A Motivated Introduction to Sheaves and Cohomology

Miika Rankaviita
Imperial College London

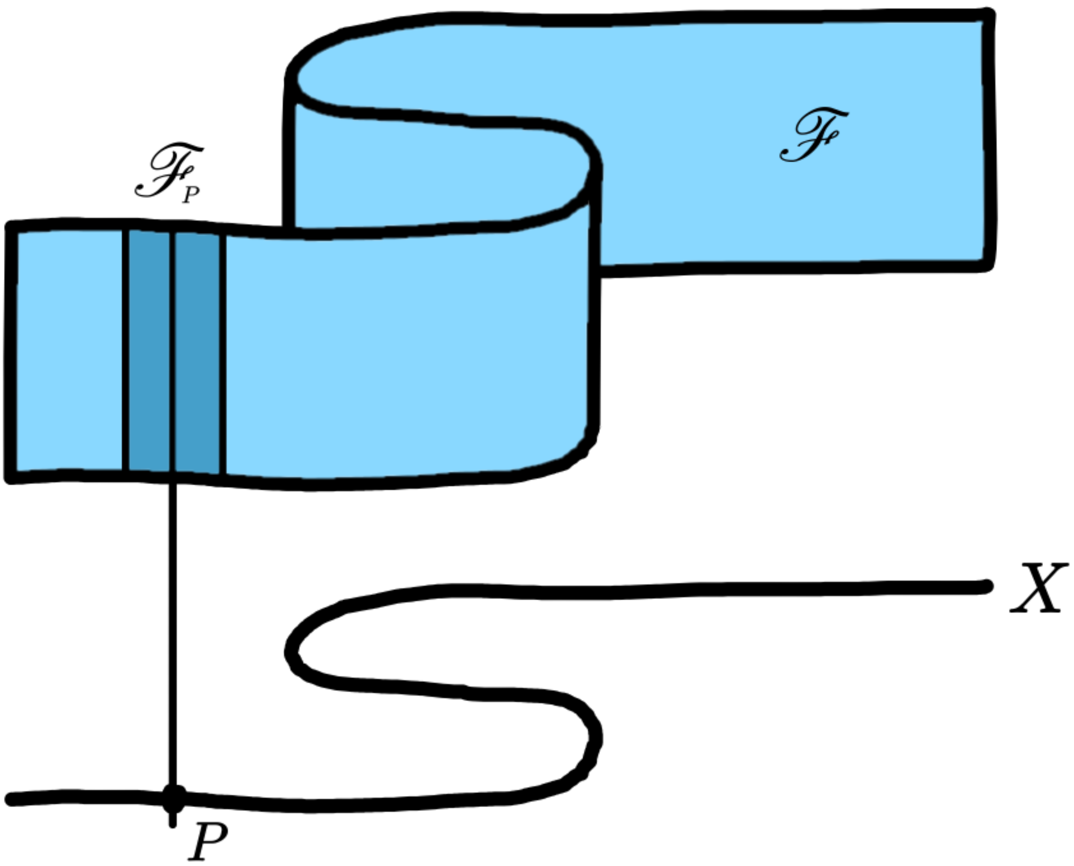

Contents

Acknowledgements

I want to thank my supervisor Professor Johannes Nicaise, whose guidance throughout the project was always invaluable. He patiently answered my numerous questions, while trusting my mathematical abilities. A huge thanks goes also to my friend Yourong Zang for giving good feedback on this article and encouraging me to publish it even when I was about to left it unfinished. Finally, I want to thank the Department of Mathematics at Imperial College London for funding this project.

Introduction

In this article, I am introducing sheaf theory and the cohomology of sheaves in the context of algebraic geometry with the motivation of proving the Riemann-Roch theorem, which is one of the fundamental results in algebraic geometry and which we shall apply to classify algebraic curves. This text is aimed at undergraduate students with basic knowledge of varieties.

Sheaves are formed by attaching algebraic data to *local* patches of a space. We usually associate data about rational functions, and the Riemann-Roch theorem will give *global* information about certain types of functions on a space. I will use the theorem to show the existence of a globally defined function on a curve X of genus 0, which defines an isomorphism with the projective line. Thus, we get a complete classification of curves of genus 0. One can use the Riemann-Roch theorem to classify curves of higher genera as well.

The main ingredient in the proof of the Riemann-Roch theorem is Serre Duality, which is a more general result. The proof will again be cohomological and is usually rather abstract. However, I will approach the proof by giving concrete interpretations of the cohomology groups by following Serre's exposition [Ser12].

The reader is assumed to know basic algebraic geometry. An excellent introduction to the topic is [Rei88], and a good set of lecture notes which goes beyond the basics is [Gat02]. I will use the language of discrete valuation rings, which is not explained in this paper. See [Ful08] for an introduction to this topic. Having familiarity with exact sequences is desirable. Throughout the paper, I will also try to point out category theoretical contexts of the concepts discussed. However, it is not necessary to know any category theory to understand this paper, and I will always place a warning for people who are not fond of abstract nonsense.

1 Sheaves

The central objects of study in this paper are *sheaves*. At first glance sheaves might seem complicated and abstract, but in fact they are inherently geometric objects. Hence, I will not start with the definition of sheaves, but instead I begin by discussing *fibre bundles*, which are more concrete objects that appear throughout geometry. This discussion should give a good motivation for the definition of sheaves and also help build geometric intuition right away.

1.1 Motivation: fibre bundles

The idea of a fibre bundle is to attach some space, called a *fibre* F, to every point of some topological space B, which will be called the *base space* [Hat01]. When we attach the fibres to the points of B, we should get a topological space E called the *total space*. More formally, a fibre bundle consists of two spaces E and B and a continuous surjection (henceforth *a projection*) $\pi : E \to B$, where the fibres $\pi^{-1}(b)$ of the projection are homeomorphic to some fixed fibre space F. We require in addition that one can find an open neighbourhood U for every point of B and a homeomorphism $h : \pi^{-1}(U) \to U \times F$ such that the following diagram commutes.

$$\begin{array}{ccc} \pi^{-1}(U) & \xleftrightarrow{\;h\;} & U \times F \\ & {\scriptstyle \pi}\searrow \quad \swarrow{\scriptstyle \mathrm{proj}_1} & \\ & U & \end{array}$$

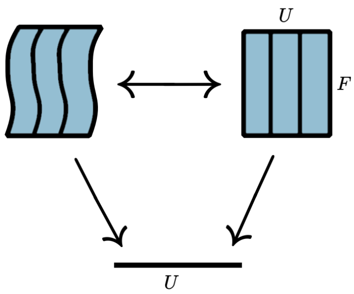

Thus, when we attach the fibres to points of B, they should be gathered so that the total space looks locally like a product space.

Let us look at some examples. Suppose the base space is the circle S^1 and the fibre space is the interval $(-1, 1)$, which is topologically speaking a line. An example of such a *line bundle* is the cylinder. The cylinder is a *trivial bundle*, because the total space can be written as a product: $S^1 \times (-1, 1)$. An example of a *non-trivial line bundle* over S^1 is the *Möbius strip*.

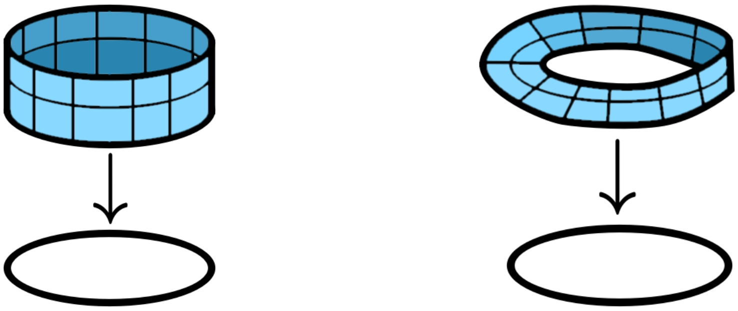

Figure 1: The two line bundles over S^1.

Similarly, we can consider *circle bundles* over S^1. A trivial circle bundle over S^1 is the torus and a non-trivial example is given by the *Klein bottle*. An important example showing up in differential and algebraic geometry is the concept of the *tangent bundle*, where the fibres are the tangent spaces at the points of the base space. Consider for example the sphere S^2, which is a 2-dimensional manifold. Then the fibres of the tangent bundle of S^2 are the planes tangent to the sphere.

The things we are usually the most interested in are the *sections* of a given fibre bundle. A section s of a fibre bundle $\pi : E \to B$ is a continuous map $s : B \to E$ such that $\forall b \in B,\ \pi(s(b)) = b$. This condition ensures that s maps every point of the base space to the corresponding fibre. A useful analogy is to think of a river, which can be represented by a curve C on the Euclidean plane. One could measure the temperature of the river at different points and get a temperature reading in $\mathbb{R}^+$ (assuming the measurement is taken in Kelvins). If one were to collect the temperature readings across all points of C, one would get a section of the $\mathbb{R}^+$-bundle over the river C. Thus, if fibre bundles are a way of attaching data to points of a space, then sections of a fibre bundle are continuous collections of data over the base space. Other interesting classes of examples of sections are vector fields and differential forms, which are sections of the tangent bundle and the cotangent bundle respectively. These are central objects of study in geometry.

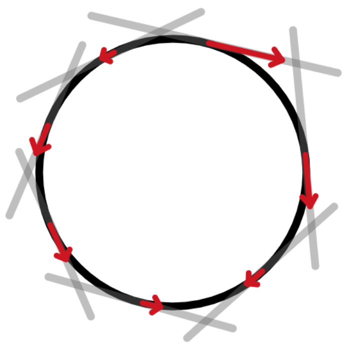

Figure 2: A section of the tangent bundle on S^1.

We often like to consider sections $s|_U : U \to E$ on some open set $U \subseteq B$. These have the following property. Given two sections $s_1 : U_1 \to E$ and $s_2 : U_2 \to E$, we can glue the two sections provided the sections agree on $U_1 \cap U_2$ to get a section $s : U_1 \cup U_2 \to E$. Then, one can piece together larger sections from smaller sections. This ability to define sections locally is very useful.

1.2 Defining sheaves

To summarise, fibre bundles are a way of attaching geometric data to points of a given space, and one can consider sections over the fibre bundle, which can be glued from small patches. We wish to extend the notion of a fibre bundle, because the requirement that the fibres must be topological spaces which form another topological space when bundled together is more strict than we want.

The notion of a fibre bundle is extended by the notion of a sheaf. Since taking sections is really the interesting part of fibre bundles, the definition of a sheaf will not be built with fibres and projection maps but with sections. For the rest of this section, I use [Gat02] as my main source for sheaf theoretic results. Next I will state the formal definition of a *pre-sheaf*. Pay attention to the properties that are motivated by fibre bundles.

Definition 1.1. A pre-sheaf $\mathscr{F}$ of sets on a topological space X associates to each open set $U \subseteq X$ a set $\mathscr{F}(U)$, which is called the set of sections over U. For every pair of nested open sets $V \subseteq U$, one defines a function $\mathrm{res}^U_V : \mathscr{F}(U) \to \mathscr{F}(V)$, called the restriction function, such that $\mathrm{res}^U_U = \mathrm{id}_U$ and for every triple $W \subseteq V \subseteq U$, we have $\mathrm{res}^V_W \circ \mathrm{res}^U_V = \mathrm{res}^U_W$. For $s \in \mathscr{F}(U)$, one usually writes $s|_V$ for $\mathrm{res}^U_V(s)$. Also, the set $\mathscr{F}(X)$ of global sections is denoted by $\Gamma(\mathscr{F})$.

One can also define pre-sheaves of rings, for example, where the sets $\mathscr{F}(U)$ are rings and the restriction functions are ring homomorphisms.

Note that pre-sheaves generalise the concept of taking sections on a topological space, but this definition doesn't capture the important property that sections should be able to be glued if they agree on the intersection. Thus, sheaves are defined in the following way.

Definition 1.2. A pre-sheaf $\mathscr{F}$ is a sheaf if the following property holds. Suppose $U \subseteq X$ is an open set with an open cover $(U_i)_{i\in I}$. If $s_i \in \mathscr{F}(U_i)$ are sections such that $s_i|_{U_i \cap U_j} = s_j|_{U_i\cap U_j}$ for all pairs i and j, then there is a unique section $s \in \mathscr{F}(U)$ such that $\forall i \in I,\ s|_{U_i} = s_i$.

⚠ Abstract nonsense ahead

Category theoretically speaking, a pre-sheaf $\mathscr{F}$ of sets is simply a functor $\mathscr{F} : \mathrm{Op}(X)^{\mathrm{op}} \to \mathbf{Set}$, where $\mathrm{Op}(X)$ is the posetal category of open sets of X. Moreover, a pre-sheaf of rings is a functor $\mathscr{F} : \mathrm{Op}(X)^{\mathrm{op}} \to \mathbf{Ring}$, and in general, a $\mathscr{C}$-valued pre-sheaf is a functor $\mathscr{F} : \mathrm{Op}(X)^{\mathrm{op}} \to \mathscr{C}$, where we define $\mathscr{F}(\emptyset)$ to be the final object of $\mathscr{C}$ [Vak17]. If $\mathscr{C}$ has limits, we can also express the sheaf axiom category theoretically [MLM92]. Namely, $\mathscr{F}$ is a sheaf if for every open set U of X and an open cover U_i of U the following diagram is an equaliser diagram.

$$\mathscr{F}(U) \longrightarrow \prod_i \mathscr{F}(U_i) \rightrightarrows \prod_{i,j} \mathscr{F}(U_i \cap U_j) , \tag{1}$$

where the first map is the product of the restrictions $\mathrm{res}^U_{U_i}$ and the pair of maps are products of the restrictions $\mathrm{res}^{U_i}_{U_i\cap U_j}$ and $\mathrm{res}^{U_j}_{U_i\cap U_j}$ respectively.

Now would be a good time to list some important examples of sheaves.

1. Any fibre bundle $\pi\colon E \to B$ defines a sheaf by
$$\mathscr{E}(U) = \{\, s : U \to E \mid s \text{ a section of } \pi \,\}.$$
2. Every affine variety X carries the structure sheaf $\mathscr{O}_X$ defined by
$$\mathscr{O}_X(U) = \{\, f \in k(X) \mid f \text{ regular on } U \,\}.$$
3. Let A be a group and $P \in X$ a point. The skyscraper sheaf A_P is defined by
$$A_P(U) = \begin{cases} A, & P \in U \\ 0, & P \notin U \end{cases}$$
4. Suppose X is any topological space and $V \subseteq X$ is an open subset. If $\mathscr{F}$ is a sheaf on X, then we can restrict it to a sheaf $\mathscr{F}|_V$ on V by defining
$$\mathscr{F}|_V(U) = \mathscr{F}(U \cap V).$$
5. If $f\colon X \to Y$ is a continuous map of topological spaces and $\mathscr{F}$ is a sheaf on X, then we can construct the so-called pushforward sheaf $f_*\mathscr{F}$ on Y by defining
$$f_*\mathscr{F}(U) = \mathscr{F}(f^{-1}(U)).$$

⚠ Abstract nonsense ahead

Let X be a topological space and $\iota\colon V \hookrightarrow X$ an open subset. The assignments
$$\iota_*\colon \mathbf{Sh}(V) \to \mathbf{Sh}(X) \quad \text{and} \quad (-)|_V\colon \mathbf{Sh}(X) \to \mathbf{Sh}(V) : \mathscr{F} \mapsto \mathscr{F}|_V$$
are functorial and define an adjoint pair $\iota_* \dashv (-)_V$.
More generally, one can define the *pullback* $f^*\mathscr{F}$ of a sheaf $\mathscr{F}$ for any continuous map $f\colon X \to Y$, which is the left adjoint to the pushforward f_*.

1.3 Stalks and sheafification

Now I will introduce two important constructions: stalks, which give a zoomed-in picture of a sheaf, and sheafification, which turns pre-sheaves into sheaves.

Suppose $\mathscr{F}$ is a pre-sheaf on some space X. We would like to understand, what the structure of $\mathscr{F}$ is at some point P by associating to the point a set $\mathscr{F}_P$ that describes this structure. Since the only information specifying a pre-sheaf are the sections, we could define $\mathscr{F}_P$ to be the set of sections over open sets containing P. However, this is most often an unimaginably large a set. Thus, we say that $\mathscr{F}_P$ is the set of equivalence classes of such sections, where two sections are equivalent if they agree on some sufficiently small open neighbourhood of P. More formally:

Definition 1.3. Given a pre-sheaf $\mathscr{F}$ on X and a point $P \in X$, the stalk $\mathscr{F}_P$ at P is defined as the set

$$\mathscr{F}_P = \{\, (s, U) \mid U \subseteq X \text{ open}, P \in U, s \in \mathscr{F}(U) \,\} / \sim,$$

where two pairs (s_1, U_1) and (s_2, U_2) are equivalent if there is an open set $V \subseteq U_1, U_2$ containing P such that

$$s_1|_V = s_2|_V.$$

The equivalence class of a pair (s, U) is called the *germ* of s at P. Given a section $s \in \mathscr{F}(U)$, its germ at P is usually denoted by s_P.

⚠️ Abstract nonsense ahead

The stalks can be written as the direct limit

$$\mathscr{F}_P = \varinjlim_{U \ni P} \mathscr{F}(U).$$

Next I will do an example computation with stalks, which will be useful later.

Proposition 1.4. *Let X be a variety. The stalk $\mathscr{O}_{X,P}$ of $\mathscr{O}_X$ at some point $P \in X$ is given by*

$$\mathscr{O}_{X,P} = \left\{ \frac{f}{g} \in k(X) \mid f, g \in k[X], g(P) \neq 0 \right\}.$$

Proof. I will prove the two inclusions separately.

$\subseteq$) Let $[(\varphi, U)] \in \mathscr{O}_{X,P}$. Since $P \in U$ and $\varphi \in \mathscr{O}(U)$, φ must be regular at P. Thus, φ has a representation $\frac{f}{g}$ as a rational function around P, where $g(P) \neq 0$. Moreover, if $[(\varphi_1, U_1)] = [(\varphi_2, U_2)]$, then there is an open set $V \subseteq U_1, U_2$ such that $\varphi_1|_V = \varphi_2|_V$. Therefore, φ_1 and φ_2 can be represented by the same rational function around P.

$\supseteq$) Now suppose $\frac{f}{g} \in k(X)$ with $g(P) \neq 0$. Then, there is an open neighbourhood U of P where g is non-zero. Clearly, $(\frac{f}{g}, U)$ defines a germ in $\mathscr{O}_{X,P}$.

□

The stalks of a sheaf are analogous to the fibres of a fibre bundle. Just as a section of a fibre bundle takes values in the fibres, a section of a sheaf takes values in the stalks.

Note however that while a stalk at P contains all the data "above" P just as a fibre at P does, stalks are larger, because they also contain local information **around** P. One can find two sections with equal values at P, but which have different germs at P.

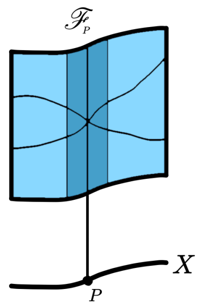

Next, we will study properties of germs that leads to the construction of the sheafification of a pre-sheaf. Fix a sheaf $\mathscr{F}$ on some space X and choose a section $s \in \mathscr{F}(U)$ on some open set $U \subseteq X$. Then, consider the totality of germs of s over U, namely the collection $(s_P)_{P \in U}$ with $s_P \in \mathscr{F}_P$. Since $\mathscr{F}$ is a sheaf, it turns out that one can recover the section s from the germs: Each germ s_P has a representative $(s|_{U_P}, U_P)$. The set U is covered by the U_P and since the $s|_{U_P}$ are restrictions of the common section s onto the sets U_P, one can use the sheaf axiom to glue the sections $s|_{U_P}$ together and obtain a section on U which must be equal to s by uniqueness. Generalising this observation, one can write the sections of a sheaf as a collection of germs:

Lemma 1.5. *Given a sheaf $\mathscr{F}$ on a space X, one can write*

$$\begin{aligned} \mathscr{F}(U) = \Big\{ (s_P)_{P \in U} \,\Big|\, & s_P \in \mathscr{F}_P \text{ such that} \\ & \forall P \in U\, \exists U_P \subseteq U \text{ an open neighbourhood of } P \\ & \exists r \in \mathscr{F}(U_P)\, \forall Q \in U_P,\ r_Q = s_Q \Big\} \end{aligned}$$

for every open set $U \subseteq X$

This statement may seem bewildering at first, but note that the important part in the above observation was that when we fixed a point P, there was an open set U_P where all the germs s_Q for $Q \in U_P$ agreed with some section $s|_{U_P}$ on U_P. We want to ensure that this is the case with the collection $(s_P)_{P \in X}$ of the proposition so that we can glue the germs in the same way as above. This statement will be essential for building theory, but it is also useful for manipulating sheaves in practise and I will routinely represent sections of sheaves as collections of germs. Therefore, make sure you understand the statement and the proof completely.

Proof. I will again break up the proof into cases.

$\subseteq \big)$ Suppose $s \in \mathscr{F}(U)$ is a section on U. Fix an arbitrary point $P \in U$ and choose a representative $(s|_{U_P}, U_P)$ of the germ s_P of s. If one takes $r = s|_{U_P} \in \mathscr{F}(U_P)$, then it follows immediately that

$$\forall Q \in U_P,\ r_Q = s_Q.$$

$\supseteq$) Now, suppose $(s_P)_{P\in U}$ is a collection of germs in the RHS set. Then, U is covered by the sets U_P associated to each germ s_P. Now, fix two points $P_1, P_2 \in U$, and consider the associated sections $r_1 \in \mathscr{F}(U_{P_1})$ and $r_2 \in \mathscr{F}(U_{P_2})$. It is clearly the case that $r_1|_{U_{P1}\cap U_{P2}} = r_2|_{U_{P1}\cap U_{P2}}$, because of the germs of r_1 and r_2 are related through the germs s_P on $U_{P_1} \cap U_{P_2}$. Therefore, we can take all the sections $r_P \in \mathscr{F}(U_P)$ associated to the germs s_P and glue them together to obtain a section s on U. The germs of this section are clearly the germs s_P.

□

Now, since one can take stalks of pre-sheaves too, one might try this construction on a pre-sheaf. If $\mathscr{F}'$ is a pre-sheaf on some space X, one can define another pre-sheaf $\mathscr{F}$ on X by

$$\begin{aligned}\mathscr{F}(U) = \Big\{ (s_P)_{P\in U} \,\Big|\, & s_P \in \mathscr{F}'_P \text{ such that} \\ & \forall P \in U\, \exists U_P \subseteq U \text{ an open neighbourhood of } P \\ & \exists r \in \mathscr{F}'(U_P)\, \forall Q \in U_P,\ r_Q = s_Q \Big\}.\end{aligned}$$

Now, I leave it as an exercise for the reader to check that $\mathscr{F}$ is in fact a sheaf. Thus, one can associate a sheaf to every pre-sheaf using this construction, which is called *sheafification*.

Remark. *Note that Lemma 1.5 is saying that sheafifying a sheaf does not do anything. One might also notice that this construction preserves the stalks of $\mathscr{F}'$, in other words, $\mathscr{F}'_P = \mathscr{F}_P$ for every $P \in X$. This is true, because every element (s, U) of $\mathscr{F}'_P$ can be associated to the element $(s_Q)_{Q\in U}$ of $\mathscr{F}(U)$, which in turn has a stalk at P in $\mathscr{F}_P$ and an element $((s_Q)_{Q\in U}, U)$ of $\mathscr{F}_P$ can be associated to $s_P \in \mathscr{F}'_P$.*

⚠️ Abstract nonsense ahead

The sheafification functor $\mathbf{PSh}(X) \to \mathbf{Sh}(X)$ from the category of pre-sheaves on X to the category of sheaves on X is the left adjoint to the forgetful functor $\mathbf{Sh}(X) \to \mathbf{PSh}(X)$ [Vak17]. In particular, there is a universal property defining the sheafification of a pre-sheaf.

A typical example illustrating sheafification is to consider the constant pre-sheaf and its sheafification. The constant pre-sheaf $\mathscr{F}'$ on a space X is defined as follows

$$\mathscr{F}'(U) = \{\, c : U \to \mathbb{R} \mid c \text{ a constant function } \}.$$

One can see that this is not a sheaf: take any two disjoint open sets $U_1, U_2 \subseteq X$ and two constant functions $c_1 : U_1 \to \mathbb{R}$ and $c_2 : U_2 \to \mathbb{R}$ with different values. Then, the two sections don't glue to form a constant function on $U_1 \cup U_2$ (all of this assumes that X is a space with at least one pair of disjoint open sets). The sheafification $\mathscr{F}$ of $\mathscr{F}'$ is given by sections $(c_P)_{P\in U}$. By the definition of sheafification, there is for every $P \in U$ an open neighbourhood $U_P \subseteq U$ of P and a constant function $d : U_P \to \mathbb{R}$ such that $\forall Q \in U_P,\ c_Q = d_Q$. Hence, the restriction of $(c_Q)_{Q\in U}$ to U_P can be identified with the constant function d. Therefore, the sheafification of the constant pre-sheaf consists of *locally constant functions*. In general, given a set A, the constant sheaf $\underline{A}$ is defined as the sheafification of the constant pre-sheaf with values in A.

1.4 Ringed spaces and algebraic varieties

Now that we have some feel for sheaves, I will briefly go over the construction of abstract algebraic varieties, which are built with the help of sheaves. Whenever I mention varieties, I refer to spaces that are defined as in this subsection. The precise details of the construction are not essential for the present article, and the reader should consult [Gat02] for a more detailed discussion.

In geometry, one is not only interested in the points of the space, but also in the functions on the space. Thus, I will extend the notion of a topological space by attaching the data of the functions on the space.

Definition 1.6. A ringed space is a pair $(X, \mathscr{O}_X)$ consisting of a topological space X and a sheaf of rings $\mathscr{O}_X$ called the structure sheaf of X.

Next, I will define morphisms between ringed spaces [Gat02].

Definition 1.7. Suppose $(X, \mathscr{O}_X)$ and $(Y, \mathscr{O}_Y)$ are two ringed spaces with structure sheaves consisting of k-valued functions for some field k. Then a morphism from $(X, \mathscr{O}_X)$ to $(Y, \mathscr{O}_Y)$ is a (set-theoretic) continuous map $f : X \to Y$ such that

$$\forall U \subseteq Y \text{ open},\ f^*\mathscr{O}_Y(U) \subseteq \mathscr{O}_X(f^{-1}(U)), \tag{2}$$

where $f^*g = g \circ f$ for $g \in \mathscr{O}_Y(U)$.

To understand the condition in (2), it helps to look at the following commutative diagram.

$$\begin{array}{ccc} X & \xrightarrow{\ f\ } & Y \\ & {\scriptstyle f^*g}\searrow \quad \swarrow{\scriptstyle g} & \\ & k & \end{array}$$

The condition requires that functions $g \in \mathscr{O}_Y(U)$ "pull back" to functions in $\mathscr{O}_X(f^{-1}(U))$ along the map $f : X \to Y$. As usual, one can define isomorphisms of ringed spaces as morphisms $f : (X, \mathscr{O}_X) \to (Y, \mathscr{O}_Y)$ such that there is a morphism $g : (Y, \mathscr{O}_Y) \to (X, \mathscr{O}_X)$ with $f \circ g = g \circ f = \text{id}$. Next I will define what it means for a ringed space to be an affine variety.

Definition 1.8 ([Gat02, Def. 2.3.15])**.** A ringed space $(X, \mathscr{O}_X)$ is an affine variety if

- X is an irreducible space,
- $\mathscr{O}_X$ is a sheaf of k-valued functions,
- There is an irreducible algebraic set $Y \subseteq \mathbb{A}^n$ such that $(X, \mathscr{O}_X) \cong (Y, \mathscr{O}_Y)$.

Note that such ringed spaces are not always zero sets of some polynomials; See for example [Gat02, Lemma 2.3.16].

Now that I have defined what it means for a ringed space to be an affine variety, I will form *varieties* by "gluing" together open affine varieties in the same way as manifolds are glued from open subsets of the Euclidean space.

Definition 1.9 ([Gat02, Def. 2.4.1])**.** A ringed space $(X, \mathscr{O}_X)$ is a pre-variety if

- X is an irreducible space,
- $\mathscr{O}_X$ is a sheaf of k-valued functions,
- X can be covered with finitely many open sets U_i such that each $(U_i, \mathscr{O}_X|_{U_i})$ is an affine variety.

The reason why we do not call these spaces varieties is because the definition allows the construction of some pathological spaces such as the *line with two origins*, see [Gat02]. To avoid such spaces, we add one more condition in the final definition of a variety.

Definition 1.10 ([Gat02, Def. 2.5.1])**.** A pre-variety X is a variety, if for every pre-variety Y and any pair of morphisms $f_1, f_2 : Y \to X$, the set

$$\{ P \in Y \mid f_1(P) = f_2(P) \}$$

is closed in Y.

The following lemma will be needed later.

Lemma 1.11 ([Vak17, Prop. 10.1.8])**.** *Suppose X is a variety and $U, V \subseteq X$ are two open affine subsets. Then, the intersection $U \cap V$ is an open affine subset of X.*

Proof. Since the intersection of two open sets is open by definition, we need to only check that the intersection is affine. Recall that the product of two affine varieties is an affine variety, and consider the projection maps $\pi_U : U \times V \to U$ and $\pi_V : U \times V \to V$. Also, denote by $\iota_U : U \to X$ and $\iota_V : V \to X$ the inclusion maps. Then, we can see that

$$U \cap V = \{ P \in U \times V \mid (\iota_U \circ \pi_U)(P) = (\iota_V \circ \pi_V)(P) \},$$

which is closed, since X is a variety. Therefore, the intersection is a closed subset of an affine variety. □

1.5 Algebraic constructions with sheaves

The aim of this subsection is to define algebraic constructions for sheaves so that we can study them using the tools of *homological algebra*. First I need to define morphisms between pre-sheaves.

Definition 1.12. Suppose $\mathscr{F}$ and $\mathscr{G}$ are pre-sheaves of sets or rings on some space X. Then, a morphism $\alpha : \mathscr{F} \to \mathscr{G}$ consists of set functions or ring homomorphisms $\alpha_U : \mathscr{F}(U) \to \mathscr{G}(U)$ which commute with restriction maps. That is, the following diagrams commute.

$$\begin{array}{ccc} \mathscr{F}(U) & \xrightarrow{\text{res}^U_V} & \mathscr{F}(V) \\ {\scriptstyle \alpha_U}\downarrow & & \downarrow{\scriptstyle \alpha_V} \\ \mathscr{G}(U) & \xrightarrow[\text{res}^U_V]{} & \mathscr{G}(V) \end{array}$$

Morphisms between sheaves are defined exactly the same way.

One might ask, why this should be the right definition. The reason is that this type of morphism **preserves the structure of the sheaf**. Compare this with group homomorphisms. Firstly, a group homomorhism $\varphi : G \to H$ associates to every element of G a corresponding element in H. In the same way a sheaf morphism α must associate to every section $s \in \mathscr{F}(U)$ a corresponding section $\alpha_U(s) \in \mathscr{G}(U)$ of $\mathscr{G}$. Secondly, if three elements of the group are related by the group structure so that $a \cdot b = c$, then the homomorphism should respect this structure so that $\varphi(a) \cdot \varphi(b) = \varphi(c)$. Similarly, the sheaf morphism α should respect the structure of the sheaves given by the restriction maps: if $s|_V = r$, then $\alpha(s)|_V = \alpha(r)$.

⚠️ Abstract nonsense ahead

As one can immediately see, morphisms of pre-sheaves and sheaves are simply natural transformations of contravariant functors.

In order to proceed to define algebraic constructions on sheaves, I must restrict to some class of sheaves with an algebraic structure. The class usually considered in algebraic geometry is the class of *sheaves of* $\mathscr{O}_X$*-modules* on a ringed space $(X, \mathscr{O}_X)$.

Definition 1.13. Given a ringed space $(X, \mathscr{O}_X)$, a sheaf $\mathscr{F}$ on X is a sheaf of $\mathscr{O}_X$-modules, if the sets $\mathscr{F}(U)$ are $\mathscr{O}_X(U)$-modules and the restriction maps satisfy the following condition. Suppose $V \subseteq U \subseteq X$ are open and $s, s_1, s_2 \in \mathscr{F}(U)$ are sections. Then,

$$(s_1 + s_2)|_V = s_1|_V + s_2|_V \quad \text{and} \quad (\lambda s)|_V = \lambda|_V \cdot s|_V$$

for all $\lambda \in \mathscr{O}_X(U)$.

I can now define the usual algebraic operations for sheaves of $\mathscr{O}_X$-modules.

Definition 1.14. Suppose $\mathscr{F}_1$ and $\mathscr{F}_2$ are sheaves of $\mathscr{O}_X$-modules and $f : \mathscr{F}_1 \to \mathscr{F}_2$ is a morphism of sheaves. Then the sheaves $\ker f$, $\operatorname{coker} f$,$\operatorname{im} f$, $\mathscr{F}_1 \oplus \mathscr{F}_2$, $\mathscr{F}_1 \otimes \mathscr{F}_2$, and $\mathscr{F}_1^\vee$ are defined as the sheafifications of the corresponding pre-sheaves given by the following:

1. $(\ker' f)(U) = \ker(f_U : \mathscr{F}_1(U) \to \mathscr{F}_2(U))$
2. $(\operatorname{coker}' f)(U) = \operatorname{coker}(f_U : \mathscr{F}_1(U) \to \mathscr{F}_2(U))$
3. $(\operatorname{im}' f)(U) = \operatorname{im}(f_U : \mathscr{F}_1(U) \to \mathscr{F}_2(U))$
4. $(\mathscr{F}_1 \oplus' \mathscr{F}_2)(U) = \mathscr{F}_1(U) \oplus \mathscr{F}_2(U)$
5. $(\mathscr{F}_1 \otimes' \mathscr{F}_2)(U) = \mathscr{F}_1(U) \otimes_{\mathscr{O}_X(U)} \mathscr{F}_2(U)$
6. $\mathscr{F}_1^{\vee\prime}(U) = \operatorname{Hom}_{\mathscr{O}_X(U)}(\mathscr{F}_1(U), \mathscr{O}_X(U))$

One can actually check that the pre-sheaves $\ker' f$ and $\mathscr{F}_1 \oplus' \mathscr{F}_2$ are in fact already sheaves, so sheafification doesn't change the definition. We also have the following usual definitions.

Definition 1.15.

1. A morphism $f : \mathscr{F}_1 \to \mathscr{F}_2$ of sheaves of $\mathscr{O}_X$-modules is injective if $\ker f = 0$ and surjective if $\operatorname{coker} f = 0$.
2. If f is injective, we write $\mathscr{F}_2/\mathscr{F}_1$ for the cokernel.
3. A sequence

$$\cdots \longrightarrow \mathscr{F}_{i-1} \longrightarrow \mathscr{F}_i \longrightarrow \mathscr{F}_{i+1} \longrightarrow \cdots$$

of sheaves of $\mathscr{O}_X$-modules is exact if

$$\ker(\mathscr{F}_i \to \mathscr{F}_{i+1}) = \operatorname{im}(\mathscr{F}_{i-1} \to \mathscr{F}_i), \ \forall i.$$

⚠️ Abstract nonsense ahead

One can show that the category $\mathscr{O}_X$-**Mod** of sheaves of $\mathscr{O}_X$-modules is an *abelian category*, which means that it is a perfect setting for homological algebra (see [Vak17]).
Furthermore, note that equaliser diagrams in abelian categories can be expressed in terms of exact sequences. Therefore, the sheaf axiom in (1) can be expressed as the exact sequence

$$0 \longrightarrow \mathscr{F}(U) \longrightarrow \bigoplus_i \mathscr{F}(U_i) \xrightarrow{\operatorname{res}^{U_j}_{U_i \cap U_j} - \operatorname{res}^{U_i}_{U_i \cap U_j}} \bigoplus_{i,j} \mathscr{F}(U_i \cap U_j) \tag{3}$$

in the category $\mathscr{O}_X$-**Mod** [MLM92].

Next I will show that the exactness of sequences of sheaves is a local property. To see this, one should first note that a morphism $f : \mathscr{F}_1 \to \mathscr{F}_2$ of sheaves of $\mathscr{O}_X$-modules induces an $\mathscr{O}_{X,P}$-module homomorphism

$$f_P : (\mathscr{F}_1)_P \to (\mathscr{F}_2)_P : [(s,U)] \mapsto [(f_U(s),U)]$$

on the stalks. Then one can show the following theorem.

Theorem 1.16 ([Gat21, Lemma 13.21]). *A sequence*

$$\cdots \xrightarrow{f_{i-2}} \mathscr{F}_{i-1} \xrightarrow{f_{i-1}} \mathscr{F}_i \xrightarrow{f_i} \mathscr{F}_{i+1} \xrightarrow{f_{i+1}} \cdots$$

of sheaves of $\mathscr{O}_X$-modules is exact if and only if the induced sequences

$$\cdots \xrightarrow{(f_{i-2})_P} (\mathscr{F}_{i-1})_P \xrightarrow{(f_{i-1})_P} (\mathscr{F}_i)_P \xrightarrow{(f_i)_P} (\mathscr{F}_{i+1})_P \xrightarrow{(f_{i+1})_P} \cdots$$

on the stalks are all exact.

Proof. I will prove the two directions of the equivalence separately.

$\Longrightarrow$)

The exactness of the first sequence implies that $\ker(f_i) = \operatorname{im}(f_{i-1})$. Thus $(\ker(f_i))_P = (\operatorname{im}(f_{i-1}))_P$, and so

$$\begin{aligned}
[(s,U)] \in \ker\left((f_i)_P\right) &\iff \exists V \subseteq U,\ s|_V \in \ker'(f_i)(V) \\
&\iff [(s|_V,V)] \in (\ker(f_i))_P \\
&\iff [(s|_V,V)] \in (\operatorname{im}(f_{i-1}))_P \\
&\iff \exists V \subseteq U,\ s|_V \in \operatorname{im}'(f_{i-1})(V) \\
&\iff [(s,U)] \in \operatorname{im}\left((f_{i-1})_P\right).
\end{aligned}$$

Therefore, $\ker\left((f_i)_P\right) = \operatorname{im}\left((f_{i-1})_P\right)$

$\Longleftarrow$)

Choose an arbitrary open set $U \subseteq X$. Then,

$$\begin{aligned}
s \in \ker(f_i)(U) &\iff \forall P \in U,\ s_P \in \ker\left((f_i)_P\right) \\
&\iff \forall P \in U,\ s_P \in \operatorname{im}\left((f_{i-1})_P\right) \\
&\iff s \in \operatorname{im}(f_{i-1})(U).
\end{aligned}$$

In the first and the last equivalence one uses Lemma 1.5 to break the section s into its germs and piece them back together. Thus, $\ker(f_i)(U) = \operatorname{im}(f_{i-1})(U)$ for all open sets $U \subseteq X$, which implies that the sheaves are equal.

□

One can prove the following results as an immediate corollary.

Corollary 1.16.1. *Suppose $f : \mathscr{F} \to \mathscr{G}$ is a morphism of sheaves of $\mathscr{O}_X$-modules, and let $f_P : \mathscr{F}_P \to \mathscr{G}_P$ be the induced homomorphism. Then,*

1. *f is an injection iff f_P is an injection,*
2. *f is a surjection iff f_P is a surjection,*
3. *f is an isomorphism iff f_P is an isomorphism.*

1.5.1 Quasi-coherent sheaves

In algebraic geometry, one further restricts attention to a certain class of sheaves of $\mathscr{O}_X$-modules called *quasi-coherent sheaves*. Indeed, most sheaves that algebraic geometers are interested in belong to this class. The definition is algebraic and makes more sense in the framework of scheme theory, but we can state it in elementary terms now. The reason we introduce quasi-coherent sheaves is that Lemma 1.20 below is what makes our definition of Čech cohomology work (Subsection 2.4).

Definition 1.17. Suppose X is an affine variety and M is a module over the coordinate ring $\mathscr{O}_X(X)$. Then, the sheaf $\widetilde{M}$ associated to the module M is defined as follows.

$$\begin{aligned}\widetilde{M}(U) = \Big\{ (\varphi_P)_{P\in U} \,\Big|\, & \varphi_P \in M_{\mathfrak{m}_P} \text{ such that} \\ & \forall P \in U \, \exists U_P \subseteq U \text{ an open neighbourhood of } P \\ & \exists m \in M \, \exists g \in \mathscr{O}_X(X) \, \forall Q \in U_P,\ \varphi_Q = \frac{m}{g} \Big\},\end{aligned}$$

where $M_{\mathfrak{m}_P}$ is the localisation of M at the maximal ideal $\mathfrak{m}_P$ of the point P.

Quasi-coherent sheaves are then constructed from such sheaves.

Definition 1.18. A sheaf $\mathscr{F}$ of $\mathscr{O}_X$-modules is quasi-coherent, if for every affine open set $U \subseteq X$, the sheaf $\mathscr{F}|_U$ is a sheaf associated to some module M over the coordinate ring $\mathscr{O}_X(U)$.

Remark. *To check quasi-coherence, it is enough to have an affine cover $(U_i)_{i\in I}$, where $\mathscr{F}|_{U_i}$ is of the form $\widetilde{M}$ for all $i \in I$, see [Gat02].*

Remark. *The constructions in Def. 1.14 preserve quasi-coherence. I refer the reader to [Gat02].*

Example 1.19. The structure sheaf $\mathscr{O}_X$ of a variety X is always quasi-coherent with $\mathscr{O}_X = \widetilde{M}$, where $M = \mathscr{O}_X(X)$.

Here are two basic properties satisfied by the construction.

Lemma 1.20. *If $\widetilde{M}$ is the sheaf on an affine variety X associated to the module M over $\mathscr{O}_X(X)$, then*

(a) $\forall P \in X, \left(\widetilde{M}\right)_P = M_{\mathfrak{m}_P}$, *and*

(b) $\Gamma\left(\widetilde{M}\right) = M$.

Proof. The proof of part (a) is basically the same as the proof of Prop. 1.4. Thus, I will only prove part (b) using ideas from the proof of [Gat21, Proposition 3.8]. Note that $M \subseteq \Gamma\left(\widetilde{M}\right)$ holds trivially, and thus I only need to show $\Gamma\left(\widetilde{M}\right) \subseteq M$. Begin by fixing an element $\varphi \in \Gamma\left(\widetilde{M}\right)$. At a point $P \in X$, the section is represented by $\varphi_P = m_P/g_P$, where $m_P \in M$ and $g_P \in \mathscr{O}_X(X)$. Furthermore, φ has this representation on an open neighbourhood U_P.

Note first that the representations of φ at any two points $P, Q \in X$ agree on the intersection $U_P \cap U_Q$. Since X is an irreducible topological space, its non-empty open subsets are dense. Therefore, the intersection $U_P \cap U_Q$ is non-empty, and we have that

$$g_Q m_P = g_P m_Q$$

for every pair of points $P, Q \in X$.

Now, since $\varphi_P \in M_{\mathfrak{m}_P}$, we may assume $g_P(P) \neq 0$ for every $P \in X$. In particular, $P \notin V(g_P)$. This implies that

$$\bigcap_{P \in X} V(g_P) = \varnothing.$$

By the Nullstellensatz, $\sqrt{(g_P \mid P \in X)} = (1)$, and therefore there are $k_P \in \mathscr{O}_X(X)$ such that

$$1 = \sum_{P \in I} k_P g_P,$$

where I is some finite subset of X. If we set

$$m = \sum_{P \in I} k_P m_P,$$

I claim that we can write $\varphi = m$ on the whole of X. Indeed, at a point Q, we have

$$\varphi_Q = \frac{m_Q}{g_Q} = \frac{\sum_{P \in I} k_P g_P \cdot m_Q}{g_Q} = \frac{\sum_{P \in I} k_P g_Q m_P}{g_Q} = \sum_{P \in I} k_P m_P = m.$$

Therefore, we can conclude that $\varphi \in M$. □

Quasi-coherent sheaves have the following nice property, which plays a major role in the following section.

Proposition 1.21 ([Gat02, Lemma 7.2.7])**.** *Given an affine variety X and a short exact sequence*

$$0 \longrightarrow \widetilde{M_1} \longrightarrow \widetilde{M_2} \longrightarrow \widetilde{M_3} \longrightarrow 0$$

of quasi-coherent sheaves on X, the induced sequence

$$0 \longrightarrow \Gamma(\widetilde{M_1}) \longrightarrow \Gamma(\widetilde{M_2}) \longrightarrow \Gamma(\widetilde{M_3}) \longrightarrow 0$$

is exact.

Proof. By Lemma 1.20(b), I need to check the exactness of the sequence

$$0 \longrightarrow M_1 \longrightarrow M_2 \longrightarrow M_3 \longrightarrow 0 \ .$$

By a result in commutative algebra, this sequence is exact if the sequences

$$0 \longrightarrow (M_1)_{\mathfrak{m}} \longrightarrow (M_2)_{\mathfrak{m}} \longrightarrow (M_3)_{\mathfrak{m}} \longrightarrow 0$$

are exact for all maximal ideals $\mathfrak{m}$ of $\boldsymbol{k}[X]$ (see Proposition 6.27 of [Gat13]). Finally, combining Thm. 1.16 with Lemma 1.20(a), we see that the sequences of the localisations are exact. □

⚠️ Abstract nonsense ahead

If X is affine with a coordinate ring $R = \mathscr{O}_X(X)$, then the functors

$$\mathrm{QCoh}(X) \to R\text{-}\mathbf{Mod} : \mathscr{F} \mapsto \Gamma(\mathscr{F})$$

and

$$R\text{-}\mathbf{Mod} \to \mathrm{QCoh}(X) : M \mapsto \widetilde{M}$$

define quasi-inverse equivalences of categories [Sta23, Tag 01IB].

2 Sheaf cohomology

Studying sheaves using homological algebra turns out to be surprisingly useful in many situations. For example, knowing that there is a short exact sequence (SES)

$$0 \longrightarrow \mathscr{F} \longrightarrow \mathscr{G} \longrightarrow \mathscr{H} \longrightarrow 0$$

lets us relate the three sheaves together. By Thm. 1.16 this information is inherently **local** since this sequence is exact if and only if the corresponding sequences on stalks are exact. Then the question is: Can we get **global** information from such exact sequences? We would hope that just as the sequence is exact on stalks, it would also be exact on global sections:

$$0 \longrightarrow \Gamma(\mathscr{F}) \longrightarrow \Gamma(\mathscr{G}) \longrightarrow \Gamma(\mathscr{H}) \longrightarrow 0.$$

Unfortunately, this is not the case. For example, let $X = \mathbb{P}^1_{\mathbb{C}}$ and consider the sheaf morphism $\mathscr{O}_X \to \mathbb{C}_{P_0} \oplus \mathbb{C}_{P_1}$, which evaluates a section of $\mathscr{O}_X$ at two distinct points $P_0, P_1 \in X$. Then, the morphism is clearly surjective on the stalks. But it is not surjective on global sections, since the global sections of $\mathscr{O}_X$ are the constant functions. In other words, the exact sequence

$$\mathscr{O}_X \longrightarrow \mathbb{C}_{P_0} \oplus \mathbb{C}_{P_1} \longrightarrow 0$$

does not yield an exact sequence on global sections. However, we have the following.

Proposition 2.1. *If the sequence*

$$0 \longrightarrow \mathscr{F} \xrightarrow{\alpha} \mathscr{G} \xrightarrow{\beta} \mathscr{H} \longrightarrow 0$$

is exact, then the sequence

$$0 \longrightarrow \Gamma(\mathscr{F}) \xrightarrow{\Gamma\alpha} \Gamma(\mathscr{G}) \xrightarrow{\Gamma\beta} \Gamma(\mathscr{H})$$

is also exact.

Proof.

Exactness at $\Gamma(\mathscr{F})$

Suppose $s \in \ker(\Gamma\alpha)$, which is to say that s is in the kernel of the component α_X of the morphism α. Since $\ker(\alpha) = \operatorname{im}(0 \to \mathscr{F})$, we have $s \in \operatorname{im}(0 \to \mathscr{F})(X)$. But $\operatorname{im}(0 \to \mathscr{F})$ is clearly the zero sheaf so that $s = 0$. Also, $\operatorname{im}(0 \to \Gamma(\mathscr{F}))$ is clearly the zero module. Therefore, $\ker(\Gamma\alpha) = 0 = \operatorname{im}(0 \to \Gamma(\mathscr{F}))$.

Exactness at $\Gamma(\mathscr{G})$

Since α is injective, $\operatorname{im}(\alpha)$ can be identified with the sheaf $\mathscr{F}$. Therefore $\ker(\beta) = \mathscr{F}$ and so $\Gamma(\mathscr{F}) = \ker(\beta)(X) = \ker(\beta_X) = \ker(\Gamma\beta)$. Since the second sequence is exact at $\Gamma(\mathscr{F})$, the image $\operatorname{im}(\Gamma\alpha)$ can also be identified with $\Gamma(\mathscr{F})$. Therefore, $\operatorname{im}(\Gamma\alpha) = \Gamma(\mathscr{F}) = \ker(\Gamma\beta)$. □

> ⚠️ Abstract nonsense ahead
>
> We say that the global sections functor $\Gamma : \mathrm{Sh}(X) \to \mathscr{O}_X(X)$-**Mod** is left-exact but not right-exact.

Although exact sequences are not completely preserved under taking global sections, we don't give up! There might still be **a way of measuring how much exactness fails**. We could measure the *obstruction* to exactness by continuing the sequence to the right so that the following sequence is exact.

$$0 \longrightarrow \Gamma(\mathscr{F}) \longrightarrow \Gamma(\mathscr{G}) \longrightarrow \Gamma(\mathscr{H}) \longrightarrow H^1(\mathscr{F}) \longrightarrow H^1(\mathscr{G}) \longrightarrow H^1(\mathscr{H}) \longrightarrow \cdots$$

$$\cdots \longrightarrow H^i(\mathscr{F}) \longrightarrow H^i(\mathscr{G}) \longrightarrow H^i(\mathscr{H}) \longrightarrow \cdots.$$

This problem of extending incomplete short exact sequences appears in other parts of mathematics and has a general solution: *derived functors*. The vector spaces $H^i(-)$ given by derived functors are then called the sheaf cohomology groups. In practise, they are difficult to compute, and thus I will need to define the *Čech cohomology* which is a tool for computing sheaf cohomology. In the next subsections I will introduce derived functors and give a complete explanation of how we arrive at Čech cohomology. The contents of these subsections will be more technical than the rest of the paper, and it is probably a good idea to skip straight to Definition 2.13, which can be taken as the definition of sheaf cohomology. The primary source I use is [Vak17].

2.1 Derived functors*

Let us consider the general problem of extending a left-exact functor $F : \mathcal{A} \to \mathcal{B}$ between abelian categories (which one may think of as categories of modules) to the right. First, I will clarify what is meant by a *left-exact functor*.

Definition 2.2. Consider a functor $F : \mathcal{A} \to \mathcal{B}$ between abelian categories. The functor is said to be left-exact if

1. it is additive: if A, B are objects of $\mathcal{A}$ and $f, g \in \mathrm{Hom}(A, B)$, then $F(f + g) = F(f) + F(g)$ and
2. given a SES
$$0 \longrightarrow A \longrightarrow B \longrightarrow C \longrightarrow 0$$
of objects of $\mathcal{A}$, the sequence
$$0 \longrightarrow F(A) \longrightarrow F(B) \longrightarrow F(C)$$

is exact.

Now, I want to find functors $R^iF : \mathcal{A} \to \mathcal{B}$ so that for objects A, B, C in $\mathcal{A}$ fitting into a SES

$$0 \longrightarrow A \longrightarrow B \longrightarrow C \longrightarrow 0 \ ,$$

the following sequence is exact.

$$\begin{array}{l} 0 \longrightarrow F(A) \longrightarrow F(B) \longrightarrow F(C) \longrightarrow \\ \quad \longrightarrow R^1F(A) \longrightarrow R^1F(B) \longrightarrow R^1F(C) \longrightarrow \cdots \\ \\ \cdots \longrightarrow R^iF(A) \longrightarrow R^iF(B) \longrightarrow R^iF(C) \longrightarrow \cdots . \end{array}$$

The functors R^iF will be called the *right derived functors* of F. The following lemma from homological algebra gives a hint as to what approach one should take to find such functors.

Lemma 2.3 (Zig-zag lemma)**.** *Suppose $A^\bullet, B^\bullet, C^\bullet$ are cochain complexes in some abelian category. If there is a SES*

$$0 \longrightarrow A^\bullet \longrightarrow B^\bullet \longrightarrow C^\bullet \longrightarrow 0 \ ,$$

then there are maps between the cohomology groups of these complexes such that the sequence

$$\begin{array}{l} H^0(A^\bullet) \longrightarrow H^0(B^\bullet) \longrightarrow H^0(C^\bullet) \longrightarrow \\ \quad \longrightarrow H^1(A^\bullet) \longrightarrow H^1(B^\bullet) \longrightarrow H^1(C^\bullet) \longrightarrow \cdots \\ \\ \cdots \longrightarrow H^i(A^\bullet) \longrightarrow H^i(B^\bullet) \longrightarrow H^i(C^\bullet) \longrightarrow \cdots \end{array}$$

is exact.

This lemma can be shown using a typical diagram chasing argument: The maps $H^i(A^\bullet) \to H^i(B^\bullet)$ and $H^i(B^\bullet) \to H^i(C^\bullet)$ are given by functoriality, and the connecting morphisms $H^i(C^\bullet) \to H^{i+1}(A^\bullet)$ are given by *the snake lemma*. Working out the details of this diagram chasing argument is a good exercise for the reader, but I will instead prove the statement using spectral sequences. Explaining the theory of spectral sequences is beyond the scope of this paper and I will refer the reader to [Vak17, section 1.7].

Proof. Define the zeroth page of a spectral sequence to be the following double complex given by the SES of complexes.

$$
\begin{array}{ccccccccc}
 & & \vdots & & \vdots & & \vdots & & \\
 & & \uparrow & & \uparrow & & \uparrow & & \\
0 & \longrightarrow & A^2 & \longrightarrow & B^2 & \longrightarrow & C^2 & \longrightarrow & 0 \\
 & & \uparrow & & \uparrow & & \uparrow & & \\
0 & \longrightarrow & A^1 & \longrightarrow & B^1 & \longrightarrow & C^1 & \longrightarrow & 0 \\
 & & \uparrow & & \uparrow & & \uparrow & & \\
0 & \longrightarrow & A^0 & \longrightarrow & B^0 & \longrightarrow & C^0 & \longrightarrow & 0 \\
 & & \uparrow & & \uparrow & & \uparrow & & \\
 & & 0 & & 0 & & 0 & &
\end{array}
$$

Since the rows are exact the first page is zero when we use the rightward orientation. Now, let us compute the first page using upward orientation. We get the following.

$$
\begin{array}{ccccccccc}
 & & \vdots & & \vdots & & \vdots & & \\
0 & \longrightarrow & H^2(A^\bullet) & \xrightarrow{\alpha_2} & H^2(B^\bullet) & \xrightarrow{\beta_2} & H^2(C^\bullet) & \longrightarrow & 0 \\
0 & \longrightarrow & H^1(A^\bullet) & \xrightarrow{\alpha_1} & H^1(B^\bullet) & \xrightarrow{\beta_1} & H^1(C^\bullet) & \longrightarrow & 0 \\
0 & \longrightarrow & H^0(A^\bullet) & \xrightarrow{\alpha_0} & H^0(B^\bullet) & \xrightarrow{\beta_0} & H^0(C^\bullet) & \longrightarrow & 0
\end{array}
$$

Finally, in the second page we have

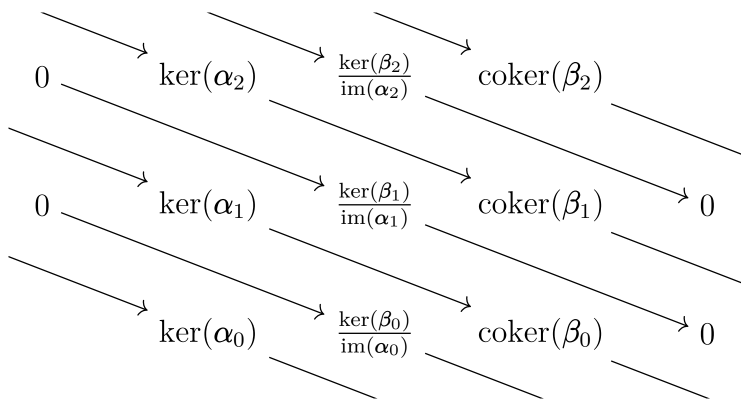

One can see that the spectral sequence will converge on the third page. Since the sequence converges to zero, the sequences

$$0 \longrightarrow \ker(\alpha_{i+1}) \longrightarrow \operatorname{coker}(\beta_i) \longrightarrow 0,$$

given by the differentials on the second page must be exact. These isomorphisms induce maps

$$\delta_i : H^i(C^\bullet) \to H^{i+1}(A^\bullet).$$

The convergence of the spectral sequence also implies that $\ker(\beta_i)/\operatorname{im}(\alpha_i) = 0$. Putting these results together, we see that the sequence

$$\begin{array}{l} H^0(A^\bullet) \xrightarrow{\alpha_0} H^0(B^\bullet) \xrightarrow{\beta_0} H^0(C^\bullet) \xrightarrow{\delta_0} \\ H^1(A^\bullet) \xrightarrow{\alpha_1} H^1(B^\bullet) \xrightarrow{\beta_1} H^1(C^\bullet) \xrightarrow{\delta_1} \cdots \end{array}$$

$$\cdots \xrightarrow{\delta_{i-1}} H^i(A^\bullet) \xrightarrow{\alpha_i} H^i(B^\bullet) \xrightarrow{\beta_i} H^i(C^\bullet) \xrightarrow{\delta_i} \cdots$$

is exact. □

Therefore, in order to extend the sequence

$$0 \longrightarrow F(A) \longrightarrow F(B) \longrightarrow F(C) \ ,$$

we wish to find cocomplexes $A^\bullet, B^\bullet, C^\bullet$ associated to A, B, C such that

1. The cochain complexes $A^\bullet, B^\bullet, C^\bullet$ fit into a SES

$$0 \longrightarrow A^\bullet \longrightarrow B^\bullet \longrightarrow C^\bullet \longrightarrow 0$$

2. The zeroth cohomology coincides with F:

$$H^0(A^\bullet) = F(A), \quad H^0(B^\bullet) = F(B), \quad H^0(C^\bullet) = F(C).$$

One way of associating a cochain complex to an object A is to take its *resolution*. In other words, by finding objects A^i and morphisms such that the sequence

$$0 \longrightarrow A \xrightarrow{\alpha} A^0 \xrightarrow{\alpha_0} A^1 \xrightarrow{\alpha_1} A^2 \xrightarrow{\alpha_2} \cdots$$

is exact. Let us concentrate on the first few terms of this sequence.

$$0 \longrightarrow A \xrightarrow{\alpha} A^0 \xrightarrow{\alpha_0} A^1 \ .$$

Since F is left-exact, we get an another exact sequence:

$$0 \longrightarrow F(A) \xrightarrow{\alpha^*} F(A^0) \xrightarrow{\alpha_0^*} F(A^1) \ .$$

Now, exactness implies that $\ker(\alpha_0^*) = \mathrm{im}(\alpha^*) = F(A)$. Therefore, if I were to replace $F(A)$ by 0, then the cohomology at $F(A^0)$ would be $F(A)$. Thus, considering the cochain complex

$$0 \longrightarrow F(A^0) \xrightarrow{\alpha_0^*} F(A^1) \xrightarrow{\alpha_1^*} F(A^2) \xrightarrow{\alpha_2^*} \cdots ,$$

we see that taking the cohomology of this complex will give

$$H^0(F(A^\bullet)) = F(A).$$

Hence, if I construct the cochain complexes $A^\bullet, B^\bullet, C^\bullet$ from resolutions of A, B, C as above, then the resulting cohomolgy will satisfy the second requirement. Next, I want to find the right type of resolution so that the cochain complexes satisfy the first requirement above.

What we have currently is the following picture.

$$\begin{array}{ccccccccc}
 & & \vdots & & \vdots & & \vdots & & \\
 & & {\scriptstyle\alpha_1}\uparrow & & {\scriptstyle\beta_1}\uparrow & & {\scriptstyle\gamma_1}\uparrow & & \\
 & & A^1 & & B^1 & & C^1 & & \\
 & & {\scriptstyle\alpha_0}\uparrow & & {\scriptstyle\beta_0}\uparrow & & {\scriptstyle\gamma_0}\uparrow & & \\
 & & A^0 & & B^0 & & C^0 & & \\
 & & {\scriptstyle\alpha}\uparrow & & {\scriptstyle\beta}\uparrow & & {\scriptstyle\gamma}\uparrow & & \\
0 & \longrightarrow & A & \xrightarrow{f} & B & \xrightarrow{g} & C & \longrightarrow & 0 \\
 & & \uparrow & & \uparrow & & \uparrow & & \\
 & & 0 & & 0 & & 0 & &
\end{array}$$

If I use *injective resolutions*, the diagram can be filled in with appropriate morphisms so that it gives a SES of complexes.

Definition 2.4. An object I of an abelian category $\mathcal{A}$ is injective, if for every injection $f : A \hookrightarrow B$ and every morphism $g : A \to I$, there is a morphism $B \to I$ such that the following diagram commutes.

$$\begin{array}{ccc}
I & & \\
{\scriptstyle g}\uparrow & \nwarrow & \\
A & \underset{f}{\hookrightarrow} & B
\end{array}$$

Then, an injective resolution of an object A is a long exact sequence

$$0 \longrightarrow A \longrightarrow I^0 \longrightarrow I^1 \longrightarrow \cdots$$

where the I^i are injective. Note that it is not obvious that an object of an abelian category should have an injective resolution in the first place. Thus, we assume that the category $\mathcal{A}$ has *enough injectives* meaning that for every object A of $\mathcal{A}$, there is an injection $A \hookrightarrow I$ into some injective object I. Then, for every object A we can construct an injective resolution inductively. The first object I^0 is given directly by the assumption. Then suppose we have constructed an exact sequence

$$0 \longrightarrow A \longrightarrow I^0 \xrightarrow{\iota_0} \cdots \longrightarrow I^{n-1} \xrightarrow{\iota_{n-1}} I^n \ .$$

Let us take I^{n+1} to be an injective object such that there is an injection

$$\operatorname{coker}(\iota_{n-1}) \hookrightarrow I^{n+1}.$$

Then we have

$$0 \longrightarrow A \longrightarrow I^0 \xrightarrow{\iota_0} \cdots \xrightarrow{\iota_{n-1}} I^n \twoheadrightarrow \operatorname{coker}(\iota_{n-1}) \hookrightarrow I^{n+1} \ .$$

When the injection is composed with the projection, we get the exact sequence

$$0 \longrightarrow A \longrightarrow I^0 \xrightarrow{\iota_0} \cdots \longrightarrow I^{n-1} \xrightarrow{\iota_{n-1}} I^n \xrightarrow{\iota_n} I^{n+1} \ .$$

Now I will quickly prove a useful lemma about injective objects:

Lemma 2.5. *A product of injective objects is injective.*

Proof. Suppose $(J_i)_{i \in I}$ is a collection of injective objects in an abelian category, and denote

$$J = \prod_{i \in I} J_i.$$

Then, suppose we have a morphism $g : X \to J$ and an injection $f : X \hookrightarrow Y$. Recall that J is injective if we can extend g along f so that the following diagram commutes.

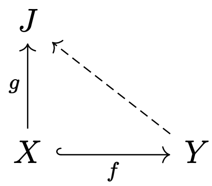

We can compose g with the projection morphisms $\pi_{J_i} : J \to J_i$, and then extend the compositions along f since the J_i are injective by assumption:

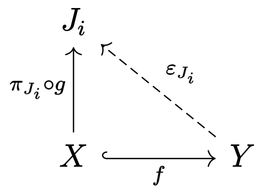

Now, by the universal property of the product J, there is a morphism $\varepsilon : Y \to J$, such that the following diagrams commute for all $i \in I$.

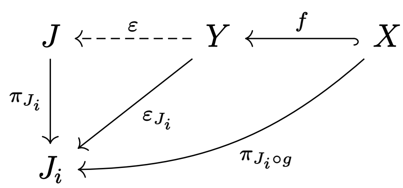

By the universal property of the product J, we also have that g is the unique morphism making the triangles

$$\begin{array}{ccc} J & \xleftarrow{\;g\;} & X \\ {\scriptstyle \pi_{J_i}}\downarrow & \swarrow {\scriptstyle \pi_{J_i}\circ g} & \\ J_i & & \end{array}$$

commute for all $i \in I$. Since these triangles are the same as the ones we get by composing ε with f, we must have $\varepsilon \circ f = g$ and therefore, ε is an extension of g along f. $\square$

Now, the definition of injective objects gives a way of constructing a SES of injective resolutions given a SES of objects in an abelian category.

Lemma 2.6. *If A, B, C are objects of an abelian category $\mathcal{A}$ with enough injectives fitting into a SES*

$$0 \longrightarrow A \xrightarrow{\;f\;} B \xrightarrow{\;g\;} C \longrightarrow 0$$

and A and C have injective resolutions $A^\bullet$ and $C^\bullet$, then there is an injective resolution $B^\bullet$ of B and maps $f^i : A^i \to B^i$ and $g^i : B^i \to C^i$ such that the following diagram commutes and has exact rows.

$$\begin{array}{ccccccccc}
 & & \vdots & & \vdots & & \vdots & & \\
 & & {\scriptstyle \alpha_1}\uparrow & & {\scriptstyle \beta_1}\uparrow & & {\scriptstyle \gamma_1}\uparrow & & \\
0 & \longrightarrow & A^1 & \xrightarrow{f^1} & B^1 & \xrightarrow{g^1} & C^1 & \longrightarrow & 0 \\
 & & {\scriptstyle \alpha_0}\uparrow & & {\scriptstyle \beta_0}\uparrow & & {\scriptstyle \gamma_0}\uparrow & & \\
0 & \longrightarrow & A^0 & \xrightarrow{f^0} & B^0 & \xrightarrow{g^0} & C^0 & \longrightarrow & 0 \\
 & & {\scriptstyle \alpha}\uparrow & & {\scriptstyle \beta}\uparrow & & {\scriptstyle \gamma}\uparrow & & \\
0 & \longrightarrow & A & \xrightarrow{f} & B & \xrightarrow{g} & C & \longrightarrow & 0 \\
 & & \uparrow & & \uparrow & & \uparrow & & \\
 & & 0 & & 0 & & 0 & &
\end{array}$$

Proof. One can take $B^i = A^i \times C^i$, and then the rows

$$\begin{array}{ccccccccc}
0 & \longrightarrow & A^i & \longrightarrow & A^i \times C^i & \longrightarrow & C^i & \longrightarrow & 0 \\
 & & a & \longmapsto & (a, 0) & & & & \\
 & & & & (a, c) & \longmapsto & c & &
\end{array}$$

are clearly exact. Moreover, the B^i are injective by Lemma 2.5, so I only need to find maps $\beta : B \to B^0$ and $\beta^i : B^i \to B^{i+1}$ making the diagram commute so that the B^i form an injective resolution of B.

Let us construct the morphism $\beta : B \to B^0$ which makes the following diagram commute.

$$\begin{array}{ccccccccc} 0 & \longrightarrow & A^0 & \xrightarrow{f^0} & A^0 \times C^0 & \xrightarrow{g^0} & C^0 & \longrightarrow & 0 \\ & & {\scriptstyle\alpha}\uparrow & & {\scriptstyle\beta}\uparrow & & {\scriptstyle\gamma}\uparrow & & \\ 0 & \longrightarrow & A & \xrightarrow{f} & B & \xrightarrow{g} & C & \longrightarrow & 0 \\ & & \uparrow & & \uparrow & & \uparrow & & \\ & & 0 & & 0 & & 0 & & \end{array}$$

Since A^0 is injective, we can extend α along the injection f:

$$\begin{array}{ccc} A^0 & & \\ {\scriptstyle\alpha}\uparrow & \nwarrow {\scriptstyle\alpha'} & \\ A & \xrightarrow[f]{} & B \end{array}$$

Then, we get the diagram

$$\begin{array}{ccccc} A^0 & \xleftarrow{\pi_{A^0}} & A^0 \times C^0 & \xrightarrow{\pi_{C^0}} & C^0 \\ & {\scriptstyle\alpha'}\nwarrow & {\scriptstyle\beta}\uparrow & \nearrow{\scriptstyle\gamma\circ g} & \\ & & B & & \end{array},$$

where β is given by the universal property of the product. The morphisms β_i are constructed in exactly the same way.

Finally, I need to check the exactness of the sequence

$$0 \longrightarrow B \xrightarrow{\beta} B^0 \xrightarrow{\beta_0} B^1 \xrightarrow{\beta_1} \cdots ,$$

but this follows by a nearly trivial spectral sequence argument. Define a spectral sequence where the 0th page is the double complex we have constructed. Computing the 1st page using rightward orientation yields 0, since all rows are exact. Thus, the spectral sequence converges to 0. One can see that the spectral sequence must converge on the 1st page also when using the upward orientation. Therefore, the column corresponding to the injective resolution of B must be exact. □

Now I only need to apply the functor F on this double complex and remove the bottom row. The only worry is that the rows don't stay exact. Thus, I will need to prove one more small result.

Lemma 2.7. *Suppose $F : \mathcal{A} \to \mathcal{B}$ is a left-exact functor between abelian categories. If I and J are injective objects of $\mathcal{A}$, then F is exact on the SES*

$$0 \longrightarrow I \longrightarrow I \times J \longrightarrow J \longrightarrow 0 .$$

Proof. In abelian categories, the object $I \times J$ is both a product and a coproduct. It follows that additive functors preserve these products. The result follows directly from this remark, because applying the functor F on the SES yields

$$0 \longrightarrow F(I) \longrightarrow F(I) \times F(J) \longrightarrow F(J) \longrightarrow 0 \ ,$$

which is clearly exact. □

In summary, given a left-exact functor $F : \mathcal{A} \to \mathcal{B}$ and an object A of $\mathcal{A}$, one constructs the ith right derived functor of F at A in the following way:

1. Find an injective resolution $0 \to A \to I^\bullet$ of A
2. Apply F on the cochain complex $0 \to I^\bullet$
3. Take the ith cohomology of the cochain complex

$$0 \longrightarrow F(I^0) \longrightarrow F(I^1) \longrightarrow F(I^2) \longrightarrow \cdots$$

We denote $R^iF(A) = H^i(F(I^\bullet))$ for the value of the derived functor.

The alert reader might have noticed that this definition depends a-priori on the injective resolution we choose. However, one can show that this is not the case, see [Vak17].

⚠ Abstract nonsense ahead

One could ask: "How do we know that derived functors give the 'correct' way of extending the left-exact functor?" This question can be formalised by considering so-called (cohomological) δ-functors [Vak17], which consist of pairs (T^i, δ^i), where

1. the $T^i : \mathcal{A} \to \mathcal{B}$ are additive functors between abelian categories with $T^i = 0$ for $i < 0$, and
2. $\delta^i : T^i(C) \to T^{i+1}(A)$ are morphisms in $\mathcal{B}$,

such that for every SES

$$0 \longrightarrow A \longrightarrow B \longrightarrow C \longrightarrow 0$$

in $\mathcal{A}$, we have a long exact sequence

$$\begin{array}{l} 0 \longrightarrow T^0(A) \longrightarrow T^0(B) \longrightarrow T^0(C) \xrightarrow{\delta_0} \\ \quad T^1(A) \longrightarrow T^1(B) \longrightarrow T^1(C) \xrightarrow{\delta_1} \cdots \end{array}$$

$$\cdots \xrightarrow{\delta_{i-1}} T^i(A) \longrightarrow T^i(B) \longrightarrow T^i(C) \xrightarrow{\delta_i} \cdots$$

In addition, we require functoriality of this construction: If

$$\begin{array}{ccccccccc} 0 & \longrightarrow & A & \longrightarrow & B & \longrightarrow & C & \longrightarrow & 0 \\ & & \downarrow & & \downarrow & & \downarrow & & \\ 0 & \longrightarrow & A' & \longrightarrow & B' & \longrightarrow & C' & \longrightarrow & 0 \end{array}$$

is a morphism of short exact sequences in $\mathcal{A}$, then the squares

$$\begin{array}{ccc} T^i(C) & \xrightarrow{\delta^i} & T^{i+1}(A) \\ \downarrow & & \downarrow \\ T^i(C') & \xrightarrow[\delta^i]{} & T^{i+1}(A') \end{array}$$

commute.
One can then define morphisms of δ-functors so that they form a category. After that, one formulates the concept of a *universal δ-functor* in this category: A δ-functor (T^i, δ^i) is universal if for every other δ-functor (S^i, γ^i) with a natural transformation $\alpha : T^0 \Rightarrow S^0$, there is a unique morphism of δ-functors $(T^i, \delta^i) \to (S^i, \gamma^i)$ extending α. One can then prove that derived functors are universal δ-functors.

Now, consider a sheaf $\mathscr{F}$ of $\mathscr{O}_X$-modules. Its *sheaf cohomology* is defined as the right derived functor of the global sections functor:

$$H^i(X, \mathscr{F}) := R^i\Gamma(\mathscr{F}) \, .$$

This definition relies on the assumption that the category of sheaves of $\mathscr{O}_X$-modules has enough injectives.

Theorem 2.8. *The category $\mathscr{O}_X$-***Mod** *has enough injectives.*

Proof. The proof of this theorem is given as a series of exercises in [Vak17]. □

2.2 Acyclic resolutions*

Derived functors provide the framework for the cohomology theory of sheaves, but working with injective resolutions is difficult in practise. Thus, we need an alternative way of constructing the cohomology groups. This is done by replacing injective resolutions by *acyclic resolutions*. If $F : \mathcal{A} \to \mathcal{B}$ is a left-exact functor between abelian categories, then an object A of $\mathcal{A}$ is acyclic (w.r.t. F) if $R^iF(A) = 0$ for $i > 0$.

Lemma 2.9. *Let $F : \mathcal{A} \to \mathcal{B}$ be a left-exact functor between abelian categories, where $\mathcal{A}$ has enough injectives. Suppose A is an object of $\mathcal{A}$ with acyclic resolution*

$$0 \longrightarrow A \longrightarrow A^0 \longrightarrow A^1 \longrightarrow \cdots .$$

Then, computing the cohomology of the cochain complex

$$0 \longrightarrow F(A^0) \longrightarrow F(A^1) \longrightarrow \cdots$$

agrees with the right derived functor of F at A.

Proof. I will prove the statement using a spectral sequence argument. Thus, I need to set up a double complex, which will be the 0th page of the sequence. First note that the long exact sequence

$$0 \longrightarrow A \xrightarrow{\alpha} A^0 \xrightarrow{\alpha_0} A^1 \xrightarrow{\alpha_1} \cdots$$

can be broken into short exact sequences

$$0 \longrightarrow A \xrightarrow{\alpha} \operatorname{im}\alpha \longrightarrow 0$$

and

$$0 \longrightarrow \ker\alpha_i \longrightarrow A^i \xrightarrow{\alpha_i} \operatorname{im}\alpha_i \longrightarrow 0 .$$

Lemma 2.6 can now be used to construct short exact sequences of injecitve resolutions, which can be combined to form a double complex.

$$\begin{array}{ccccccccccc}
 & & \vdots & & \vdots & & \vdots & & \vdots & & \\
 & & \uparrow & & \uparrow & & \uparrow & & \uparrow & & \\
0 & \longrightarrow & I^2 & \longrightarrow & I^{0,2} & \longrightarrow & I^{1,2} & \longrightarrow & I^{2,2} & \longrightarrow & \cdots \\
 & & \uparrow & & \uparrow & & \uparrow & & \uparrow & & \\
0 & \longrightarrow & I^1 & \longrightarrow & I^{0,1} & \longrightarrow & I^{1,1} & \longrightarrow & I^{2,1} & \longrightarrow & \cdots \\
 & & \uparrow & & \uparrow & & \uparrow & & \uparrow & & \\
0 & \longrightarrow & I^0 & \longrightarrow & I^{0,0} & \longrightarrow & I^{1,0} & \longrightarrow & I^{2,0} & \longrightarrow & \cdots \\
 & & \uparrow & & \uparrow & & \uparrow & & \uparrow & & \\
0 & \longrightarrow & A & \longrightarrow & A^0 & \longrightarrow & A^1 & \longrightarrow & A^2 & \longrightarrow & \cdots \\
 & & \uparrow & & \uparrow & & \uparrow & & \uparrow & & \\
 & & 0 & & 0 & & 0 & & 0 & &
\end{array}$$

Removing the bottom row and the left-most column and applying the functor F gives the following double complex.

$$
\begin{array}{ccccccccc}
 & & \vdots & & \vdots & & \vdots & & \\
 & & \uparrow & & \uparrow & & \uparrow & & \\
0 & \longrightarrow & F(I^{0,2}) & \longrightarrow & F(I^{1,2}) & \longrightarrow & F(I^{2,2}) & \longrightarrow & \cdots \\
 & & \uparrow & & \uparrow & & \uparrow & & \\
0 & \longrightarrow & F(I^{0,1}) & \longrightarrow & F(I^{1,1}) & \longrightarrow & F(I^{2,1}) & \longrightarrow & \cdots \\
 & & \uparrow & & \uparrow & & \uparrow & & \\
0 & \longrightarrow & F(I^{0,0}) & \longrightarrow & F(I^{1,0}) & \longrightarrow & F(I^{2,0}) & \longrightarrow & \cdots \\
 & & \uparrow & & \uparrow & & \uparrow & & \\
 & & 0 & & 0 & & 0 & &
\end{array}
$$

Now, I take this double complex to be the 0th page of the spectral sequence. Note that computing the cohomology groups of the rows and columns is the same as computing the values of the right derived functors of F; The cohomology groups of the rows correspond to the values of the right derived functor at the objects I^i and the cohomology groups of the columns correspond to the values at A^i. Note also that injective objects are acyclic, because an injective object I has the trivial resolution $0 \to I \to I \to 0$. Thus, when one computes the first page of the sequence starting with the rightward orientation, the only non-zero column will be the first column:

$$
\begin{array}{c}
\vdots \\
\uparrow \\
F(I^1) \\
\uparrow \\
F(I^0) \\
\uparrow \\
0
\end{array}
$$

Then, the entries on the second page will be equal to the values of the right derived functors R^iF at the object A.

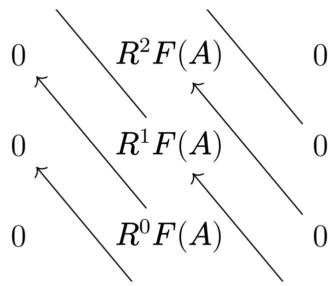

The sequence collapses at the 2nd step and one can see that the total cohomology of the complex corresponds to the values of the right derived functor at A. Similarly, computing the first page starting with the upward orientation gives only one non-zero row, since the A^i are acyclic:

$$0 \longrightarrow F(A^0) \longrightarrow F(A^1) \longrightarrow F(A^2) \longrightarrow \cdots$$

Computing the second page gives the cohomology groups of this complex. Again, the sequence collapses and therefore the values of the right derived functor agree with the cohomology groups of the complex. □

Finally, let us prove the analogue of Lemma 2.5 for acyclic objects.

Lemma 2.10. *Suppose $F : \mathcal{A} \to \mathcal{B}$ is a left-exact functor between abelian categories. If A_1 and A_2 are F-acyclic objects of $\mathcal{A}$, then the product $A_1 \times A_2$ is F-acyclic.*

Proof. Suppose A_1 and A_2 have injective resolutions

$$0 \longrightarrow A_1 \longrightarrow I_1^0 \xrightarrow{d_1^0} I_1^1 \xrightarrow{d_1^1} I_1^2 \xrightarrow{d_1^2} \cdots$$

and

$$0 \longrightarrow A_2 \longrightarrow I_2^0 \xrightarrow{d_2^0} I_2^1 \xrightarrow{d_2^1} I_2^2 \xrightarrow{d_2^2} \cdots .$$

Then,

$$0 \longrightarrow A_1 \times A_2 \longrightarrow I_1^0 \times I_2^0 \xrightarrow{d_1^0 \times d_2^0} I_1^1 \times I_2^1 \xrightarrow{d_1^1 \times d_2^1} \cdots$$

is an injective resolution of $A_1 \times A_2$ by Lemma 2.5. Applying the functor F to the cochain complex $0 \to I_1^\bullet \times I_2^\bullet$ yields

$$0 \longrightarrow F(I_1^0) \times F(I_2^0) \xrightarrow{F(d_1^0) \times F(d_2^0)} F(I_1^1) \times F(I_2^1) \xrightarrow{F(d_1^1) \times F(d_2^1)} \cdots , \tag{4}$$

as additive functors preserve finite products. Since A_1 and A_2 are F-acyclic, the cohomologies of the cochain complexes $0 \to F\left(I_1^\bullet\right)$ and $0 \to F\left(I_2^\bullet\right)$ vanish in positive degrees. Therefore, one can immediately see that the cohomology of the cochain complex in (4) vanishes in positive degrees. □

2.3 Discovering the Čech complex*

In this subsection, we will construct acyclic resolutions for sheaves, so that we can compute their sheaf cohomology. We end up discovering the Čech complex, which gives the Čech cohomology groups. To start, recall Prop. 1.21, which states that if we have a SES of quasi-coherent sheaves on an affine variety, then the global sections functor is exact on this SES. This leads us to suspect that quasi-coherent sheaves on affine varieties are acyclic, which is indeed the case. The proof of this statement relies on flasque sheaves, which I have not introduced.

Lemma 2.11 ([Har97, Thm. 3.5]). *If X is an affine variety and $\mathscr{F}$ is a quasi-coherent sheaf, then $H^i(X, \mathscr{F}) = 0$ for all $i > 0$.*

Proof. If I denote $M = \Gamma(\mathscr{F})$, then $\mathscr{F} = \widetilde{M}$. Now, the module M has an injective resolution $0 \to M \to I^\bullet$. This gives an exact sequence $0 \to \widetilde{M} \to \widetilde{I}^\bullet$. The sheaves $\widetilde{I}^i$ are acyclic by [Har97, Prop. 3.4] and [Har97, Prop. 2.5] (the sheaves $\widetilde{I}^i$ are flasque and flasque sheaves are acyclic). Therefore, the sequence $0 \to \widetilde{M} \to \widetilde{I}^\bullet$ is an acyclic resolution of $\widetilde{M} = \mathscr{F}$. Applying the global sections functor to this resolution returns the original sequence $0 \to M \to I^\bullet$, which is exact. After removing the M term, the cohomology groups at I^0 is $\Gamma(\mathscr{F})$, but the cohomology groups at I^i are zero for $i > 0$. □

Equipped with this key insight, consider now a quasi-coherent sheaf $\mathscr{F}$ on a general variety X, and recall the exact sequence (3). If one substitutes $U = X$ and the U_i form an affine open cover of X, then the sequence reads

$$0 \longrightarrow \mathscr{F}(X) \longrightarrow \bigoplus_{i_0} \mathscr{F}(U_{i_0}) \longrightarrow \bigoplus_{i_0,i_1} \mathscr{F}(U_{i_0} \cap U_{i_1}) \ ,$$

which looks awfully like the beginning of an acyclic resolution in light of the above lemma. Thus, I would like to find sheaves with global sections that match the ones in the exact sequence. But this is quite simple in fact; I can use the constructions from Subsection 1.2 to define

$$\mathscr{F}_{i_0,\ldots,i_k} := \iota_*^{i_0,\ldots,i_k}\left(\mathscr{F}|_{U_{i_0,\ldots,i_k}}\right),$$

where $U_{i_0,\ldots,i_k} = U_{i_0} \cap \cdots \cap U_{i_k}$ and $\iota^{i_0,\ldots,i_k} : U_{i_0,\ldots,i_k} \hookrightarrow X$. I can then define an exact sequence of these sheaves following the sequence (3):

$$0 \longrightarrow \mathscr{F} \longrightarrow \bigoplus_{i_0} \mathscr{F}_{i_0} \xrightarrow{d^0} \bigoplus_{i_0,i_1} \mathscr{F}_{i_0,i_1} \ .$$

Now, I want to continue this sequence to the right by constructing a morphism d^1 from $\bigoplus \mathscr{F}_{i_0,i_1}$ such that $\ker d^1 = \operatorname{im} d^0$. Thus, fix an open set $V \subseteq X$ and consider a section ψ of the pre-sheaf $\operatorname{im}' d^0(V)$. Such a section consists of components

$$\psi_{i_0,i_1} = \varphi_{i_1}|_{V \cap U_{i_0,i_1}} - \varphi_{i_0}|_{V \cap U_{i_0,i_1}},$$

where φ is a section of $\bigoplus \mathscr{F}_{i_0}$ with $\varphi_{i_0} \in \mathscr{F}(V \cap U_{i_0})$ and $\varphi_{i_1} \in \mathscr{F}(V \cap U_{i_1})$. The aim now is to construct a morphism, which sends ψ to zero. This can be done by considering one more index i_2 and taking an alternating sum of sections restricted to $V \cap U_{i_0,i_1,i_2}$. For conciseness, I omit the restriction symbols in the following computation, because every section is restricted to a common set.

$$\begin{aligned} \psi_{i_1,i_2} - \psi_{i_0,i_2} + \psi_{i_0,i_1} &= (\varphi_{i_2} - \varphi_{i_1}) - (\varphi_{i_2} - \varphi_{i_0}) + (\varphi_{i_1} - \varphi_{i_0}) \\ &= (\varphi_{i_2} - \varphi_{i_2}) + (\varphi_{i_1} - \varphi_{i_1}) + (\varphi_{i_0} - \varphi_{i_0}) = 0. \end{aligned}$$

Thus, it would be natural to extend the sequence by a morphism

$$d^1 : \bigoplus_{i_0,i_1} \mathscr{F}_{i_0,i_1} \to \bigoplus_{i_0,i_1,i_2} \mathscr{F}_{i_0,i_1,i_2}$$

defined component-wise by

$$\left(d^1_V(\psi)\right)_{i_0,i_1,i_2} = \psi_{i_1,i_2} - \psi_{i_0,i_2} + \psi_{i_0,i_1}.$$

Now, it is an easy task to check the exactness of the sequence

$$\bigoplus_{i_0} \mathscr{F}_{i_0} \xrightarrow{d^0} \bigoplus_{i_0,i_1} \mathscr{F}_{i_0,i_1} \xrightarrow{d^1} \bigoplus_{i_0,i_1,i_2} \mathscr{F}_{i_0,i_1,i_2}$$

by checking exactness on the stalks at points $P \in X$. Note that $\operatorname{im} d^0_P \subseteq \ker d^1_P$ follows by exactly the same computation as above. Thus, I will quickly demonstrate $\ker d^1_P \subseteq \operatorname{im} d^0_P$. Suppose $\psi \in \ker d^1_P$ and fix an index i. Then, take an arbitrary germ φ_i in $(\mathscr{F}_i)_P$. Furthermore, if j is any index, define the germ φ_j in $(\mathscr{F}_j)_P$ as $\psi_{i,j} + \varphi_i$. Now, for any pair of indecies i_0, i_1, we have

$$\begin{aligned} \psi_{i_0,i_1} &= \psi_{i,i_1} - \psi_{i,i_0} \\ &= (\varphi_{i_1} - \varphi_i) - (\varphi_{i_0} - \varphi_i) \\ &= \varphi_{i_1} - \varphi_{i_0}. \end{aligned}$$

Therefore, one can see that the germ φ formed from these components φ_j map to ψ under d^0_P.

One can continue this sequence in the same fashion by taking alternating sums of sections, thus obtaining a resolution

$$0 \longrightarrow \mathscr{F} \longrightarrow \bigoplus_{i_0} \mathscr{F}_{i_0} \longrightarrow \bigoplus_{i_0,i_1} \mathscr{F}_{i_0,i_1} \longrightarrow \bigoplus_{i_0,i_1,i_2} \mathscr{F}_{i_0,i_1,i_2} \longrightarrow \cdots .$$

Remark. *The above computations might seem familiar if you have seen homology or cohomology before. In fact, Čech cohomology can be seen as a singular cohomology of the **nerve** of the open covering (U_i) of the space [Hat01]. The nerve of a covering is an abstract simplicial complex, where the sets U_i are the points and a non-empty intersection of k of the sets is a k-simplex.*

The last thing I need to show is that the resolution is acyclic.

Lemma 2.12. *The sheaves $\bigoplus_{i_0,\ldots,i_k} \mathscr{F}_{i_0,\ldots,i_k}$ are acyclic.*

Proof. By Lemma 2.10, I only need to check that the sheaf $\mathscr{F}_{i_0,\ldots,i_k}$ is acyclic for an arbitrary index $i_0, \ldots, i_k$. Thus, fix an index and denote $V = U_{i_0,\ldots,i_k}$ and $\iota \colon V \hookrightarrow X$. Then, suppose $\mathscr{F}|_V$ has an injective resolution

$$0 \longrightarrow \mathscr{F}|_V \longrightarrow \mathscr{I}^0 \longrightarrow \mathscr{I}^1 \longrightarrow \mathscr{I}^2 \longrightarrow \cdots .$$

This is now a sequence of sheaves on V, and one can push them forward to X using the construction from Subsection 1.2. It is easy to see that the resulting sequence

$$0 \longrightarrow \iota_*\left(\mathscr{F}|_V\right) \longrightarrow \iota_*\mathscr{I}^0 \longrightarrow \iota_*\mathscr{I}^1 \longrightarrow \iota_*\mathscr{I}^2 \longrightarrow \cdots$$

is exact. If we can show that the sheaves $\iota_*\mathscr{I}^i$ are injective, then we have constructed an injective resolution of $\mathscr{F}_{i_0,\ldots,i_k}$.

Thus, suppose there is an injection $\mathscr{A} \hookrightarrow \mathscr{B}$ of sheaves on X and a morphism $g : \mathscr{A} \to \iota_*\mathscr{I}^i$. After restricting to V, we can extend $g|_V$ due to the injectivity of $\mathscr{I}^i$.

$$\begin{array}{ccc} \mathscr{I}^i & & \\ {\scriptstyle g|_V}\uparrow & \nwarrow{\scriptstyle u} & \\ \mathscr{A}|_V & \hookrightarrow & \mathscr{B}|_V \end{array}$$

This can be used to define an extension $u^+ \colon \mathscr{B} \to \iota_*\mathscr{I}^i$ of g. Fix an open set $U \subseteq X$. Then, we can simply define

$$u_U^+ \colon \mathscr{B}(U) \to \mathscr{I}^i(U \cap V) : \sigma \mapsto u_{U\cap V}(\sigma|_{U\cap V}).$$

It is clear that u^+ fits into the following commutative triangle.

$$\begin{array}{ccc} \iota_*\mathscr{I}^i & & \\ {\scriptstyle g}\uparrow & \nwarrow{\scriptstyle u^+} & \\ \mathscr{A} & \hookrightarrow & \mathscr{B} \end{array}$$

We can conclude that $\iota_*\mathscr{I}^i$ is injective.

Note that the sequence

$$0 \longrightarrow \Gamma\left(\iota_*\mathscr{I}^0\right) \longrightarrow \Gamma\left(\iota_*\mathscr{I}^1\right) \longrightarrow \Gamma\left(\iota_*\mathscr{I}^2\right) \longrightarrow \cdots .$$

is simply equal to

$$0 \longrightarrow \Gamma\left(\mathscr{I}^0\right) \longrightarrow \Gamma\left(\mathscr{I}^1\right) \longrightarrow \Gamma\left(\mathscr{I}^2\right) \longrightarrow \cdots .$$

This shows that the cohomology of $\mathscr{F}_{i_0,\ldots,i_k}$ coincides with that of $\mathscr{F}|_V$. But $\mathscr{F}|_V$ is quasi-coherent since $\mathscr{F}$ is and V is an affine open set by Lemma 1.11. Thus, $\mathscr{F}|_V$ is acyclic by Lemma 2.11. Therefore, so is $\mathscr{F}_{i_0,\ldots,i_k}$. $\square$

2.4 Čech cohomology

We finally arrive at the definition of Čech cohomology.

Definition 2.13. The Čech cohomology of a quasi-coherent sheaf $\mathscr{F}$ on a variety X is the cohomology of the cochain complex

$$\bigoplus_{i_0} \mathscr{F}(U_{i_0}) \xrightarrow{d^0} \bigoplus_{i_0,i_1} \mathscr{F}(U_{i_0,i_1}) \xrightarrow{d^1} \bigoplus_{i_0,i_1,i_2} \mathscr{F}(U_{i_0,i_1,i_2}) \xrightarrow{d^2} \cdots ,$$

where the U_i form an affine open cover of X. I denote

$$U_{i_0,\dots,i_k} = U_{i_0} \cap \cdots \cap U_{i_k}$$

and the differentials are defined as the products of the maps

$$d^k_{i_0,\dots,i_{k+1}} : \bigoplus_{j_0,\dots,j_k} \mathscr{F}(U_{j_0,\dots,j_k}) \to \mathscr{F}(U_{i_0,\dots,i_{k+1}})$$

given by

$$d^k_{i_0,\dots,i_{k+1}}(\alpha) = \sum_{j=0}^{k+1} (-1)^j \alpha_{i_0,\dots,\hat{i}_j,\dots,i_{k+1}}|_{U_{i_0,\dots,i_{k+1}}}.$$

More specifically, the ith cohomology group $H^i(X, \mathscr{F})$ is defined as the k-vector space $\ker(d^i)/\operatorname{im}(d^{i-1})$ if we take d^{-1} to be the zero-morphism.

I will now use this new tool to compute some cohomology groups.

Proposition 2.14. *If $\underline{A}$ is a constant sheaf that is quasi-coherent on some irreducible variety X, then*

$$H^0(X, \underline{A}) = A \quad \textit{and} \quad H^i(X, \underline{A}) = 0 \textit{ for } i > 0.$$

Proof. The statement $H^0(X, \underline{A}) = A$ follows directly from the purely topological fact that locally constant functions on a connected space are constant: A global section of $\underline{A}$ can be represented by a locally constant function $f : X \to A$. Fix an arbitrary value $a \in A$ in the image of f and consider the inverse image $f^{-1}(a)$. For an arbitrary point $P \in f^{-1}(a)$, there is an open neighbourhood of P, where f is constant. This immmediately implies that $f^{-1}(a)$ is open, since every point $P \in f^{-1}(a)$ has an open neighbourhood contained in $f^{-1}(a)$. But if $f^{-1}(a)$ is not the whole of X, then $f^{-1}(A \setminus \{ a \})$ is a non-empty open set which disconnects X together with $f^{-1}(a)$. Hence, we arrive at a contradiction.

I will now show $H^1(X, \underline{A}) = 0$, which follows from the exactness of the sequence

$$\bigoplus_{i_0} \mathscr{F}(U_{i_0}) \xrightarrow{d^0} \bigoplus_{i_0,i_1} \mathscr{F}(U_{i_0,i_1}) \xrightarrow{d^1} \bigoplus_{i_0,i_1,i_2} \mathscr{F}(U_{i_0,i_1,i_2}) .$$

Since this is a part of a cochain complex, we know that $\operatorname{im}(d^0) \subseteq \ker(d^1)$ and I only need to show $\ker(d^1) \subseteq \operatorname{im}(d^0)$. Thus, suppose $b \in \ker(d^1)$ and fix an index i. Then, let a_i be an arbitrary constant on $\underline{A}(U_i)$. Now, for an arbitrary index j, one can define a section a_j of $\underline{A}(U_j)$ as $b_{i,j} + a_i$ and check that the condition $b \in \ker d^1$ implies that the components a_j form a section of $\bigoplus \mathscr{F}(U_{i_0})$ which is mapped onto b by d^0 and we can conclude that $a \in \operatorname{im}(d^0)$. The computation for the higher cohomology groups is exactly the same, but there are more indices to juggle. □

I leave the following proposition as an exercise for the reader.

Proposition 2.15. *If A_P is a skyscraper sheaf that is quasi-coherent on some irreducible variety X, then cohomology groups $H^i(X, A_P)$ are zero for $i > 0$.*

3 The Riemann-Roch theorem

Equipped with sheaf cohomology, I will prove the Riemann-Roch theorem using the methods we have learnt.

The rest of this article will be concerned with algebraic curves. Thus, from now on X will always denote an irreducible, non-singular, projective curve over an algebraically closed field k. We will often use the fact that in this case the field of rational functions $k(X)$ consists of functions f/g, where f and g are homogeneous polynomials over k with $\deg f = \deg g$.

3.1 Divisors and differentials

Before tackling the Riemann-Roch theorem, I will quickly review two constructions that we will need in the last two sections of this paper: *divisors* and *differentials*. I use [Gat02] and [Ser12] as my sources.

Like a sheaf, a divisor is an object that associates additional data to an algebraic variety. Unlike sheaves, divisors contain **discrete** data. More specifically, a divisor D on X associates an integer to each point of X and only finitely many of these integers are non-zero. Then, we can represent the divisor as a formal linear combination

$$D = \sum_{P \in X} n_P P,$$

where the integer n_P is the value associated to the point P. Now, two divisors can be added component by component so that divisors on X form an abelian group $\operatorname{Div} X$. Given a divisor D, I write $D(P)$ for the value n_P associated to the point $P \in X$.

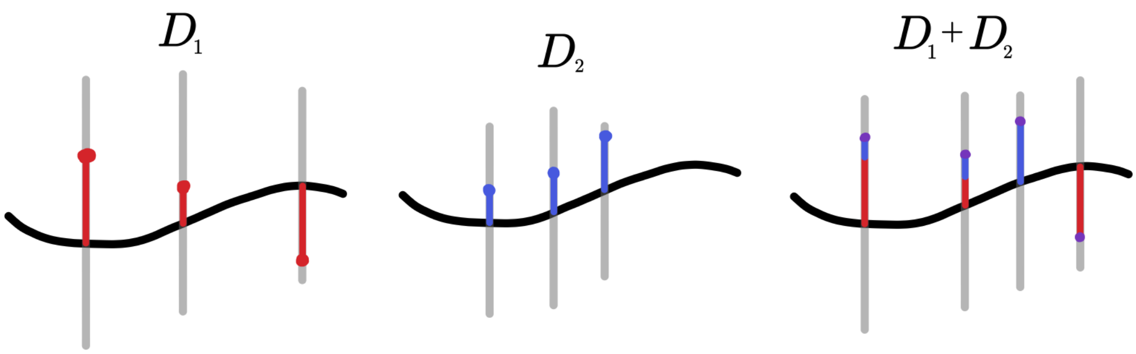

Figure 3: Visualising divisors and their sums.

I will now define the *degree* of a divisor and an order relation on the group of divisors. The degree of a divisor D is simply the integer

$$\deg(D) = \sum_{P \in X} D(P).$$

Next, I say D is *effective* and write $D \geq 0$ if $D(P) \geq 0$ for all $P \in X$. Then, given two divisors D, D', I can define $D \geq D'$ if the divisor $D - D'$ is effective.

One can associate a divisor to a non-zero rational function $f \in k(X)^\times$, and these form an important class of divisors. For a point $P \in X$, the ring $\mathscr{O}_{X,P}$ is a discrete valuation ring (DVR) with valuation ord_P. Then, the divisor of f is defined as follows.

$$(f) = \sum_{P \in X} \operatorname{ord}_P(f) P.$$

Example 3.1. Consider $X = \mathbb{P}^1$ and

$$f(X_0, X_1) = \frac{X_1 - X_0}{X_1} \in k(\mathbb{P}^1).$$

Since $f(X_0, X_1)$ is a unit in the the local rings $\mathscr{O}_{X,P}$ for $P \neq [1:0], [1:1]$, we have $\operatorname{ord}_P(f) = 0$ at those points P. The points $[1:0]$ and $[1:1]$ are contained in the affine piece $\mathbb{A}^1_0 = \{\, [X_0, X_1] \in \mathbb{P}^1 \mid X_0 \neq 0 \,\}$. The dehomogenised version of f on $\mathbb{A}^1_0$ is given by $f(x) = \frac{x-1}{x}$. One can immediately see that $\operatorname{ord}_{P_0}(f) = -1$ for $P_0 = [1:0]$ and $\operatorname{ord}_{P_1}(f) = 1$ for $P_1 = [1:1]$. Therefore,

$$(f) = P_1 - P_0.$$

Remark. *The divisor of a rational function should be thought of as counting the orders of zeros and poles of the function.*

It is useful to consider a non-constant element f of $k(X)$ as a morphism $f : X \to \mathbb{P}^1_k$, because then pre-composition by f defines a homomorphism of fields

$$k\left(\mathbb{P}^1_k\right) \to k\left(X\right) : g \mapsto g \circ f.$$

Since homomorphisms of fields are always injective, this map defines the inclusion $k\left(\mathbb{P}^1_k\right) \subseteq k\left(X\right)$. This fact can be used to define a homomorphism $f^* : \operatorname{Div}\left(\mathbb{P}^1_k\right) \to \operatorname{Div}\left(X\right)$ in the following way [Har97]. Firstly, suppose Q is a point of $\mathbb{P}^1_k$ and let t be a local uniformiser of $\mathscr{O}_{\mathbb{P}^1,Q}$. Then, t can be seen as a function in $k(X)$ and one can define

$$f^*(Q) = \sum_{f(P)=Q} \operatorname{ord}_P(t)\, P.$$

This definition is independent of the choice of local uniformiser: Suppose t' is another local uniformiser of $\mathscr{O}_{\mathbb{P}^1,Q}$. Then, there is a unit u of $\mathscr{O}_{\mathbb{P}^1,Q}$ such that $t' = ut$. Since u is a unit, its denominator and numerator do not vanish at Q. Thus, the denominator and the numerator of the corresponding function $u \circ f$ in $k(X)$ do not vanish at P, as $f(P) = Q$. Therefore, $u \circ f$ is a unit of $\mathscr{O}_{X,P}$, and so, $\operatorname{ord}_P(t) = \operatorname{ord}_P(t')$. Furthermore, this definition can be linearly extended for an arbitrary divisor $D \in \operatorname{Div}(\mathbb{P}^1_k)$. Now, the following lemma will be useful.

Lemma 3.2. *If $f \in k(X)$ is a non-constant function, then for the induced morphism $f^* : \operatorname{Div}\left(\mathbb{P}^1_k\right) \to \operatorname{Div}\left(X\right)$ we have*

$$\deg\left(f^*(D)\right) = [k(X) : k(\mathbb{P}^1_k)]\, \deg(D).$$

Proof. See Proposition 6.9 of [Har97] chapter II. $\square$

Corollary 3.2.1. *Suppose $f \in k(X)^\times$. Then,* $\deg((f)) = 0$.

Proof. If f is a constant, then the statement is immediate. Thus, assume f is not a constant. Then, one can write $(f) = f^*(0) - f^*(\infty)$ and compute the degree:

$$\deg((f)) = \deg(f^*(0)) - \deg(f^*(\infty)) = [k(X) : k(\mathbb{P}^1_k)] - [k(X) : k(\mathbb{P}^1_k)] = 0.$$

$\square$

Remark. *Note that divisors of rational functions form a group since $(f) + (g) = (fg)$. Then, we can take the quotient of* $\operatorname{Div} X$ *by this group. The quotient is called the Picard group* $\operatorname{Pic} X$ *and its elements are called* ***divisor classes****. Two elements of the same class are said to be* ***linearly equivalent*** *and one has*

$$D \sim D' \iff \exists f \in k(X)^\times,\ D' = D + (f).$$

In this article, I will use divisors to control the "order of vanishing" of a rational function $f \in k(X)^\times$. For example, if I want to allow f to have a pole only at some point $P \in X$ with order at most 2, I can express this requirement in the following way. Define a divisor $D = -2P$ and require that $(f) \geq D$. Conventionally, we would actually set $D = 2P$ and require that $(f) \geq -D$. This leads us to define the sheaf $\mathscr{O}_X(D)$ of such functions:

$$(\mathscr{O}_X(D))(U) = \{\, f \in k(X) \mid \forall P \in U,\ \operatorname{ord}_P(f) \geq -D(P) \,\}$$

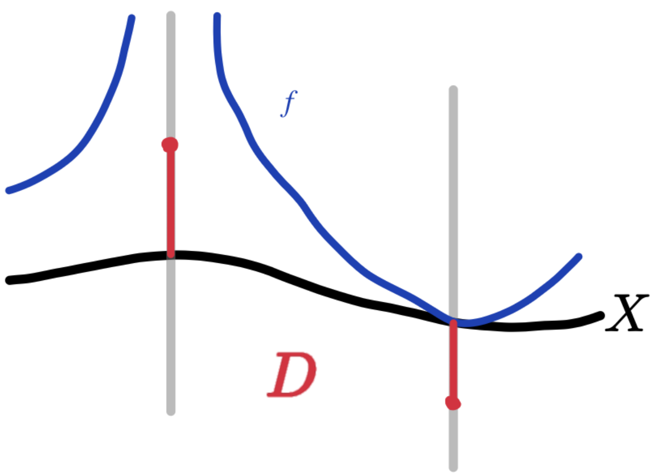

Figure 4: A section of the sheaf $\mathscr{O}_X(D)$.

Proposition 3.3. *The sheaves $\mathscr{O}_X(D)$ are quasi-coherent for every divisor D.*

Proof. Firstly, $\mathscr{O}_X(D)$ is a sheaf of $\mathscr{O}_X$-modules: If $f, g \in \mathscr{O}_X(D)(U)$, then it is clear that

$$\forall P \in U,\ \operatorname{ord}_P(f + g) \geq -D(P).$$

Moreover, if $h \in \mathscr{O}_X(U)$, then multiplying f by h only increases the order at all points, since $\operatorname{ord}_P(h)$ is never negative on U.

To finish the proof, I will show that $\mathscr{O}_X(D)$ is locally isomorphic to the structure sheaf $\mathscr{O}_X$. Thus, fix a point $P \in X$ and choose a local uniformiser $t \in k(X)$ at P so that $\operatorname{ord}_P(t) = 1$. The function t has finitely many zeroes and poles, so we can find an affine open neighbourhood $U \subseteq X$ of P such that t has order 0 on $U \setminus \{ P \}$. Moreover, one may restrict U further so that P is the only point of U where D is non-zero. Next, note that the function $\varphi = t^{D(P)}$ satisfies

$$\operatorname{ord}_Q(\varphi) = \begin{cases} D(P), & Q = P \\ 0, & Q \neq P \end{cases}.$$

Since the sections of $\mathscr{O}_X(D)|_U$ consist of functions f such that $\operatorname{ord}_P(f) \geq -D(P)$ and $\operatorname{ord}_Q(f) \geq 0$ for $Q \neq P$, multiplying by φ yields regular functions. Hence, it is easy to see that multiplication by φ defines an isomorphism $\mathscr{O}_X(D)|_U \to \mathscr{O}_U$. Since quasi-coherence can be checked locally, we are done. □

The following proposition will also be useful later.

Proposition 3.4. *If a divisor D has negative degree, then $H^0(X, \mathscr{O}_X(D)) = 0$.*

Proof. Suppose $f \in H^0(X, \mathscr{O}_X(D))$. Then, $(f) \geq -D$ so that $\deg((f)) \geq -\deg(D)$, but this implies $\deg(D) \geq 0$ by Prop. 3.2.1, which contradicts the assumption on D. □

I will now turn to discussing differentials. In geometry, spaces are often studied locally by looking at their tangent spaces. But the notion of the tangent space does not translate directly to algebraic geometry, and we would like to have a more algebraic alternative. It turns out that it is easier to give an algebraic definition of the *cotangent space*, the dual space of the tangent space. The cotangent space can be defined abstractly as a module of differentials, which makes it nice to work with from an algebraic perspective. I will not attempt to make the connection to geometry more apparent since giving geometric motivation for the definitions would take us too far from the focus of the article, and thus I leave it out. For a soft exposition of differential forms in analysis, I recommend reading Terence Tao's excellent article [Tao].

Definition 3.5. For a commutative algebra F over a field k, the module of k-differentials of F is the free F-module $\Omega(F)$ generated by the symbols df for $f \in F$ with the following rules.

1. $d(f + g) = df + dg$ for $f, g \in F$,
2. $d(fg) = f\,dg + g\,df$ for $f, g \in F$,
3. $da = 0$ for $a \in k$.

Note that this definition implies that the differential map $d : F \to \Omega(F)$ is k-linear, as $d(af) = a\,df + f\,da = a\,df$ for $a \in k$. I will write Ω for the module $\Omega(k(X))$.

Proposition 3.6. *Suppose t is a local uniformiser of $\mathscr{O}_{X,P}$ at some point $P \in X$. Then, Ω is generated by the differential dt as a $k(X)$-module.*

The proof relies on the following lemma from commutative algebra.

Lemma 3.7 (Nakayama's lemma, [AM69, Proposition 2.8]). *Let M be a finitely generated module over a local ring $(A, \mathfrak{m})$. If $x_1, \ldots, x_n$ are elements of M such that their images form a k-vector basis of $M/\mathfrak{m}M$, where $k = A/\mathfrak{m}$, then they generate M.*

Proof of Proposition. I will first prove that $\Omega(\mathscr{O}_{X,P})$ is generated by dt by showing that dt spans the k-vector space $\Omega(\mathscr{O}_{X,P})/\mathfrak{m}_P\Omega(\mathscr{O}_{X,P})$, where $\mathfrak{m}_P$ is the maximal ideal of $\mathscr{O}_{X,P}$. To apply Nakayama's lemma, I need to first check that $\Omega(\mathscr{O}_{X,P})$ is finitely generated. Recall that $\mathscr{O}_{X,P}$ consists of rational functions in finitely many variables $X_1, \ldots, X_n$. Using the rules of differentiation, any such rational function can be written as a $\mathscr{O}_{X,P}$-linear combination of the differentials $dX_1, \ldots, dX_n$. In other words, $\Omega(\mathscr{O}_{X,P})$ is finitely generated as an $\mathscr{O}_{X,P}$-module.

Now, take an element $\sum f_i\,dg_i \in \Omega(\mathscr{O}_{X,P})$. One can write $f_i = a_i + m_i$ with $a_i \in k$ and $m_i \in \mathfrak{m}$, so

$$f_i\,dg_i = a_i\,dg_i + m_i\,dg_i \equiv a_i\,dg_i.$$

Thus, it suffices to show that dg_i is in the k-linear span of dt. We can now write $g_i = a_i' + m_i'$, where $a_i' \in k$ and $m_i' \in \mathfrak{m}$. Since $\mathfrak{m} = (t)$, we can further write $m_i' = (a_i'' + m_i'')t$, where $a_i'' \in k$ and $m_i'' \in \mathfrak{m}$. Then,

$$dg_i = dm_i' = t\,da_i'' + a_i''\,dt + t\,dm_i'' + m_i''\,dt.$$

Note that the first term vanishes, since $a_i'' \in k$. Also, the third and fourth terms vanish modulo $\mathfrak{m}\mathscr{O}_{X,P}$, because $t, m_i'' \in \mathfrak{m}$. Therefore, $dg_i = a_i''\,dt$, and Nakayama's lemma implies $\mathscr{O}_{X,P}$ is generated by dt.

Now, suppose $f \in k(X)$. Such a function has an expansion

$$f = a_{-m}t^{-m} + a_{-m+1}t^{-m+1} + \cdots + a_{-1}t^{-1} + ut^{\ell},$$

where $a_i \in k$, u is a unit in $\mathscr{O}_{X,P}$ and $m, \ell \geq 0$. Then,

$$df = -ma_{-m}t^{-m-1}\,dt + \cdots - a_{-1}t^{-2}\,dt + d\left(ut^{\ell}\right)$$

Since $ut^{\ell} \in \mathscr{O}_{X,P}$, it is in the $\mathscr{O}_{X,P}$-linear span of dt by what we proved above. This shows the generators of Ω lie in the $k(X)$-linear span of dt. $\square$

This proposition lets us to define the order of a differential $\omega \in \Omega$ as follows. Write $\omega = f\,dt$ for $f \in k(X)$. Then,

$$\mathrm{ord}_P(\omega) = \mathrm{ord}_P(f).$$

Now, one can define the divisor (ω) of ω in the same way as the divisor of a rational function:

$$(\omega) = \sum_{P \in X} \mathrm{ord}_P(\omega) P.$$

It turns out that divisors of this kind are all linearly equivalent: suppose $\omega_1 = f_1\,dt$ and $\omega_2 = f_2\,dt$. Then,

$$\omega_1 = \frac{f_1}{f_2} f_2\,dt = \frac{f_1}{f_2}\omega_2 \implies (\omega_1) = \left(\frac{f_1}{f_2}\right) + (\omega_2).$$

Thus, the divisors of differentials lie in the same divisor class called the *canonical class*. Any representative of the class is called the *canonical divisor* K_X.

Next I define a module of differentials related to a divisor D on X in the same way as we defined the sheaf $\mathscr{O}_X(D)$:

$$\Omega(D) = \{\, \omega \in \Omega \mid \forall P \in X,\ \mathrm{ord}_P(\omega) \geq D(P) \,\}.$$

(in the modern literature one requires $\mathrm{ord}_P(\omega) \geq -D(P)$ to match the definition of $\mathscr{O}_X(D)$, but here I follow [Ser12] with the notation). Lastly, I will define the *residue* of a differential, which will be the main ingredient in the proof of Serre Duality in the next section.

Definition 3.8. Let $\omega = f\,dt \in \Omega$, where t is a local uniformiser of $\mathscr{O}_{X,P}$ for some point $P \in X$. Then, f can be embedded in the ring $k((t))$ of formal series over k, where it has a series expansion in terms of t:

$$f = \sum_{i \geq n} a_i t^i,$$

where $n \in \mathbb{Z}$ and $a_i \in k$. Then, the residue of ω at P is defined as $\mathrm{Res}_P(\omega) = a_{-1}$.

After this point, it is easy to get lost in all the preliminary results we need to prove about residues. It is not forbidden to skip to Subsection 3.2 if this happens.

A priori, this definition depends on the local uniformiser t, and thus I need to show that the definition is indeed independent of the choice of a local uniformiser. I will only give a proof sketch, but for a more detailed proof, see [Ser12].

Lemma 3.9. *For a non-zero function f, we have* $\mathrm{Res}_P(df/f) = \mathrm{ord}_P(f)$.

Proof. If t is a local uniformiser of $\mathscr{O}_{X,P}$, then f can be written as $f = ut^n$, where $n = \text{ord}_P(f)$. Then,

$$df/f = \frac{t^n\,du + nut^{n-1}\,dt}{ut^n} = du/u + n\,dt/t.$$

Thus, $\text{Res}_P(df/f) = \text{Res}_t(du/u) + n$, but since u is a unit in $\mathscr{O}_{X,P}$, the residue $\text{Res}_t(du/u)$ is clearly zero. □

Proposition 3.10. *Fix a point $P \in X$ and let t and u be two local uniformisers of $\mathscr{O}_{X,P}$. Denote by Res_t and Res_u the function Res_P calculated using t and u respectively. Then, $\text{Res}_t(\omega) = \text{Res}_u(\omega)$ for all differentials $\omega \in \Omega$.*

Proof sketch. Suppose $f \in k(X)$ is a rational function with a series expansion

$$f = \sum_{i \geq n} a_i u^i,$$

in $k((u))$. Then, it is possible to construct a module of differentials, where

$$df = \left(\sum_{i \geq n} i a_i u^{i-1}\right) du.$$

This is probably the most non-trivial statement of the proof, and the construction is laid out by Serre [Ser12]. Now, we can write a differential ω in this module as

$$\omega = \sum_{n \geq 0} a_n\,du/u^n + \omega_0,$$

where ω_0 is a differential with $\text{ord}_P(\omega_0) \geq 0$. Then, $\text{Res}_u(\omega) = a_1$ and $\text{Res}_t(\omega) = \sum a_n \text{Res}_t(du/u^n)$. Now, concentrate first on the term $a_1 \text{Res}_t(du/u)$. I can apply Lemma 3.9 to get $\text{Res}_t(du/u) = \text{ord}_P(u) = 1$. Therefore,

$$\text{Res}_t(\omega) = a_1 + \sum_{n > 0} \text{Res}_t(du/u^n).$$

Hence, it is enough to show that $\text{Res}_t(du/u^n) = 0$ for $n > 0$.

In characteristic zero, we can write

$$du/u^n = d\left(-\frac{1}{(n-1)u^{n-1}}\right).$$

But this immediately implies that $\text{Res}_t(du/u^n) = 0$ since differentiating a series can never result in a term of the form $a_{-1}t^{-1}$. The proof for positive characteristic follows from the statement in zero characteristic by an argument by Serre [Ser12]. □

Now, I will prove the *residue formula*, which is used in the proof of Serre Duality.

Theorem 3.11 (Residue Formula)**.** *For every differential $\omega \in \Omega$, we have that*

$$\sum_{P \in X} \text{Res}_P(\omega) = 0.$$

First I prove the theorem for the case when $X = \mathbb{P}^1_k$.

Lemma 3.12. *The residue formula holds for $X = \mathbb{P}^1_k$.*

Proof. Fix a differential $\omega = f\,dt$ on X. For convenience, I work with dehomogenised representation, and take f to be a rational function in one variable t. This function has a partial fractions decomposition, which is a linear combination of terms of the form listed below. I will consider each type separately.

Term of type $\omega = t^n\,dt$:

There are no poles at finite points, so the only pole could be at infinity. Changing to $u = 1/t$, we have $dt = -u^{-2}\,du$ and

$$\omega = (u^{-n})(-u^{-2}\,du) = \frac{du}{u^{n+2}}.$$

Then, the residue clearly vanishes: $\text{Res}_\infty(\omega) = 0$.

Term of type $\omega = \frac{dt}{t-a}$:

Clearly $\text{Res}_a(\omega) = 1$, and there are no other poles at finite points. But there is also a pole at infinity. Again, changing to $u = 1/t$, we get

$$\begin{aligned} f(u) &= \left(\frac{u}{1-au}\right)(-u^{-2}\,du) = -\frac{1}{u}\cdot\frac{1}{1-au}\,du \\ &= -\frac{1}{u}\left(1 + au + (au)^2 + \cdots\right)\,du. \end{aligned}$$

Therefore, $\text{Res}_a(\omega) = -1$ and the residues of the two points cancel.

Term of type $\omega = \frac{dt}{(t-a)^n}$ for $n > 1$:

Following a similar argument as above, one can check that in this case the residue is zero also at a and ∞.

□

Let X be a curve as before. Again, a non-constant function $\varphi \in k(X)$ induces an embedding $k(\mathbb{P}^1_k) \hookrightarrow k(X)$, and I hope to use this embedding to apply the above lemma in the case of an arbitrary curve X. Now, one can consider the trace map $\text{Tr}_{k(X)/k(\mathbb{P}^1_k)}$ defined as follows [Mil22]. Multiplication by an element $\alpha \in k(X)$ defines a $k(\mathbb{P}^1_k)$-linear map $k(X) \to k(X) : f \mapsto \alpha \cdot f$. Then, we simply define $\text{Tr}_{k(X)/k(\mathbb{P}^1_k)}(\alpha)$ to be the usual trace of this linear transformation. This definition translates to differentials on $k(X)$. Any differential $\omega \in \Omega(k(X))$ can be written as $\omega = f\,d\varphi$ and one can make the following definition.

$$\text{Tr} : \Omega(k(X)) \to \Omega(k(\mathbb{P}^1_k)) : f\,d\varphi \mapsto \left(\text{Tr}_{k(X)/k(\mathbb{P}^1_k)}(f)\right)\,d\varphi$$

Finally, the residue formula is implied by the following lemma [Ser12].

Lemma 3.13. *For every point $P \in \mathbb{P}^1_k$, we have*

$$\sum_{Q \in \varphi^{-1}(P)} \text{Res}_Q(\omega) = \text{Res}_P(\text{Tr}(\omega)).$$

This finishes the subsection on divisors and differentials, and we are now ready to move on to discussing the Riemann-Roch theorem and Serre Duality.

3.2 Proof of Riemann-Roch

I am now able to state and prove an "incomplete" version of the Riemann-Roch theorem, which I will make complete after proving Serre Duality.

Theorem 3.14 (Riemann-Roch, cohomology version). *For every divisor D on X,*

$$h^0(X, \mathscr{O}_X(D)) - h^1(X, \mathscr{O}_X(D)) = \deg(D) + 1 - g,$$

where $g = h^1(X, \mathscr{O}_X)$ and $h^i(X, \mathscr{F})$ denotes $\dim(H^i(X, \mathscr{F}))$.

Proof. One can use an induction argument, because any divisor D is obtained from the zero divisor by adding and subtracting finitely many points.

base case)
Since $\mathscr{O}_X(0) = \mathscr{O}_X$ and $\deg(0) = 0$, I need to verify that

$$h^0(X, \mathscr{O}_X) - h^1(X, \mathscr{O}_X) = 1 - g.$$

But note that the only globally defined regular functions on X are constant and thus they form a one-dimensional vector space. Moreover, $h^1(X, \mathscr{O}_X) = g$ by definition so that the equality holds.

induction step)
In the induction step I want to relate the 0th and the 1st cohomology groups of $\mathscr{O}_X(D)$ to the 0th and 1st cohomology groups of $\mathscr{O}_X(D+P)$, where P is some point. To do this, first note that $\mathscr{O}_X(D)$ is a subsheaf of $\mathscr{O}_X(D+P)$, since the orders of the sections of $\mathscr{O}_X(D+P)$ at P are allowed to be smaller than the orders of the sections of $\mathscr{O}_X(D)$ at P. Thus, there is an exact sequence

$$0 \longrightarrow \mathscr{O}_X(D) \longrightarrow \mathscr{O}_X(D+P) \longrightarrow Q \longrightarrow 0,$$

where Q is the quotient sheaf. The stalks of Q are clearly zero away from P. The stalk at P consists of zero and elements of the form u/t^{n+1}, where t is the local uniformiser of $\mathscr{O}_{X,P}$, u is a unit in $\mathcal{O}_{X,P}$, and n is the order of P in D. As $\mathcal{O}_{X,P}/(t) = k$, we can write $u = vt + r$, where $v \in \mathcal{O}_{X,P}$ and $r \in k$. Then,

$$\frac{u}{t^{n+1}} = \frac{v}{t^n} + \frac{r}{t^{n+1}}.$$

Since v/t^n is an element of $\mathscr{O}_X(D)_P$, we conclude that every element of $\mathscr{O}_X(D+P)_P$ is equivalent to an element r/t^{n+1} modulo $\mathscr{O}_X(D)_P$ for some $r \in k$. Therefore, $Q_P \cong k$ and Q is the skyscraper sheaf k_P.

Now we apply our cohomology machinery on the SES

$$0 \longrightarrow \mathscr{O}_X(D) \longrightarrow \mathscr{O}_X(D+P) \longrightarrow k_P \longrightarrow 0$$

to get the following exact sequence (using Prop. 2.15).

$$0 \longrightarrow H^0(X, \mathscr{O}_X(D)) \longrightarrow H^0(X, \mathscr{O}_X(D+P)) \longrightarrow H^0(X, k_P)$$

$$\longrightarrow H^1(X, \mathscr{O}_X(D)) \longrightarrow H^1(X, \mathscr{O}_X(D+P)) \longrightarrow 0.$$

This exact sequence of vector spaces implies the following equality.

$$h^0(X, \mathscr{O}_X(D)) - h^0(X, \mathscr{O}_X(D+P)) + 1 - h^1(X, \mathscr{O}_X(D)) + h^1(X, \mathscr{O}_X(D+P)) = 0.$$

Therefore,

$$\begin{aligned} h^0(X, \mathscr{O}_X(D+P)) - h^1(X, \mathscr{O}_X(D+P)) \\ &= \left(h^0(X, \mathscr{O}_X(D)) - h^1(X, \mathscr{O}_X(D+P))\right) + 1 \\ &= \deg(D) + 1 - g + 1 \quad \text{(by induction hypothesis)} \\ &= \deg(D+P) + 1 - g. \end{aligned}$$

This is exactly the induction step we wanted to prove. We also need to prove

$$h^0(X, \mathscr{O}_X(D-P)) - h^1(X, \mathscr{O}_X(D-P)) = \deg(D-P) + 1 - g,$$

but we can run the same argument starting with the SES

$$0 \longrightarrow \mathscr{O}_X(D-P) \longrightarrow \mathscr{O}_X(D) \longrightarrow k_P \longrightarrow 0.$$

□

One would need to also prove that the cohomology groups are finite dimensional to begin with. See [Ser12] for the details.

This form of the theorem is not the most useful one for applications, because computing $h^1(X, \mathscr{O}_X(D))$ is not easy. However, one can get a nice expression for this space using the Serre Duality, which I will prove in the next section:

Theorem 4.7 (Serre Duality)**.** If X is an algebraic curve as before and D is a divisor on X, there is an isomorphism

$$H^1(X, \mathcal{L}(D))^\vee \cong \Omega(D).$$

of k-vector spaces.

Using the below lemma, we get a nicer looking version of the Riemann-Roch theorem.

Lemma 3.15. *The following isomorphism holds.*

$$\Omega(D) = H^0(X, \mathscr{O}_X(K_X - D))$$

Proof. First, let t be a local uniformiser at some point $P \in X$, so one can write $K_X = (dt)$. Then, define

$$\Phi \colon \Omega(D) \to H^0(X, \mathscr{O}_X(K_X - D)) : f\,dt \mapsto f.$$

The map is well-defined, because

$$f\,dt \in \Omega(D) \iff (f) + (dt) \geq D \iff (f) \geq D - K_X.$$

Since the above is not only a chain of implications but a chain of equivalences, we immediately see that Φ is bijective. Since it is trivially linear, we are done. □

Finally, we can write the "complete" form of the Riemann-Roch theorem.

Theorem 3.16 (Riemann-Roch)**.** *For every divisor D,*

$$h^0(X, \mathscr{O}_X(D)) - h^0(X, \mathscr{O}_X(K_X - D)) = \deg(D) + 1 - g,$$

where $g = h^1(X, \mathscr{O}_X)$.

3.3 An application to the classification of curves

Before proving Serre Duality, I want to take some time to look at an application of the Riemann-Roch theorem using [Har97] as my source. A major project in algebraic geometry is to give a classification of different algebraic varieties. The Riemann-Roch theorem helps us prove statements about curves based solely on topological data, namely the genus of the curve. First, I give a formula for the degree of the canonical divisor on a curve.

Lemma 3.17. *The degree of the canonical divisor K_X is $2g - 2$, where g is the genus of X.*

Proof. Set $D = K_X$ so that the Riemann-Roch theorem yields

$$\begin{aligned} h^0(X, \mathscr{O}_X(K_X)) - h^0(X, \mathscr{O}_X) &= \deg(D) + 1 - g \\ h^1(X, \mathscr{O}_X) - 1 &= \deg(D) + 1 - g \\ 2g - 2 &= \deg(D), \end{aligned}$$

where $h^0(X, \mathscr{O}_X(K_X)) = h^1(X, \mathscr{O}_X)$ is given by Serre Duality. □

Now I can give a classification of curves of genus 0.

Theorem 3.18. *Any curve X of genus 0 is isomorphic to $\mathbb{P}^1_k$.*

Proof. Fix two point $P, Q \in X$ and consider the divisor $D = P - Q$. The Riemann-Roch theorem implies

$$h^0(X, \mathscr{O}_X(D)) - h^0(X, \mathscr{O}_X(K_X - D)) = \deg(D) + 1 = 1.$$

Note that

$$\deg(K_X - D) = -2 - 0 = -2 < 0$$

by the above lemma. Thus, $h^0(X, \mathscr{O}_X(K_X - D)) = 0$ by Prop. 3.4 and hence, the space $H^0(X, \mathscr{O}_X(D))$ is non-empty. Now, let $f \in H^0(X, \mathscr{O}_X(D))$. Then, there is an effective divisor D' on X such that $D' = (f) + D$. Then, we have $\deg(D') = \deg((f)) + \deg(D) = 0$. But the only effective divisor with degree 0 is the zero divisor. Therefore, $(f) = Q - P$. One can see that $f^*(0) = Q$, where f^* is the homomorphism induced by f. Applying Lemma 3.2, we get the following.

$$\begin{aligned}[k(X) : k(\mathbb{P}^1)] \cdot 1 &= \deg(f^*(0)) \\ &= \deg(Q) = 1.\end{aligned}$$

Therefore, the function fields $k(X)$ and $k(\mathbb{P}^1_k)$ are isomorphic. Since X and $\mathbb{P}^1_k$ are non-singular, it follows by [Har97, Proposition 6.7] that the two curves are isomorphic. □

4 Serre duality

The rest of this paper is devoted to proving the Serre duality. I will prove the theorem by first finding a more concrete representation of $H^1(X, \mathscr{O}_X(D))$ and then constructing a perfect pairing between $H^1(X, \mathscr{O}_X(D))$ and $\Omega(D)$, which will give us the isomorphism.

4.1 Concrete representation of 1st cohomology

To prove the Serre duality, we do not want to directly work with the Čech cohomology definition of $H^1(X, \mathscr{O}_X(D))$. Instead we want to give a more concrete description of $H^1(X, \mathscr{O}_X(D))$ by finding some SES involving $\mathscr{O}_X(D)$ and then taking the cohomology sequence of the SES. Since $\mathscr{O}_X(D)$ is a subsheaf of the constant sheaf $\underline{k(X)}$, one can simply consider the following SES.

$$0 \longrightarrow \mathscr{O}_X(D) \longrightarrow \underline{k(X)} \longrightarrow \underline{k(X)}/\mathscr{O}_X(D) \longrightarrow 0,$$

which yields the following exact sequence

$$H^0(X, \underline{k(X)}) \longrightarrow H^0(X, \underline{k(X)}/\mathscr{O}_X(D)) \longrightarrow H^1(X, \mathscr{O}_X(D)) \longrightarrow H^1(X, \underline{k(X)}).$$

But Prop. 2.14 implies that $H^0(X, \underline{k(X)}) = k(X)$ and $H^1(X, \underline{k(X)}) = 0$ so that the exact sequence simplifies to

$$k(X) \longrightarrow H^0(X, \underline{k(X)}/\mathscr{O}_X(D)) \longrightarrow H^1(X, \mathscr{O}_X(D)) \longrightarrow 0. \tag{5}$$

Let us first try to understand the space $H^0(X, \underline{k(X)}/\mathscr{O}_X(D))$. An element of the space is of the form $([f_P])_{P\in X}$, where $[f_P]$ is the equivalence class of some $f_P \in k(X)$ modulo $\mathscr{O}_X(D)_P$. These elements are actually quite simple, because one can show that the components $[f_P]$ are zero almost everywhere. In other words, we can write $H^0(X, \underline{k(X)}/\mathscr{O}_X(D))$ as a direct sum of stalks.

Lemma 4.1. *For a divisor D on a curve X, the following equality holds.*

$$H^0(X, \underline{k(X)}/\mathscr{O}_X(D)) = \bigoplus_{P\in X} k(X)/\mathscr{O}_X(D)_P.$$

Proof. I will first show the inclusion in forward direction:

$$H^0(X, \underline{k(X)}/\mathscr{O}_X(D)) \subseteq \bigoplus_{P\in X} k(X)/\mathscr{O}_X(D)_P.$$

I need to check that the components $[f_P]$ of some $([f_P])_{P\in X} \in H^0(X, \underline{k(X)}/\mathscr{O}_X(D))$ are non-zero for only finitely many $P \in X$. Thus, suppose $([f_P])_{P\in X} \in H^0(X, \underline{k(X)}/\mathscr{O}_X(D))$. Firstly, the divisor D is non-zero only at finitely many points. The components $[f_P]$ may or may not be zero at those points, but I can ignore those points in any case since there are finitely many of them. Thus, assume P is a point where $D(P) = 0$. At such a point, the stalk $\mathscr{O}_X(D)_P$ is equal to the ring of germs with non-negative order at P, which is of course the ring $\mathscr{O}_{X,P}$. Now, it is a basic result in algebraic geometry that a rational function $f \in k(X)$ on a curve X has negative order only at finitely many points so that $f_P \in \mathscr{O}_{X,P}$ for almost all points P. Therefore, we see that $[f_P] = 0$ for almost all points P.

Next I show the inclusion in the reverse direction:

$$H^0(X, \underline{k(X)}/\mathscr{O}_X(D)) \supseteq \bigoplus_{P \in X} k(X)/\mathscr{O}_X(D)_P.$$

Thus, let $([f_P]_{P\in X}) \in \bigoplus k(X)/\mathscr{O}_X(D)_P$. I want to show that the components $[f_P]$ form a global section of $\underline{k(X)}/\mathscr{O}_X(D)$. Let us fix an arbitrary point $P \in X$. I want to find an open neighbourhood $U \ni P$ and a section $g \in k(X)$ such that $\forall Q \in U,\ [f_Q] = [g]$. If I denote by $P_1, \ldots, P_r$ the points where D is non-zero and by $Q_1, \ldots, Q_s$ the points where the f_{Q_i} have negative order, then there are two cases:

$P \notin \{ P_1, \ldots, P_r, Q_1, \ldots, Q_s \}$)
: Since P is none of the points $P_1, \ldots, P_r$, $\mathscr{O}_X(D)_P$ consists of all functions $f \in k(X)$ with non-negative order at P: $\operatorname{ord}_P(f) \geq 0$. Since P is non of the points $Q_1, \ldots, Q_s$, $f_P \in \mathscr{O}_{X,P}$. Therefore, $f_P \in \mathscr{O}_X(D)_P$ so that $[f_P] = [0]$. Now, if we let U be the complement of the set $\{ P_1, \ldots, P_r, Q_1, \ldots, Q_s \}$, we see that $[f_Q] = [0]$ for every r_Q on U so that we can simply choose $0 \in k(X)$ as the section on the open neighbourhood U.

$P \in \{ P_1, \ldots, P_r, Q_1, \ldots, Q_s \}$)
: First, denote $Y = \{ P_1, \ldots, P_r, Q_1, \ldots, Q_s \} \setminus \{ P \}$ and $g = f_P \in k(X)$. Next, let $S_1, \ldots, S_t$ be the points where the g_{S_i} have negative order. Then, let U be the complement of $Y \cup \{ S_1, \ldots, S_t \}$. As above, $[f_Q] = [0]$ for all $Q \in U$ except for $Q = P$. But since the points S_i are also included in the complement, $[g] = [0]$ away from P. Thus, $[f_Q] = [g]$ on U.

□

This direct sum can be expressed in the following way. Consider the vector space R of families $\{r_P\}_{P\in X}$, where $r_P \in k(X)$ and $r_P \in \mathscr{O}_{X,P}$ for almost all points $P \in X$. (Serre calls such a family a *répartition* [Ser12]). Then, define $R(D) = \{ \{r_P\}_{P\in X} \mid \operatorname{ord}_P(r_P) \geq -D(P) \}$. It is clear that we have the following isomorphism.

$$\bigoplus_{P\in X} k(X)/\mathscr{O}_X(D)_P \cong R/R(D).$$

Note that répartitions are easier to work with, because the components f_P of some section $(f_P)_{P\in X}$ of $\underline{k(X)}/\mathscr{O}_X(D)$ must be related together so that the sections satisfy sheaf axioms, whereas there is no such requirement for répartitions. Therefore, I will make use of répartitions in the rest of the section.

Now we can return to the SES (5) derived above and replace $H^0(X, \underline{k(X)}/\mathscr{O}_X(D))$ by $R/R(D)$:

$$k(X) \longrightarrow R/R(D) \longrightarrow H^1(X, \mathscr{O}_X(D)) \longrightarrow 0.$$

This exact sequence finally gives us the representation of the first cohomology group: By exactness, the second map is a surjection such that its kernel is the image of $k(X)$ in $R/R(D)$. By the first isomorphism theorem, we have

$$H^1(X, \mathscr{O}_X(D)) \cong R/\left(R(D) + k(X)\right).$$

The dual space $H^1(X, \mathcal{L}(D))^\vee$ is simply the space of linear functionals on R, which vanish on $R(D)$ and $k(X)$.

4.2 Constructing a pairing

Next I will construct a bilinear form

$$\langle -, - \rangle : \Omega(D) \times H^1(X, \mathscr{O}_X(D)) \to k.$$

Recall that the space $\Omega(D)$ consists of differential forms ω such that $(\omega) \geq D$. Now, define the bilinear form as follows.

$$\langle \omega, r \rangle = \sum_{P \in X} \mathrm{Res}_P(r_P \omega),$$

where $r = [\{r_P\}_{P \in X}] \in R/\left(R(D) + k(X)\right)$. I must check that the map is well-defined: only finitely many terms of the sum can be non-zero, and its value must be independent of the representation modulo $R(D) + k(X)$.

Firstly, $\mathrm{Res}_P(r_P\omega)$ can be non-zero only when P is a point such that $D(P) \neq 0$ or $r_P \notin \mathscr{O}_{X,P}$. Otherwise, $r_P, f \in \mathscr{O}_{X,P}$, if we write $\omega = f\,dt$ where t is a local uniformiser at P. This clearly implies that the coefficients of the terms of negative degree in the series expansion of $r_P f$ are all zero. Secondly, given a repartition $r \in R(D)$, we have $(r_P\omega) = (r_P) + (\omega) \geq -D + D = 0$ and thus $\mathrm{Res}_P(r_P\omega) = 0$ by the same argument as above. Also, if $r \in k(X)$, then $\langle \omega, r \rangle = 0$ by the residue formula (Thm. 3.11).

Now, Serre duality will follow if we can show that the map

$$\iota_D : \Omega(D) \to H^1(X, \mathcal{L}(D))^\vee : \omega \mapsto \langle \omega, - \rangle$$

is a bijection. Every differential $\omega \in \Omega(D)$ is indeed mapped to elements of $H^1(X, \mathcal{L}(D))^\vee$ since they vanish on $R(D)$ and $k(X)$ by the arguments in the previous paragraph. In the proof of the bijectivity of ι_D, I will diverge little bit from the proof given by Serre, and I mix in ideas from [For81].

I will proceed to prove the bijectivity of ι_D. The proof relies on the observation that the spaces $H^1(X, \mathcal{L}(D))^\vee$ and $\Omega(D)$ form so called *filtered families*, which are tied together by ι_D. In the proofs of injectivity and surjectivity we will "move through the filtration" to derive the wanted results. Thus, let us make the following observation. For any two divisors D_1, D_2 such that $D_1 \geq D_2$, we have $R(D_1) \supseteq R(D_2)$. Therefore, a linear functional vanishing on $R(D_1)$ will also vanish on $R(D_2)$ so that $H^1(X, \mathscr{O}_X(D_1))^\vee \subseteq H^1(X, \mathscr{O}_X(D_2))^\vee$. One can also see that if $D_1 \geq D_2$, then $\Omega(D_1) \subseteq \Omega(D_2)$. Moreover, these inclusions trivially commute with $\iota_\bullet$ so that the following square is commutative.

$$\begin{array}{ccc} H^1(X, \mathcal{L}(D_1))^\vee & \overset{i_{D_1}^{D_2}}{\hookrightarrow} & H^1(X, \mathcal{L}(D_2))^\vee \\ {\scriptstyle \iota_{D_1}} \uparrow & & \uparrow {\scriptstyle \iota_{D_2}} \\ \Omega(D_1) & \hookrightarrow & \Omega(D_2) \end{array} \tag{6}$$

Abstract nonsense ahead

The commutativity of this square shows that the maps $\iota_\bullet$ define a natural transformation between the contravariant functors

$$\Omega(-), H^1(X, \mathcal{L}(-))^\vee : \mathrm{Div}(X) \to k\text{-}\mathbf{Vect},$$

where $\mathrm{Div}(X)$ is the posetal category of divisors on X. Thus, not only do we get an isomorphism of the vector spaces $\Omega(D)$ and $H^1(X, \mathcal{L}(D))^\vee$, but we also get a natural isomorphism between the functors $\Omega(-)$ and $H^1(X, \mathcal{L}(-))^\vee$.

Now, the following lemma will let us "transport the problem along the filtration".

Lemma 4.2. *Suppose D_1 and D_2 are two divisors on X such that $D_1 \geq D_2$. Furthermore, let $\lambda \in H^1(X, \mathcal{L}(D_1))^\vee$ and $\omega \in \Omega(D_2)$. If $i_{D_1}^{D_2}(\lambda) = \iota_{D_2}(\omega)$, then $\omega \in \Omega(D_1)$ and $\iota_{D_1}(\omega) = \lambda$.*

This lemma is effectively saying that if we can invert λ along ι_{D_2}, then we can invert it along ι_{D_1}, when $D_1 \geq D_2$.

Proof. Assume to the contrary that $\omega \notin \Omega(D_1)$. Then, there is a point $P \in X$ such that $\mathrm{ord}_P(\omega) < D_1(P)$. Now we can construct a répartition $r = \{r_Q\}_{Q \in X}$ such that $r_Q = 0$ when $Q \neq P$ and $r_P = 1/t^{\mathrm{ord}_P(\omega)+1}$, where t is a local uniformiser at P. We have $r \in R(D_1)$, because

$$\mathrm{ord}_P(r_P) = -\mathrm{ord}_P(\omega) - 1 > -D_1(P) - 1 \implies \mathrm{ord}_P(r_P) \geq -D_1(P).$$

Then,

$$\lambda(r) = i_{D_1}^{D_2}(r) = \iota_{D_2}(\omega)(r) = \sum_{Q \in X} \mathrm{Res}_Q(r_Q \omega) = \mathrm{Res}_P(r_P \omega).$$

Since $\mathrm{ord}_P(r_P \omega) = \mathrm{ord}_P(r_P) + \mathrm{ord}_P(\omega) = -\mathrm{ord}_P(\omega) - 1 + \mathrm{ord}_P(\omega) = -1$, we have that $\lambda(r)$ is non-zero. Therefore, λ doesn't vanish on $R(D_1)$, which contradicts the assumption that $\lambda \in H^1(X, \mathcal{L}(D_1))^\vee$. $\square$

Injectivity of ι_D follows easily from this lemma.

Proposition 4.3. *The map* $\iota_D : \Omega(D) \to H^1(X, \mathcal{L}(D))^\vee$ *is an injection.*

Proof. The map is injective if its kernel is trivial. Thus, suppose $\iota_D(\omega) = 0$. Since $0 \in H^1(X, \mathcal{L}(D'))^\vee$ for every divisor D', $\omega \in \Omega(D')$ for every divisor D' such that $D' \geq D$ by the above lemma. This clearly implies that ω must be zero. □

Proving surjectivity is not quite as easy. Let us fix an element $\lambda \in H^1(X, \mathcal{L}(D))^\vee$. I want to find a suitable divisor D' with $D' \leq D$ such that it is easy to invert λ along $\iota_{D'}$. But first I will introduce an extra degree of freedom, which I can work with. Namely, I consider an arbitrary element $\psi \in H^0(X, \mathcal{O}_X(\Delta))$ for some divisor Δ. This section induces a map

$$H^1(X, \mathcal{O}_X(D - \Delta)) \to H^1(X, \mathcal{O}_X(D)) : [\{r_P\}_{P \in X}] \mapsto [\{\psi r_P\}_{P \in X}].$$

It is easy to check that this map is well-defined. Then, the dual map

$$H^1(X, \mathcal{L}(D))^\vee \to H^1(X, \mathcal{L}(D - \Delta))^\vee$$

is defined so that $(\psi f)(r) = f(\psi r)$. Note that $\frac{1}{\psi} \in H^0(X, \mathcal{O}_X((\psi)))$, and it induces a map $H^1(X, \mathcal{L}(D - \Delta))^\vee \to H^1(X, \mathcal{L}(D - \Delta - (\psi)))^\vee$ in the same way. If I now take $D' = H^1(X, \mathcal{O}_X(D - \Delta - (\psi)))$, the inclusion $i_D^{D'} : H^1(X, \mathcal{L}(D))^\vee \to H^1(X, \mathcal{L}(D'))^\vee$ can be factored as follows.

$$\begin{array}{ccc} H^1(X, \mathcal{L}(D))^\vee & \xrightarrow{\psi\cdot} & H^1(X, \mathcal{L}(D - \Delta))^\vee \\ & \searrow^{i_D^{D'}} & \downarrow \tfrac{1}{\psi}\cdot \\ & & H^1(X, \mathcal{L}(D - \Delta - (\psi)))^\vee \end{array}$$

Now, I will prove a result analogous to Lemma 4.2, which lets us "move along these maps" induced by elements of $H^0(X, \mathcal{O}_X(D))$.

Lemma 4.4. *Suppose D_1 and D_2 are two divisors on X and $\psi \in H^0(X, \mathcal{O}_X(D_2))$. Then, the following square commutes.*

$$\begin{array}{ccc} H^1(X, \mathcal{L}(D_1))^\vee & \xrightarrow{\psi\cdot} & H^1(X, \mathcal{L}(D_1 - D_2))^\vee \\ \uparrow \iota_{D_1} & & \uparrow \iota_{D_2} \\ \Omega(D_1) & \xrightarrow[\psi\cdot]{} & \Omega(D_2) \end{array}$$

Proof. Let $\omega \in \Omega(D_1)$ and $r \in H^1(X, \mathcal{O}_X(D_1 - D_2))$. Then,

$$\begin{aligned} (\psi \circ \iota_{D_1})(\omega)(r) &= \langle \omega, \psi r \rangle \\ &= \sum_{P \in X} \operatorname{Res}(\psi r_P \cdot \omega) \\ &= \langle \omega\psi, r \rangle = (\iota_{D_2} \circ \psi)(\omega)(r). \end{aligned}$$

Since this equality holds for every ω and r, we have $\psi \circ \iota_{D_1} = \iota_{D_2} \circ \psi$. □

Now I will invert $\psi\lambda$ along $\iota_{D-\Delta}$, which will finally let us prove the surjectivity of ι_D.

Lemma 4.5. *Let* $\lambda \in H^1(X, \mathcal{L}(D))^\vee$*. Then, there is a divisor* Δ*, a section* $\psi \in H^0(X, \mathscr{O}_X(\Delta))$ *and a differential* $\omega \in \Omega(D-\Delta)$ *such that* $\psi\lambda = \iota_{D-\Delta}(\omega)$*.*

Proof. Let Δ be an arbitrary divisor. The elements $\psi\lambda$ form a subspace

$$\Lambda = \left\{ \psi\lambda \mid \psi \in H^0(X, \Delta) \right\}$$

of $H^1(X, \mathcal{L}(D-\Delta))^\vee$. I will use a dimensional argument to show that Λ must interesect with $\mathrm{im}(\iota_{D-\Delta})$. Therefore, let us first find a bound for the dimension of Λ. Consider the map

$$H^0(X, \Delta) \to H^1(X, \mathcal{L}(D-\Delta))^\vee : \psi \mapsto \psi\lambda.$$

Claim. *This map is an injection.*

Proof. Assume for a contradiction that the kernel of this map is non-trivial so that there is some non-zero $\psi \in H^0(X, \Delta)$ such that $\psi\lambda = 0$. Note that the multiplication map

$$H^1(X, \mathscr{O}_X(D-\Delta)) \to H^1(X, \mathscr{O}_X(D)) : [\{r_P\}_{P\in X}] \mapsto [\{\psi r_P\}_{P\in X}]$$

is clearly surjective since for a repartition $[\{r_P\}_{P\in X}] \in H^1(X, \mathscr{O}_X(D))$, we have $[\{1/\psi \cdot r_P\}_{P\in X}] \in H^1(X, \mathscr{O}_X(D-\Delta))$. Thus, the dual map

$$\psi\cdot : H^1(X, \mathcal{L}(D))^\vee \to H^1(X, \mathcal{L}(D-\Delta))^\vee$$

is an injection. But this is a contradiction, because we also have that $\psi \cdot 0 = 0$. Therefore, the kernel of the map $H^0(X, \Delta) \to H^1(X, \mathcal{L}(D-\Delta))^\vee$ is trivial, and hence the map is injective. ■

Now, this claim implies that $\dim\Lambda = h^0(X, \Delta)$, and I can apply the cohomological version of the Riemann-Roch theorem to get a bound

$$\dim\Lambda \geq \deg\Delta - g + 1. \tag{7}$$

Similarly, since $\iota_{D-\Delta} : \Omega(D-\Delta) \to H^1(X, \mathcal{L}(D-\Delta))^\vee$ is an injection, we have that

$$\dim(\mathrm{im}(\iota_{D-\Delta})) = \dim(\Omega(D-\Delta)).$$

Using Lemm. 3.15, I can apply Riemann-Roch again:

$$\dim(\mathrm{im}(\iota_{D-\Delta})) = h^0(X, \mathscr{O}_X((\omega) - D + \Delta)) \geq \deg((\omega) - D + \Delta) - g + 1. \tag{8}$$

Of course, I can also apply Riemann-Roch to $\mathscr{O}_X(D-\Delta)$ to get $h^0(X, \mathscr{O}_X(D-\Delta)) - h^1(X, \mathscr{O}_X(D-\Delta)) = \deg(D-\Delta) - g + 1$. Since Δ was an arbitrary divisor, I can choose it so that $\deg \Delta > \deg D$. Then the zeroth cohomology group vanishes and we are left with

$$h^1(X, \mathscr{O}_X(D - \Delta)) = \deg(\Delta) - \deg D + g - 1. \tag{9}$$

Combining inequalities (7) and (8), we get the following:

$$\dim \Lambda + \dim\left(\operatorname{im}(\iota_{D-\Delta})\right) \geq 2\deg(\Delta) - \deg D + \deg(\omega) - 2g + 2.$$

Again, since Δ was arbitrary, we can make its degree large enough so that

$$\dim \Lambda + \dim\left(\operatorname{im}(\iota_{D-\Delta})\right) \geq h^1(X, \mathscr{O}_X(D - \Delta)).$$

Now, since the sum of the dimensions of these two subspaces is larger than the total space, they must intersect. ☐

Combining all these lemmas, the surjectivity of ι_D follows.

Proposition 4.6. *The map* $\iota_D : \Omega(D) \to H^1(X, \mathcal{L}(D))^\vee$ *is a surjection.*

Proof. Fix an element $\lambda \in H^1(X, \mathcal{L}(D))^\vee$. Then, by Lemma 4.5, there is a section $\psi \in H^0(X, \mathscr{O}_X(\Delta))$ and a differential $\omega \in \Omega(D - \Delta)$ such that $\psi\lambda = \iota_{D-\Delta}(\omega)$. Then,

$$\lambda = \frac{1}{\psi}(\psi\lambda) = \frac{1}{\psi}\iota_{D-\Delta}(\omega),$$

and by Lemma 4.4, this is

$$\lambda = \iota_{D-\Delta-(\psi)}(\frac{1}{\psi}\omega).$$

Finally, Lemma 4.2 implies that $\frac{1}{\psi}\omega \in H^1(X, \mathcal{L}(D))^\vee$ and that $\iota_D(\frac{1}{\psi}\omega) = \lambda$, concluding the proof. ☐

⚠️ Abstract nonsense ahead

This proof boils down to chasing the following diagram.

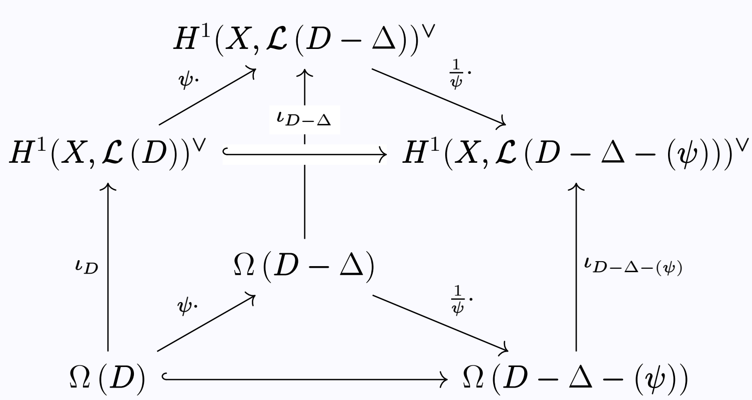

We finally arrive at Serre Duality!

Theorem 4.7 (Serre Duality). *If X is an algebraic curve as in the previous section and D is a divisor on X, there is an isomorphism*

$$H^1(X, \mathcal{L}(D))^\vee \cong \Omega(D)$$

of $\boldsymbol{k}$-vector spaces.

Proof. Combining Propositions 4.3 and 4.6 shows that the linear map

$$\iota_D : \Omega(D) \to H^1(X, \mathcal{L}(D))^\vee$$

is a bijection so that it defines an isomorphism between the two spaces. □

References

[AM69] Michael Atiyah and Ian Grant Macdonald. *Introduction to Commutative Algebra*. Addison-Wesley Publishing Company, 1969.

[For81] Otto Forster. *Lectures on Riemann Surfaces*, volume 81. Springer-Verlag, 1981.

[Ful08] William Fulton. Algebraic curves: An introduction to algebraic geometry. https://dept.math.lsa.umich.edu/~wfulton/CurveBook.pdf, 2008.

[Gat02] Andreas Gathmann. Algebraic geometry. https://www.mathematik.uni-kl.de/~gathmann/class/alggeom-2002/alggeom-2002.pdf, 2002.

[Gat13] Andreas Gathmann. Commutative algebra. https://www.mathematik.uni-kl.de/~gathmann/class/commalg-2013/commalg-2013.pdf, 2013.

[Gat21] Andreas Gathmann. Algebraic geometry. https://www.mathematik.uni-kl.de/~gathmann/class/alggeom-2021/alggeom-2021.pdf, 2021.

[Har97] Robin Hartshorne. *Algebraic Geometry*, volume 52. Springer-Verlag, 1997.

[Hat01] Allen Hather. Algebraic topology. https://pi.math.cornell.edu/~hatcher/AT/AT.pdf, 2001.

[Mil22] James S. Milne. Flieds and galois theory (v5.10). https://www.jmilne.org/math/CourseNotes/FT.pdf, 2022.

[MLM92] Saunders Mac Lane and Ieke Moerdijk. *Sheaves in Geometry and Logic: A First Introduction to Topos Theory*. Springer-Verlag, 1992.

[Rei88] Miles Reid. *Undergraduate algebraic geometry*, volume 12. Cambridge University Press, 1988.

[Ser12] Jean-Pierre Serre. *Algebraic groups and class fields*, volume 117. Springer-Verlag, 2012.

[Sta23] The Stacks project authors. The stacks project. https://stacks.math.columbia.edu, 2023.

[Tao] Terence Tao. Differential forms and integration. https://www.math.ucla.edu/~tao/preprints/forms.pdf.

[Vak17] Ravi Vakil. The rising sea: Foundations of algebraic geometry. https://math.stanford.edu/~vakil/216blog/FOAGnov1817public.pdf, November 2017.