemma 4.2 $i_{x*}E(\mathcal{F}_x)$ is again injective. Moreover one can easily check that any product of injective objects in injective.

Thus for any $\mathcal{F}$ we consider this composition of injections :

$$\mathcal{F} \to \prod_{x \in X} i_{x*}\mathcal{F}_x \to \prod_{x \in X} i_{x*}E(\mathcal{F}_x)$$

The first map is an injection as in the proof of Lemma 2.18, and for the second map we note that i_{x*} as a left exact functor takes injections to injections.

This provides the desired injection to an injective sheaf. □

By a similar argument the category of sheaves of $\mathcal{R}$-modules on a ringed space also has enough injectives.

The injective replacement is an extra level of nuisance and it can be avoided.

Definition 4.4. A sheaf $\mathcal{F}$ is called *flasque* or *flabby* if for any open sets $V \subset U$ the restriction map $r_{UV} : \mathcal{F}(U) \to \mathcal{F}(V)$ is surjective.

Example 4.5. Let F_x be an abelian group for any x in some subset S of X. Then the sheaf $\prod i_{x*}F_x$ is flabby by definition.

Lemma 4.6. *Consider a short exact sequence of sheaves*

$$0 \to \mathcal{F} \xrightarrow{i} \mathcal{I} \xrightarrow{q} G \to 0$$

such that $\mathcal{F}$ is flasque and $\mathcal{I}$ is injective. Then for any open set U there is a surjection $\mathcal{I}(U) \to \mathcal{G}(U)$.

Proof. Take $\sigma \in \mathcal{G}(U)$. We consider the set S of all pairs (W, τ) with $W \subset U$ open, $\tau \in \mathcal{I}(W)$ and $q(\tau) = \sigma|_W$. The set is naturally partially ordered by inclusion of open sets and restriction of sections. As any chain has an upper bound (provided by the union of all opens in the chain) there is a maximal element of S. We will show this maximal element must be (U, τ_U) where τ_U is a lift of σ.

Note that for any two pairs (W_1, τ_1) and (W_2, τ_2) in S we may consider $\tau_1|_{W_1 \cap W_2} - \tau_2|_{W_1 \cap W_2}$ which lies in the image of $\mathcal{F}(W_1 \cap W_2)$ and by flasqueness of $\mathcal{F}$ it may be extended to a section $\rho \in \mathcal{F}(W_1)$. Now $\tau_1' = \tau_1 - i\rho$ is in $\mathcal{I}(W_1)$ and its restriction agrees with τ_2 on $W_1 \cap W_2$ and we may thus glue τ_1' and τ_2 to a section τ in $\mathcal{I}(W_1 \cup W_2)$ with image $\sigma|_{W_1 \cup W_2}$.

The upshot is that $(W_1 \cup W_2, \tau)$ is in S.

So assume (W_{max}, τ_{max}) is the maximal element of S but $W_{max} \neq U$. Then let x be a point in $U \setminus W_{max}$. By exactness there is (W_x, τ_x) in S with x contained in W_x. By what we have shown above $(W_{max} \cup W_x, \tau)$ is in S for some τ, contradicting maximality. □

Lemma 4.7. *Any flasque sheaf is Γ-acyclic.*

Proof. Let $\mathcal{F}$ be flabby and take an injective $\mathcal{I}$ as in the proof of Lemma 4.3 with $\mathcal{F} \subset \mathcal{I}$. As in Example 4.5 $\mathcal{I}$ is flasque. (In fact it's not hard to show any injective sheaf is flasque.)

We complete to a short exact sequence

$$0 \to \mathcal{F} \to \mathcal{I} \xrightarrow{q} \mathcal{G} \to 0$$

By Lemma 4.6 for any open U of X the map $\mathcal{I}(U) \to \mathcal{G}(U)$ is surjective.

Thus we see $0 \to \mathcal{F}(X) \to \mathcal{I}(X) \to \mathcal{G}(X) \to 0$ is exact. As $\mathcal{I}$ is injective the long exact sequence from Corollary 3.28 shows $H^1(X, \mathcal{F}) = 0$ and $H^k(X, \mathcal{F}) = H^{k-1}(X, \mathcal{G})$ for $k \geq 2$.

Lemma 4.6 also shows that $\mathcal{G}$ is flasque: Consider any $U \subset V$ and $\sigma \in \mathcal{G}(U)$. We may lift to $\mathcal{I}(U)$ by the lemma and to a section τ in $I(V)$ by flasqueness of $\mathcal{I}$. But then the image of τ in $\mathcal{G}(V)$ lifts σ. Thus by induction $H^{k-1}(X, \mathcal{G}) = 0$ and all higher cohomology of $\mathcal{F}$ vanishes. □

4.2. Mayer-Vietoris sequences

Theorem 4.8. *Let $X = U \cup V$ with U, V open and let $\mathcal{F}$ be a sheaf on X. Then there is a long exact sequence of homologoy groups*

$$\cdots \to H^k(X, \mathcal{F}) \xrightarrow{a} H^k(U, \mathcal{F}) \oplus H^k(V, \mathcal{F}) \xrightarrow{b} H^k(U \cap V, \mathcal{F}) \to H^{k+1}(X, \mathcal{F}) \to \cdot$$

with the maps $a : s \mapsto (s|_U, -s|_V)$ and $b : (s, t) \mapsto s|_{U\cap V} + t|_{U\cap V}$.

Proof. By the sheaf condition we the following sequence which is exact on the left:

$$0 \to \Gamma(X, \mathcal{F}) \to \Gamma(U, \mathcal{F}) \oplus \Gamma(V, \mathcal{F}) \to \Gamma(U \cap V, \mathcal{F})$$

Note that if we assume that $\mathcal{F}$ is flasque, then the sequence is also exact on the right. So we replace $\mathcal{F}$ be a flasque resolution $F^\bullet$ and obtain the short exact sequence of complexes

$$0 \to \Gamma(X, F^\bullet) \to \Gamma(U, F^\bullet) \oplus \Gamma(V, F^\bullet) \to \Gamma(U \cap V, F^\bullet) \to 0$$

By the Snake Lemma 3.27 and the fact flasque resolutions compute cohomology (Lemma 4.7) we obtain the desired long exact sequence. □

Sheaf cohomology also admits a Mayer-Vietoris sequence for closed subsets which has a quite different proof.

Lemma 4.9. *Let $Z \subset X$ be closed. Then $R\Gamma(Z, \mathcal{F}) \cong R\Gamma(X, i_*\mathcal{F})$.*

Proof. We choose an injective I resolution of $\mathcal{F}$ on Z. By Lemma 4.2 i_*I is injective on Z. As i_* is exact by Example 3.11 we have $i_*\mathcal{F} \simeq i_*I$ by Lemma 3.47, thus i_*I is an injective resolution of $i_*\mathcal{F}$ and $\Gamma(X, i_*I) = \Gamma(Z, I)$ computes both $R\Gamma(X, i_*\mathcal{F})$ and $R\Gamma(Z, \mathcal{F})$. □

Theorem 4.10. *Let $X = A \cup B$ with A, B closed subsets of X and let $\mathcal{F}$ be a sheaf on X (we also write $\mathcal{F}$ for the restriction to A, B).. Then there is a long exact sequence*

$$\cdots \to H^k(X, \mathcal{F}) \xrightarrow{a} H^k(A, \mathcal{F}) \oplus H^k(B, \mathcal{F}) \xrightarrow{b} H^k(A \cap B, \mathcal{F}) \to H^{k+1}(X, \mathcal{F}) \to \cdot$$

with the maps $a : s \mapsto (s|_A, -s|_B)$ and $b : (s, t) \mapsto s|_{A\cap B} + t|_{A\cap B}$.

Proof. We denote the inclusions by $i : A \to X$, $j : B \to X$ and $h : A \cap B \to X$. Then there is a short exact sequence of sheaves

$$0 \to \mathcal{F} \to i_* i^{-1}\mathcal{F} \oplus j_* j^{-1}\mathcal{F} \to h_* h^{-1}\mathcal{F} \to 0$$

where the maps are induced by the (sums and differences of) the adjunctions $i^{-1} \dashv i_*$ etc.

The exactness may be checked at stalks and we may compute $(i_* i^{-1}\mathcal{F})_x = \mathcal{F}_x$ etc. Thus we reduce to the exactness of $\mathcal{F}_x \to \mathcal{F}_x \to 0$ respectively $\mathcal{F}_x \to \mathcal{F}_x \oplus \mathcal{F}_x \to \mathcal{F}_x$ (depending which of $A, B, A \cap B$ the point lies in).

Finally we take the long exact sequence on cohomology from Corollary 3.28 and the result follows from Lemma 4.9. □

Remark 4.11. From algebraic topology you may remember the Mayer-Vietoris sequence for open decompositions, but not for closed decompositions. This is an indication that for some topological spaces there is a difference between singular cohomology and sheaf cohomology with constant coefficients.

We would now like to compute some cohomology of basic topological spaces, but the Mayer-Vietoris sequence is no good if we don't know the cohomology of A and B. And so far we don't know any cohomology, except for injective and flabby sheaves.

4.3. Cohomology with constant coefficiets

We will now do an honest computation of some sheaf cohomology from first principle.

Lemma 4.12. *Let $I = [0, 1]$ with its Euclidean topology and let $\mathcal{F}$ be a sheaf on I such that for all $x \in I$ there is a surjection $\mathcal{F}(I) \to \mathcal{F}_x$. Then $H^j(I, \mathcal{F}) = 0$ for $j > 0$.*

In particular $H^j(I, \underline{R}) = 0$ for $j > 0$.

Proof. Fix $j > 0$. Let $s \in H^j(I; \mathcal{F})$ and for any $t \leq t'$ we consider restriction maps $r_{t,t'} : H^j(I; \mathcal{F}) \to H^j([t, t'], \mathcal{F})$.

We now look at the set $J = \{t \in [0, 1] \mid r_{0,t}(s) = 0\}$.

As the point has no higher cohomology we have $0 \in J$. Moreover as we can factor restriction maps it follows that for $0 \leq t \leq t'$ from $t' \in J$ we can deduce $t \in J$. So J is an interval. Next it is true that for all $t \in I$ we have

$$H^j([0, t]; \mathcal{F}) = \varinjlim_{t' > t} H^j([0, t']; \mathcal{F}).$$

(This is a (hard) exercise on the 6th example sheet, a proof will be added to the notes later.)

This equation follows from the definition of the restriction of a sheaf to a closed subset (applied to the flabby resolution F of $\mathcal{F}$) by replacing the colimit on the right hand side by a colimit of cohomology groups of open intervals (the diagram of open and of closed intervals are both cofinal in the diagrom of all intervals, thus they give the same colimit).

It follows that $r_{0,t}(s) = 0$ implies that $r_{0,t'}(s) = 0$ for some $t < t'$ and J is open.

We choose $t_0 = \sup J$ and some $t \leq t_0$. Then we consider the Mayer-Vietoris sequence for a closed cover by writing $[0, t_0] = [0, t] \cup [t, t_0]$ giving

$$\cdots \to H^j([0, t_0]; \mathcal{F}) \to H^j([0, t]; \mathcal{F}) \oplus H^j([t, t_0]; \mathcal{F}) \to H^j(\{t\}; \mathcal{F}) \to \cdots$$

which gives

$$H^j([0, t_0]; \mathcal{F}) \cong H^j([0, t]; \mathcal{F}) \oplus H^j([t, t_0]; \mathcal{F}).$$

Note that for $j = 1$ this needs the assumption that $\mathcal{F}(I) \to \mathcal{F}_t$ is surjective.

As $\varinjlim_{t<t_0} H^j([t, t_0]; \mathcal{F}) = 0$ we find $t < t_0$ with $r_{t,t_0}(s) = 0$. By construction of J we also have $r_{0,t} = 0$. Together this gives $r_{0,t_0}(s) = 0$. Thus the interval J is also closed and we have $J = I$, proving the vanishing of $H^j(I; \mathcal{F})$.

The application follows as $\underline{R}(I) \to \underline{R}_x$ is just the identity map. □

Another example to which this lemma applies is smooth functions on the interval: any germ comes from a smooth function defined on the interval. It does not apply to analytic functions: The germ of an analytic function is a Taylor series with positive radius of convergence which may not define an analytic function on all of I.

We constructed here in an ad hoc way a map $H^*(Y; \mathcal{F}) \to H^*(X; r^{-1}\mathcal{F})$ for a subspace $r : X \to Y$.

For later use we note the following more general statement.

Lemma 4.13. *For any $f : X \to Y$ and $\mathcal{F} \in \mathsf{Sh}(Y)$ and $\mathcal{G} \in \mathsf{Sh}(X)$ equipped with a map $f^{-1}\mathcal{F} \to \mathcal{G}$ there is a natural map $f^\# : H^*(Y; \mathcal{F}) \to H^*(X; \mathcal{G})$.*

Proof. We define the map in the derived category $D(\mathsf{Ab})$ where we have

$$Rp_{Y*}\mathcal{F} \to Rp_{Y*}Rf_*f^{-1}\mathcal{F} \cong Rp_{X*}f^{-1}\mathcal{F} \to Rp_{X*}\mathcal{G}$$

where the first map is induced by the unit of the adjunction $f^{-1} \vdash Rf_*$, which is the derived version of the adjunction $f^{-1} \dashv f_*$, see Lemma 4.14 below. □

In particular the lemma applies to the case $\mathcal{G} = f^{-1}\mathcal{F}$ and to constant sheaves as $\underline{R}_X = p_X^{-1}R = (p_Y \circ f)^{-1}R = f^{-1}\underline{R}_Y$.

Note that equivalent to a map $f^{-1}\mathcal{F} \to \mathcal{G}$ is a map $\mathcal{F} \to f_*\mathcal{G}$.

Lemma 4.14. *There is an adjunction $f^{-1} \dashv Rf_* : D(Y) \to D(X)$.*

Proof. The adjunction $f^{-1} \dashv f_*$ of sheaves from Theorem 2.42 lifts to an adjunction of categories of chain complexes of sheaves and indeed of homotopy categories: In fact there is an isomorphism of hom complexes $\underline{\mathrm{Hom}}_{\mathsf{Ch}(\mathsf{Sh}(Y))}(f^{-1}A, B) \cong \underline{\mathrm{Hom}}_{\mathsf{Ch}(\mathsf{Sh}(Y))}(f^{-1}A, B)$ as the adjunction isomorphism is by naturality compatible with differentials. Next we note that for an injective resolution I of B we have

$$\mathrm{Hom}_{D(X)}(f^{-1}A, B) \cong \mathrm{Hom}_{D(X)}(f^{-1}A, I) \cong \mathrm{Hom}_{D(Y)}(A, f_*I) \cong \mathrm{Hom}_{D(Y)}(A, Rf_*B)$$

using the fact that f_*I is still injective by Lemma 4.2 andwe can compute homs in the derived category by Theorem 3.42. This completes the proof. □

At this point we would like to deduce homotopy invariance of sheaf cohomolgy with constant coefficients, but this will need some more ingredients that we will only introduce later.

Example 4.15. We can compute $H^(S^1, \underline{R})$ from the closed cover consisting of two intervals A, B meeting in a point. As all higher cohomology of $A, B, A \cap B$ vanishes the long exact sequence of Theorem 4.10 reduces to

$$0 \to H^0(S^1, \underline{R}) \to R \oplus R \xrightarrow{\alpha} R \oplus R \to H^1(S^1, \underline{R}) \to 0$$

where $\alpha : (s, t) \mapsto (s + t, s + t)$ is given by the restriction maps. We obtain $H^0(S^1, \underline{R}) \cong H^1(S^1, \underline{R}) \cong R$ as expected, with all other cohomology vanishing.

4.4. Čech Cohmology

The Mayer-Vietoris sequence allows us to compute cohomology of a union $A \cup B$ from the cohomolies of A, B and $A \cap B$. It is natural to ask what happens for a triple union $A \cup B \cup C$. If the pieces themselves have higher cohomology this gets computationally very complicated. Imagine instead that all the pieces have no higher cohomology.

Then Mayer-Vietoris simplifies to

$$R\Gamma(A \cup B, \mathcal{F}) \simeq (\Gamma(A, \mathcal{F}) \oplus \Gamma(B, \mathcal{F}) \to \Gamma(A \cap B, \mathcal{F}))$$

where we suppress the notation for restriction of sheaves when they are obvious from context.

And for $X = \cup U_i$ with $R^{>0}\Gamma(\cap_{i \in I} U_i, \mathcal{F}) = 0$ for any subset I of $\{1, \dots, n\}$ we may expect

$$R\Gamma(X, \mathcal{F}) \simeq \left(\oplus_i \Gamma(U_i, \mathcal{F}) \to \oplus_{i,j} \Gamma(U_i \cap U_j, \mathcal{F}) \to \cdots \Gamma(U_1 \cap \cdots \to \cap U_n, \mathcal{F})\right)$$

This is indeed true.

Definition 4.16. We fix a topological space X and a cover $\mathfrak{U} = \{U_i\}_{i \in I}$ which we equip with a well-ordering. Given a presheaf $\mathcal{P}$ on X we define the *Čech complex* $\check{C}^\bullet(\mathfrak{U}, \mathcal{P})$ by

$$\check{C}^p(\mathfrak{U}, \mathcal{P}) = \prod_{|J|=p} \mathcal{P}(\cap_{j \in J} U_j)$$

with differential given by the alternating sum of the restriction maps

$$(\delta f)_{i_0 \dots i_{p+1}} = \sum_{k=0}^{p+1} (-1)^k f_{i_0 \dots \widehat{i_k} \dots i_{p+1}} |_{U_{i_0 \dots i_{p+1}}}$$

where we write $U_{i_0 \dots i_{p+1}}$ for $U_{i_0} \cap \cdots U_{i_{p+1}}$ and assume the indices are ordered $i_0 < i_1 < \cdots < i_{p+1}$.

The cohomology of the Čech complex is denoted by $\check{H}(\mathfrak{U}, \mathcal{P})$ and is called the *Čech cohomology* of $\mathcal{P}$ on X with respect to the cover $\mathfrak{U}$.

This is indeed a complex as $d^2 = 0$ by observing each restriction occurs twice with opposite signs.

In particular for a two-element cover we recover the Mayer-Vietoris complex from above.

Lemma 4.17. *If $\mathcal{F}$ is a sheaf we have $\check{H}^0(\mathfrak{U}, \mathcal{F}) = \Gamma(X, \mathcal{F})$.*

Proof. for $f \in \check{C}^0$ to be in the kernel of δ means $f|_{U_j} - f|_{U_j} = 0$ on each intersection $U_i \cap U_j$. By the sheaf axioms this means exactly that f is a global section. □

Definition 4.18. For a sheaf $\mathcal{F}$ on a space X with cover $\mathfrak{U}$ we define the sheaf version of the Čech complexby

$$\check{\mathcal{C}}^p(\mathfrak{U}, \mathcal{F}) = \prod_{|J|=p} \iota_{J*}\iota_J^{-1}\mathcal{F}$$

where we write $\iota_J : \cap_J U_j \to X$ for any inclusion of an intersection of U_i into X. The differential is given by the same formula as in Definiton 4.16.

We elaborate a bit on the differential. For any inclusion $f : I \to J$ from a set with p elements into a set with $p + 1$ elements there is an associatedd inclusion $\kappa : U_J \to U_I$ and $\iota_I \circ \kappa = \iota_J$. The unit $\mathbf{1} \to \kappa_*\kappa^{-1}$ induces a map $\iota_{I*}\iota_I^{-1} \to \iota_{J*}\iota_J^{-1}$ which is just the restriction along κ. Note that we have $\iota_{I*}\iota_I^{-1}\mathcal{F}(V) = (\iota_I^{-1}\mathcal{F})(V \cap U_I) = \mathcal{F}(V \cap U_I)$ etc. which simplifies things. The differential is the alternating sum of these maps.

Note we have $\Gamma(X, \check{\mathcal{C}}(\mathfrak{U}, \mathcal{P})) = \check{C}(\mathfrak{U}, \mathcal{P})$ by construction.

Lemma 4.19. *$H^i(\mathcal{C}(\mathfrak{U}, \mathcal{F})) = 0$ for $i > 0$. In particular the natural map $\epsilon : \mathcal{F} \to \check{\mathcal{C}}(\mathfrak{U}, \mathcal{F})$ is a quasi-isomorphism.*

Proof. We define the map ϵ as a product of the units $\mathcal{F} \to \iota_{i*}\iota_i^*\mathcal{F}$. If $\mathcal{F}$ is a sheaf this is an isomorphism on H^0 by Lemma 4.17.

We now check exactness at $p > 0$, which as usual we do on stalks. For $x \in X$ we find U_j containing x. Then an element in $\check{\mathcal{C}}_x^p$ is represented by (V, s) and we may assume $V \subset U_j$. For any p-tuple $i_0 < \cdots < i_{p-1}$ we define

$$(hs)_{i_0,\ldots,i_{p-1}} := s_{ji_0\ldots i_{p-1}}$$

where we define for an unordered index set with repetitions $s_{j_0\cdots j_p} = 0$ and without repetitions

$$s_{j_0\cdots j_p} = \epsilon(\sigma)s_{j_{\sigma^{-1}(0)}\cdots j_{\sigma^{-1}(p)}}$$

with σ the permutation putting the j_i in order and $\epsilon(\sigma)$ its sign. Note that this definition does indeed make sense as $V \subset U_j$ and thus $V \cap U_{i_0\cdots i_{p-1}} = V \cap U_{ji_0\cdots i_{pi-1}}$.

Then we check $(\delta h + h\delta)(s) = s$ on $\mathcal{C}^p$ for $p \geq 1$. This shows that all cohomology groups vanish. □

We may consider Čech resolutions not just for single sheaves but for complexes of sheaves. Then we have a complex with two distinct gradings and two distinct kinds of differentials.

Definition 4.20. A *double complex* in an abelian category $\mathcal{A}$ is a collection of objects L^{pq} in $\mathcal{A}$ for $p, q \in \mathbb{Z}$ equipped with horizontal differentials $d_h : L^{pq} \to L^{pq+1}$ and vertical differentials $d_v : L^{pq} \to L^{p+1q}$ such that $(d_h + d_v)^2 = 0$. As this equality has to hold in each degree it may translated into $d_h^2 = 0$, $d_v^2 = 0$ and $d_h d_v = -d_v d_h$.

The associated *total complex* $\mathrm{Tot}(L)$ is defined by $\mathrm{Tot}(L)^i = \oplus_{p+q=i} L^{pq}$ with differential $d = d_h + d_v$.

We say an element of L^{pq} has *bidegree* (p, q). Note that in the definition of the total complex we could have just as well used the direct product on each anti-diagonal. These two constructions are both useful in their own right, denoted $\mathrm{Tot}^{\oplus}(A)$ and $\mathrm{Tot}^{\prod}(A)$.

We will be mainly focussed on *first quadrant double complexes* with $L^{pq} = 0$ if $p < 0$ or $q < 0$, and for these the two totalizations agree.

We introduce the following very useful. lemma. It may be viewed as a special case of the spectral sequence of a double complex.

Lemma 4.21. *Let $(A^{\bullet\bullet}, d_h, d_v)$ be a double complex with $A^{ij} = 0$ for $i < 0$ or $j < 0$. Assume that each row $A^{i\bullet}$ and each colum $A^{\bullet j}$ is exact except in degree 0. Define $B^{\bullet}$ be the complex of kernels of d_h^0 and let $C^{\bullet}$ be the complex of kernels of the d_v^0.*

Then the natural maps $B \to \mathrm{Tot}(A)$ and $C \to \mathrm{Tot}(A)$ are quasi-isomorphism.

Proof. By symmetry it suffices to consider the case of $B \to \mathrm{Tot}(A)$.

$$\begin{array}{ccccccccc}
B^2 & \longrightarrow & A^{20} & \longrightarrow & A^{21} & \longrightarrow & A^{22} & \longrightarrow & \cdots \\
\uparrow & & \uparrow & & \uparrow & & \uparrow & & \\
B^1 & \longrightarrow & A^{10} & \longrightarrow & A^{11} & \longrightarrow & A^{12} & \longrightarrow & \cdots \\
\uparrow & & \uparrow & & \uparrow & & \uparrow & & \\
B^0 & \longrightarrow & A^{00} & \longrightarrow & A^{01} & \longrightarrow & A^{02} & \longrightarrow & \cdots
\end{array}$$

Denote the map $B \to \mathrm{Tot}(A)$ by ι. Let us be explicit about the differnital: Given $b \in B^i$ we have $d_v(\iota b) \in \ker(d_h)$, thus it is of the form $\iota b'$ for some b' (which is unique by injectivitiy of ι), and this is our definition of the differenial db.

THis inclusion is indeed a chain map as $\iota(db) = d_v \iota b$ by definition and furthmer equals $d(\iota b)$ as $d_h \iota(b) = 0$. In particular the map is well-defined on cohomlogy.

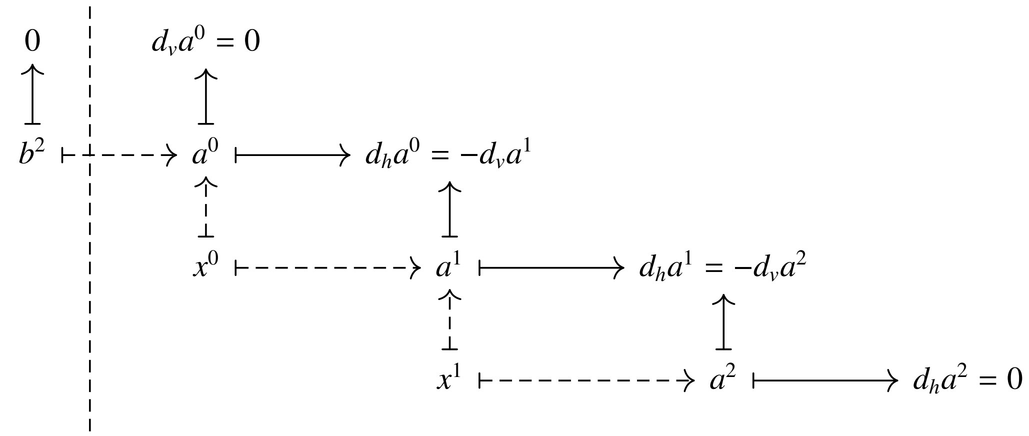

We check surjectivity on cohomology. Consider an cycle $a = \sum_{i=0}^{k} a^i$ of degree m in $\mathrm{Tot}(A)$. Here $k \leq m$ is maximal such that a^k is nonzero. As this is a cycle we have $d_h a^i = -d_v a^{i+1}$ and $d_h a^k = 0$. By exactness of the bottom column there must be x^{k-1} with $d_h x^{k-1} = a^k$. We now consider $\tilde{a} = a - dx^{k-1}$, which is homologous to a but the last nonzero entry in horizontal degree $k-1$ or smaller. Proceeding by induction we obtain that a is equivalent to a' in bidegree $(m, 0)$. As $d_v a' = 0$ this lies in the image of B.

Similarly we check injectivity on cohomology. Consider b with $db = 0$ such that $\iota b = dx$ with $x = \sum x^i$. As dx is concentrated in bidegree $(m, 0)$ we have $d_h x^i = -d_v x^{i+1}$ and $d_h x^k = 0$ for k large enough. Thus $x^k = d_h y^{k-1}$ and $x^{k-1} - d_v y^{x-1} \in \ker(d_h)$. Continuing inductively we have $x' = x + dy$ concentrated in bidegree $(m-1, 0)$ with $d_h x' = 0$ and $x' = ib'$ for a suitable b' which also satisfies $d_v b' = b$, showing $[b] = 0 \in H^m(B)$. □

Theorem 4.22. *Let $\mathcal{F}$ be a sheaf on a topological space X and $\mathfrak{U}$ a cover of X such that on each intersection of opens $H^{>0}(U_I, \mathcal{F}) = 0$. Then $H^*(X, \mathcal{F}) \cong \check{H}^*(\mathfrak{U}, \mathcal{F})$.*

Proof. We resolve $\mathcal{F}$ by a complex of flasque sheaves F and consider the double complex $\check{\mathcal{C}}^p(\mathfrak{U}, F^q)$ with horizontal differential given by the Čech differential and vertical differential induced by the differential $F^q \to F^{q+1}$.

This gives us a double complex of sheaves to which we will apply global sections.

The key claim now is that for any p, q the sheaf $\check{\mathcal{C}}^p(\mathfrak{U}, F^q)$ is again flasque: Restriction to opens, pushforward and product all preserve flasqueness. (This is true in general and particularly easy to check for flasque sheaves that are products of skyscraper sheaves, and that suffices for our puproses.)

Thus the row $\check{\mathcal{C}}^p(\mathfrak{U}, F^*)$ (fixed p) is a flasque resolution of $\check{\mathcal{C}}^p(\mathfrak{U}, \mathcal{F})$ as each $\mathcal{C}(U_I, F)$ resolves $\mathcal{C}(U_I, \mathcal{F})$. We take global sections, which computes the cohomology groups $\prod_I H^*(U_I, F)$. By assumption this is exact except in degree 0, where we obtain $C^*(\mathfrak{U}, F)$.

Next we consider the columns $\check{C}^*(\mathfrak{U}, F^q)$ (fixed q), which are Čech resolutions of the sheaves F^q. They are exact by Lemma 4.19, thus we have flabby resolutions of each F^q. But F^q is already flabby so the higher cohomology vanishes and in degree zero we obtain $\Gamma(X, F^*)$.

Thus the double complex satisfies the assumptions of Lemma 4.21 and We have a quasi-isomorphism

$$C^*(\mathfrak{U}, F) \to \operatorname{Tot} C^*(\mathfrak{U}, F^*) \leftarrow \Gamma(X, F^*).$$

Taking cohomology we obtain the desired result. □

Thus Čech cohomology reduces the problem of computing sheaf cohomology to explicit combinatorics, provided we can find small pieces of our topological space on which there is no higher cohomology.

For constant (or more generally locally constant) coefficients we will show that contractible sets serve this purpose.

In algebraic geometery cohomology of quasi-coherent sheaves can be computed by reducing to affine sets which have no higher cohomology.

Cohomology of coherent sheaves on a complex manifold can be computed by decomposing the space into Stein submanifolds (in particular a decomposition into polydiscs works).

Remark 4.23. This is a very useful result, but the dependence on an open cover is a somewhat unsatisfactory feature of Čech cohomology. As long as there are small enough pieces on which cohomology vanishes, Čech cohomology computes what we want. If not it does not even seem to be an invariant of the space and sheaf! One may define Čech cohomology in general as follows: Define a partial ordering on all open covers of X by saying $\mathfrak{U} < \mathfrak{V}$ if $\mathfrak{U}$ is a refinement of $\mathfrak{V}$, i.e. each open in $\mathfrak{U}$ is contained in some open in $\mathfrak{V}$. Then define $\check{H}^*(X, \mathscr{P}) = \operatorname{colim}_{\mathfrak{U}} \check{H}^*(\mathfrak{U}, \mathscr{P})$.

One can show this agree with sheaf cohomology in degrees 0 and 1, but not necessarily in general.

5. Sheaves on locally compact spaces

5.1. Compact support

We will assume in this chapter that X is locally compact Hausdorff. This is true for all subsets of $\mathbb{R}^n$, all CW complexes and all manifolds. It is not true for algebraic varieties with the Zariski topology.

Recall that for $\mathcal{F}$ an $\mathcal{A}$-valued sheaf on X we write $\Gamma(X, \mathcal{F})$ for $\mathcal{F}(X)$, which is nothing but $\pi_*(\mathcal{F})$ for $\pi : X \to *$ the projection.

We will now be interested in a variation, the space of sections of $\mathcal{F}$ with *compact support*.

To do this we restrict our attention to reasonably nice space:

Convention 5.1. From now on let all our spaces be locally compact and Hausdorff unless stated otherwise.

Locally compact means every point has a compact neighbourhood. (There are variations of this definition if X is not Hausdorff, but for Hausdorff spaces they all agree.)

Definition 5.2. Define the *support* of a section $s \in \mathcal{F}(U)$, denoted $\mathsf{supp}(s)$, to be the set of all $x \in U$ such that $s_x \neq 0$ in $\mathcal{F}_x$.

Define the *support* of a sheaf $\mathcal{F}$, denoted $\mathsf{supp}(\mathcal{F})$, to be the closure of the set of all $x \in X$ with $\mathcal{F}_x \neq 0$.

The support of a section is always closed: Take $x \notin \mathsf{supp}(s)$, then $s_x = 0$, thus for some open neighbourhoof V of x we have $s|_V = 0$ and then $V \cap \mathsf{supp}(s) = \emptyset$.

A map $f : X \to Y$ is called *proper* if the preimage of any compact set is compact. (If our spaces are not nice then there are several competing definitions of properness.)

Definition 5.3. Define $\Gamma_c(X, \mathcal{F}) \subset (X, \mathcal{F})$ to be space of *sections with compact support* of $\mathcal{F}$.

Remark 5.4. While $\Gamma(-, \mathcal{F})$ is a sheaf, we cannot define a sheaf with sections $\Gamma_c(-, \mathcal{F})$ as there is no restriction map for sections with compact support.

Instead $\Gamma_c(-, \mathcal{F})$ gives a *cosheaf*, the dual notion of a sheaf, i.e. a covariant functor $\mathrm{Op}(X) \to \mathcal{A}$. For any inclusion $U \subset V$ of open sets there is an inclusion map $\Gamma_c(U, \mathcal{F}) \to \Gamma_c(V, \mathcal{F})$.

Example 5.5. For $X = (0, 1)$ we have $\Gamma(X, \underline{\mathbb{R}}) = \mathbb{R}$ but $\Gamma_c(X, \underline{\mathbb{R}}) = 0$ (there support of a locally constant funtion is open as well as closed).

If on the other hand X is compact we have $\Gamma(X, \underline{\mathbb{R}}) = \mathbb{R} = \Gamma_c(X, \underline{\mathbb{R}})$.

Definition 5.6. For any function $f : X \to Y$ we can consider the functor $f_! : \mathsf{Sh}(X, \mathcal{A}) \to \mathsf{Sh}(Y, \mathcal{A})$, called *proper pushforward* or *direct image with proper support* defined by

$$f_!\mathcal{F}(V) = \{s \in \Gamma(f^{-1}(V), \mathcal{F}) \mid f|_{\mathsf{supp}(f)} : \mathsf{supp}(f) \to V \text{ is proper}\}$$

It is now easy to see that for $\pi : X \to *$ we have $\pi_! = \Gamma_c$.
Straight from the construction we can show:

Lemma 5.7. *1. There is a natural injection $f_!\mathcal{F} \to f_*\mathcal{F}$.*

2. The presheaf $f_!\mathcal{F}$ is indeed a sheaf.

3. The functor $f_! : \mathsf{Sh}(X) \to \mathsf{Sh}(Y)$ is left exact.

Proof. 1. By construction.

2. The condition of having proper support is local (this is an exercise in topology), so the sub-presheaf $f_!\mathcal{F} \subset f_*\mathcal{F}$ is also a sheaf.

3. This follows by combining (a) with left exactness of f_*.

□

Example 5.8. If f is proper then $f_! = f_*$. This is the case for example if f is a closed inclusion (as closed subset of a compact set is compact).

Example 5.9. Let $j : X \to Y$ a locally closed inclusion and $\mathcal{F}$ a sheaf on K. Then $j_!\mathcal{F}$ is the *extension by zero* satisfying $(j_!\mathcal{F})_x = \mathcal{F}_x$ if $x \in X$ and 0 otherwise. (This follows immediately by Proposition 5.14 below.) In particular $j_!$ is exact in this case.

As we have another left exact functor we can compute its right derived functors using injective resolutions.

Definition 5.10. We define the *cohomology with compact support* of a sheaf $\mathcal{F}$ on a space X as $H_c^*(X, \mathcal{F}) := H^*(R\Gamma_c(X, \mathcal{F}))$.
We also consider $Rf_!(\mathcal{F})$ for any $f : X \to Y$ and $\mathcal{F} \in \mathsf{Sh}(X)$.

5.2. Base change

We now carefully prove the general version of the result from Question 6.4. Recall that a Hausdorff space is *paracompact* if every open cover has a *locally finite* subcover, i.e. there is a subcover such that each point has a neighbourhood that meets only finitely many sets of the subcover.

Lemma 5.11. *Let K be a subset of a space X with a fundamental system of paracompact open neighbourhoods, or let K be a compact subset of an arbitrary space X*
Write $K = \cap_i U_i$ as the intersection of all of its open neighbourhoods $U_i \subset X$. Write $\mathcal{F}|_K$ for the pullback of $\mathcal{F}$ to K etc. Then $\Gamma(K, \mathcal{F}|_K) = \operatorname{colim}_i \Gamma(U_i, \mathcal{F}_{U_i})$.

The condition is in particular satisfied if X is hereditarily paracompact, i.e. every subspace is paracompact. This is visibly the case for $\mathbb{R}^n$. It is also the case for every metric space and every CW complex.
The proof uses the following two facts from point set topology:

1. Any paracompact space is normal.

2. A locally finite cover U_i of a normal space has a locally finite cover V_i with $\overline{V_i} \subset U_i$.

The (harder) second statement is sometimes called the shrinking lemma, see [Wil12, Theorem 15.10] for a proof.

Proof. We just write $\mathcal{F}$ for all of its restrictions. The natural maps $\Gamma(U_i, \mathcal{F}_{U_i}) \to \Gamma(K, \mathcal{F}|_K)$ induce a map

$$f : \operatorname*{colim}_i \Gamma(U_i, \mathcal{F}) \to \Gamma(K, \mathcal{F})$$

It is clear that f is injective: If $f(s) = 0$ then each point in K has a neighbourhood (in X) on which s vanishes, and s is zero on their union and has image equal to 0 in the colimit.

For surjectivity consider a section of $\Gamma(K, \mathcal{F}|_K)$, given by a collection $s_\alpha \in \mathcal{F}(U_\alpha)$ for $\alpha \in I$ where $K \subset U_\alpha$ and s_α and s_β agree on $U_\alpha \cap U_\beta \cap K$. We need to find for each point an open neighbourhood in X where the sections agree. By assumption we may assume that I is locally finite: If K is compact we may even assume that I is finite. If on the other hand K has a system of paracompact neighbourhoods we find a paracompact open B such that $K \subset B \subset \cup U_\alpha$ and we replace $\{U_\alpha\}$ by a locally finite subcover of $\{B \cap U_\alpha\}$.

By the facts (a) and (b) recalled above we can use paracompactness to further refine to a covering of K consisting of V_α with $\overline{V}_\alpha \subset U_\alpha$.

We define $W = \{x \in X \mid x \in \overline{V}_\alpha \cap \overline{V}_\beta \Rightarrow s_\alpha(x) = s_\beta(x)\}$. This is going to be the open neighbourhood on which we extend s. For each $x \in X$ define $J(x) = \{\alpha \mid x \in \overline{V}_\alpha\}$. For each x there is a neighbourhood $N(x)$ such that $y \in N(x)$ implies $J(y) \subset J(x)$, i.e. a neighbourhood not meeting any $\overline{V}_\alpha$ except the ones containing x. To find $N(x)$ we use local finiteness of the cover to find $N'(x)$ that meets only finitely many $\overline{V}_\alpha$, and then we remove all $\overline{V}_\alpha$ not containing x to obtain $N(x)$ as claimed. As we removed finitely many closed sets this is indeed a neighboourhood of x.

By construction $K \subset W$.

Moreover, W is open: Since $J(x)$ is finite the sections s_α for $\alpha \in J(x)$ which agree in $\mathcal{F}_x$ must agree in a neighbourhood of x, and the intersection of this neighbourhood with $N(x)$ is contained in W.

We define $t \in \mathcal{F}(W)$ by $t(x) = s_\alpha(x)$ for $x \in V_\alpha \cap W$, by definition of W this is well-defined. Now the image of (W, t) in the colimit is sent by f to s and we have completed the proof. □

Corollary 5.12. *Let K be a subset of X with a fundamental sytem of paracompact neighbourhoods. The restriction of a flabby sheaf from X to K is flabby.*

Proof. Consider open subsets $W \cap K \subset V \cap K \subset K$ and the restriction map $r : \Gamma(V \cap K, \mathcal{F}) \to \Gamma(W \cap K, \mathcal{F})$ for a flabby sheaf $\mathcal{F}$. Take $s \in \Gamma(W \cap K, \mathcal{F})$. By Lemma 5.11 we find some open neighbourhood W' of $W \cap K$ with s represented by (W', s'). Then s' extends to a section $\tilde{s}$ on $W' \cup V$ by flabbiness of $\mathcal{F}$ and the restriction of $\tilde{s}$ to $V \cap K$ is a preimage of s under r. □

Corollary 5.13. *Let K be a subset of X with a fundamental sytem of paracompact neighbourhoods. Let $\mathcal{F} \in \mathsf{Sh}(X)$. Then $H^*(K, \mathcal{F}|_K) = \operatorname{colim}_i H^*(U_i, \mathcal{F}_{U_i})$*

Proof. We resolve $\mathcal{F}$ by a flabby complex F. By Corollary 5.12 the restriction to K is a flabby resolution of $\mathcal{F}|_K$. We now apply Lemma 5.11 to find $\Gamma(K, F|_K) = \mathrm{colim}_i \Gamma(U_i, F_{U_i})$, and taking cohomology proves the corollary. (Cohomology commutes with colimits of complexes.) □

We now turn to a key application, comparing the stalk of the proper pushforward with global sections of the fiber.

Proposition 5.14. *Let $f : Y \to X$, $x \in X$ and $\mathcal{F}$ a sheaf on Y be given. Then there is an isomorphism $(f_!\mathcal{F})_x \cong \Gamma_c(f^{-1}(x), \mathcal{F})$.*

Proof. There is indeed a canonical morphism α, for every $x \in V \subset X$ we have $\Gamma(V, f_!\mathcal{F}) \to \Gamma(f^{-1}(V), \mathcal{F}) \to \Gamma(f^{-1}(x), \mathcal{F}|_{f^{-1}(x)})$ which factors through Γ_c.

First we check injectivity: Take (V, s) representing an element t in $(f_!\mathcal{F})_x$, i.e. V an open neighbourhood of x and $s \in \Gamma(V, f_!\mathcal{F})$ with $\mathsf{supp}(s) \to V$ proper. We assume $\alpha(t) = 0$. Thus $\mathsf{supp}(s) \cap f^{-1}(x) = \emptyset$, i.e. $x \notin f(\mathsf{supp}(s))$. But the latter set is compact and thus closed, so x has an open neighbourhood on which t vanishes.

For surjectivity we fix $s \in \Gamma_c(f^{-1}(x); \mathcal{F}|_{f^{-1}(x)})$ and denote its (compact) support by K. We have to produce a neighbourhood W of x together with a section $\tilde{s}$ in $\mathcal{F}(f^{-1}(W))$ such that f is proper on the support of $\tilde{s}$.

By Lemma 5.11 there is an open neighbourhood U of K in Y and $t \in \Gamma(U; \mathcal{F})$ with $t|_K = s|_K$. We can moreover assume $t|_{f^{-1}(x)\cap U} = s|_{f^{-1}(x)\cap U}$ by shrinking U if necessary (so it doesn't meet $\mathsf{supp}(t) \setminus \mathsf{supp}(s)$, we can subtract the closed set $\mathsf{supp}(t - s)$).

Now we let V be a relatively compact open neighbourhood of K in U, i.e. we have $K \subset V \subset \overline{V} \subset U$ with $\overline{V}$ compact.

We consider $V' = \overline{V} \cap \mathsf{supp}(t)$, which is compact. Then $V' \setminus V$ is closed and its image $f(V' \setminus V)$ is closed in X as proper maps between Hausdorff spaces are closed. We also have

$$f^{-1}(x) \cap V' \subset f^{-1}(x) \cap U \cap \mathsf{supp}(t) \subset K \subset V$$

so $x \notin f(V' \setminus V)$ and we can find an open neighbourhood W of x with $W \cap f(V' \setminus V) = \emptyset$, i.e. $f^{-1}(W) \cap V' \subset V$.

We now define $\tilde{s} \in \Gamma(f^{-1}(W); \mathcal{F})$ by setting

$$\tilde{s}|_{f^{-1}(W)\setminus V'} = 0$$
$$\tilde{s}|_{f^{-1}(W)\cap V} = t|_{f^{-1}(W)\cap V}$$

This is well-defined as the support of t does not meet the intersection $f^{-1}(W) \setminus V' \cap f^{-1}(W) \cap V$ as $\mathsf{supp}(t) \cap V \subset V'$.

Moreover $\mathsf{supp}(\tilde{s})$ is contained in $f^{-1}(W) \cap \bar{V}$ and the map $f^{-1}(W) \cap \bar{V} \to W$ is proper: The preimage of any compact in W is a closed subset of the compact set $\bar{V}$. Thus $f|_{\mathsf{supp}(\tilde{s})} : \mathsf{supp}(\tilde{s}) \to W$ is proper and we have found $(W, \tilde{s}) \in (f_!\mathcal{F})_x$. the restriction to $f^{-1}(x)$ is equal to s by construction. □

Note that the analogue of the proposition does not hold for the regular pushforward, see the example sheet!

The proposition is the special case of the key property called *base change*.

Theorem 5.15. *Given a pullback square of locally compact spaces*

$$\begin{array}{ccc} \tilde{Y} & \xrightarrow{\tilde{f}} & Z \\ \downarrow{\scriptstyle \tilde{g}} & & \downarrow{\scriptstyle g} \\ Y & \xrightarrow{f} & X \end{array}$$

there is a canonical isomorphism of functors

$$g^{-1} \circ f_! \simeq \tilde{f}_! \circ \tilde{g}^{-1}.$$

Proof. We first define a morphism

$$\phi : f_! \circ \tilde{g}_* \to g_* \circ \tilde{f}_!.$$

Take any $\mathcal{F} \in \mathsf{Sh}(\tilde{Y})$ and consider an open subset V of X. Then we claim that $t \in f_!\tilde{g}_*\mathcal{F}$ is represented by a section $s \in \Gamma((f \circ \tilde{g})^{-1}(V), \mathcal{F})$ with $\mathsf{supp}(s) \subset \tilde{g}^{-1}(T)$ for some subset $T \subset f^{-1}V$ that is proper over V. Let us justify this claim. By definiton $t \in f_!\tilde{g}_*\mathcal{F}$ is given by some $s' \in \tilde{g}_*\mathcal{F}(f^{-1}V)$ with $\mathsf{supp}(s') \to V$ proper. Let $T = \mathsf{supp}(s')$. The claim now is that $\mathsf{supp}(s) \subset g^{-1}(\mathsf{supp}(s'))$, and indeed if some $y \notin \mathsf{supp}(s')$ then there is some neighbourhood W of y with s' represented by $(W, 0)$. But then on $(g^{-1}(W), 0)$ also represents s and $g^{-1}(y)$ does not meet $\mathsf{supp}(s)$.

(Here t, s, s' all live in the same space, namely $\mathcal{F}(g^{-1}f^{-1}(V))$, but are sections of different sheaves!)

If $T \to V$ is proper then so is the pullback $\tilde{g}^{-1}(T) \to g^{-1}(V)$. This follows from the fact that proper maps (of locally compact Hausdorff spaces) are (exactly) universally closed maps, i.e. maps all of whose pullbacks send closed sets to closed sets. See [Wed16, Definition and Theorem 1.30, Problem 1.20].

Thus s also defines a section of $g_*f_!\mathcal{F}$ as $\mathsf{supp}(s) \to g^{-1}(V)$ is proper as a composition of two proper maps (closed inclusions are proper).

Now we use Theorem 2.42, specifically the unit $\tilde{\eta}$ of $\tilde{g}^{-1} \dashv \tilde{g}_*$ and the counit ϵ of $g^{-1} \dashv g_*$ to define

$$g^{-1}f_! \xrightarrow{\tilde{\eta}} g^{-1}f_!\tilde{g}_*\tilde{g}^{-1} \xrightarrow{\phi} g^{-1}g_*\tilde{f}_!\tilde{g}^{-1} \xrightarrow{\epsilon} \tilde{f}_!\tilde{g}^{-1}.$$

To prove we have an isomorphism we consider the fiber at $z \in Z$.

$$\begin{aligned} (g^{-1}f_!\mathcal{F})_z &\cong (f_!\mathcal{F})_{g(z)} \\ &\cong \Gamma_c(f^{-1}(g(z)); \mathcal{F}) \\ &\cong \Gamma_c(\tilde{f}^{-1}(z), \tilde{g}^{-1}\mathcal{F}) \\ &\cong (\tilde{f}_!\tilde{g}^{-1}\mathcal{F})_z \end{aligned}$$

Here we used Lemma 2.41, Proposition 5.14 twice, and in between we use that in a pullback square there is a homeomorphism of fibers $\tilde{g} : f^{-1}(gz) \cong \tilde{f}^{-1}(z)$ which identifies the categories of sheaves and their cohomology with compact support.

We are not quite finished with the proof though, as we have produced a canonical map and an isomorphism of stalks, but we need to see that the map of stalks (that is an isomorphism) is indeed induced by the map θ.

It suffices to show that $\epsilon \circ \phi$ induces an isomorphism on stalks, as $\tilde{\eta}$ certainly does: By Proposition 5.14 we may check on the fiber, on which $\tilde{g}$ is a homeomorphism, thus the unit map is an isomorphsm.

So let us take a representative (S, t) for an element of $(g^{-1} f_! \tilde{g}_* \tilde{g}^{-1} \mathcal{F})_z$. Now $t \in g^{-1} f_! \tilde{g}_* \tilde{g}^{-1} \mathcal{F}(S)$ may be further represented (by definition of the pullback) by (S', t') with t' in $f_! \tilde{g}_* \tilde{g}^{-1} \mathcal{F}(S')$ and $g(S) \subset S'$. By Lemma 2.41 we may consider this (S', t') as an element of the stalk $(f_! \tilde{g}_* \tilde{g}^{-1} \mathcal{F})_{gz}$ with $gz \in S' \subset X$ and $t' \in f_! \tilde{g}_* \tilde{g}^{-1} \mathcal{F}(S')$.

Now explicitly t' is a certain element in $\tilde{g}^{-1} \mathcal{F}(\tilde{g}^{-1} f^{-1}(S'))$. (A properness condition holds, but here we are just identifying a particular element.) All the map ϕ does now is recast (S', t') by considering $t' \in \tilde{g}^{-1} \mathcal{F}(\tilde{f}^{-1} g^{-1}(S'))$ which is just another name for $\tilde{g}^{-1} \mathcal{F}(\tilde{g}^{-1} f^{-1}(S'))$.

We finally unravel definitions for the counit map which maps

$$(g^{-1} g_* \tilde{f}_! \tilde{g}^{-1} \mathcal{F})_z \cong \operatorname*{colim}_{gz \in U} \tilde{f}_! \tilde{g}^{-1} \mathcal{F}(g^{-1}(U))$$

to

$$(\tilde{f}_! \tilde{g}^{-1} \mathcal{F})_z \cong \operatorname*{colim}_{z \in W} \tilde{f}_! \tilde{g}^* \mathcal{F}(W)$$

by sending the class of (S', t') to the class of $(g^{-1}(S'), t')$ where we reinterpret the element t' as living in $\tilde{f}_! \tilde{g}^{-1} \mathcal{F}(g^{-1}(S'))$. (You can check this by unravelling the proof of Theorem 2.42.)

In other words, after applying the fiberwise isomorphism $\tilde{g}$ our canonical map is the identity on representatives. We just reinterpret the section t' as living in different incarnations of the same abelian group (viewed as spaces of sections of different sheaves). In this sense the map is canonical, and on stalks it induces the isomorphism induced by $\tilde{g}$, just like the map on stalks we considered above. □

5.3. Derived base change

We would like to derive the results that we worked hard for in the last section.

To do this, we will need a new class of Γ_c-acyclic sheaves. It's not just that injective sheaves are big and cumbersome, they do not have the formal properties we would expect.

For example, we would like to know $R(g_! \circ f_!) = Rg_! \circ Rf_!$. For the usual pushforward this follows immediately by replacing everything injectively as injectives are preserved by pushforward. Proper pushforward does not preserve injectives, so we need a different class of Γ_c acyclic sheaves.

Definition 5.16. A sheaf $\mathcal{F}$ on a locally compact space X is *c-soft* if for any compact subset $K \subset X$ the restriction map $\Gamma(X, \mathcal{F}) \to \Gamma(K, \mathcal{F}|_K)$ is surjective.

The c here stands for compact and we will later meet soft sheaves as well. We collect the key properties of c-soft sheaves in the following proposition. We leave out the proofs, as they are technical, not very hard, and somewhat similar to the proofs of similar properties for flasque sheaves.

Proposition 5.17. *For X, Y locally compact spaces and $f : X \to Y$ continuous we have*

1. *Any flabby sheaf on X (in particular any injective sheaf) is c-soft.*
2. *$\mathcal{F}$ is c-soft if and only if for any closed subset $Z \subset X$ there is a surjective restriction morphism $\Gamma_c(X, \mathcal{F}) \to \Gamma_c(Z, \mathcal{F}|_Z)$.*
3. *The restriction of a c-soft sheaf to a locally closed subset is c-soft.*
4. *For a continuous map $f : X \to Y$ and $\mathcal{F} \in \mathsf{Sh}(x)$ c-soft the sheaf $f_!\mathcal{F}$ is c-soft, too.*
5. *c-soft sheaves are Γ_c-acyclic and $f_!$-acyclic.*

Proof.

1. See example sheet. 5.11.
2. [KS90, Proposition 2.5.6].
3. [KS90, Proposition 2.5.7(i)]
4. [KS90, Proposition 2.5.7(ii)]
5. [KS90, Proposition 2.5.8 & Corollary 2.5.9]

□

There is also a compactly supported version of MayerVietoris, see III.1.8 in [Ive86].

We now obtain the following key result, *derived proper base change*.

Theorem 5.18. *Given a pullback square of locally compact spaces*

$$\begin{array}{ccc} \tilde{Y} & \xrightarrow{\tilde{f}} & Z \\ \downarrow{\scriptstyle \tilde{g}} & & \downarrow{\scriptstyle g} \\ Y & \xrightarrow{f} & X \end{array}$$

for any sheaf $\mathcal{F}$ on Y there is a canonical quasi-isomorphism

$$g^{-1} \circ Rf_!\mathcal{F} \simeq R\tilde{f}_! \circ \tilde{g}^{-1}\mathcal{F}.$$

Proof. This is the derived version of Theorem 5.15. As g^{-1} is exact we see that $g^{-1} \circ Rf_!$ is the derived functor of $g^{-1} \circ f_!$. It remains to show that $R\tilde{f}_! \circ \tilde{g}^{-1}$ is the derived functor of $\tilde{f}_! \circ \tilde{g}^{-1}$. We need a class of sheaves with which we can resolve sheaves on Y and which is $\tilde{f}_! \circ \tilde{g}^{-1}$-acyclic.

Our first guess would be injective sheaves, but as g^{-1} does not preserve injectives it is not clear that this computes $R\tilde{f}_! \circ \tilde{g}^{-1}$.

We will consider the class $\mathcal{I}_Y$ on Y of sheaves $\mathcal{G}$ satisfying $\mathcal{G}|_{f^{-1}(x)}$ is c-soft for any $x \in X$.

We define similarly the class $\mathcal{I}_{\tilde{Y}}$ on $\tilde{Y}$ of sheaves $\mathcal{G}$ satisfying $\mathcal{G}|_{f^{-1}(z)}$ is c-soft for any $z \in Z$. Then $\tilde{g}^{-1}$ takes $\mathcal{I}_Y$ to $\mathcal{I}_{\tilde{Y}}$ as $\tilde{g}$ induces an isomorphism on the fibers.

By Proposition 5.17 any injective sheaf lies in $\mathcal{I}_Y$ thus we may resolve in $\mathcal{I}_Y$.

The class $I_{\tilde{Y}}$ is $\tilde{f}_!$-acyclic: To check it sends exact sequences to exact sequences we check on stalks by Proposition 5.14, where it is true precisely by definition of $\mathcal{I}_{\tilde{Y}}$: Consider an exact sequence of sheaves $0 \to A \to B \to C \to 0$ on $\tilde{Y}$, the stalks of the pushforwards form an exact sequence if $\tilde{f}_* B_z \to \tilde{f}_* C_z$ is surjcetive, i.e. if $\Gamma_c(f^{-1}z, B) \to \Gamma_c(f^{-1}(z), C)$ is surjective, which is true as $R^1\Gamma(f^{-1}(z), A) = 0$.

As g^{-1} is exact (on all sheaves) the composition $f_! \circ g^{-1}$ is also exact on $\mathcal{I}_Y$ and has no higher derived functors, cf. Lemma 3.47. □

Remark 5.19. We may also consider a bounded below complex of sheaves F instead of the sheaf $\mathcal{F}$ in the proof, we just need the existence of injective resolutions of bounded below complexes, compare the proof of Corollary 3.45.

5.4. Homotopy invariance

We can use Lemma 4.12 prove homotopy invariance of certain sheaf cohomology groups.

Theorem 5.20. *We write $p : X \times I \to X$ for the natural projection and fix a sheaf $\mathcal{F}$ on X. We have $H^*(X; \mathcal{F}) \cong H^*(X \times I; p^{-1}\mathcal{F})$.*

Proof. We show $p^\# : Rp_{X*}\mathcal{F} \to Rp_{X*}Rp_*p^{-1}\mathcal{F} = (Rp_{X\times I})_*\mathcal{F}$ from Lemma 4.13 is a quasi-isomorphism where we write $p_Y : Y \to *$ for $Y = X, X \times I$ and (later) I.

It suffices to check that $\mathcal{F} \to Rp_*p^{-1}\mathcal{F}$ is a quasi-isomorphism, which we may check at a fiber $i : * \to X$. Thus we consider

$$i^{-1}\mathcal{F} \to i^{-1}Rp_*p^{-1}\mathcal{F}$$

where the left hand side is just $\mathcal{F}_x$. By compactness of I we see that p is proper and we have $p_* = p_!$, thus we may apply base change Theorem 5.18. Writing $j : I \to X \times I$ for the natural inclusion of $p^{-1}(x)$ we have

$$i^{-1}Rp_!p^{-1}\mathcal{F} \simeq Rp_{I!}j^{-1}p^{-1}\mathcal{F} \simeq Rp_{I*}p_I^{-1}i^{-1}\mathcal{F} \simeq R\Gamma(I; p_I^{-1}\mathcal{F}_x) \cong \mathcal{F}_x$$

where we used Lemma 4.12 to conclude. So the unit of the adjunction induces a canonical map $\mathcal{F}_x \to \mathcal{F}_x$. We can even see it explicitly: it sends an object $s \in \mathcal{F}_x$ to the global section of $p^{-1}\mathcal{F}_x$ that is constantly s, and this is an isomorphism. □

Corollary 5.21. *Let $f, g : X \to Y$ be homotopic and $\mathcal{F} \in \mathsf{Sh}(Y)$. Then the induced maps on cohomology agree, i.e. $f^\# = g^\# : H^*(Y, \underline{R}) \to H^*(X, \underline{R})$.*

Proof. The homotopy equivalence of f and g means there is $h : X \times I \to Y$ with $h \circ i_0 = f$ and $h \circ i_1 = g$ where $i_j : x \mapsto (x, j)$.

It is clear from the definition of $-^{\#}$ that it is functorial, thus it suffices to show that $i_0^{\#} = i_1^{\#}$.

By definition $p \circ i_j = \mathbf{1}_X$ and by Theorem 5.20 $p^{\#}$ is invertible. Thus its right inverse is unique and $i_0^{\#} = i_1^{\#}$. □

Corollary 5.22. *Let X, Y be homotopy equivalent. Then $H^*(X, \underline{R}) \cong H^*(Y, \underline{R})$.*

Proof. This follows directly from Corollary 5.21 as the homotopy equivalences $f : X \to Y$ and $g : Y \to X$ induce inverse maps on cohomology. □

Corollary 5.23. *Let $f : X \to Y$ be a proper map with contractible fibers. Then $H^*(X, \underline{R}) \cong H^*(Y, \underline{R})$.*

This is sometimes called the *Vietoris-Begle theorem.*

Proof. We take an injective resolution I of $\underline{R}_X$. Then $H^*(X, \underline{R})$ is the cohomology of $\Gamma(X, I) \cong Gamma(Y, f_* I) \cong R\Gamma(Y, Rf_* \underline{R})$ using that f_* preserves injectives. But the natural comparison $\underline{R}_Y \to Rf_* \underline{R}_X$ is an isomorphism in the derived category by checking on stalks since $(Rf_* \underline{R})_y = (Rf_! \underline{R})_y \cong R\Gamma_c(f^{-1}(y), \underline{R}) \cong R$ where we first used that f is proper, followed by Theorem 5.18 applied to the inclusion of a point, and finally Corollary 5.22. □

Example 5.24. For $n > 0$ we have $H^k(S^n, \underline{R}) = R$ if $k = 0, n$ and 0 otherwise.

We proceed by induction, and may even start with $n = 0$ if we formulate our result as $H^*(S^0, \underline{R}) = R[0] \oplus R[-k]$.

For $n = 0$ the constant sheaf on a discrete space is easily seen to be flasque and we are done.

For general n we cover S^n by $U = S^n \setminus \{N\}$ and $V = S^n \setminus \{S\}$ where N, S are distinct points. Then U, V are contractible and $U \cap V$ is homotopy equivalent to S^{n-1}, so we can proceed by induction knowing Corollary 5.22.

The Mayer Vietoris sequence (suppressing coefficients $\underline{R}$) is

$$\cdots \to H^i(S^n) \to H^i(U) \oplus H^i(V) \to H^i(S^{n-1}) \to H^{i+1}(S^n) \to H^{i+1}(U) \oplus H^{i+1}(V) \to \cdots$$

In the case $n = 1$

For all $i \geq 1$ we have $H^{i+1}(S^n) \cong H^i(S^{n-1})$ by vanishing of the middle term above degree 0.

For $i = 0$ we have the exact sequence

$$R \to R \oplus R \to R$$

if $n \geq 2$ and we have the sequence with cokernel R

$$R \to R \oplus R \to R \oplus R$$

for $n = 1$. Together with vanishing of the middle term in degree 1 this gives $H^1(S^1) = R$. and $H^1(S^n) = 0$ if $n \geq 2$.

5.5. A fundamental short exact sequence

Given an open subset $j : U \to X$ we recall from Example 5.9 that $j_!$ is the exact functor given by extension by 0.

From Exercise 8.3 we recall that $j_! j^{-1}\mathcal{F}(V)$ is given by all sections of $\mathcal{F}(V$ with support contained in $U \cap V$, and in particular there is a monomorphism $j_! j^{-1}\mathcal{F} \to \mathcal{F}$. (We'll give a proof below.)

Theorem 5.25. *Let $j : U \to X$ be an open inclusion and $i : Z = X \backslash U \to X$ the complementary closed inclusion.*

For any sheaf $\mathcal{F}$ on X there is a short exact sequence

$$0 \to j_! j^{-1}\mathcal{F} \to \mathcal{F} \to i_* i^{-1}\mathcal{F} \to 0.$$

Proof. For any open $W \subset X$ we know that $j_! j^{-1}\mathcal{F}(W)$ is exactly the subset of sections of $\mathcal{F}(W)$ whose support is contained in U. For completeness we recall the proof: If $s \in \mathcal{F}(W)$ has support contained in U then the map $\mathsf{supp}(s) \subset W \cap U \to W$ is proper and $s \in j_! j^{-1}(W)$. Conversely take a section s in $j_! j^{-1}(W)$. We have to show it extends to a section $s' \in \mathcal{F}(W)$. The idea is to define s' on the cover $W = (W \cap U) \cup (W \setminus \mathsf{supp}(s))$ by 0 on $(W \setminus \mathsf{supp}(s))$ and s otherwise. This is well-defined as long as this is indeed an open cover of W. This is clear for $W \cap U$ but we have to check for $W \setminus \mathsf{supp}(s))$ as $\mathsf{supp}(s)$ is a prioir only closed in U. But we know that $\mathsf{supp}(s) \subset X$ is an inclusion that is a proper map so it is in particular closed. This completes the proof.

We denote the natural injection $\alpha : j_! j^{-1}\mathcal{F} \to \mathcal{F}$.

We may compose with the unit map of the adjunction unit $\mathcal{F} \to i_* i^{-1}\mathcal{F}$ to obtain the sequence and need to check exactness. At a point $x \in U$ the sequence of stalks reduces to

$$0 \to j_! j^{-1}\mathcal{F}_x \to \mathcal{F}_x \to 0 \to 0$$

which is exact as $j_! j^{-1}\mathcal{F}_x = j^{-1}\mathcal{F}_x = \mathcal{F}_x$ (and α induces the canonical isomorphism of stalks). At a point $x \in Z$ we have

$$0 \to 0 \to \mathcal{F}_x \to i_* i^{-1}\mathcal{F}_x \to 0$$

which is exact by Example 3.11. □

This short exact sequence is very useful for computation:

Corollary 5.26. *For $Z \subset X$ closed and R is an abelian group then there is a long exact sequence of cohomology*

$$\ldots \to H_c^{k-1}(Z, \underline{R}) \to H_c^k(X \setminus Z; \underline{R}) \to H_c^k(X; \underline{R}) \to H_c^k(Z, \underline{R}) \to \ldots$$

In particular if Z is compact $H_c^{k-1}(Z, \underline{R})$ may be replaced by $H^{k-1}(Z, \underline{R})$.

Proof. We apply $R\Gamma_c$ to the exact sequence

$$0 \to j_! \underline{R} \to \underline{R} \to i_* \underline{R} \to 0$$

on X. Then the result follows if we can show $R\Gamma_c(X, j_!\underline{R}) \simeq R\Gamma_c(X \setminus Z, \underline{R})$ and $R\Gamma_c(X, i_!\underline{R}) \simeq R\Gamma_c(Z, \underline{R})$. The first result follows as $j_!$ is exact and preserves c-soft sheaves this is immediate. The second is similar. □

Example 5.27. Consider $* \subset S^n$ with complement homeomorphic to $\mathbb{R}^n$. We get

$$\ldots \to H^{k-1}(*) \to H^k_c(\mathbb{R}^n, \underline{R}) \to H^k(S^n; \underline{R}) \to H^k(*, \underline{R}) \to \ldots$$

and we can read off from Example 5.24 that $H^i_c(\mathbb{R}^n, \underline{R}) \cong H^i(S^n, \underline{R})$ for $i > 0$ while $H^0_c(\mathbb{R}^n) = 0$. In other words, $H^*_c(\mathbb{R}^n, \underline{R})$ is R in degree n and 0 elsewhere.

The example shows in particular that compactly supported cohomology is not homotopy invariant.

To understand the theoretical importance of Theorem 5.25 we continue with the derived version, which is immediate as all functors are already exact.

For simplicity we write $D^+(X)$ for $D^+(\mathsf{Sh}(X))$ in the following theorem.

Corollary 5.28. *Let $j : U \subset X$ be open and let $i : Z = X \setminus U \to Z$ have a paracompact neighbourhood basis. For any complex of sheaves C in $D^+(X)$) there is an exact triangle $j_! j^{-1} C \to C \to i_* i^{-1} G$.*

Theorem 5.29. *Given a locally compact topological space X with an open subset U and complement Z. There are fully faithful functors $j_! : D^+(U) \to D^+(X)$ and $i_* : D^+(Z) \to D^+(X)$. Moreover any object F fits in an exact triangle $j_1 A \to F \to i_* B$ with $A \in D^+(U)$, $B \in D^+(Z)$. For any $A \in D^+(U)$, $B \in D^+(Z)$ we have $R\operatorname{Hom}(j_! A, i_* B) = 0$.*

Proof. The full embedding of $D(Z)$ is exercise 4.2. The full embedding of $D(U)$ follows as we obtain a map $j^{-1} j_! \to \mathbf{1}_{D(U)}$ from $j_! \to j_*$ via adjunction. The map is an isomorphism as we can see comparing stalks.

The existence of exact triangles is Corollary 5.28.

The vanishing of homs follows by using the adjunction $i^{-1} \dashv Ri_* = i_*$ from Lemma 4.14 and observing $i^{-1} j_!$ is equivalent to the zero functor by checking $(i^{-1} j_! F)_z = (j_! F)_z = 0$. □

We summarise the situation of Theorem 5.29 by saying $D(X)$ has a *semi-orthogonal decomposition* into $D(U)$ and $D(X)$. Much like F decomposes into $j_! j^{-1} F$ and $i_* i^{-1} F$ we think of the whole derived category decomposing into the two smaller categories. Semi-orthogonal refers to the fact that in one direction there are no maps between the two factors. One can reconstruct the category $D(X)$ from $D(U)$, $D(Z)$ and some gluing information that can be encoded in the functor $i^{-1} j_*$.

Warning: If we are studying sheaves of modules (as in algebraic geometry) there is a similar theorem, but we have to replace $D(Z)$ by the derived category of sheaves on X with support on Z.

5.6. Relative cohomology and comparison with singular cohomology

We have now seen that $R\Gamma(-, \underline{R})$ satisfies functoriality (Lemma 4.13), homotopy invariance (Corollary 5.21) and the Mayer-Vietoris theorem 4.8.

These are key properties of ordinary cohomology, and in fact we can show that sheaf cohomology with constant coefficients agrees (on CW complexes) with singular cohomology by showing it satisfies the *Eilenberg-Steenrod axioms*.

We quote the following result from algebraic topology

Definition 5.30. An *unreduced cohomology theory* $E^\bullet$ is a collection of cotravariant functors $E^n : (X, A) \to E^n(X, A)$ from the category of CW-pairs to abelian groups together with natural transformations $E^n(A) \to E^{n+1}(X, A)$ where we we write $E^n(A, \emptyset)$ as $E^n(A)$

1. $E^\bullet$ is homotopy invariant
2. There is a long exact sequence

$$\cdots \to E^n(X, A) \to E^n(X) \to E^n(A) \xrightarrow{\delta} E^{n+1}(X, A) \to \cdots$$

3. (Excision) For $W \subset A \subset X$ with $\bar{W} \subset A^\circ$ the inclusion : $(X \setminus W, A \setminus W) \to (X, A)$ induces an isomorphism on $E^\bullet$.
4. (Additivity)If $(X, A) = \amalg_i (X_i, A_i)$ the canonical comparison map

$$E^n(X, A) \cong \prod_i E^n(X_i, A_i)$$

   is an isomorphism

If moreover $E^n(*)$ vanishes for $n \neq 0$ we say $E^\bullet$ is *ordinary*.

To show sheaf cohomology defines a cohomology theory we need to define it on pairs. We can use a trick:

Theorem 5.31. *Let R be an abelian group. Sheaf cohomology with coefficients in constant sheaves defines an ordinary additive unreduced cohomology theory by setting $E^n(X, A) = H^n(X, j_!\underline{R})$ for $j : X \setminus A \to X$. In particular $E^n(X) = H^n(X, \underline{R})$*

Proof. $R\Gamma(-, \underline{R})$ is functorial by Lemma 4.13 and homotopy invariant by Corollary 5.21. To check functoriality for relative cohomology we fix $f : (X, A) \to (Y, B)$ and denote the open inclusions by $i : A \to X$, $j : X \setminus A \to X$ and $h : B \to Y$ and $k : Y \setminus B \to Y$. We need to produce a map $k_!\underline{R} \to f_* j_!\underline{R}$. We note that the middle and right vertical map in the below diagram exist by adjunctions (and are the maps inducing the map on cohomology). They are compatible by construction using $f \circ i = h \circ f|_A$ and induce the desired map $k_!\underline{R} \to f_* j_!\underline{R}$ on kernels.

$$
\begin{array}{ccccc}
f_* j_! \underline{R}_{X\setminus A} & \longrightarrow & f_* f^{-1}\underline{R} & \longrightarrow & f_* i_* i^{-1} f^{-1}\underline{R} \\
\uparrow & & \uparrow & & \uparrow \\
k_! \underline{R}_{Y\setminus B} & \longrightarrow & \underline{R}[r] & \longrightarrow & h_* h^{-1}\underline{R}
\end{array}
$$

The long exact sequence of pairs arises by applying $R\Gamma$ to the short exact sequence obtained by putting $\underline{R}$ into Theorem 5.25.

To obtain homotopy invariance for pairs we use homotopy invariance for singletons together with the long exact sequence of pairs and the five-lemma.

We know $R\Gamma(X, i_*\underline{R}) \simeq R\Gamma(A, \underline{R})$ from Lemma 4.9.

Additivity is Exercise 5.3 and the value at a point is immediate. □

Theorem 5.32. *Any ordinary additive unreduced cohomology theory E^* on the category of CW pairs is naturally isomorphic to singular cohomology with coefficients $E^0(*)$.*

Proof. See [Hat02, Theorem 4.59]. □

Putting the two theorems together we have:

Corollary 5.33. *Let X be a CW complex and R an abelian group. Then $H^n(X, \underline{R}) \cong H^n_{sing}(X, \underline{R})$.*

Remark 5.34. This is not the strongest possible result: The result holds not just for CW complexes but for all locally contractible spaces (in fact, for all spaces with locally trivial sheaf cohomology). And our proof is very indirect, relying on a powerful axiomatic characterization singular cohomology.

The most natural proof would be to present singular cochains as global sections of some complex of sheaves that is a Γ-acyclic resolution of the constant sheaf.

The first problem is that singular cochains do not glue: we would have to determine the value of a simplex $\Delta^n \to A \cup B$ that is not contained in either A or B from the value on simplices with values in A and in B. This is reminiscent of problems in algebraic toplogoy, and the solution is barycentric subdivision. However, to build a sheaf we need to take the limit over all barycentric subdivisions. This is possible, but it gives a somewhat unwieldy object and it is in particular not obvious that it is Γ-acyclic. A careful proof along these ilines can be found in [Sel16].

Instead we consider a fixed cover $\mathfrak{U}$ of X by contractibles and take Čech cohomology with coefficients in the the presheaf associated to the singular cochains $S^*_{\mathfrak{U}}(X, A)$. Čech cohomology works with coefficients in any presheaf, but for presheaves we may not just use Lemma 4.19. But one can prove by hand that $S^*_{\mathfrak{U}}(U_i, A)$ is exact. This argument is spelled out in the book [BT82].

The final and arguably most natural proof is to observe that while singular cochains do not form a sheaf, they form a *hypersheaf*. That means they satisfy the gluing axiom not on the nose but in a derived or homotopical sense (one has to derive the limit functor in the definition of sheaves). This is a very good observation, but turning it into a proof still needs some non-trivial input from homotopy theory, the details are in the note [Pet21].

6. Sheaves on manifolds

6.1. Tangent and cotangent sheaf

We now consider topological spaces with extra structure.

Recall that a ringed space $(X, \mathcal{R})$ is a topological space X equipped with a sheaf of rings $\mathcal{R}$.

We note that if $\mathcal{R}$ is a sheaf of rings then the stalk $\mathcal{R}_x$ naturally has a ring structure, and for any continuos map the sheaves $f^{-1}\mathcal{R}$ and $f_*\mathcal{R}$ are also sheaves of rings. (The key ingredient here is that both limits and filtered colimits of rings are given by the natural ring structure on the limits and filtered colimits of underlying abelian groups.)

It is called *locally ringed* if all the stalks $\mathcal{R}_x$ are local rings (i.e. they have a unique maximal ideal).

$(\mathbb{R}^n, \mathcal{C}^\infty)$ is an example of a locally ringed space, as the the ring of germs $\mathcal{C}^\infty_x$ has a natural evaluation map to $\mathbb{R}$ whose kernel is the unique maximal ideal. In fact the existence of evaluation maps at stalks is equivalent to the existence of a unique maximal ideal and the reason we care about local rings.

A map of ringed spaces $(f, f^\#) : (X, \mathcal{R}) \to (Y, \mathcal{S})$ consists of a continuous map $f : X \to Y$ together with a map $f^\# : f^{-1}\mathcal{S} \to \mathcal{R}$ of sheaves of rings on X. Instead of $f^\#$ we could equivalently ask for a map $\mathcal{S} \to f_*\mathcal{R}$. Such a map $(f, f^\#) : (X, \mathcal{R}) \to (Y, \mathcal{S})$ is a map of locally ringed spaces if X and Y are locally ringed and the induced map $f^\#_x$ is local, i.e. the pre-image of the maximal ideal is maximal.

Thus a map of ringed spaces $(f, f^\#)$ is an isomorphism if f is a homeomorphism and $f^\#$ is a sheaf isomorphism.

Definition 6.1. A locally ringed space $(X, \mathcal{R})$ is a *smooth manifold* if it is Hausdorff, second countable and locally isomorphic to $(\mathbb{R}^n, \mathcal{C}^\infty)$, i.e. there is a cover $\cup U_i \to X$ such that every $(U_i, \mathcal{O}|_{U_i})$ is isomorphic as a ringed space to $(\mathbb{R}^n, \mathcal{C}^\infty)$.

A locally ringed space $(X, \mathcal{R})$ is a *complex manifold* if it is Hausdorff, second countable and locally isomorphic to $(\mathbb{C}^n, \mathcal{O})$ where $\mathcal{O}$ is the sheaf of holomorphic functions on $\mathbb{C}^n$.

We often denote the sheaf $\mathcal{R}$ by $\mathcal{C}^\infty$ and call its sections *smooth functions*.

For reference we recall the usual definitions: A *chart* on a topological space X is an open subset $U \subset X$ together with a homeomorphism ϕ from U to an open set $D \subset \mathbb{R}^n$. A *smooth atlas* on X is a collection of charts $U_i \subset X$, $\phi_i : U_i \to D_i$ such that for every pair U_i, U_j the *transition functino*, i.e. the composition $\phi_i \cap \phi_j^{-1} : \phi_j(U_i \cap U_j) \to \phi_i(U_i \cap U_j)$ is smooth (and thus a diffeomorphism). A smooth atlas $\mathcal{U}$ is *maximal* if it contains all compatible chart, i.e. if any chart ϕ on X is such that $\mathcal{U} \cup \{\phi\}$ is also an atlas then $\phi \in \mathcal{U}$.

A *smooth manifold* M is a topological space that is Hausdorff and second countable (i.e. the topology has a countable basis) together with a maximal smooth atlas.

If M is connected there is an n called the *dimension* of M such that all charts map to $\mathbb{R}^n$.

A maximal atlas is by definition very large, it is usually enough to consider an atlas on X such that any point in X is contained in a chart. This determines a maximal atlas.

A *holomorphic atlas* on a topological space M is a collection of charts $U \to D \subset \mathbb{C}^n$ such that all transition functions are biholomorphisms.

A *complex manifold* is a second countable Hausdorff space equipped with a maximal holomorphic atlas.

A continuous function $f : V \to \mathbb{R}$ on an open subet V of a manifold M is called *smooth* if for any chart $\phi : U \to D$ the function $f \circ \phi^{-1}|_{U \cap V}$ is smooth.

We then define the *sheaf of smooth functions* $\mathbb{C}^\infty$ on M as the subsheaf of the sheaf of continuous functions whose sections are smooth.

Remark 6.2. A cover U_i with $(U_i, \mathbb{R}_{U_i}) \cong (\mathbb{R}^n, \mathcal{C}^\infty)$ provides an atlas and vice versa we can glue a locally ringed space from an atlas. A detailed comparision can be found in in [Wed16, Section 4.2].

On a smooth manifold we can do calculus. The main ingredient for this is a special sheaf:

Definition 6.3. Let M be a smooth manifold. We define the *tangent sheaf* $\mathcal{T}_M$ as follows:

$$\mathcal{T}_M(U) = \{D : \mathcal{C}^\infty(U) \to \mathcal{C}^\infty(U) \mid D \text{ is a derivation}\}$$

where a derivation is a map D satisfying the *Leibniz rule* $D(fg) = D(f)g + fD(g)$.

Note that the tangent sheaf is indeed a sheaf: It is a sub-presheaf of the hom sheaf $\mathcal{H}om(\mathcal{C}^\infty, \mathcal{C}^\infty)$ defined by the local condition of being a derivation.

It is moreover a sheaf of modules over the sheaf of rings $\mathcal{C}$, as fD is a derivation for any smooth function f and derivation D.

In fact this sheaf is the sheaf of sections of a rank n vector bundle over M, the *tangent bundle*. We just check this locally.

We call smooth section $v : U \to E \times V$ a *smooth vector field* on U and define the *directional derivative* D_v as the map sending $f \in \mathcal{C}^\infty(U)$ to

$$D_v(f) := \lim_{h \to 0} \frac{f(x + hv(x)) - f(x)}{h}$$

Lemma 6.4. *Let $E = \mathbb{R}^n$. Then on $U \subset E$ the direction derivative gives an isomorphism beween smooth vector fields on U and the sheaf $\mathcal{T}_U$.*

Proof. Given a smooth vector field $v : U \to E \times V$ we can compute D_v via the chain rule by $f \mapsto \sum_i v_i \frac{\partial f}{\partial v_i}$.

Thus D_v is only 0 if all v_i are 0 and D_- is an injection.

Consider now an arbitrary derivation $D \in \mathcal{T}_E(U)$. We claim that D is equal to the directional derivative $\sum_i D(x_i) \frac{\partial}{\partial x_i}$.

We consider $P = D - \sum_i D(x_i)\frac{\partial}{\partial x_i}$ which is a derivation, and $P(x_i) = 0$ for all x_i. We will show P is in fact 0 to show our map is surjectiev. It is enough to check this locally. Let now f be an arbitrary smooth function on U and fix $u \in U$. In a neighbourhood of u we can write $f(x) = f(u) + \sum_i (x_i - u_i)h_i(x)$ for smooth functions h_i. To see this write

$$f(x) - f(u) = \int_0^1 \frac{d}{dt} f(t(x-u)+u)dt = \sum_i (x_i - u_i) \int_0^1 \frac{\partial f}{\partial x_i}(t(x-u)+u)dt$$

using the chain rule.

But then

$$Pf(x) = P(f(u)) + P\left(\sum_i (x_i - u_i)h_i(x)\right) = 0 + \sum_i (0 + (x_i - u_i)Ph_i(x))$$

and it follows that $(Pf)(u) = 0$. □

Definition 6.5. We define the *bracket* of derivations by $[v, w](f) = v(w(f) - w(v(f))$

In fact, this makes $\mathcal{T}_M$ into a sheaf of Lie algebras.

Definition 6.6. The *cotangent sheaf* Ω_M of a manifold M is defined as the dual of the tangent sheaf $\mathcal{H}om_{\mathcal{C}^\infty}(\mathcal{T}_M, \mathcal{C}^\infty)$.

On $\mathbb{R}^n$ the global sections of $\mathcal{T}$ are spanned (over $\mathcal{C}$) by the derivations $\frac{\partial}{\partial x_i}$ and the global section of Ω are spanned by their dual differential forms dx_i.

6.2. De Rham cohomology

We now show that sheaves of smooth sections of any vector bundle are Γ-acyclic.

Definition 6.7. A sheaf $\mathcal{F}$ on X is *fine* if it is a module over a sheaf of rings $\mathcal{A}$ with the property that for any open cover $\{U_i\}$ of X there is a set of sections $\phi_i \in \mathcal{A}$ such that

1. the support of ϕ_i is contained in U_i ,
2. any point has a neighbourhood where only finitely many ϕ_i are nonzero,
3. $\sum_i \phi_i = 1_{\mathcal{A}}$.

In other words, $\mathcal{A}$ admits partitions of unity. In particular, in the definition above $\mathcal{A}$ itself is fine as it is a module over itself.

Example 6.8. The sheaf of smooth functions on a manifold is fine as partitions of unity exist. One uses a chart and the existence of bump functions.

It follows that all sheaves of $\mathcal{C}$-modules, like $\mathcal{T}_M$ and Ω_M are fine.

Lemma 6.9. *Any fine sheaf is* Γ*-acyclic.*

Proof. Let $\mathcal{F}$ be a fine sheaf of $\mathcal{A}$-modules. We use the Godemenet resolution I as a flasque resolution of $\mathcal{F}$ and observe that I also consists of $\mathcal{A}$-modules in each degree and the differential is $\mathcal{A}$-linear.

We compute $R^k\Gamma(X, \mathcal{F}) \cong H^k(\Gamma(X, I^\bullet))$. Assume this is nonzero for some $k > 0$ and take $\alpha \neq 0 \in H^k(\Gamma(X, I^\bullet))$. As I is locally exact there is a cover $\{U_i\}$ of X such that $\alpha|U_i = d\beta_i$ for some $\beta_i \in I^{k-1}(U_i)$. Choosing a partition of unity ϕ_i subordinate to the cover we define

$$\beta = \sum_i \phi_i \beta_i.$$

As the support of ϕ_i is contained in U_i we may view $\phi_i\beta_i$ as a global section of I^{k-1}. As the sum is locally finite we indeed have $\beta \in I^{k-1}(X)$. We compute

$$d\beta = \sum_i \phi_i d(\beta_) = \sum \phi_i \alpha|_{U_i} = \alpha$$

using $\mathcal{A}$-linearity. This completes the proof. □

We thus have a large and well-understoof supply of Γ-acyclic sheaves for the first time. However, fine sheaves will almost always take coifficients in $\mathbb{R}$ or $\mathbb{C}$ and cannot be used to compute cohomology with coefficients in $\mathbb{Z}$.

We now consider the *exterior powers* of the cotangent sheaf of an n-manifold M, defined as

$$\Omega_M^p \coloneqq \bigwedge^p \Omega_M$$

where the p-th exterior power is defined as the quotient of the p-fold tensor product $\Omega_M^{\otimes p}$ by the ideal generated by all $\alpha \otimes \alpha$ for $\alpha \in \Omega$.

Remark 6.10. To avoid the quotient one may also (if the characteristic is 0) identify the exterior algebra with the subsheaf consisting of section on which the symmetric group S_p acts via the sign representation.

Then the sections of Ω^p are maps $\alpha : \mathcal{T}^{\otimes p} \to \mathcal{C}^\infty$ satisfying $\alpha(v_1, \ldots, v_p) = \epsilon(\sigma)\alpha(v_{\sigma(1)}, \ldots, v_{\sigma(p)})$.

Definition 6.11. The sheaf of *differential p-forms* on M is defined as Ω_M^p.

Ω^p is the sheaf of sections of a vector bundle of rank $\binom{n}{p}$.

Lemma 6.12. *There is a differential* $d : \Omega^p \to \Omega^{p+1}$.

Proof. We define $d; \mathcal{C}^\infty \to \Omega$ by $f \mapsto (v \mapsto v(f))$ recalling that Ω is the dual of derivations.

For $\alpha \in \Omega^p$ we define

$$d\alpha(v_1, \ldots, v_p) = \sum_i (-1)^{i+1} v_i \alpha(v_1, \ldots, \widehat{v_i}, \ldots, v_{p+1}) + \sum_{i<q} (-1)^{i+j} \alpha([v_i, v_j], v_1, \ldots, \widehat{v_i}, \ldots, \widehat{v_j}, \ldots v_{p+1})$$

and one may check that $d^2 = 0$. □

Definition 6.13. We call $(\Omega^\bullet(M), d)$ the *de Rham complex* of M and its cohomology the *de Rham cohomology*, denoted $H^i_{dR}(M)$.

Theorem 6.14. *The complex of sheaves $(\Omega^\bullet, d)$ is a soft resolution of the constant sheaf $\underline{\mathbb{R}}$. In particular de Rham cohomology computes sheaf cohomology, $R^i\Gamma(M, \underline{\mathbb{R}}) \cong H^i_{dR}(M)$.*

Sketch of proof. As $\Omega^\bullet$ is fine by Example 6.8 it is Γ-acyclic by Lemma 6.9.

It remains to show exactness, which we can check for each point $x \in M$ in a neighbourhood. Thus it suffices to check on an open disc $D \subset \mathbb{R}^n$ with centre the origin. This is the *Poincaré lemma*:

We define a contraction operator $h^p : \Omega^p(D) \to \Omega^{p-1}(D)$ by sending $\alpha = f dx_{i_1} \wedge \cdots \wedge dx_{i_p}$ to

$$h^p\alpha = x_{i_1}\left(\int_0^1 f(0, \ldots, 0, tx_{i_1}, x_{i_1+1}, \ldots, x_n)dt\right) dx_{i_2} \wedge \ldots dx_{i_p}$$

and have to verify the formula $h^{p+1}d^p + d^{p-1}h^p = \mathbf{1}$ for $p \geq 1$. This shows that inclusion and evaluation at 0 define a homotopy equivalence of complexes between $\mathbb{R}$ and $(\Omega^\bullet(D), d)$. This is a computation. Details for a slightly different phrasing of the same proof can be found in [BT82, I.4]. □

Example 6.15. We know that $H^*_{dR}(S^2) \cong \mathbb{R}[0] \oplus \mathbb{R}[-2]$ from Theorem 6.14.

It is clear that the constant functions are in the kernel of d, so the constant function 1 generates $H^0_{dR}(S^2)$.

To exhibit a generator of H^2_{dR} we need some analysis. Consider $S^2 = \{(x_1, x_2, x_3) \mid x_1^2 + x_2^2 + x_3^2 = 1\} \subset \mathbb{R}^3$ and consider ω as the restriction of the form $x_1 dx_2 \wedge dx_2 - x_2 dx_1 \wedge dx_3 + x_3 dx_1 \wedge dx_2$.

For degree reasons $d\omega = 0$. But ω cannot be exact as it as a volume form: One can compute. $\int_{S^2} \omega = 4\pi$. If $\omega = d\beta$ then by Stokes theorem $\int_{S^2} \omega = \int_{\partial S^2} \beta = 0$ and we have a contradiction.

Definition 6.16. The *exterior product* is the natural product $\Omega^p \times \Omega^q \to \Omega^{p+q}$ induced by the concatenation of tensor product. We write it as

$$(\alpha, \beta) \mapsto \alpha \wedge \beta.$$

ON

Lemma 6.17. *The exterior product is bilinear and satisfies*

$$\alpha \wedge \beta = (-1)^{pq} \beta \wedge \alpha$$

(it is graded commutative*) and*

$$d(\alpha \wedge \beta) = d\alpha \wedge \beta + (-1)^p \alpha \wedge d\beta$$

(it satisfies the graded Leibniz rule*). Here $\alpha \in \Omega^p$ and $\beta \in \Omega^q$.*

Sketch of proof. Bilinearity is clear. For $p = q = 1$ the graded commutativitiy follows from the relation $(\alpha+\beta)\wedge(\alpha+\beta) = 0$. From this one may deduce $\alpha_1\wedge\cdots\wedge\alpha_p = \epsilon(\sigma)\alpha_{\sigma(1)}\wedge\cdots\wedge\alpha_{\sigma(p)}$. As any p-form is the exterior product of p 1-forms the result follows.

The Leibniz rule follows by computation. □

The Leibniz rule ensures that the exterior product descends to cohomology, thus de Rham cohomology comes with a natural graded commutative product structure.

Remark 6.18. You also know that singular cohomology has a product. So what about sheaf cohomology? The category of rings is not abelian, so we cannot take derived global sections of a sheaf of rings as rings.

But we may rewrite the sheaf cohomology $R^\bullet\Gamma(X, \underline{R})$ as $\mathrm{Ext}^\bullet_{\mathsf{Sh}(X)}(\underline{R}, \underline{R})$ which has a natural product structure!

This product agrees with the exterior product on de Rham cohomology (and any other reasonable product), see [Ive86, Section II.10].

6.3. Line bundles

Let M be a real or complex manifold, so let $k = \mathbb{R}$ or $\mathbb{C}$ and let $\mathcal{C}$ denote either $\mathcal{C}^\infty$ or the sheaf of holomorphic functions.

Recall that a smooth vector bundles of rank n is a space E with a surjection $p : E \to X$ such that all fibers are n-dimensional k-vector spaces and X has a cover U_i and each $p^{-1}(U_i)$ is diffeomorphic via some h_i to $U_i \times k^n$ that restricts to a linear isomorphism on each fiber.

The composition (of restrictions) $h_j \circ h_i^{-1}$ defines a *transition function* ϕ_{ij} on each $U_i \cap U_j$ which is a function in $GL_n(\mathcal{C})$, i.e. a smooth/holomorphic function with values in $GL_n(k)$.

Compatibility demands that $\phi_{ij}\circ\phi_{jk} = \phi_{ik}$ on $U_i\cap U_j\cap U_k$. One may check that whenever this compatibility is satisfied we can glue a vector bundle from the data of the transition functions: We define a quotient by $\amalg_i U_i \times k$ with the relation $(u^i, x) \sim (u^j, \phi_{ij}(x))$ where $u^i \in U_i \cap U_j \subset U_i$ and $u^j \in U_i \cap U_j \subset U_j$ denote the same element in $U_i \cap U_j$. The check that this is indeed a vector bundle is left to the reader/material for a differential geometry course.

On the other hand our choice of trivializations was far from unique, we can postcompose with any set $\{g_i\}$ of smooth/holomorphic $GL_n(k)$-valued function on U_i. Thus two sets of transition functions ϕ_{ij} and ϕ'_{ij} are equivalent if there are $g_i \in \mathcal{C}(U_i, GL_n(k))$ with $\phi'_{ij} = g_j \circ \phi_{ij} \circ g_i^{-1}$

We restrict attention to $n = 1$. A *line bundle* is a vector bundle of rank 1.

Lemma 6.19. *Fix a manifold X and a cover $\mathfrak{U}$. The set of line bundles on X that trivialise on $\mathfrak{U}$ up to isomorphism is given by the first Čech cohomology group $\check{H}^1(\mathfrak{U}, \mathcal{C}^\times)$ where $\mathcal{C}^\times$ is the sheaf of functions not taking the value 0.*

We give a detailed outline of the proof:

Proof. We define a map Ψ from line bundles to Čech cohomology as follows. Given a line bundle L we pick transition functions ϕ_{ij} as above and define the Čech cochain $\Psi(L)$ by ϕ_{ij} on U_{ij}. It is a cocycle by the compatibility condition.

If we pick different transition functions then the associated Čech cocycle will differ by $g_j g_i^{-1}$ for some Čech 0-cochain, so Ψ is well-defined into $\check{H}^1(\mathfrak{U}, \mathcal{C}^\times)$. (We use multilpicative notation for familiarity, but of course the multiplication is commutative.)

As we can glue a line bundle form every Čech cocycle this is a surjective map.

Assume now that $\Psi(L)$ and $\Psi(M)$ are cohomologous via a 1-cochain g. Then $g|_{U_i} : L|_{U_i} \to M|_{U_i}$ assemble into a line bundle isomorphism.

Conversely, assume two line bundles are isomorphic via $f : L \to M$ then $f|_{U_i}$ defines an element $f \in \check{C}^0(\mathfrak{U}, \mathcal{C}^\times)$ with $\delta f = \Psi(M)\Psi(L)^{-1}$. □

Corollary 6.20. *Assume X has set $\mathcal{T}$ of open subsets that is closed under intersection and forms a neighbourhood basis for every point und such that $\mathcal{C}^\times$ is Γ-acyclic on each $T \in \mathcal{T}$. Then the set of all line bundles on X up to isomorphism is given by the sheaf cohomology $H^1(X, \mathcal{C}^\times)$.*

The condition is satisfied on a smooth manifold by taking geodesically convex open sets, it will follow from the considerations below that $\mathcal{C}^\times$ is Γ-acyclic on contractible sets. The condition is also satsified on a complex manifold but that is harder to check.

Proof. The isomorphism classes of line bundles are given by the colimit over all covers, thus by the Čech cohomology group as in Remark 4.23. If any poin has a neighbourhood basis of opens on which $\mathcal{C}^\times$ is Γ-acyclic then the diagram of Čech cohomolgy groups is constant on a cofinal subdiagram by Theorem 4.22 and we obtain sheaf cohomology. □

Remark 6.21. Note that even if $n > 1$ we may view the transition functions as a Čech cocycle with coefficients in the sheaf of smooth functions with values in $GL_n(k)$. This is certainly not a sheaf of abelian groups in general, thus we have to specify the order of the group operation in the differential carefully, but this can be done.

More seriously we do not have a comparision with sheaf cohomology (as we have no derived global sections for a sheaf of nonabelian groups). Nevertheless, it can be useful to consider $\check{H}^1(X, GL_n(\mathcal{C}^\infty))$.

Example 6.22. We consider a variation of Example 2.29, considering the sheaf of smooth complex-valued functions $\mathcal{C}^\infty$ on a smooth manifold X and the sheaf $(\mathcal{C}^\infty)^\times$ of smooth nonzero functions. There is a short exact sequence of sheaves:

$$0 \to \underline{\mathbb{Z}} \to \mathcal{C}^\infty \xrightarrow{\exp} (\mathcal{C}^\infty)^\times \to\to 0$$

with an associated long exact sequence in cohomology.

But now we know that $\mathcal{C}^\infty$ is fine and thus the long exact sequence gives

$$0 \to H^1(X, (\mathcal{C}^\infty)^\times) \to H^2(X, \underline{\mathbb{Z}}) \to 0$$

As $\mathbb{Z}$ and $\mathcal{C}^\infty$ are locally Γ-acyclic so is $(\mathcal{C}^\infty)^\times$, thus we may apply Corollary 6.20 and the complex line bundles on X are classified exactly by the easily computed group $H^2(X, \mathbb{Z})$! Geometrically the boundary map providing the isomorphism is know as the *first Chern class*.

The addition operation on $H^2(X, \mathbb{Z})$ corresponds precisely to the tensor product on line bundles.

Example 6.23. Let us now consider Example 2.29 proper. The long exact sequence on cohomology associated with the short exact sequence of sheaves on a complex manifold X

$$0 \to \underline{\mathbb{Z}} \to \mathcal{O} \to \mathcal{O}^\times \to 0$$

contains

$$\cdots \to H^1(X, \mathbb{Z}) \to H^1(X, \mathcal{O}) \to H^1(X, \mathcal{O}^\times) \to H^2(X, \mathbb{Z}) \to H^2(X, \mathcal{O}) \to \cdots$$

As $\mathcal{O}$ is not fine (or flabby, or soft, or injective) we cannot simplify as in the smooth case.

The set of holomorphic line bundles on a complex manifold, up to isomorphism is given by an extension of $\ker(H^2(X, \mathbb{Z}) \to H^2(X, \mathcal{O}))$ by $H^1(X, \mathcal{O})/H^1(X, \mathbb{Z})$.

In many cases this is still very instructive.

Take for example $X = \mathbb{C}^\times$. It is clear that $H^2(X, \mathbb{Z}) = 0$. Moreover, one can use the resolution $\mathcal{O} \to \mathcal{C}^\infty \xrightarrow{\bar{\partial}} \mathcal{C}^\infty d\bar{z}$ to show that $\mathcal{O}$ is Γ-acylic on $\mathbb{C}^*$. The fact that this is a resolution, i.e. that $\bar{\partial}$ is locally surjective is the so-called Dolbeault-Grothendieck Lemma or $\bar{\partial}$-Poincaré lemma.

The fact that the $\mathcal{C}^\infty(\mathbb{C}^*) \xrightarrow{\bar{\partial}} \mathcal{C}^\infty(\mathbb{C}^*)d\bar{z}$ is surjective and there is no higher cohomology follows by a variation of the same argument. Details can be found in Taylor, *Several complex variables*, Proposition 1.4.3. In summary, all holomorphic line bundles on $\mathbb{C}^\times$ are trivial.

A non-trivial example is given by an elliptic curve X, i.e. the set of solutions of a cubic equation in $\mathbb{C}P^2$ (assuming this set is non-singular and we fix a point). One can compute $H^1(E, \mathcal{O}) \cong \mathbb{C}$ and $t : H^1(E, \mathbb{Z}) \cong \mathbb{Z}^2 \to H^1(E, \mathcal{O})$ is an injection. Then the set of holomorphic line bundles on E given by an extension of $\mathbb{Z}$ (the set of topological line bundles) by $\mathbb{C}/\mathrm{Im}(t)$, the set of holomorphic structures on a topologically trivial line bundle, which does in fact have a holomorphic structure isomorphic to E, i.e. the map t encodes all the information about E.

Example 6.24. Now we let $\mathcal{C}^\infty$ and $(\mathcal{C}^\infty)^\times$ denote real-valued smooth functions on a manifold (arbitrary and nonzero respectively) X. We still have an exponential map on sheaves, but now it is injective with cokernel the sheaf $\underline{\mathbb{Z}/2}$. As $\mathcal{C}^\infty$ is fine we obtain $H^1(X, (\mathcal{C}^\infty)^\times) \cong H^1(X, \underline{\mathbb{Z}/2})$ and this is the group classifying real line bundles.

7. Verdier duality and $f^!$

7.1. Projection formula and Künneth

This section will be a bit compressed due to time reasons. Let $f : X \to Y$ be map of locally compact spaces and let $\mathcal{R}$ be a sheaf of rings on Y.

We can then consider the tensor product over $\mathcal{R}$ and over $f^{-1}\mathcal{R}$.

It is not hard to see that for a left and right $\mathcal{R}$-module $\mathcal{F}$ and $\mathcal{G}$ we have $f^{-1}\mathcal{F} \otimes_{f^{-1}\mathcal{R}} f^{-1}\mathcal{G} \cong f^{-1}(\mathcal{F} \otimes_{\mathcal{R}} \mathcal{G})$ [KS90, Proposition II.2.3.5].

How the tensor product interacts with pushforwards is less clear, but there is the following powerful result.

Proposition 7.1. *Let $X, Y, \mathcal{R}$ be as above and let $\mathcal{G} \in \mathsf{Sh}(f^{-1}\mathcal{R}^{op})$ and $\mathcal{F} \in \mathsf{Sh}(\mathcal{R})$ be sheaves.*

There is a natural map $f_!\mathcal{G} \otimes_{\mathcal{R}} \mathcal{F} \to f_!(\mathcal{G} \otimes_{f^{-1}\mathcal{R}} f^{-1}\mathcal{F})$. If $\mathcal{F}$ is flat over $\mathcal{R}$ this is an isomorphism.

Sketch of proof. The map is induced by a natural map

$$f_*\mathcal{G} \otimes_{\mathcal{R}} \mathcal{F} \to f_*(\mathcal{G} \otimes_{f^{-1}\mathcal{R}} f^{-1}\mathcal{F})$$

which comes via adjunction from $f^{-1}(f_*\mathcal{G} \otimes \mathcal{F}) \cong f^{-1}f_*\mathcal{G} \otimes f^{-1}\mathcal{F} \to \mathcal{G} \otimes f^{-1}\mathcal{F}$, one needs to check it descends to proper pushforward (and also check that $f_!\mathcal{G}$ is a $\mathcal{R}$-module as $\mathcal{G}$ is a $f^{-1}\mathcal{R}$-module).

Then we show there is an isomorphism by base change, but we need the non-trivial lemma that $\Gamma_c(X, \mathcal{F}) \otimes_R M \cong \Gamma_c(X, F \otimes_{\underline{R}} \underline{M})$. □

We want to derive this statement and drop the assumption that $\mathcal{F}$ is flat. To this end we may derive the tensor product by choosing a flat resolution of $\mathcal{F}$. Flat sheaves (i.e. those tensoring with which preserves injections) are an adapted class for the tensor product. (This statement in particular means there are enough flats to resolve arbitrary sheaves, see [KS90, Proposition 2.4.12].)

To make sure our resolutions stay finite we will typically assume that the sheaf of rings $\mathcal{R}$ has *finite weak global dimension*, i.e. there is some n such that the weak global dimension of each stalk $\mathcal{R}_x$ is at most n, i.e. each $\mathcal{R}_x$-module has a flat resolution of length at most n. This is clearly the case for $\underline{\mathbb{Z}}$.

Theorem 7.2. *Let $X, Y, \mathcal{R}$ be as above and let $G \in D^+(f^{-1}\mathcal{R}^{op})$ and $F \in D^+(\mathcal{R})$. Let $\mathcal{R}$ have finite weak global dimension. There is a natural quasi-isomorphism $Rf_!G \otimes^L_{\mathcal{R}} F \to Rf_!(G \otimes^L_{f^{-1}\mathcal{R}} f^{-1}F)$.*

Sketch of proof. We prove the theorem if F is a flat sheaf $\mathcal{F}$. First show that $-\otimes_{f^{-1}\mathcal{R}} f^{-1}F$ sends c-soft sheaves to $f_!$-acyclic sheaves. Then $Rf_!(-\otimes^L_{f^{-1}\mathcal{R}} f^{-1}\mathcal{F})$ is the derived functor of $f_!(-\otimes^L_{f^{-1}\mathbb{R}} f^{-1}\mathcal{F})$. Since $-\otimes_{)} \mathcal{R}F$ is exact we also have that $Rf_!(-)\otimes^L_{\mathcal{R}} F$ is the derived functor of $f_!()\otimes_{\mathcal{R}}F$. Thus the result follows by Proposition 7.1.

We extend by flatly resolving arbitrary F, noting that the resolution is still bounded above has $\mathcal{R}$ has finite weak global dimension. □

More details on the projection formula can be found in [Ive86, VII.2.4] and [KS90, Proposition II.2.5.13]. The following is a special case of a more general statement.

Corollary 7.3 (Künneth theorem)**.** *Let $(S,\mathcal{R})$ be a ringed space such that $\mathcal{R}$ has finite weak global dimension. Let $p_X : X \to S$ and $p_Y : Y \to S$ be maps of l of locally compact spaces. Let $\mathcal{F} \in D^+(p_X^{-1}\mathcal{R})$ and $\mathcal{G} \in D^+(p_Y^{-1}\mathcal{R})$ be sheaves on and $\mathcal{G}$. Consider the projections $q_X : X\times_S Y \to X$ and $q_Y : X\times_S Y \to Y$ and the map $f = p_X q_X = p_Y q_Y : X\times_S Y \to S$. Define the exterior product $\mathcal{F}\boxtimes\mathcal{G} = \pi_X^*\mathcal{F}\otimes_{f^{-1}\mathcal{R}}\pi_Y^*\mathcal{G}$ on $X\times Y$. Then*

$$Rf_!(X\times_S Y, \mathcal{F}\boxtimes\mathcal{G}) \simeq Rf_!(X,\mathcal{F})\otimes^L_R Rf_!(Y,\mathcal{G})$$

in particular if $S = *$ and $\mathcal{R}$ is a field k we have

$$H^n_c(X\times Y,\underline{k}) = \oplus_{i+j=n} H^i_c(X,\underline{k})\otimes H^j_c(Y,\underline{k})$$

Proof. We compute

$$\begin{aligned} Rf_!(X\times Y,\mathcal{F}\boxtimes\mathcal{G}) &\simeq Rp_{X,!}Rq_{X,!}(p_X^{-1}\mathcal{F}\otimes_{f^{-1}\mathcal{R}} p_Y^{-1}\mathcal{G}) \\ &\simeq Rp_{X,!}(\mathcal{F}\otimes_{p_X^{-1}\mathcal{R}} Rq_{X,!}q_Y^{-1}\mathcal{G}) \\ &\simeq Rp_{X,!}(\mathcal{F}\otimes_{p_X^{-1}\mathcal{R}} p_X^{-1}Rq_{Y,!}\mathcal{G}) \\ &\simeq Rp_{X,!}\mathcal{F}\otimes_{\mathcal{R}} Rp_{Y,!}\mathcal{G} \end{aligned}$$

Where we used the projection formula from Theorem 7.2 twice and base change (Theorem 5.18) in between.

For the last part just note that working over a field the cohomology of a tensor product of bounded below complexes is the tensor product of the cohomologies. □

7.2. Verdier duality

Given a map of spaces $f : X \to Y$ and sheaves on X and Y we now have five useful functors (namely $f_*, f^{-1}, f_!, \otimes$ and $\underline{\mathrm{Hom}}$) satisfying some relus. In particular $f^{-1} \dashv f_*$ and $-\otimes M \dashv \underline{\mathrm{Hom}}(M,-)$ are adjunctions. There is no adjoint for $f_!$. We have seen that it is not in general right exact, so it cannot be a left adjoint. But it is not a right adjoint either, global sections with compact support does not preserve products. (Consider for example $\prod \mathbb{R}_n$ on $\mathbb{R}$, see also Exercise 11.4.)

However, the derived functor $Rf_!$ is exact and has a right adjoint, which will be denoted by $f^!$. The notation has no R or L as this is not the derived functor of anything, but it is a functor that is only defined on the level of derived categories!

We will state the theorem and its main consequences before turning to the proofs.

Theorem 7.4 (Verdier Duality)**.** *Let $f : X \to Y$ be a map of locally compact spaces that has finite cohomological dimension, i.e. there is some n such that $R^i f_!(\mathcal{F}) = 0$ for all $i > n$ and all $\mathcal{F} \in \mathsf{Sh}(X)$. Let R be commutative ring of finite weak global dimension. Then there is an adjunction*

$$Rf_! \dashv f^! : D^+(\underline{R}_X) \leftrightarrows D^+(\underline{R}_Y)$$

where $f^!$ is called the exceptional inverse image.

Poincaré duality is a direct consequence of Verdier duality once we compute one very special case of the exceptional inverse image:

Proposition 7.5. *Let M be an n-manifold and k a field and $p_M : M \to *$ the natural projection. Then the shifted complex $p_M^! k$ is quasi-isomorphic to a shifted sheaf ω_M that is locally isomorphic to the constant sheaf $\underline{k}[n]$. If M is orientable $\omega_M \simeq \underline{k}$.*

The proposition also holds if k is a commutative ring with finite global dimension.

Corollary 7.6. *Let M be an orientable n-manifold. Then for each i we have $H_c^{n-i}(M, k)^\vee \cong H^i(M, k)$.*

Proof. We apply Theorem 7.4 to the projection map p_M to get

$$\mathrm{Hom}_{D(k)}(Rp_!\underline{k}, k) \simeq \mathrm{Hom}_{D^+(M)}(\underline{k}, p_M^! k)$$

In fact, just as for the adjunction $f^{-1} \dashv Rf_*$ the adjunction extends to a quasi-isomorphism on hom complexes. Thus by Proposition 7.5 we have $(Rp_!\underline{k})^* \simeq R\Gamma(\underline{k}[n])$ and taking cohomology we obtain the desired result. □

Example 7.7.

1. For $M \cong S^n$ we have $H_c^i \cong H^i$ and Poincaré duality swaps H^0 and H^n, all other cohomolgy groups being zero.
2. For $\mathbb{R}^n$ this matches with our previous computations, in particular we obtain another proof that $H^i(\mathbb{R}^n, k) \cong k$ if and only if $i = n$ and 0 otherwise.
3. Take now a non-orientable manifold like $\mathbb{R}P^2$. Theorem 7.4 and Proposition 7.5 together give us that there is a locally constant sheaf $\omega \simeq p^! k$ on $\mathbb{R}P^2$ such that $R\Gamma(\mathbb{R}P^2, \omega[2])$ is dual to $R\Gamma(\mathbb{R}P^2, \underline{\mathbb{Z}})$, thus

$$R\Gamma(\mathbb{R}P^2, \omega) \simeq \mathrm{Hom}\left(\mathbb{Z} \xrightarrow{0} \mathbb{Z} \xrightarrow{2} \mathbb{Z}, \mathbb{Z}\right)[-2]$$

which has cohomology $\mathbb{Z}$ in degree 2 and $\mathbb{Z}/2$ in degree 1 and 0 elsewhere. So we learn about the sheaf ω which is not $\underline{k}_{\mathbb{R}P^2}$. Note that we may just compute $R\Gamma(\mathbb{R}P^2, \mathbb{Z})$ by any complex with the correct cohomology because $\mathbb{Z}$ has global dimension 1 (we also say it is hereditary), i.e. any module has a resolution of length 2, and it follows that any complex is quasi-isomorphic to its cohomology.

Before embarking on a proof we wil consider one special case, that we will appeal to in the course of the proof.

Lemma 7.8. *Let $j : U \to X$ be an open inclusion. Then $j_! \dashv j^{-1} : \mathsf{Sh}(U) \to \mathsf{Sh}(X)$.*

In particular in this case $j^! = j^{-1}$ is an (exact) functor on the level of abelian categories.

Proof. We have already constructed a map $j_! j^{-1}\mathcal{F} \to \mathcal{F}$ for $\mathcal{F}$ a sheaf on X given by inclusion of those sections whose support is contained in U, see Theorem 5.25. Given a sheaf $\mathcal{G}$ on U we can check on each open $W \subset U$ that $j^{-1}j_!\mathcal{G}(W) = j_!\mathcal{G}(W) = \mathcal{G}(W)$. Thus we have a candidate counit and unit for our natural transformation and the unit is the identity.

The triangle identities follow: The composition $j_!\mathcal{G} \to j_!(j^{-1}j_!)\mathcal{G} = (j_!j^{-1})j_!\mathcal{G} \to j_!\mathcal{G}$ is just the composition of two identity maps.

The composition $j^{-1}\mathcal{F} \to (j^{-1}j_!)j^{-1}\mathcal{F} \to j^{-1}(j_!j^{-1})\mathcal{F} \to j^{-1}\mathcal{F}$ is also the identity as we can see for example on stalks. □

7.3. Proof of Verdier duality

We begin by fixing the idea of the proof. Given $f : X \to Y$ for any open $j : V \subset X$ and $F \in D^+\underline{R}_Y$ we expect to have

$$\begin{aligned} R\Gamma(V, f^! F) &\simeq R\operatorname{Hom}(j_!\underline{R}_V, f^! F) \\ &\cong R\operatorname{Hom}(Rf_! j_!\underline{R}_V, F) \end{aligned}$$

and if F is a complex of injectives we may choose a suitable resolution K of $j_!\underline{R}_V$ and get

$$R\Gamma(V, f^! F) \simeq \underline{\operatorname{Hom}}(f_! K, F).$$

We will now make all of this precise.

We will fix a commutative ring R of finite weak global dimension and work with sheaves of $\underline{R}$-modules. There is a version of Verdier duality for more general $\mathcal{R}$-modules, see [KS90, Remark III.3.1.6].

We introduce a notion that we briefly met in the proof of Theorem 5.18:

Definition 7.9. A sheaf $\mathcal{F}$ on X is called *f-soft* if for any $y \in Y$ the sheaf $\mathcal{F}|_{f^{-1}(y)}$ is c-soft.

Definition 7.10. A map $f : X \to Y$ of locally compact spaces that has *cohomological dimension n*, if $R^i f_!(\mathcal{F}) = 0$ for all $i > n$ and all $\mathcal{F} \in \mathsf{Sh}(X)$ and n is maximal with this property.

As in the statement of the theorem we will from now on assume that $f : X \to Y$ has finite cohomological dimension.

Lemma 7.11. *If f has cohomological dimension n and $\mathcal{F}^0 \to \mathcal{F}^1 \cdots \to \mathcal{F}^r$ is an exact sequence with all $\mathcal{F}^i$ for $i < n$ f-soft then $\mathcal{F}^n$ is also n-soft.*

Proof. See Exercise 12.3. □

Definition 7.12. We introduce the following notation: For $\mathcal{F}$ a sheaf on X and $j : A \to X$ the inclusion of a locally closed subset we denote by $\mathcal{F}_A$ the sheaf $j_! j^{-1}\mathcal{F}$.

We now define for any $\underline{\mathbb{Z}}_X$ module $\mathcal{K}$ and $\underline{R}_X$-module $\mathcal{F}$ an auxiliary presheaf $f^!_{\mathcal{K}}\mathcal{F}$. We define

$$V \mapsto f^!_{\mathcal{K}}\mathcal{F}(V) = \operatorname{Hom}_{\underline{R}_Y}\left(f_!\left(\underline{R}_X \otimes_{\underline{\mathbb{Z}}_X} \mathcal{K}_V\right), \mathcal{F}\right).$$

For any inclusion of opens $W \subset V$ with $k : W \to X$ we have a natural extension map $\mathcal{K}_W = k_! k^{-1}\mathcal{K} \to j_! j^{-1}\mathcal{K} = \mathcal{K}_V$ induced by the natural transformation $i_! i^{-1} \to \mathbf{1}$ as in the proof of Theorem 5.25. This induces $f_!\left(\underline{R}_X \otimes_{\underline{\mathbb{Z}}_X} \mathcal{K}_W\right) \to f_!\left(\underline{R}_X \otimes_{\underline{\mathbb{Z}}_X} \mathcal{K}_V\right)$ and thus the restriction map $f^!_{\mathcal{K}}\mathcal{F}(V) \to f^!_{\mathcal{K}}\mathcal{F}(W)$.

We next have to show that what we have is a sheaf. We need some preliminary results, some of which we could have considered earlier.

Lemma 7.13. *The proper pushforward functor $f_!$ commutes with arbitrary direct sums.*

Proof. See Exercise 12.4. □

Lemma 7.14. *Let $\mathcal{F}$ be a an $\mathcal{R}$-module on a ringed space $(X, \mathcal{R})$. Then there is an epimorphism $\mathcal{P} \to \mathcal{F}$ from an $\mathcal{R}$-flat sheaf $\mathcal{P}$.*

Proof. We consider the indexing set I consisting of all pairs $(U \subset X, s \in \mathcal{F}(U))$ for U open in X. Then we set $\mathcal{P} = \oplus_{(j,s)\in I}\mathcal{R}_U$. Each section $S \in \mathcal{F}(U)$ defines a map $\mathcal{R}_U \to \mathcal{F}_U \to \mathcal{F}$ sending $1 \in \Gamma(U, \mathcal{R}_U)$ to $s \in \Gamma(U, \mathcal{F}_U)$ and applying the natural inclusion $\mathcal{F}_U \to \mathcal{F}$. This is an epimorphism by construction and as each stalk is a free $\mathcal{R}_x$ module $\mathcal{P}$ is flat. □

Lemma 7.15. *Let $\mathcal{K}$ be a flat and f-soft $\underline{\mathbb{Z}}_X$-module. Then for any sheaf $\mathcal{F}$ on X the sheaf $\mathcal{F} \otimes_{\mathbb{Z}_X} \mathcal{K}$ is f-soft. In particular $\mathcal{F} \mapsto f_!(\mathcal{F} \otimes_{\mathbb{Z}_X} \mathcal{K})$ is exact.*

Proof. By the proof of Lemma 7.14 we have a resolution

$$\cdots F^{-r} \cdots F^{-1} \to F^0 \to \mathcal{F}$$

where each F^i is a direct sum of sheaves $\underline{\mathbb{Z}}_V$ for $V \subset X$ open. The tensor product of f-softs is again f-soft as tensor product commutes with pullback and by definition a tensor product of c-soft sheaves is c-soft.

Thus we have an exact complex of sheaves all but the last of which are f-soft sheaves

$$\cdots \to F^{-r} \otimes_{\underline{\mathbb{Z}}_X} \mathcal{K} \to \cdots F^{-1} \otimes_{\underline{\mathbb{Z}}_X} \mathcal{K} \to F^0 \otimes_{\underline{\mathbb{Z}}_X} \mathcal{K} \to \mathcal{F} \otimes_{\underline{\mathbb{Z}}_X} \mathcal{K}$$

Truncating at a large enough r we may apply Lemma 7.11. (Note that we are switching from this complex as a left resolution to viewing it as a right resolution!)

It follows from base change (5.18) that f-soft sheaves are $f_!$-acyclic and as $-\otimes\mathcal{K}$ is exact the second result follows. □

We are now approaching the heart of the proof.

Lemma 7.16. *Let $\mathcal{K}$ be a flat and f-soft $\underline{\mathbb{Z}}_X$-module and $\mathcal{G}$ an $\underline{R}_Y$-module. Then the presheaf $f^!_{\mathcal{K}}\mathcal{G}$ is a sheaf. The presheaf $f^!_{\mathcal{K}}\mathcal{G}$ is a sheaf.*

Proof. We simplify notation and write $\otimes$ for $\otimes_{\underline{\mathbb{Z}}_X}$.

We prove first that $f^!_{\mathcal{K}}\mathcal{G}$ is a sheaf. So take an open subset V of X and choose an open covering $\{U_i\}$. Slightly generalising the proof of the Mayer Vietoris theorem with compact support (Exercise 8.3) we have an exact sequence

$$\bigoplus_{j<k}\underline{R}_{U_j\cap U_k}\to\bigoplus_i\underline{R}_{U_i}\to\underline{R}_V\to 0$$

(note that this sequence is not exact on the left if we have more than two opens in our cover) and by Lemma 7.15 this gives an exact sequence

$$f_!\left(\bigoplus_{j,k}\underline{R}_{U_j\cap U_k}\otimes\mathcal{K}\right)\to f_!\left(\bigoplus_i\underline{R}_{U_i}\otimes\mathcal{K}\right)\to f_!(\underline{R}_V\otimes\mathcal{K})\to 0.$$

Now this gives us an exact sequence

$$0\to\mathrm{Hom}_{\underline{R}_Y}(f_!(\underline{R}_V\otimes\mathcal{K}),\mathcal{G})\to\mathrm{Hom}_{\underline{R}_Y}(f_!(\oplus_i\underline{R}_{U_i}\otimes\mathcal{K}),\mathcal{G})\to\mathrm{Hom}_{\underline{R}_Y}(f_!(\oplus_{j,k}\underline{R}_{U_j\cap U_k}\otimes\mathcal{K}),\mathcal{G})$$

(In a previous version we used here an extra assumption that $\mathcal{G}$ is injective (as do [**KS**]) but as pointed out by Muskan Abbas this is not needed. One can check that any left exact functor F sends any exact sequence $0\to A\to B\to C$ to $0\to FA\to FB\to FC$.)

By Lemma 7.13 we can pull the direct sum past $f_!$. Next we observe that $\mathbb{R}_V\otimes K$ is fact isomorphic to $\mathbb{R}_X\otimes K_V$. We write $j_!j^{-1}\mathbb{R}_X\otimes K\cong j_!(j^{-1}\mathbb{R}_X\otimes j^{-1}K)\cong\mathbb{R}_X\otimes j_!j^{-1}K$ using the projection fromula 7.1 twice.

Then by the definition of $f^!_{\mathcal{K}}$ our sequence is isomorphic to

$$0\to f^!_{\mathcal{K}}\mathcal{F}(V)\to\prod_i f^!\mathcal{F}_{\mathcal{K}}(U_i)\to\prod_{j,k}f^!\mathcal{F}_{\mathcal{K}}(U_j\cap U_k)$$

which precise says that $f^!_{\mathcal{K}}\mathcal{F}$ is a sheaf! □

Lemma 7.17. *Let again $\mathcal{K}$ be a flat and f-soft $\underline{\mathbb{Z}}_X$-module and $\mathcal{G}$ an $\underline{R}_Y$-module. Then there is a canonical isomorphism*

$$\mathrm{Hom}_{\underline{R}_Y}(f_!(\mathcal{F}\otimes_{\underline{\mathbb{Z}}_X}\mathcal{K}),\mathcal{G})\cong\mathrm{Hom}_{\underline{R}_X}(\mathcal{F},f^!_{\mathcal{K}}\mathcal{G})$$

Proof. The first task is to define the map. We are trying to build $\alpha_{\mathcal{F}} : \mathrm{Hom}_{\underline{R}_Y}(f_!(\mathcal{F}\otimes\mathcal{K}),\mathcal{G}) \to \mathrm{Hom}_{\underline{R}_X}(\mathcal{F}, f^!_{\mathcal{K}}\mathcal{G})$ so let us begin with $\phi : f_!(\mathcal{F}\otimes\mathcal{K}) \to \mathcal{G}$ and now for any open set $V \subset X$ we want to construct, functorially in V, a map

$$\alpha_V : \mathcal{F}(V) \to (f^!_{\mathcal{K}}\mathcal{G})(V) \cong \mathrm{Hom}_{\underline{R}_Y}(f_!(\underline{R}_X \otimes \mathcal{K}_V), \mathcal{G}).$$

We now apply Corollary 2.44 to the projection $p : Y \to *$ and then Corollary 2.45 to obtain

$$\begin{aligned}\mathrm{Hom}_R(\mathcal{F}(V), \mathrm{Hom}_{\underline{R}_Y}(f_!(\underline{R}_X\otimes\mathcal{K}_V),\mathcal{G})) &\cong \mathrm{Hom}_R(\mathcal{F}(V), p_*\,\mathcal{H}om_{\underline{R}_Y}(f_!(\underline{R}_X\otimes\mathcal{K}_V,\mathcal{G})))\\ &\cong \mathrm{Hom}_{\underline{R}_Y}(p^{-1}(\mathcal{F}(V)), \mathcal{H}om_{\underline{R}_Y}(f_!(\underline{R}_X\otimes\mathcal{K}_V),\mathcal{G}))\\ &\cong \mathrm{Hom}_{\underline{R}_Y}(p^{-1}(\mathcal{F}(V))\otimes_{\underline{R}_X} f_!(\underline{R}_X\otimes\mathcal{K}_V),\mathcal{G})).\end{aligned}$$

So we have to use ϕ to find an element in $\mathrm{Hom}_{\underline{R}_Y}(\underline{\mathcal{F}(V)}\otimes_{\underline{R}_X}(f_!(\underline{R}_X\otimes\mathcal{K}_V),\mathcal{G}))$.

$$\begin{aligned}\underline{\mathcal{F}(V)}\otimes_{\underline{R}_X} f_!(\underline{R}_X\otimes\mathcal{K}_V) &\to f_!(\underline{\mathcal{F}(V)}\otimes\mathcal{K}_V) && \text{(A)}\\ &\xrightarrow{\cong} f_!((pf)^{-1}(pfj)_*j^{-1}\mathcal{F}\otimes\mathcal{K}_V)\\ &\to f_!(j_*j^{-1}\mathcal{F}\otimes j_!j^{-1}\mathcal{K})\\ &\to f_!(j_!(j^{-1}\mathcal{F}\otimes j^{-1}\mathcal{K}) && \text{(B)}\\ &\xrightarrow{\cong} f_!(j_!j^{-1}(\mathcal{F}\otimes\mathcal{K})\\ &\to f_!(\mathcal{F}\otimes\mathcal{K})\\ &\xrightarrow{\phi} \mathcal{G}\end{aligned}$$

Where we applied the projection formula 7.1 in (A). Then we unravelled the meaning of $\underline{\mathcal{F}(V)}$ where $j : V \to X$ is the inclusion and applied the counit of the adjunction $(pf)^{-1} \dashv (pf)_*$. For (B) we use the projection formula and adjunction again to deduce that for any two sheaves $\mathcal{F}, \mathcal{K}$ and any map g there is a natural map

$$g_*\mathcal{F}\otimes g_!\mathcal{K} \to g_!(g^{-1}g_*\mathcal{F}\otimes\mathcal{K}) \to g_!(\mathcal{F}\otimes\mathcal{K}).$$

Intuitively speaking a section of the tensor product has proper support if one of the two factors does. We conclude with the compatibility of pullback and tensor product followed by the natural transformation $j_!j^{-1} \to \mathbf{1}_{\mathsf{Sh}(X)}$.

As one can check all our constructions are functorial with respect to V, i.e. $\alpha_W \circ r_{VW} \cong r_{VW}\alpha_V$ for $W \subset V$. This completes the construction of α.

It is also clear that our constructions are natural with respect to $\mathcal{F}$ and $\mathcal{G}$.

To check $\alpha_{\mathcal{F}}$ is an isomorphism we will proceed in three steps. First let us assume $\mathcal{F} = \underline{R}_U$ for some open U. Then the map $\alpha_{\underline{R}_U}$ simplifies considerably, we have

$$\begin{aligned}\mathrm{Hom}_{\underline{R}_Y}(f_!(\mathcal{F}\otimes\mathcal{K}),\mathcal{G}) &\cong \mathrm{Hom}_{\underline{R}_Y}(f_!(\underline{R}_U\otimes\mathcal{K}),\mathcal{G})\\ &\cong (f^!_{\mathcal{K}}\mathcal{G})(U)\\ &\cong \mathrm{Hom}_{\underline{R}_U}(\underline{R}_U, j^{-1}(f^!_{\mathcal{K}}(\mathcal{G})))\\ &\cong \mathrm{Hom}_{\underline{R}_X}(\underline{R}_U, f^!_{\mathcal{K}}\mathcal{G}),\end{aligned}$$

where we used Lemma 7.8 in the last step.

One can check that this is indeed the map we constructed above.

Secondly, we deduce the result for $\mathcal{F} = \oplus_i \mathbb{R}_{U_i}$ for some collection of opens $\{U_i\}$. Here $\alpha_{\oplus \underline{R}_i}$ is just the product $\prod_i \alpha_{\underline{R}_i}$, using that $f_!$ commutes with direct sums (Lemma 7.13).

Finally we consider an arbitrary $\underline{R}_X$-module $\mathcal{F}$ and resolve it as in the proof of Lemma 7.14 with

$$0 \to \mathcal{F}' \to \oplus_j \underline{R}_{U_j} \to \mathcal{F} \to 0$$

Here $\alpha_{\oplus \underline{R}_{U_i}}$ is an isomorphism by the above.

We consider the diagram

$$\begin{array}{ccccccc}
0 & \longrightarrow & \mathrm{Hom}(f_!(\mathcal{F} \otimes \mathcal{K}), \mathcal{G}) & \longrightarrow & \mathrm{Hom}(f_!(\oplus_j \underline{R}_{U_j} \otimes \mathcal{K}), \mathcal{G}) & \longrightarrow & \mathrm{Hom}(f_!(\mathcal{F}' \otimes \mathcal{K}), \mathcal{G}) \\
\downarrow & & \downarrow{\scriptstyle \alpha_{\mathcal{F}}} & & \downarrow{\scriptstyle \alpha_{\oplus \underline{R}_{U_i}}} & & \downarrow{\scriptstyle \alpha_{\mathcal{F}'}} \\
0 & \longrightarrow & \mathrm{Hom}(\mathcal{F}, f^!_{\mathcal{K}} \mathcal{G}) & \longrightarrow & \mathrm{Hom}(\oplus_j \underline{R}_{U_j}, f^!_{\mathcal{K}} \mathcal{G}) & \longrightarrow & \mathrm{Hom}(\mathcal{F}', f^!_{\mathcal{K}} \mathcal{G})
\end{array}$$

where both rows are exact by Lemma 7.15. Moreover we have just shown $\alpha_{\oplus \underline{R}_{U_i}}$ is an isomorphism. This forces $\alpha_{\mathcal{F}}$ to be injective. But by the same reasining $\alpha_{\mathcal{F}'}$ is injective and this implies (by the refined 5-lemma) that $\alpha_{\mathcal{F}}$ is also surjective. □

Corollary 7.18. *Let again $\mathcal{K}$ be a flat and f-soft $\underline{\mathbb{Z}}_X$-module and let now $\mathcal{G}$ an injective $\underline{R}_Y$-module. The sheaf $f^!_{\mathcal{K}} \mathcal{G}$ is an injective $\underline{R}_X$-module.*

Proof. Let $0 \to \mathcal{A} \to \mathcal{B} \to \mathcal{C} \to 0$ be a short exact sequence of sheaves on X. Then by Lemma 7.15 and injectivity of $\mathcal{G}$ we have a short exact sequence

$$0 \to \mathrm{Hom}(f_!(\mathcal{A} \otimes \mathcal{K}), \mathcal{G}) \to \mathrm{Hom}(f_!(\mathcal{B} \otimes \mathcal{K}), \mathcal{G}) \to \mathrm{Hom}(f_!(\mathcal{C} \otimes \mathcal{K}), \mathcal{G}) \to 0$$

And by Lemma 7.17 this shows $f^!_K(\mathcal{G})$ is injective. □

Next we will apply Lemma 7.17 to suitable resolutions to deduce Theorem 7.4.

Lemma 7.19. *Let f have cohomological dimension n. The sheaf $\underline{\mathbb{Z}}_X$ admits a resolution $0 \to \underline{\mathbb{Z}}_X \to K^0 \to K^1 \to \ldots \to K^n \to 0$ such that all K^i are flat and f-soft $\underline{\mathbb{Z}}_X$-modules.*

Proof. We use a Godement resolution: Let $\delta : X^\delta \to X$ the identity map from X with the discrete topology and let $K^0 = \delta_* \delta^{-1} \underline{\mathbb{Z}}_X$, $K^1 = \delta_* \delta^{-1}(K^0 / \underline{\mathbb{Z}}_X)$ and so on with $K^j = \delta_* \delta^{-1}(\mathrm{coker}\, K^{j-2} \to K^{j-1})$ for $1 < j < r$ and $K^r = \mathit{coker}(K^{r-2} \to K^{r-1})$.

Then as the K^j for $j < r$ are flabby they are rc-soft by Proposition 5.17.1 and f-soft by Proposition 5.17.3. It follows that K^r is also f-soft by Lemma 7.11.

To show that the K^i are flat we show that for any flat shaf $\mathcal{G}$ both $\delta_* \delta^{-1} \mathcal{G}$ and $\delta_* \delta^{-1} \mathcal{G} / G$ are flat. For any point $x \in X$ we take the stalks:

$$(\delta_* \delta^{-1} \mathcal{G})_x = \operatorname*{colim}_{x \in U} \prod_{x' \in U} \mathcal{G}_{x'}$$

and

$$(\delta_*\delta^{-1}\mathcal{G}/\mathcal{G})_x = \operatorname*{colim}_{x \in U} \prod_{x' \in U \setminus \{x\}} \mathcal{G}_{x'}$$

by definition. Both of these are torsion free abelian groups and thus flat. □

Definition 7.20. We define the functor $f^! : D^+(\underline{R}_Y) \to D^+(\underline{R}_X)$ by choosing a resolution K of $\mathbb{Z}_X$ as in Lemma 7.19 and an injective resolution I of a complex G representing an object in $D^+(\underline{R}_X)$ by Corollary 3.45. Then we let $f^!(G)$ be the class of the total complex associated to the double complex $f^!_K(I) = \mathrm{Hom}_{\underline{R}_Y}\left(f_!\left(\underline{R}_X \otimes_{\underline{\mathbb{Z}}_X} K^\bullet_V\right), I^\bullet\right)$. We call this functor the *exceptional inverse image*.

We note that the double complex is concentrated in finitely many columns so we can choose sum or product totalization. It is indeed bounded below. The complex $f^!_K(I)$ consists of injectives by Corollary 7.18. The construction sends morphisms homotopy to zero to morphisms homotopic to zero (since $\mathrm{Hom}(f_!(F \otimes K), -)$ preserves homotopies) thus it is indeed well-defined on homotopy categories and thus on the derived category.

Proof of Theorem 7.4. Consider objects $F \in D^+(\underline{R}_X)$ and $G \in D^+(\underline{R}_Y)$ and choose injective resolutions $F \simeq I$ and $G \simeq J$. We choose K as in the definition of $f^!$. We compute

$$\begin{aligned}
\mathrm{Hom}_{D^+(\underline{R}_Y)}(Rf_!F, G) &\cong \mathrm{Hom}_{K^+(\underline{R}_Y)}(Rf_!I, J) \\
&\cong \mathrm{Hom}_{K^+(\underline{R}_Y)}(f_!(I \otimes_{\underline{\mathbb{Z}}_X} K, J) \\
&\cong \mathrm{Hom}_{K^+(\underline{R}_X)}(I, f^!_K J) \\
&\cong \mathrm{Hom}_{D^+)\underline{R}_X}(F, f^! G)
\end{aligned}$$

where we used Corollary 3.45, the fact that $I \otimes K$ is f-soft by Lemma 7.15 and most importantly Lemma 7.17 in each bidegree. We finish with another use of Corollary 3.45 as $f^I G$ is an complex of injectives by Corollary 7.18. □

Remark 7.21. The condition of finite cohomological dimension is not unreasonable. Thanks to base change it suffices to check the cohomological dimension of the fibers and as you saw in Exercise 9.2 any n-manifold and any locally closed subspace of $\mathbb{R}^n$ has cohomological dimension n. This is explained in [Ive86, Section III.9].

We have the following refinement:

Corollary 7.22. *Let $f : X \to Y$ be a map of locally compact spaces that has finite cohomological dimension. Then for F in $D^b(\underline{R}_X)$ and G in $D^+(\underline{R}_Y)$ we have natural quasi-isomorphisms*

$$R\mathcal{H}om(Rf_!F, G) \simeq Rf_*R\mathcal{H}om(F, f^!G)$$

where we use $\mathcal{H}om$ to denote the sheaf hom complex, and

$$R\,\underline{\mathrm{Hom}}(Rf_!F, G) \simeq R\,\underline{\mathrm{Hom}}(F, f^!G).$$

Note that we restrict F to be bounded her, if F is only bounded below then the sheaf hom complex $\mathcal{H}om(F, f^!G)$ could be unbounded in both directions

Proof. The second statement follows from the first by taking global sections. We first note that there is a map

$$Rf_*R\mathcal{H}om(F, f^!G) \to R\mathcal{H}om(Rf_!F, Rf_!f^!G)$$

To show this we use, as in the proof of Lemma 7.17, that for any two sheaves (or complexes of sheaves) A and B there is a map $f_!A \otimes f_*B \to f_!(A \otimes B)$ and then set $A = F$ and $B = \mathcal{H}om(F, f^!G)$ to conclude using the natural map $F \otimes \mathcal{H}om(F, f^!G) \to f^!G$.

The adjunction from Theorem 7.4 provides the counit $Rf_!f^!G \to G$ and together we have a map

$$R\mathcal{H}om(Rf_!F, G) \to R\mathcal{H}om(F, f^!G).$$

It remains to check this is a quasi-isomorphism on each open $U \subset Y$. We compute using Theorem 7.4:

$$\begin{aligned} H^j(R\Gamma(V, Rf_*R\mathcal{H}om(F, f^!G))) &\cong \mathrm{Hom}_{D^+(\underline{R}_V)}(F|_{f^{-1}(V)}, f^!G[j]|_{f^{-1}(V)}) \\ &\cong \mathrm{Hom}_{D^+(\underline{R}_V)}(Rf_!F|_V, G[j]|_V) \\ &\cong H^j(R\Gamma(V, R\mathcal{H}om(Rf_!F, G))) \end{aligned}$$

where we also used $R\Gamma(V, -) \circ Rf_* \simeq R\Gamma(f^{-1}(V), -)$. □

We obtain some useful properties for $f^!$ by its adjoint property, for example $(g \circ f)^! = f^! \circ g^!$ as both of these are righ adjoint to $R(g \circ f)_! \simeq Rg_! \circ Rf_!$.

Remark 7.23. There are different approaches to proving Verdier duality. One can use an adjoint functoror representability theorem to show that $f_!$ must have a right adjoint because it commutes with direct sums. This is the approach in [GM03, Section III.8].

To avoid the long chain of maps we used to define α one may try to define α on sums of $\underline{R}_U$ and then extend to cokernels. In our particular situation this works quite smoothly: The resolutions by $\underline{R}_U$ we have are not just canonical but also functorial: any map $f : \mathcal{F} \to \mathcal{G}$ of functors induces maps $(U, s \in \mathcal{F}(U)) \to (U, f(s) \in \mathcal{G}(U))$ and thus a map on resolutions inducing f. (We can truncate resolutions in sufficiently low degree and consider the induced map on kernels.) The interested reader is welcome to work out the details!

In general, some care is needed when extending a functor from generators to a category.

7.4. Submersions

To deduce Poincaré duality form Verdier duality we have to compute $p^!R$ for the projection to the point. In fact, the exceptional preimage of the constant sheaf is a key input in many computations.

A *topological manifold* is just a second countable Hausdorff space that is locally homeomorphic to $\mathbb{R}^n$ (or a locally ringed space locally isomorphic to $(\mathbb{R}^n, C^0)$).

A map $f : X \to Y$ between topological manifolds is called a *submersion* of fiber dimension n if each point $x \in X$ has a neighbourhood U such that $V = f(U) \subset Y$ is open and there is a homeomorphism $h : U \cong V \times \mathbb{R}^n$ compatible with the maps to U.

One can also consider smooth submersions by asking f and h are smooth maps.

Definition 7.24. Let $f : X \to Y$ have finite cohomological dimension and assume a commutative ring R is fixed. We define the *relative dualizing complex* $\omega_{X/Y}$ (or ω_f) to be $f^! \underline{R}_Y$. For $Y = *$ we call $\omega_{X/Y}$ the *dualizing complex* and denote it by ω_X.

Lemma 7.25. *Let $f : X \to Y$ be a submersion of fiber dimension n. Then $f^! \underline{R}_Y$ is quasi-isomorphic to a locally constant sheaf $\omega_{X/Y}$ with fiber $R[n]$ on X.*

This means $H^\bullet(f^! \underline{R}_Y)$ is concentrated in degree $-n$. Over $\mathbb{Z}$ it is not hard to see that any sheaf F that is locally isomorphic to F satisfies $F \otimes F \cong \underline{\mathbb{Z}}$.

Proof. We consider $f^! \underline{R}_Y$ locally on U where we may assume U is homeomorphic to $\mathbb{R}^{\dim Y}$ as open balls are cofinal in the topology of a manifold. As f is a submersion we may also assume $U \cong V \times D$ for some open ball in Y and some disk of dimension $\dim X - \dim Y$. We have $R\Gamma(U, f^! \underline{R}_Y) \cong R\operatorname{Hom}_{\underline{R}_U}(\underline{R}_U, f^! \underline{R}_Y) \simeq R\operatorname{Hom}_{\underline{R}_V}(Rf_! \underline{R}_U, \underline{R}_Y)$. But the restriction of f is just projection from a product, $fj = pr_1 : U \cong V \times D \to V$ so we can apply base change 5.18 to the diagram

$$\begin{array}{ccc} V \times D & \xrightarrow{p_2} & D \\ \downarrow{\scriptstyle p_1 = fj} & & \downarrow{\scriptstyle p_D} \\ V & \xrightarrow{p_V} & * \end{array}$$

and find $Rp_{1,!} p_2^{-1} \underline{R} \simeq p_V^{-1} Rp_{D,!}(D, \underline{R})$, thus $Rf_! \underline{R}_U \simeq R\Gamma_c(\mathbb{R}^n, \underline{R}) \otimes \underline{R}_V \simeq \underline{R}_V[n]$. □

If the dualizing complex ω_X is quasi-isomorphic to a shifted sheaf $\mathfrak{or}[k]$, for example if X is a manifold, we call the sheaf the *orientation sheaf.*

Definition 7.26. An *R-orientation* on an n-manifold X is a quasi-isomorphism $\underline{R}_X[n] \simeq \omega_X$.

Such a quasi-isomorphism is nothing but a compatible choice of generators of $R\Gamma_C(U, R) \cong H^n(U, U \setminus \{x\})$ for a cover of X by open balls. So this agrees with notion of orientation you may have met in previous courses on topology.

Proof of Proposition 7.5. This is a direct corollary of Lemma 7.25, together with Definition 7.26 □

In fact, the dualizing complex allows us to find $f^!$ for any submersion. We skip the proof of the next theorem for reasons of time.

Theorem 7.27. *Let $f : X \to Y$ be a submersion of manifolds. Then $f^! G \simeq \omega_{X/Y} \otimes f^{-1}(G)$ for $G \in D^+(\underline{R}_Y)$.*

Proof. You construct the natural map in Exercise 13.3. See [KS90, Proposition III.3.3.2] for the proof it is a quasi-isomorphism. □

7.5. Immersions and cohomology with support

Let $Z \subset X$ be closed and $j : U \to X$ the complement.

For a sheaf $\mathcal{F}$ on X we define the subsheaf with support in Z by $\Gamma_Z\mathcal{F} := \ker(\mathcal{F} \to j_*j^{-1}\mathcal{F})$, so that $\Gamma_Z(U, \mathcal{F})$ consists of sections of $\mathcal{F}$ on U whose support is contained in Z.

We consider this globally and derived global sections define cohomology with support: $H^i_Z(X, \mathcal{F}) := R^i\Gamma_Z(X, \mathcal{F})$

Lemma 7.28. *There are adjunctions* $i_* \dashv i^{-1}\Gamma_Z : \mathsf{Sh}(Z) \rightleftarrows \mathsf{Sh}(X)$ *and* $i_* \dashv i^{-1}R\Gamma_Z : D(Z) \rightleftarrows D(X)$.

Moreover there is an exact sequence

$$i_!i^!\mathcal{F} \to \mathcal{F} \to Rj_*j^{-1}\mathcal{F}$$

in $D(X)$.

Proof. We first note that for any sheaf $\mathcal{A}$ on X with support contained in Z and any other sheaf $\mathcal{F}$ on X we have $\mathrm{Hom}(\mathcal{A}, \mathcal{F}) = \mathrm{Hom}(\mathcal{A}, \Gamma_Z\mathcal{F})$ as $\mathrm{Hom}(\mathcal{A}, j_*j^{-1}\mathcal{F}) = \mathrm{Hom}(j^{-1}\mathcal{A}, j^{-1}\mathcal{F}) = 0$. Thus we compute

$$\begin{aligned}\mathrm{Hom}(i_*\mathcal{F}, \mathcal{G}) &\cong \mathrm{Hom}(i_*\mathcal{F}, \Gamma_Z\mathcal{G})\\ &\cong \mathrm{Hom}(i^{-1}i_*\mathcal{F}, i^{-1}\Gamma_Z\mathcal{G})\\ &\cong \mathrm{Hom}(\mathcal{F}, i^{-1}\Gamma_Z\mathcal{G})\end{aligned}$$

where we used the equivalence $i_{-1} \dashv i_*$ between sheaves on Z and sheaves on X supported on Z, see Exercise 4.2.

As a right adjoint $i^{-1}\Gamma_Z$ is left exact and may be derived. Since i^{-1} is exact the derived functor is given by $i^{-1}R\Gamma_Z$ (i^{-1} which right adjoint to i_* (which is also exact), this is the second statement.

As $i_* = i_!$ for closed immersions we have $i^{-1}R\Gamma_Z \simeq i^!$ by uniqueness of adjoints. As i^{-1} and i_* are inverse to each other on $D(Z)$ it follows that $R\Gamma_Z \simeq i_!i^!$.

Thus to deduce the exact sequence we consider the short exact sequence $0 \to \Gamma_Z I \to I \to j_*j^{-1}I$ from the definition of Γ_Z for an injective complex I on X. As injectives are flabby the map $I \to j_*j^{-1}I$ is now a surjection and we obtain a short exact sequence. Replacing any complex injectively the desired exact sequence in the derived category follows. □

Example 7.29. Let $i : * \to \mathbb{R}^n$ be the inclusion of a point. With the short exact sequence from Lemma 7.28 we have a long exact sequence on cohomology

$$\cdots \to H^i(i^!\underline{R}) \to H^i(\mathbb{R}^n, R) \to H^i(\mathbb{R}^n \setminus \{*\}, R) \to H^{i+1}(j^!\underline{R}) \to \cdots$$

(using i_*, j_* are exact).

Thus $i^!\underline{R}$ is given by a copy of R concentrated in degree n.

In some cases we do not even need this computation.

Lemma 7.30. *Let $i : X \to Y$ be a closed inclusion of orientable manifolds with $m = \dim X$ and $n = \dim Y$. Then $\omega_{X/Y} := i^! \underline{R}_Y \simeq \underline{R}_X[m-n]$.*

Proof. We know $i^! p_Y^! R \simeq p_X^! R \simeq \underline{R}_X[m]$. Together with $p_Y^! R \simeq \underline{R}_Y[n]$ the result follows. □

We can put together our previous results to get a more concrete description of $f^!$ for manifolds:

Theorem 7.31. *Let $f : X \to Y$ a map of manifolds. Denote by $i : X \to X \times Y$ the inclusion $x \mapsto (x, f(x))$ and by $p : X \times Y \to Y$ the projection. Then*

$$f^!(-) \simeq i^{-1} R\Gamma_X(p^{-1}(-) \otimes \omega_{X \times Y/Y})$$

.

Note that here $\omega_{X \times Y/Y}$ is quasi-isomorphic to a shifted sheaf.

Proof. This is Theorem 7.27 and Lemma 7.28 together with the fact $(p \circ i)^! \simeq i^! \circ p^!$ □

8. Constructible and perverse sheaves

We conclude with a rough outlook of where the story might go next. We fix $R = \mathbb{Z}$ just for convenience of notation. TA reference for this section is [Dim04, Chapter 4], or rather the references therein.

We have been working with the category of all sheaves on X, but this is really a rather large and unwieldy category. We have also met smaller categories like the category of locally constant sheaf, but this is not closed under the six functors we have constructed, so that is a disadvantage.

A compromise is to look at sheaves which are locally constant on some stratification of X into locally closed subsets. This works in different settings, we will use the following:

Definition 8.1. Let X be a complex analytic space (i.e. a subset of a complex manifold cut out by holomorphic equations, possibly singular).

An *admissible stratification* of X is a locally finite partition $X = \amalg_i S_i$ into non-empty, connected, locally closed subsets which we call *strata* sucht that

1. each boundary $\partial S_i = \bar{S}_i \setminus S_i$ is union of strata,
2. all $\bar{S}_i$ and ∂S_i are closed complex analytic subspaces.

One often considers further restrictions on stratifications to exclude pathological behaviour, for example we might ask that all S_i are smooth. Such good properties can always be achieved by refining the stratification, and this definition is good enough for our purposes.

We will now restrict our attention to analytic spaces X.

Definition 8.2. A sheaf $\mathcal{F}$ on X is *constructible* if X has a stratification that the restrictions $\mathcal{F}|_{S_i}$ are locally constant with fibers of finite rank.

A complex $F \in D^b(X)$ is *constructible* if all the cohomology sheaves $\mathcal{H}^i F$ are constructible. The derived category of all bounded constructible complexes on X is denoted by $D^b_c(X)$.

The finiteness of the stalks is sometimes relaxed to speak of weakly constructible sheaves.

It is not in general true (although it is true if X is algebraic) that the derived category of the category of constructible sheaves agrees with the subcategory of the derived category consisting of constructibel complexes. We always mean the latter.

We next state the difficult theorem that constructible complexes are (essentially) closed under the six operations $Rf_*, Rf_!, f^{-1}, f^!, \otimes^L, R\mathcal{H}om$.

Theorem 8.3. *Let $f : X \to Y$ be a morphism of analytic spaces. Then*

1. *$f^!$ and f^{-1} send $D^b_c(Y)$ to $D^b_c(X)$,*

2. *for $F, G \in D^b_c(X)$ we have $F \otimes^L G$ and $R\mathcal{H}om(F, G)$ in $D^b_c(X)$,*

3. *given $F \in D^b_c(X)$ if the restriction of f to $\mathsf{supp}(F)$ is proper then Rf_*F and $Rf_!G$ are constructible.*

It is not surprising we need some restriction on f for the pushforward to be well-defined. Consider for example the case $\iota : \mathbb{C}^* \to \mathbb{C}$ with the sheaf $\mathcal{F} = \oplus_n \mathbb{C}^n_{1/n}$ whose proper pushforward will be poorly behaved at the origin. We can drop the condition if we work instead in an algebraic setting where X and Y are algebraic varieties, f is an algebraic map and we consider stratificaitons such that all closures and boundaries of strata are closed algebraic subsets.

Another convenient construction is the *dualization functor* $D : D^b(X)^{\mathrm{op}} \to D^b(X)$ given by $F \mapsto R\mathcal{H}om(F, \omega_X)$ (whenever ω_X is bounded).

Lemma 8.4. *For a bounded constructible complex F the complex DF is also bounded constructible and moreover we have $DDF \simeq F$.*

Lemma 8.5. *For a constructible complex F we have*

1. $Df_!F \simeq f_*DF$
2. $Df_*F \simeq f_!DF$
3. $Df^!G \simeq f^{-1}DG$
4. $Df^{-1}G \simeq f^!DG$

We prove this lemma because the result is striking but the proof is very easy using our previous results:

Proof. The first statement follows directly from Verdier duality, Corollary 7.22 gives

$$R\mathcal{H}om(f_!F, \omega_Y) \simeq f_*R\mathcal{H}om(F, \omega_X) \simeq f_*DF.$$

Then the second statement follows from Lemma 8.4 and the first, third and fourth statement follow from invertibility of D together with uniqueness of adjoints. □

There is a subtlety we have not properly considered so far. We have defined the derived category $D(\mathcal{A})$ of an abelian category $\mathcal{A}$. Any object in $D(\mathcal{A})$ has associated to it homology objects in $\mathcal{A}$. We say $\mathcal{A}$ is a *heart* of $D(\mathcal{A})$ (I'm skipping the precise list of axioms).

Now two surprising things happen: A triangulated category can have more than one heart. And the derived category of a heart need not be the original derived category.

Let us turn to the second phenomen first. We have seen in example sheets, that locally constant sheaves are equivalent to representations of the fundamental group. But in the derived category of locally constant sheaves on X, denoted $LC(X)$, we have $\mathrm{Ext}^i_{LC(X)}(\underline{\mathbb{Z}}, \underline{\mathbb{Z}}) \cong H^i(X, \mathbb{Z})$, which depends on more than the fundamental group! The derived category of the heart of $LC(X)$ is $LC(K(\pi_1(X), 1)) \neq LC(X)$!

For the first phenomenon we can consider another way in which constructible sheaves appear. Consider a partial differential equation on a manifold X, which may be encoded in a certain sheaf $\mathcal{P}$ called a *(left) $\mathcal{D}$-module*, i.e. a sheaf of modules over the sheaf of (noncommutative) rings of differential operators generated by holomorphic functions and vector fields. $\mathcal{D}$ on the line $\mathbb{C}$ is thus generated holomorphic functions in z and ∂ with the rule $[\partial, z] = 1$. A module over $\mathcal{D}$ module could have the form $\mathcal{D}/\mathcal{D}.P$ on $\mathbb{C}$ where $P = \sum a_i(x)\partial^i$ is a differential operator of degree n. Then we think of the sheaf as locally generated by solutions of the equation $Pu = 0$. The $\mathcal{D}$-module has an underlying $\mathcal{O}$-module, which here would be $\mathcal{O}^n$ on the subset where $a_n(x)$ does not vanish.

There is a natural way of taking the local solutions of the PDE by taking a derived hom space $R\mathcal{H}om_{\mathcal{D}}(\mathcal{P}, \mathcal{O})$, where $\mathcal{O} = \mathcal{D}/(\partial)$ is a $\mathcal{D}$-module. For example

$$R\mathcal{H}om_{\mathcal{D}}(\mathcal{D}/\partial, \mathcal{O}) \simeq \mathcal{H}om_{\mathcal{D}}(\mathcal{D} \xrightarrow{\partial} \mathcal{D}, \mathcal{O}) \simeq \mathcal{O} \xrightarrow{\partial} \mathcal{O} \simeq \underline{\mathbb{C}}[1]$$

is a shifted locally constant sheaf corresponding to the constant functions satisfying $\partial f = 0$ while

$$R\mathcal{H}om_{\mathcal{D}}(\mathcal{D}/z, \mathcal{O}) \simeq \mathcal{H}om_{\mathcal{D}}(\mathcal{D} \xrightarrow{z} \mathcal{D}, \mathcal{O}) \simeq \mathcal{O} \xrightarrow{z} \mathcal{O} \simeq \mathbb{C}_0$$

is a skyscraper that might remind you of the Dirac δ-function (or rather distribution).

One can prove the solution complex is a bounded complex of constructible sheaves for any bounded complex of $\mathcal{D}$-modules. For a nice $\mathcal{D}$-module the solution complex need not be a sheaf, but it satisfies a number of strict conditions (the restrictions respectively restrictions with compact support to the strata live in certain degrees) making it a *perverse sheaf*.

Theorem 8.6. *The perverse sheaves on a complex manifold X are an abelian category forming a heart of the derived category of constructible sheaves. It is equivalent to the abelian category of (regular holonomic) $\mathcal{D}$-modules on X.*

Here regular means that the $\mathcal{D}$-module does not have bad singularities as we approach the boundary of X and holonomic means the $\mathcal{D}$-module is not too large (for example $\mathcal{D}$ itself is not holonomic, but $\mathcal{O}$ and $\mathcal{D}/(z)$ are.

This is called the *Riemann-Hilbert correspondence*, it generalises the classical statement that on a manifold local systems correpsond to finite-dimensional vector bundles equipped with a flat connection.

Perverse sheaves have many nice properties that the abelian category of constructible sheaves does not. For example the dualizing complex on a manifold is not a sheaf (because of the shift) but it is a perverse sheaf.

Another natural appearance of perverse sheaves is when studying spaces, say algebraic varieties via singular fibrations. If there is a fibration $F \to X \to Y$ it is often possibly to deduce properties of X from those of F and Y quite easily. But often this is not possible. Instead we might find a singular fibrations where all but finitely many fibers are the same.

Then to study properties of X from those of the base Y and the generic fiber F and the special fiber F_0 we need to see how F and F_0 are related. Two functors called the *vanishing cycles functor* and *nearby cycles functor* are used to measure the difference between F and F_0 and they take values in perverse sheaves. The stalk of the sheaf of vanishing cycles corresponds to the cohomology of the *Milnor fiber*, a geometric construction that was used earlier to study the nature of the singular fibers.

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* $\mathcal{C}$ consists of the following data:

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

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 $f \in \mathrm{Hom}_{\mathcal{C}}(X, Y)$.

Given f in $\mathrm{Hom}_{\mathcal{C}}(X, Y)$ 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 $\mathcal{C}$ with an object C there is a category whose objects are arrows $f : C \to D$ in $\mathcal{C}$, 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 $\mathcal{C}$ there is an *opposite category* $\mathcal{C}^{op}$ with the same objects, $\mathrm{Hom}_{\mathcal{C}^{op}}(A, B) = \mathrm{Hom}_{\mathcal{C}}(B, A)$ and $f \circ_{\mathcal{C}^{op}} g := g \circ_{\mathcal{C}} f$. Thus we obtain the opposite category $\mathcal{C}^{op}$ from $\mathcal{C}$ by turning around all arrows.

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

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 $\mathcal{C}$ and $\mathcal{D}$ consists of the following data:

- a map that associates to any $X \in \mathrm{Ob}(\mathcal{C})$ an object $F(X) \in \mathrm{Ob}(\mathcal{D})$.
- for each pair of objects $X, Y \in \mathrm{Ob}(\mathcal{C})$ a map from $\mathrm{Hom}_{\mathcal{C}}(X, Y)$ to $\mathrm{Hom}_{\mathcal{D}}(F(X), F(Y))$ 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 $\mathcal{C}$ there is an identity functor $\mathbf{1}_{\mathcal{C}}$ that does nothing on objects and morphisms.
2. Let $\mathcal{C}$ and $\mathcal{D}$ be categories and D an object of $\mathcal{D}$. Then there is a constant functor $c_D : \mathcal{C} \to \mathcal{D}$ that sends every object of $\mathcal{C}$ to D and any morphism of $\mathcal{C}$ 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 $F, G : \mathcal{C} \to \mathcal{D}$ be two functors. A *natural transformation* α from F to G consists of maps $\alpha_C : FC \to GC$ for every $C \in \mathcal{C}$ such that for every map $f : C \to C'$ in $\mathcal{C}$ 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 $F : \mathcal{C} \to \mathcal{D}$ there is the identity natural transformation $\mathbf{1}_F$ defined by $(\mathbf{1}_F)_C = \mathbf{1}_{FC}$ for every $C \in \mathcal{C}$.

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

   We will consider a functor $F : I \to \mathcal{C}$ as a *diagram* in $\mathcal{C}$, 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 $\mathcal{C}$ determines a constant functor $c_C : I \to \mathcal{C}$ that sends any i to C and any $f : i \to j$ to $\mathbf{1}_C$.

   Then natural transformation from c to another functor $F : I \to \mathcal{C}$ 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.

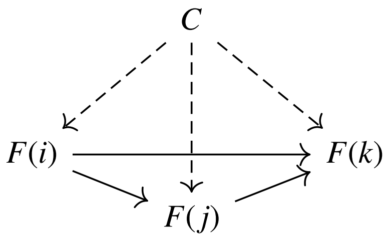

4. For every $n \geq 1$ the Hurewicz homomorphism $h_n : \pi_n(X, *) \to H_n(X, \mathbb{Z})$ from homotopy to homologoy of path connected spaces is a natural transformation. (To be precise it is a natural transformation from π_n to the composition of homology with the functor forgetting basepoints. If $n = 1$ 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* $\mathrm{Fun}(\mathcal{C}, \mathcal{D})$ 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 $F : \mathcal{C} \to \mathcal{D}$ and $G : \mathcal{D} \to \mathcal{C}$ such that $F \circ G$ is naturally isomorphic to $\mathbf{1}_\mathcal{D}$ and $G \circ F$ is naturally isomorphic to $\mathbf{1}_\mathcal{C}$.

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

Definition A.13. functor $F : \mathcal{C} \to \mathcal{D}$ is *full* if it induces surjections on all hom sets, i.e. every $g : FC \to FC'$ in $\mathcal{D}$ 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 $\mathcal{D}$ is isomorphic to some object FC in the image of F.

Then one can prove that $F : \mathcal{C} \to \mathcal{D}$ 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 $\mathcal{C}$ be any category. Then its *opposite category* $\mathcal{C}^{\mathrm{op}}$ is defined to have the same objects as $\mathcal{C}$ but $\mathrm{Hom}_{\mathcal{C}^{\mathrm{op}}}(C, D) := \mathrm{Hom}_{\mathcal{C}}(D, C)$ and $f \circ_{\mathcal{C}^{\mathrm{op}}} g := g \circ_{\mathcal{C}} f$.

In words $\mathcal{C}$ is obtained by turning around all the arrows in $\mathcal{C}$.
Clearly any functor $F : \mathcal{C} \to \mathcal{D}$ induces an opposite functor $F^{\mathrm{op}} : \mathcal{C}^{\mathrm{op}} \to \mathcal{D}^{\mathrm{op}}$.
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 $\mathcal{C}^{\mathrm{op}} \to \mathcal{D}$ a *contravariant functor* from $\mathcal{C} \to \mathcal{D}$.

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 $\mathcal{C}$ 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 $\mathcal{C}$ 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 $\mathcal{C}$.

Definition A.17. The *hom-functor*, denoted $h_C : \mathcal{C} \to \mathsf{Set}$, sends any object D to $\mathrm{Hom}_{\mathcal{C}}(C, D)$ and any morphism $f : D \to D'$ to the map $f_* : \mathrm{Hom}_{\mathcal{C}}(C, D)$ to $\mathrm{Hom}_{\mathcal{C}}(C, D')$ 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 $h^C : \mathcal{C}^{\mathrm{op}} \to \mathsf{Set}$ which is defined by $D \mapsto \mathrm{Hom}_{\mathcal{C}}(D, C)$ and $f \mapsto f^*$, where $f^*(g) = g \circ f$.
For another level of abstraction, $h_{(-)}$ defines a functor from $\mathcal{C}^{\mathrm{op}}$ to the category of functors $\mathrm{Fun}(\mathcal{C}, \mathsf{Set})$. 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 : $\mathsf{Group} \to \mathsf{Set}$ 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 $F : \mathcal{C} \to \mathsf{Set}$ 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 $\mathcal{C}$ any category. A *diagram of shape I* in $\mathcal{C}$ is just a functor $D : I \to \mathcal{C}$.

A *cone over D* is an object C in $\mathcal{C}$ 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 $F : I \to \mathcal{C}$ 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.

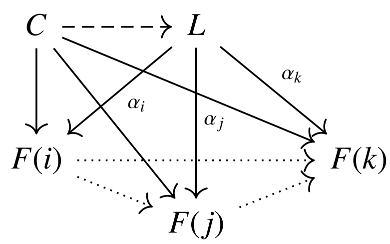

This universal property (like all universal property) ensures that if there are two limits L and L' there is a unique isomorphism between them: As L is a limit there is a unique map

$g : L' \to L$ 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 $\mathcal{C}$ we say that $\mathcal{C}$ *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 $F : I \to \mathcal{C}$ 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 $I \to \mathcal{C}$ is called the *terminal* object of $\mathcal{C}$, often written $*$. It has the property that for every $C \in \mathcal{C}$ 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 $F : I \to \mathcal{C}$ is called *equalizer*.

Definition A.26. Let I be the category with three objects $\bullet \to \bullet \leftarrow \bullet$. The limit of $F : I \to \mathcal{C}$ 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 $\mathcal{C}$, 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 $F : I \to \mathcal{C}$, denoted by $\operatorname{colim}_I F$, is an object D of $\mathcal{C}$ 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:

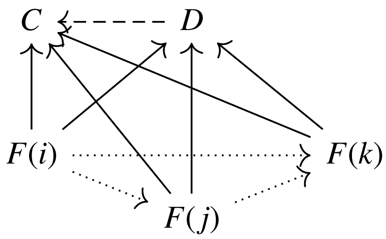

Remark A.29. To make the duality of limit and colimit more precise we can observe that (D, α) is a colimit of the diagram $F : I \to \mathcal{C}$ exactly if $(D, \alpha^{\mathrm{op}})$ is a limit of the diagram $F^{\mathrm{op}} : I^{\mathrm{op}} \to \mathcal{C}^{\mathrm{op}}$. Here $\alpha^{\mathrm{op}}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

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 $0 \leftarrow V \to W$ 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 $F : I \to \mathcal{C}$ and $G : J \to \mathcal{C}$ be diagrams. Then we have natural isomorphisms*

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

and

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

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 $I \to \mathcal{C}$ 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: $\mathbb{R} = \mathrm{colim}_{a \in \mathbb{N}}(-a, a)$ 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 $i \overset{f_1}{\longleftarrow} \cdots \overset{f_n}{\longrightarrow} i1$ making the natural diagram commute:

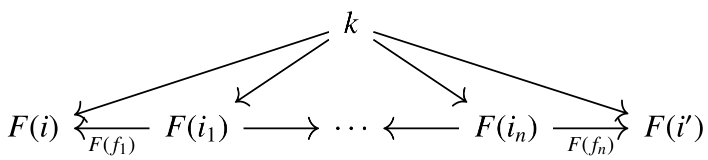

Note that the second condition is automatic if J is filtered.

Lemma A.35. *Let $F : I \to J$ be a final functor and $G : J \to \mathcal{C}$ 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 $\mathcal{C}$ *has all small limits* or *is complete* if every diagram $I \to \mathcal{C}$ has a limit. Similarly we say $\mathcal{C}$ has all small colimits or is cocomplete if every diagram $I \to \mathcal{C}$ 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 $\mathcal{C}$ has all equalizers if any equalizer diagram has a limit, and similarly for products (and other shapes of diagrams).

Lemma A.38. *A category $\mathcal{C}$ 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 $F : \mathcal{C} \to \mathcal{D}$ is *left adjoint* to $G : \mathcal{D} \to \mathcal{C}$, in symbols $F \dashv G$ if for all $C \in \mathcal{C}$ and $D \in \mathcal{D}$ there are natural isomorphisms

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

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

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

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

If $\mathcal{C}$ and $\mathcal{D}$ are locally small we can also phrase naturality as saying that the two functors $\operatorname{Hom}_{\mathcal{C}}(-, G(-))$ and $\operatorname{Hom}_{\mathcal{D}}(F(-), -)$ from $\mathcal{C}^{\text{op}} \times \mathcal{D}$ 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 $F : \mathcal{C} \to \mathcal{D}$ is left adjoint to G, then $F^{\mathrm{op}} : \mathcal{C}^{\mathrm{op}} \to \mathcal{D}^{\mathrm{op}}$ is right adjoint to G^{op}.

Let $F \dashv G : \mathcal{C} \rightleftarrows \mathcal{D}$ and $C \in \mathcal{C}$. 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 $\epsilon : \mathbf{1}_{\mathcal{C}} \Rightarrow GF$. This is called the *unit* of the adjunction.

Simlarly there is a natural transformation $\eta : FG \Rightarrow \mathbf{1}_{\mathcal{D}}$, 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 $C \in \mathcal{C}$ we have*

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

and for every $D \in \mathcal{D}$ 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 $G : I \to \mathcal{C}$ then $(FD, F\alpha)$ is a colimit for $F \circ G : I \to \mathcal{D}$.*

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

Remark A.44. Under some assumption on the categories $\mathcal{C}$ and $\mathcal{D}$ 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* $\mathrm{Fun}(I, \mathcal{C})$ *of I-shaped diagrams in $\mathcal{C}$. There is a diagonal functor $\Delta : \mathcal{C} \to$* $\mathrm{Fun}(I, \mathcal{C})$ *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.*

Bibliography

[BT82] R. Bott and L.W. Tu. *Differential Forms in Algebraic Topology*. Graduate Texts in Mathematics. Springer New York, 1982.

[Cis19] Denis-Charles Cisinski. *Higher Categories and Homotopical Algebra*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 2019. DOI: 10.1017/9781108588737.

[Dim04] Alexandru Dimca. *Sheaves in topology*. Universitext. Berlin: Springer, 2004.

[GM03] Gelfand, S.I. and Y. Manin. *Methods of Homological Algebra: Springer monographs in mathematics*. Springer Monographs in Mathematics. Springer, 2003.

[Hat02] Allen Hatcher. *Algebraic topology*. Cambridge: Cambridge University Press, 2002, pp. xii+544.

[Huy06] Daniel Huybrechts. *Fourier-Mukai transforms in algebraic geometry*. Oxford Mathematical Monographs. Oxford: Oxford Univ. Press, 2006.

[Ive86] Birger Iversen. *Cohomology of sheaves*. Universitext. Springer, 1986.

[KS90] M. Kashiwara and P. Schapira. *Sheaves on Manifolds:* Grundlehren der mathematischen Wissenschaften in Einzeldarstellungen. Springer, 1990.

[Pet21] Dan Petersen. *A remark on singular cohomology and sheaf cohomology*. 2021. arXiv: `2102.06927 [math.AT]`.

[Sel16] Y. Sella. “Comparison of sheaf cohomology and singular cohomology”. In: *ArXiv e-prints* (Feb. 2016). arXiv: `1602.06674 [math.AT]`.

[Wed16] T. Wedhorn. *Manifolds, Sheaves, and Cohomology*. Springer Studium Mathematik - Master. Springer Fachmedien Wiesbaden, 2016.

[Wei95] C.A. Weibel. *An introduction to homological algebra*. Cambridge University Press, 1995.

[Wil12] S. Willard. *General Topology*. Dover Books on Mathematics. Dover Publications, 2012.